

Multinational Expansion in Time and Space*

Stefania Garetto Xiao Ma Lindsay Oldenski Natalia Ramondo[†]
BU, CEPR, and NBER Peking University Georgetown University BU, CEPR, and NBER

February 6, 2026

Abstract

This paper studies the expansion patterns of multinational enterprises (MNEs) in time and space. Informed by a set of facts documented using a panel of US MNEs, we develop a multi-country general equilibrium dynamic model featuring a rich structure of costs to firm expansion into host and export markets. We introduce a novel compound-option formulation for the firm problem, which captures the spatial heterogeneity observed in the data. Using the calibrated model, our quantitative applications reveal that both the spatial and time dimensions of MNE activities matter 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 Spasi, 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 Andrés Rodríguez-Clare, as well as seminar participants at various conferences and institutions. We thank Juan Ignacio Jacobian for outstanding research assistance. The views expressed in this paper are those of the authors and do not necessarily represent the US Bureau of Economic Analysis (BEA) or the US Department of Commerce. BEA has reviewed this paper for unauthorized disclosure of confidential information and has approved the disclosure avoidance practices applied to this release. The statistical analysis of firm-level data was conducted at the Bureau of Economic Analysis (BEA), Department of Commerce, under arrangements that maintain legal confidentiality requirements. BEA Approval Numbers BEA-FY21-E0011-R0002 to R0004, BEA-FY23-E0011-R0001, BEA-FY24-E0011-R0001 to R0002, BEA-FY25-E0011-R0001, and other unnumbered approved disclosure requests.

[†]E-mail: garetto@bu.edu; xiaoma@phbs.pku.edu.cn; lindsay.oldenski@georgetown.edu; nramondo@bu.edu.

1 Introduction

Many questions in international economics involve the complex activities of multinational enterprises (MNEs) over time and space. Episodes such as the United Kingdom abandoning the European Union (“Brexit”), the US-China trade war in 2018, and the 2025 US tariff war, involve not only the re-arrangement of trade flows, but also of MNE activities. Undoubtedly, these shocks have global consequences. For instance, Brexit affected not only the behavior of MNEs operating in the United Kingdom, but also the behavior and expansion strategies of MNEs operating in countries nearby — or connected to — the United Kingdom. The Brexit shock also affected MNEs differently in the short and long run, as these firms make long-lasting investments in productive capacity. Answering these and other similar questions requires an understanding of the MNE expansion over time and across space, as well as of the nature of the costs these firms face.

Despite their importance for the global economy and in the policy arena, the behavior of MNEs and their affiliates over time *and* space has received relatively scarce 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.

In this paper, we introduce a multi-country general equilibrium dynamic model, which is informed by a new set of facts on the behavior of foreign affiliates of US MNEs over time and space. The rich spatial and dynamic structure of the model is aimed at answering quantitatively questions about the effects of globalization shocks on MNE expansion.

Our empirical analysis uses a long panel of US MNEs and their foreign affiliates from the Bureau of Economic Analysis (BEA).² First, we document that virtually all US MNE affiliates are born with sales to the host market, which is the main destination of sales over their entire life; exports start later in 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 already entered similar markets (e.g. markets in the same continent or with a similar income level). Fourth, sales of existing affiliates to their host market do not change when sibling affiliates (i.e. affiliates of the

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

²Not only is the United States the main source of MNEs in the world, but MNE affiliates are also 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; 40 percent of those affiliates’ sales were exports, i.e., sales to customers outside the affiliate’s host market (Yeaple, 2013).

same MNE) are created in a different market, when sibling affiliates start exporting to the affiliate's host market, or when sibling affiliates start exporting to the same export market of the affiliate. Finally, aggregate shocks and firm-level time-invariant characteristics account for more than eighty percent of the variation in affiliate sales.

Guided by these facts, we build a multi-country general equilibrium dynamic model of firm expansion. Firms decide whether, when, and where to open foreign affiliates, which, in turn, can sell both to their host market and to any other market, 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 shocks, and the array of costs.

While the static components of the model are standard and follow Melitz (2003), the formulation of the dynamic problem of the MNE is new to the international trade literature. We build on insights from the literature on real options on how to solve models of investment under uncertainty (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. This 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 affect the MNE entry decisions into that market.

We couple this dynamic structure with a standard Armington type assumption by which goods are firm- and location-specific. This assumption implies that the MNE entry decision into a market is separable across locations, avoiding a high-dimensional combinatorial problem.

In order to implement an algorithm for the quantitative model, we assume that a firm's behavior is myopic with respect to the general equilibrium effect of shocks on aggregate variables. This approach mutes general equilibrium effects on the firm's entry and exit decisions, but still captures how aggregate prices, quantities, and wages respond to firms' dynamic decisions on which markets to serve and from which locations. By enabling a guess-and-verify procedure to solve for the value functions, our approach delivers a fast algorithm for calibration and quantitative analysis. At the same time, it provides a clear theoretical intuition for the firm problem. Overall, we balance computational feasibility with the richness of the MNE problem and the inclusion of general equilibrium responses.

We calibrate the model to match static and dynamic moments related to the behavior of US MNE affiliates located in the top-ten host countries for US Foreign Direct Investment (FDI), over

thirty years. Our calibration implies that opening and operating foreign 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 countries, sales type, and type of frictions.

Armed with the calibrated model, we perform various quantitative exercises that demonstrate the importance of including both dynamics and spatial heterogeneity for evaluating the reallocation of MNE activities after a globalization shock. First, we study the effects of a change in the costs of entering and operating in a country. Including the possibility of affiliate exports—captured by the compound option structure—dramatically increases the MNE’s incentives to operate in a country, and delivers different long-run responses compared to what a dynamic model with no export platforms predicts. Second, our model predicts that an increase in export costs between the United Kingdom and EU countries —mimicking “Brexit” — would have a static effect due to the higher costs of accessing the export network, a dynamic effect on entry and exit due to the presence of sunk costs, and a general equilibrium effect. The strength of each effect on the reallocation of MNE activities depends on the nature of the export cost. For instance, an increase in per-period fixed costs would decrease the share of UK-based affiliates selling to the EU by twice as much as a similar increase in sunk export costs (12 vs 6 percentage points). 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. Finally, after a globalization shock, sunk costs create slower dynamics than a static model without sunk costs. Put together, our quantitative exercises highlight the importance of considering both the time *and* space dimensions of MNE expansion decisions, without neglecting general equilibrium effects.

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), Arkolakis et al. (2018), Alviarez (2019), Head and Mayer (2019), Fan (2024), Oberfield et al. (2024), Antràs et al. (2024), Castro-Vincenzi (2024), Arkolakis et al. (2025), and Castro-Vincenzi et al. (2025), have made substantial progress in building static models with a rich geography, in particular, by allowing firms to set up affiliates in locations — countries or regions— that differ from the destinations of their sales.³ Making progress in dynamic setups, while keeping the spatial complexity of the static models in general equilibrium, requires restricting the problem of the MNE. The sharp patterns that we document from observing US MNE affiliates over time guide us on how to simplify this problem, and how to reduce the choice set of firms in a way that is consistent with the data. By adopting a compound-option structure of

³In the context of export markets, two exceptions featuring dynamic models with interdependencies are Morales et al. (2019) and Alfaro-Ureña et al. (2024).

the dynamic problem, we are able to make substantial progress towards modeling the dynamics of MNE expansion, without sacrificing the spatial richness of static models and general equilibrium responses.

Second, there is a 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.⁴ Fillat and Garetto (2015) and Fillat et al. (2015) introduce the idea that MNE activities can be treated as a real option that gets exercised once an affiliate is opened abroad.⁵ However, both papers assume that the activities of affiliates are restricted to their market of operation, in a partial equilibrium setting. Our model treats MNE activities as a compound, rather than a simple, option. In this way, we can expand on the spatial dimension of the problem by separating the locations of MNEs' production and sales, while embedding these choices in a dynamic general equilibrium multi-country framework.⁶ Finally, McGrattan and Waddle (2020) use a neoclassical growth model augmented by technology capital (McGrattan and Prescott, 2010) to analyze the effects of Brexit. While their paper analyzes similar quantitative exercises, the technology-capital growth model is quite different from ours, notably, because we include the extensive margin of both export and MNE decisions, a rich spatial export network, and uncertainty — all features dictated by the data, and relevant when analyzing counterfactual scenarios.

Third, our paper is naturally related to the large literature on export dynamics, which has been primarily concerned with quantifying the various costs of export activities and their welfare implications (see Alessandria et al., 2021, for a review).⁷ The nature of the MNE problem, however, is more

⁴Our facts on US MNEs complement the facts in Gumpert et al. (2020), who 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 US MNE affiliates abroad, they do not provide any information about (non-MNE) US exporters. In contrast, the French and Norwegian data contain detailed information on all firms operating in those economies, while they lack detailed information on the activity of MNEs outside their home country.

⁵Impullitti et al. (2013) use a real option model to study the entry and exit patterns of exporters.

⁶Other papers in the MNE literature limit both the spatial and dynamic dimension of the analysis by considering only MNE sales to the host market and only two periods (see, for instance, Ramondo et al., 2013; Egger et al., 2014; Conconi et al., 2016).

⁷Earlier 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. (2021) 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. (2024) show that the life-cycle growth patterns of export prices and quantities are quite different.

complex than the nature of the exporter problem: MNEs choose not only which markets to serve, as an exporter does, but also the location from where 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. Admittedly, to be able to solve the multi-country model in general equilibrium, we must restrict the model to include only aggregate shocks, and exclude firm-level shocks, which are instead used in most of the literature on exporter dynamics. Overall, our analysis complements the literature on export dynamics by quantifying the frictions to MNE expansion and their implications in terms of aggregate firm dynamics.

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). While our framework can be applied more generally to dynamic problems of multi-establishment firms, our facts suggest that the dynamics of MNE affiliates are different from the dynamics of domestic firms. These differences may be indicative of domestic firms facing frictions that are different and of different magnitude in domestic and foreign markets (see e.g. Hsieh and Klenow, 2014; Kueng et al., 2014; Foster et al., 2016; Saborowski and Misch, 2019, for explanations related to development, financial frictions, learning, and organizational capital).

2 Evidence on US MNE Expansion

Our empirical evidence focuses on the expansion of affiliate activities over time, the expansion of the MNE over space, and the nature of the shocks that MNEs are exposed to. We start with a description of the data.

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. Additionally, we consolidate affiliates belonging to the same parent and operating in the same country and 3-digit industry. Online Appendix O.1 provides more details on the data coverage and sample construction.

Crucially, the BEA data break down affiliate sales by destination: the host market of operation (horizontal sales) and other markets (export sales). The data distinguish between affiliate exports

Table 1: Summary Statistics.

	Horizontal sales	Export sales
	(1)	(2)
No. of observations	132,493	132,493
with positive sales	127,220	88,060
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%

Notes: 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.

to the United States and to other markets. Every five years, the data further distinguish affiliate exports to Canada, the United Kingdom, and Japan.

Table 1 shows the number of observations with 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.

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.

2.2 Affiliate activities over time

In this section we present evidence on the composition of affiliate sales over time. To establish the evidence from time series data, we need to observe affiliates 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. Online Appendix O.1 provides a detailed description of this sample.

I. MNE affiliates' activities start in their host market and expand into export markets.

Figure 1 shows the evolution of the intensive and extensive margins of horizontal and export sales of US 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. 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.⁸ 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. The patterns in Figure 1 are confirmed by OLS regressions that include a battery of fixed effects (see Online Appendix Table O.4).

Together, these findings suggest that, for horizontal activities, changes in sales shares are coming from 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.

These findings motivate the following assumption of our model: all affiliates sell their output in the host market, starting at birth and throughout their life, while they may start exporting at birth or later in life.

II. Affiliate sales, as a share of parent's, grow only at the time of export entry.

In Figure 2a, we show that affiliate horizontal sales, relative to parent sales, are close to their long-run value at birth and remain stable over the affiliate life span, growing only ten percent in their first four years. At the same time, Figure 2b shows that there is a jump in affiliate-to-parent sales in the year when the affiliate starts exporting.

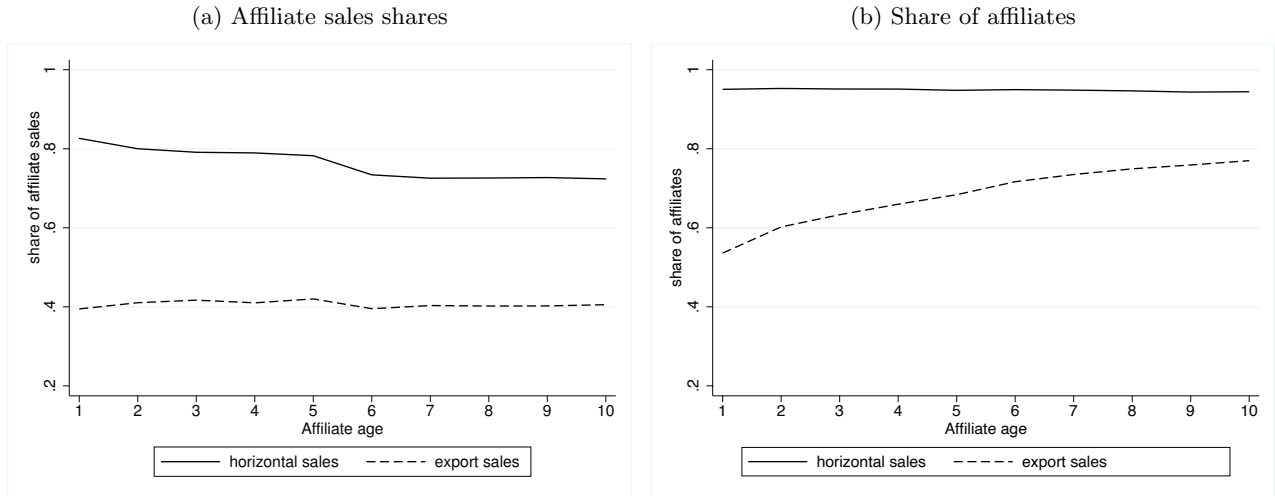
While Figure 2 shows the raw data, these patterns are robust to including country-year and affiliate fixed effects in an OLS regression (see Online Appendix Table O.5).⁹

Taken together, the results in Figure 2 show that affiliates are born with sales, relative to the

⁸ 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 number of affiliates with positive horizontal sales is larger than the ones with positive export sales. Including zeros, the share of export sales increases from 21 percent at age one to almost 32 percent at age 10.

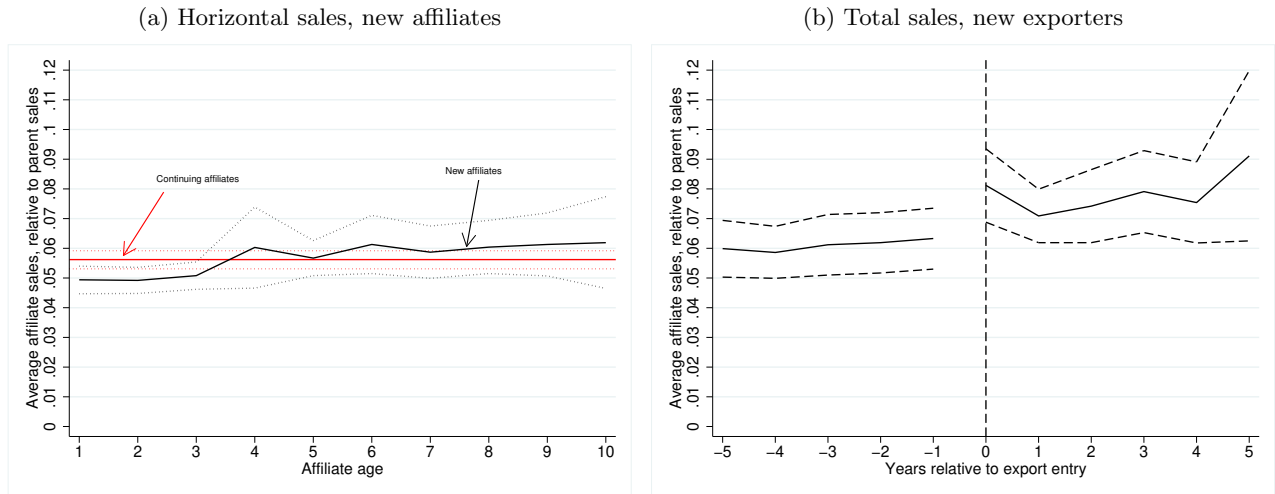
⁹Results are robust to using different subsamples of affiliates (see Online Appendix Table O.6): first, initial affiliates do not grow faster than subsequent affiliates of the same parent, suggesting that the age of the firm at the time an affiliate opens is not driving the results; second, affiliates 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 or 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 1: Affiliate sales and number of affiliates: horizontal vs export sales.



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.

Figure 2: Affiliate sales relative to parent sales.



Notes: Figure 2a reports the average value of affiliate horizontal sales relative to the domestic sales of the US parent, for new affiliates surviving at least ten years (black) and for all affiliates of all ages (red). 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. Dashed lines denote 95% confidence intervals.

parent's, which are as large as incumbent affiliate sales, and that grow significantly only when the affiliate expands to export markets. This fact suggests the presence of entry costs, both in the host and export markets, and the importance of the extensive margin for MNE expansion: (relative)

growth occurs discretely when entering new markets, including the host market.¹⁰ Our modeling strategy will focus on these margins, while our quantitative exercises will evaluate the magnitude of their responses to trade shocks.

2.3 Geography, affiliate entry, and affiliate sales

In this section, we document the expansion of US MNEs over space. Using the full sample of affiliates of US MNEs (not only affiliates surviving for at least ten years), we show that the location of a new affiliate barely depends on the location of preexisting affiliates of the same MNE, and that the sales of affiliates in a host country do not change with local and export activities 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. Figure 3a shows that, 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. This is particularly true for richer host economies. Differences are slightly larger, but still small in magnitude, for less developed host economies, which are typically more engaged in global value chains (GVCs).¹¹

To lend further credibility to our results, we implement an instrumental variable approach similar to the one used by Kovak et al. (2021). In the first stage, we predict the existence of a “sibling” affiliate in other countries using the existence of a bilateral tax treaty (BTT) between the United States and the host country of the affiliate at time t .¹² We define a sibling affiliate as an affiliate of an MNE operating in the same three-digit industry as another affiliate of the same firm. Figure

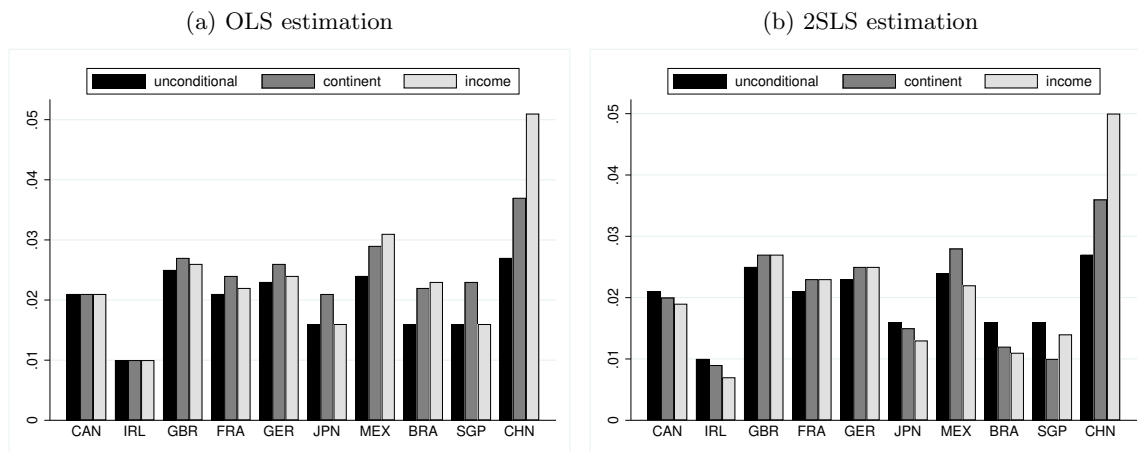
¹⁰Online Appendix Table O.14 presents additional evidence supporting the importance of sunk costs. Analogously to the literature on trade dynamics (e.g. Bernard and Jensen, 1999; Roberts and Tybout, 1997; Alessandria and Choi, 2007), we document that MNEs’ horizontal and export activities are persistent, and that US parents are larger at the time of affiliate entry than at the time of affiliate exit.

¹¹Online Appendix Table O.7 reports the estimated coefficients, also including as “similarity” variables sharing a border and sharing a language; these conditional probabilities are also extremely similar — and in most cases statistically equal — to unconditional probabilities. Furthermore, Online Appendix Table O.9 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 exports). However, differences between conditional and unconditional entry probabilities are still small and often insignificant for both GVC and non-GVC affiliates. Finally, Online Appendix Table O.10 shows that differences between unconditional and conditional probabilities become smaller when we restrict the sample to MNEs with more than five, and more than ten, affiliates.

¹²Kovak et al. (2021) show that BTTs are exogenous to MNE activities and can predict employment at foreign affiliates of US MNEs. Online Appendix O.3 presents a description of the first-stage results.

3b shows that results are virtually unchanged relative to OLS.¹³

Figure 3: Unconditional and conditional probability of affiliate entry.



Notes: 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.

Our findings are in stark contrast with the analogous findings 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*.¹⁴

The pattern of affiliate entry documented in this section motivates a key assumption of our model: affiliate entry decisions in each market are independent across host markets.¹⁵

IV. Sales of existing affiliates do not change with sibling activities.

¹³Online Appendix Table O.8 includes other “similarity” variables using the IV strategy.

¹⁴This finding does not contradict the fact that US MNEs open affiliates in closer and larger markets first, as documented in Egger et al. (2014) for German MNEs, and as we show in Online Appendix Table O.13 for US MNE affiliates.

¹⁵In recent work, Antràs et al. (2024) use newly merged data on MNE production and trade activities to document a rich set of static facts that suggest the existence of MNE-wide fixed costs to serve (or source from) a market. In contrast, the evidence in Section 2.2 (i.e. exports are incorporated slowly into the affiliates’ activities) suggests the presence of affiliate-level frictions.

Our last fact shows that horizontal 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. For this analysis, we restrict affiliate horizontal sales to *unaffiliated* parties to make sure that our findings are not contaminated by GVCs within the MNE.¹⁶

First, an affiliate may shrink, or expand, in its country of operation when a sibling affiliate opens in a different host country. Figure 4a shows 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 in 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.

Second, an MNE may serve a destination from sibling affiliates operating in the same 3-digit industry and located in different host countries. The BEA collects affiliate export sales data by destination for various geographic areas covering all countries in the benchmark surveys conducted during our sample period. However, because Canada, United Kingdom, and Japan are the only individual countries among the geographic areas for which the data are collected, for this exercise we restrict our sample to affiliates that are located in, or export to, those three countries. Figure 4b reports the average horizontal sales 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 country j .¹⁷ As the figure shows, the horizontal sales of an affiliate do not significantly change when a sibling affiliate in the same industry starts exporting to that market from a different location.


Lastly, Figure 4c reports the average affiliate export sales to country j before and after an affiliate of the same parent operating in the same industry but located in a different country starts exporting to j . These data are available for benchmark survey years so that we show exports 5, 10, 15, and 20 years before and after a sibling starts exporting to j . Also in this case, export sales of the affiliates do not significantly change when a sibling affiliate in the same industry but located in a different country starts exporting into the preexisting affiliate's export market.

The lower panels of Figure 4 shows results using the same instrumental variable approach as in Figure 3b. Once again, results are virtually unchanged relative to OLS.

¹⁶ Online Appendix Figure O.1 shows results for total horizontal sales, both to unaffiliated and affiliated parties. Results barely change.

¹⁷Of the 20,359 affiliates that exported to Canada in 1999, 64 percent have a sibling affiliate located in Canada; 70 percent of the 5,017 affiliates that exported to the United Kingdom have siblings located there; for Japan, this share

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



Notes: Horizontal sales refer to sales to unaffiliated parties in the host market of the affiliate. Export sales refer to exports to unaffiliated parties. Figure 4a (4d) includes all affiliates of US MNEs in manufacturing. Figure 4b (4e) includes affiliates located in Canada, United Kingdom, and Japan. Figure 4c (4f) includes affiliates exporting to Canada, United Kingdom, and Japan. Figures 4a (4d) and 4b (4e) show coefficients (solid line) from regressing the log of horizontal sales for an affiliate in country j and 3-digit industry h belonging to parent p at time $t \in \{-5, \dots, 5\}$ on a set of dummies indicating: in Figure 4a (4d), time from opening an affiliate in country $j' \neq j$ belonging to the same parent p and industry h ; and in Figure 4b (4e), time from export entry into j of an affiliate located in country $j' \neq j$ belonging to the same parent p and industry h , with $j = \text{Canada, UK, Japan}$. Figure 4c (4f) shows coefficients (solid line) from regressing the log of export sales to country j by affiliates located in country $k \neq j$ and 3-digit industry h belonging to parent p at time $t \in \{-20, -15, \dots, 15, 20\}$ on a set of dummies indicating time from export entry into j of an affiliate located in country $k' \neq j$ and $k' \neq k$ belonging to the same parent p and industry h as the affiliate located in k and exporting to j , with $j = \text{Canada, UK, Japan}$. 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.

The patterns of MNE entry and sales in a market documented in this section reinforce the model's assumption that both entry and sales decisions into a market (host and export) are independent across markets.

was 47 percent, out of the 5,224 affiliates that exported to this market in 1999.

2.4 MNE shocks

An important feature of our modeling strategy is the inclusion of different types of shocks that the MNE is subject to. In order to help assess the relative importance of aggregate versus firm-level uncertainty, and time-varying versus time-invariant shocks, we examine the types of shocks that drive variation in affiliate sales. Concretely, in Table 2, we show OLS estimates from regressing the log of horizontal sales of an affiliate belonging to parent p , located in country j and industry h , at time t , on different combinations of fixed effects. We analyze the increase in the R-squared as we introduce these fixed effects. Aggregate uncertainty, captured by host country-year and industry-year fixed effects, accounts for almost 20 percent of the variation in horizontal affiliate sales, while adding a parent fixed effect brings the R-squared to 0.44.¹⁸ Adding parent sales to capture time-varying parent-level shocks barely changes the R-squared. In contrast, adding affiliate fixed effects, which capture invariant characteristics of the affiliate, doubles the R-squared relative to columns 2 and 3, bringing it to 0.88.

Guided by this evidence, our model includes aggregate time-varying uncertainty and a firm-specific invariant productivity term subsuming both parent and affiliate fixed effects. We leave out time-varying firm-level shocks, which according to the R-squared of our estimates account for around 12-15 percent of the variation in affiliate sales.

Table 2: MNE shock structure. OLS.

Dependent variable	log of horizontal affiliate sales			
	(1)	(2)	(3)	(4)
Host-year fixed effect	yes	yes	yes	yes
Industry-year fixed effect	yes	yes	yes	yes
Parent fixed effect	no	yes	yes	no
Parent sales	no	no	yes	yes
Affiliate fixed effect	no	no	no	yes
R-squared	0.19	0.44	0.44	0.88
Adjusted R-squared	0.17	0.41	0.41	0.85

Notes: Full sample. Number of observations: 132,493.

¹⁸The evidence in Table 2 complements the facts documented by Cravino and Levchenko (2017). Using data containing several MNE source countries, they can distinguish between the contribution of host and source country shocks on the affiliate total sales growth, which they estimate to be 18 and 10 percent, respectively.

3 A Dynamic Model of MNE Expansion

We build a dynamic model where firms open affiliates, locally and abroad, over time. Affiliates sell in their host markets of operations, 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 introduce a compound option formulation, which allows us to characterize the richness of the decisions of MNEs in time and space. This formulation is novel to the international trade literature.

3.1 Preferences and technology

The economy consists of N countries. In each country k , a fixed mass of L_k consumers have preferences over a differentiated good,

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

where ρ denotes the subjective discount rate, and $Q_k(t)$ aggregates a continuum of tradable varieties, indexed by v ,

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

with $\eta > 1$. The variable $q_{ijk}(v, t)$ denotes consumption of variety $v \in \Omega_{ijk}(t)$, where $\Omega_{ijk}(t)$ is the set of varieties sold to country k and produced by affiliates located in j belonging to firms from i , at time t .

Assumption 1 (Armington). *Varieties are firm and 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.¹⁹

Each country is populated by a continuum of firms. Firms decide where to establish affiliates, including their own local market, and whether to export or not from each of them.

Labor is the only factor of production. Each firm produces with a linear technology, characterized by a productivity parameter φ . When a firm operates an affiliate in a foreign country, the affiliate inherits the home-country firm productivity φ . Under monopolistic competition, the profit-

¹⁹ With CES preferences, the implication of this assumption for firm entry is isomorphic to assuming that firm productivity is a random variable drawn from an independent Pareto distribution over varieties and production locations. In the static model in Arkolakis et al. (2018), this assumption implies that the probability of opening an affiliate in market j is independent from the probability of opening an affiliate in market j' .

maximizing price of a variety produced in country j destined to k is given by a mark-up over the unit cost of production,

$$p_{ijk}(\varphi, t) = \frac{\eta}{\eta - 1} \frac{\tau_{ijk} W_j(t)}{\varphi}. \quad (3)$$

The variable $W_j(t)$ denotes the wage in country j , at time t , while $\tau_{ijk} \geq 1$ is the iceberg-type cost of producing goods in j with a technology from i and shipping them to k . This cost includes the efficiency loss of using a technology for production in a different location from its origin (Ramondo and Rodríguez-Clare, 2013), the standard trade cost of shipping goods internationally, and losses from selling in a different market from the headquarter country of the MNE (Head and Mayer, 2019; Wang, 2021).

When a firm from home country i starts serving market k from production location j , it must pay a sunk cost, F_{ijk} , in units of market- j labor. For $j = k$, the cost F_{ijj} refers to the cost of setting up an affiliate in country j , a cost that must be incurred if the firm wants to sell to market j —and eventually to other markets k from j . Additionally, firms from i face a per-period (marketing) fixed cost, f_{ijk} , to serve market k , also paid in units of labor in j .

Profits from sales to destination k for a firm from i with productivity φ producing in location j at time t are

$$\pi_{ijk}(\varphi, t) = \frac{1}{\eta} x_{ijk}(\varphi, t) - W_j(t) f_{ijk}, \quad (4)$$

where

$$x_{ijk}(\varphi, t) = p_{ijk}(\varphi, t)^{1-\eta} P_k^{\eta-1}(t) X_k(t) \quad (5)$$

are revenues, $P_k(t)$ denotes the CES price index, and $X_k(t)$ denotes total expenditure in k .

For $j = k$, $\pi_{ijk}(\varphi, t)$ denotes profits from domestic sales by local firms ($i = j$) or local sales by foreign affiliates ($i \neq j$). For $j \neq k$, $\pi_{ijk}(\varphi, t)$ denotes profits from export sales, either by local firms ($i = j$) or foreign affiliates ($i \neq j$).

Our setup has the following implications. First, Assumption 1, coupled with CES preferences and monopolistic competition, implies that there is no cannibalization of sales within the MNE, since goods produced in different locations are different goods. This implication is consistent with our fourth fact in Section 2, as well as with evidence from the automobile industry in Head and Mayer (2019) where 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 in the next section, entry decisions into a location depend on other locations through the compound-option formulation of the dynamic problem and general equilibrium effects.

3.2 The MNE dynamic problem: the compound option

At each point in time, a firm decides whether and where to open an affiliate, and whether and where to export from existing affiliates. Affiliate exports may include exporting back to the home market. A firm may also decide to shut down affiliates, or to exit any of its export markets.

Guided by the empirical observations in Section 2.2, we make the following assumption.

Assumption 2 (Sequential affiliate activities). *A new affiliate must sell first to the market where its production activities are located.*

Assumption 2 leads to the notion of a *compound option* to model the dynamic problem of the MNE. Opening an affiliate 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, introducing dynamic interdependence in the MNE location choices.

The compound option structure allows us to solve the firm’s problem backwards, as suggested by Dixit and Pindyck (1994, chap. 10). Conditional on the firm having production operations 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 sales to the host market, 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.

Summing up, Assumptions 1 and 2 combine to avoid the complex combinatorial problem that the MNE dynamic problem in a multi-country setting with sunk costs would typically involve. Because of Assumption 1, we can solve the affiliate export problem market by market. Thanks to Assumption 2, we can sequentially solve the problem of opening an affiliate and exporting from it.

Aggregate shocks. Following Ghironi and Melitz (2005), we assume that firm productivity φ is the product of a time-invariant firm-specific component, z , and a time-varying home country-specific component, Z_i so that $\varphi \equiv z Z_i$. The term z is drawn from a time-invariant distribution, $G(z)$, as in Melitz (2003). The term Z_i follows a geometric Brownian motion,

$$\frac{dZ_i}{Z_i} = \mu_i^Z dt + \sigma_i^Z dW_i^Z, \tag{6}$$

for $\mu_i^Z \in \mathbb{R}$, $\sigma_i^Z > 0$, dW_i^Z denoting a standard Wiener process, where we allow for correlation across countries. When a firm operates an affiliate in a foreign country, it transfers both the aggregate and the idiosyncratic components of productivity 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 (6) is equivalent to assuming that aggregate productivity growth behaves as a random walk, and that productivity growth is independently and identically distributed. We allow, however, for non-zero correlation of productivity shocks across countries, denoted by ρ . Computationally, relying only on aggregate shocks makes feasible the aggregation of individual firms' decisions and the computation of equilibrium variables for many countries; 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. Empirically, this shock specification is consistent with the main sources of variation of US MNE affiliate sales documented in Section 2.4, whereby only a small share of that variation is driven by parent- and affiliate-level time-varying shocks. Moreover, the persistence of the aggregate productivity shock, together with positive growth rates ($\mu_i^Z \geq 0$), gives rise to the dynamic patterns documented in Figure 1: affiliates start serving their host market, and later on, they start expanding to export markets.

It is useful to re-define the productivity shock as

$$Y_i \equiv Z_i^{\eta-1}, \quad (7)$$

which is also a geometric Brownian motion, with drift and standard deviation given by $\mu_i = (\eta - 1)\mu_i^Z + \frac{(\eta-1)(\eta-2)}{2}\sigma_i^{Z^2}$, and $\sigma_i = (\eta - 1)\sigma_i^Z$. We assume that $\mu_i < \rho$, for all i , to ensure that the expected present discounted value of profits is finite.

We also define the aggregate (endogenous) destination-specific term $\Psi_k \equiv P_k^{\eta-1} X_k$, to which we refer as the destination demand shock.

3.3 The problem of the myopic firm

While the combination of Assumptions 1 and 2 avoids the combinatorial problem typically present in multi-country MNE problems with sunk costs, the inclusion of aggregate shocks coupled with general equilibrium responses makes solving for the expected discounted stream of profits of a forward-looking firm, and hence, applying standard guess-and-verify procedures to solve for the value functions, unfeasible, as we show below.

To make progress in solving the quantitative model, we assume that the firm is myopic.²⁰ We define

²⁰ Using the problem of the myopic firm is inspired by the analysis in Leahy (1993), who shows that, in a closed

a myopic firm as a firm that only takes into account the direct effect of its own home productivity shock on the future stream of profits, but ignores the future changes in aggregate variables driven by the shocks. That is, a myopic firm assumes that aggregate variables will remain at their current values in future periods. In this way, expectations driving firms' decisions are taken exclusively over the home country shock Y_i . In contrast, if the firm were forward looking, this expectation would be taken over $\mathbf{Y} = \{Y_1, \dots, Y_N\}$, since the firm recognizes the dependence of aggregate variables on the entire vector of shocks. This problem retains the convenience of the compound-option formulation, and will allow us to solve for the value functions.

The expected net present value of a firm z from country i that follows an optimal policy is given by

$$\mathcal{V}_i(z; Y_i, \mathbf{W}, \mathbf{\Psi}) = \sum_{j=1}^N \max \{V_{ij}^o(z; Y_i, W_j, \mathbf{\Psi}), V_{ij}^a(z; Y_i, W_j, \mathbf{\Psi})\}. \quad (8)$$

where $\mathbf{W} = \{W_1, \dots, W_N\}$ and $\mathbf{\Psi} = \{\Psi_1, \dots, \Psi_N\}$.²¹ The function V_{ij}^o is the option value of an affiliate in country j , and V_{ij}^a is the value of an existing affiliate in country j , regardless of the destination of its sales. In turn, the value of an affiliate in country j is given by

$$V_{ij}^a(z; Y_i, W_j, \mathbf{\Psi}) = V_{ij}^h(z; Y_i, W_j, \mathbf{\Psi}) + \sum_{k \neq j} \max \{V_{ijk}^o(z; Y_i, W_j, \Psi_k), V_{ijk}^e(z; Y_i, W_j, \Psi_k)\}. \quad (9)$$

The function V_{ij}^h captures the value of sales to the host market—i.e. sales to country j by a firm located in j , and indicates that all affiliates sell in the country where they are located. The value of an affiliate also includes the value of its exports: V_{ijk}^o is the option value of exporting to country k for a firm located in j , and V_{ijk}^e is the value of exports to country k from j .²²

Equations (8) and (9) reflect the sequential nature of the problem in Assumption 2. The problem is formulated as a compound option because opening an affiliate is equivalent to exercising an option that gives access to another set of options — the options to export to any other country. In this way, the value of an affiliate can be written as a separable additive function of the value at each location, either realized or option.

economy (i.e. partial) equilibrium with aggregate shocks, the solution to the problem of a myopic firm (i.e. a firm that ignores the evolution of investment) and the one of a forward-looking firm coincide and are efficient. The Leahy equivalence does not apply in our general equilibrium model with firms that can produce in one location and export from there to another one.

²¹ Formally, for comparison, the aggregate state space for the forward-looking firm would be $\{Y_i, W_j(\mathbf{Y}), \mathbf{\Psi}(\mathbf{Y})\}$.

²² Although V_{ij}^h represents the value associated with serving the host market, it depends on other destination states $\mathbf{\Psi}$ through the affiliate's shutdown option. See (13) below.

3.4 Solution method.

The problem of the firm can be solved first for V_{ijk}^o and V_{ijk}^e , conditional on the firm having an affiliate in country j , and then for V_{ij}^a and V_{ij}^o . Next, we write down the firm optimization problem by making explicit the Bellman equations that each of the value functions satisfies.

Export entry and exit. The Bellman equation describing the value of the option to export to country k for an affiliate in country j of a firm from i is given by

$$V_{ijk}^o(z; Y_i, W_j, \Psi_k) = \max \left\{ \frac{1}{1 + \rho\Delta t} \mathbb{E}[V_{ijk}^o(z; Y'_i, W_j, \Psi_k) | Y_i, W_j, \Psi_k]; V_{ijk}^e(z; Y_i, W_j, \Psi_k) - W_j F_{ijk} \right\}. \quad (10)$$

While the expectation of the value function is over the shock next period, Y'_i , given our definition of a myopic firm, the aggregate states are kept as in the current period, W_j and Ψ_k .²³

The expression in (10) indicates that 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 $W_j F_{ijk}$, in units of labor in j , and gets the value of exporting to k from j , V_{ijk}^e . In turn, this value is given by

$$V_{ijk}^e(z; Y_i, W_j, \Psi_k) = \max \left\{ \frac{1}{1 + \rho\Delta t} [\pi_{ijk}(z; Y_i, W_j, \Psi_k)\Delta t + \mathbb{E}[V_{ijk}^e(z; Y'_i, W_j, \Psi_k) | Y_i, W_j, \Psi_k]]; V_{ijk}^o(z; Y_i, W_j, \Psi_k) \right\} \quad (11)$$

An affiliate may keep exporting to country k —and get the export profit flow plus its continuation value—or may stop exporting to country k , in which case it gets the value of the option of exporting to k from j , V_{ijk}^o . For $i = j$, (10) and (11) describe the value of exports from the firm's home market. For $j = k$, the problem is specified next.

Host-country entry and exit. We next define the Bellman equations associated with a firm's decision to open an affiliate. If a firm has not yet opened an affiliate in country j , all the value of its operations in j is option value—i.e., the value of the possibility of entering j in the future,

$$V_{ij}^o(z; Y_i, W_j, \Psi) = \max \left\{ \frac{1}{1 + \rho\Delta t} \mathbb{E}[V_{ij}^o(z; Y'_i, W_j, \Psi_j) | Y_i, W_j, \Psi_j]; V_{ij}^a(z; Y_i, W_j, \Psi) - W_j F_{ijj} \right\}, \quad (12)$$

A firm may keep the option of entering market j , in which case it gets the continuation value of that option, or may enter country j by opening an affiliate there, in which case it pays the entry cost, $W_j F_{ijj}$, in units of labor in j , and gets the value of having an affiliate in country j , V_{ij}^a . Thanks to

²³ It is worth noting here the difference with a forward-looking firm. Following the notation in Footnote 21, the forward-looking firm would have expectations given by $\mathbb{E}[V_{ijk}^o(z; \mathbf{Y}', W_j(\mathbf{Y}'), \Psi_k(\mathbf{Y}')) | \mathbf{Y}, W_j(\mathbf{Y}), \Psi_k(\mathbf{Y})]$.

Assumption 1, a firm evaluates entry into each location separately.

The value of an affiliate is given by the value of its sales to the host and export markets, as defined in (9). Since all affiliates sell in the market where they are located, the value of sales to country j for an affiliate located in j of a firm from i is given by

$$V_{ij}^h(z; Y_i, W_j, \Psi) = \max \left\{ \frac{1}{1 + \rho \Delta t} \left[\pi_{ijj}(z; Y_i, W_j, \Psi) \Delta t + \mathbb{E}[V_{ij}^h(z; Y_i', W_j, \Psi) | Y_i, W_j, \Psi] \right]; V_{ij}^o(z; Y_i, W_j, \Psi_j) \right\}. \quad (13)$$

An affiliate located in j makes profits from selling to j , or may shut down and get the value of the option of opening an affiliate in j in (12); for $i = j$, (13) describes the value of sales to the firm's home market. This structure deliberately imposes that the exit margin depends on the value of sales to the host market exclusively, and does not involve the value of exports; once host market sales are not valuable enough, the affiliate would exit. The problem formulated in this way is justified by the fact that only an extremely small fraction of affiliates only export and do not sell to their location of production.

Writing the Bellman equations in their continuation region and applying Ito's lemma, one obtains the standard no-arbitrage conditions and can guess-and-verify the value functions (see Appendix A for details). For the export entry and exit problems, we guess that

$$V_{ijk}^o(z; Y_i, W_j, \Psi_k) = B_{ijk}^o(z; W_j, \Psi_k) Y_i^{\beta_i} \quad (14)$$

$$V_{ijk}^e(z; Y_i, W_j, \Psi_k) = \frac{\tilde{\pi}_{ijk}(z; W_j, \Psi_k)}{\rho - \mu_i} Y_i - \frac{W_j f_{ijk}}{\rho} + A_{ijk}^e(z; W_j, \Psi_k) Y_i^{\alpha_i}, \quad (15)$$

while for the host-country entry and exit problems we similarly guess that

$$V_{ij}^o(z; Y_i, W_j, \Psi) = B_{ij}^o(z; W_j, \Psi) Y_i^{\beta_i} \quad (16)$$

$$V_{ij}^h(z; Y_i, W_j, \Psi) = \frac{\tilde{\pi}_{ijj}(z; W_j, \Psi_j)}{\rho - \mu_i} Y_i - \frac{W_j f_{ijj}}{\rho} + A_{ij}^h(z; W_j, \Psi) Y_i^{\alpha_i}, \quad (17)$$

Here, $A_{ijk}^e(z; W_j, \Psi_k)$, $B_{ijk}^o(z; W_j, \Psi_k)$, $A_{ij}^h(z; W_j, \Psi)$, and $B_{ij}^o(z; W_j, \Psi)$ are firm-specific terms to be determined, the parameters $\beta_i > 1$ and $\alpha_i < 0$ are the roots of the fundamental quadratic equation, a standard object in problems of investment under uncertainty formulated as real options, and

$$\tilde{\pi}_{ijk}(z; W_j, \Psi_k) \equiv \eta^{-1} \left(\frac{\eta}{\eta - 1} \right)^{1-\eta} (\tau_{ijk} W_j)^{1-\eta} \Psi_k z^{\eta-1}. \quad (18)$$

The expression in (14) captures the option value of affiliate exports. The first two terms of (15)

are the expected present discounted value of export profits, while the third term represents the option value of exit from the export market. While the option value of exporting is increasing in the realization of the shock, the option value of exiting that export market is decreasing in the realization of the home productivity shock — i.e., the option of entering (exiting) an export market has a larger value in good times (bad times). The expressions in (16) and (17) have a similar interpretation for entry and exit into the host market j .

The myopic-firm assumption is crucial to solve the value functions. The guesses in (14)-(17) are feasible because, since the myopic firm only considers the effects of the Home shock when forming expectations, the expected present discounted value of profits is a linear function of the current shock. Hence, we can verify that the flow term of the value function corresponds to the expected present discounted value of profits.

Combining the solution for V_{ij}^h and the solution of the affiliate export problem determines V_{ij}^a in (9).

We can then solve for the policy functions, which are the export entry and exit thresholds, $\bar{Y}_{ijk}(z; W_j, \Psi_k)$ and $\underline{Y}_{ijk}(z; W_j, \Psi_k)$, as well as for the terms $A_{ijk}^e(z; W_j, \Psi_k)$ and $B_{ijk}^o(z; W_j, \Psi_k)$, from the system of value-matching and smooth-pasting conditions,

$$V_{ijk}^e(z; \bar{Y}_{ijk}, W_j, \Psi_k) - W_j F_{ijk} = V_{ijk}^o(z; \bar{Y}_{ijk}, W_j, \Psi_k), \quad \frac{\partial V_{ijk}^e(z; \bar{Y}_{ijk}, W_j, \Psi_k)}{\partial Y_i} = \frac{\partial V_{ijk}^o(z; \bar{Y}_{ijk}, W_j, \Psi_k)}{\partial Y_i},$$

$$V_{ijk}^e(z; \underline{Y}_{ijk}, W_j, \Psi_k) = V_{ijk}^o(z; \underline{Y}_{ijk}, W_j, \Psi_k) \quad \text{and} \quad \frac{\partial V_{ijk}^e(\underline{Y}_{ijk}, W_j, \Psi_k)}{\partial Y_i} = \frac{\partial V_{ijk}^o(\underline{Y}_{ijk}, W_j, \Psi_k)}{\partial Y_i}.$$

We can then solve for entry and exit thresholds into host market j , $\bar{Y}_{ij}(z; W_j, \Psi)$ and $\underline{Y}_{ij}(z; W_j, \Psi)$, as well as for the option-value terms, $A_{ij}^h(z; W_j, \Psi)$ and $B_{ij}^o(z; W_j, \Psi)$, from the corresponding system of value-matching and smooth-pasting conditions for the host-country problem,

$$V_{ij}^a(z; \bar{Y}_{ij}, W_j, \Psi) - W_j F_{ijj} = V_{ij}^o(z; \bar{Y}_{ij}, W_j, \Psi) \quad V_{ij}^h(z; \underline{Y}_{ij}, W_j, \Psi) = V_{ij}^o(z; \underline{Y}_{ij}, W_j, \Psi),$$

$$\frac{\partial V_{ij}^a(z; \bar{Y}_{ij}, W_j, \Psi)}{\partial Y_i} = \frac{\partial V_{ij}^o(z; \bar{Y}_{ij}, W_j, \Psi)}{\partial Y_i} \quad \text{and} \quad \frac{\partial V_{ij}^h(\underline{Y}_{ij}, W_j, \Psi)}{\partial Y_i} = \frac{\partial V_{ij}^o(\underline{Y}_{ij}, W_j, \Psi)}{\partial Y_i}.$$

Properties of the solution. The combination of aggregate shocks and constant firm-level productivity implies that the problem of the myopic firm displays intuitive sorting properties, which we exploit to derive additional testable implications, solve the model's general equilibrium, and implement the numerical algorithm.

First, at each point in time, more productive firms (higher z) are more likely to enter both a host

market and an export market. Under the assumptions that $\bar{Y}_{ij}(z) < \bar{Y}_{ijk}(z)$ (we ease the notation by skipping the dependence on aggregates), for all $k \neq j$, and $\mu_i > 0$, this result implies that affiliates that are exporters from birth have larger sales to their host market than non-exporting affiliates, and that affiliates that start exporting later in life have lower host-market sales than affiliates that start earlier. Second, the same property implies that MNEs with larger parent sales enter more foreign markets through FDI; and the number of MNEs with affiliates in n host markets is decreasing in n , so that there is a negative correlation between the number of firms with affiliates in n markets and their parent sales. Third, for a firm with productivity z , the affiliate entry threshold, both to host and export markets, is increasing in the iceberg-type cost and sunk entry cost. Under the assumption that $\mu_i > 0$, the model predicts that an MNE opens first its largest affiliates, and conditional on the size of the host market, affiliates in markets that are less costly to enter.²⁴ In our quantitative analysis, we compare these model predictions with the data.

The myopic firm policy functions are invertible, and can hence be written as productivity-threshold functions of the shock: $\bar{z}_{ij}(Y_i)$ ($\underline{z}_{ij}(Y_i)$) are the minimum (maximum) productivity for firms to open (shut down) an affiliate in j ; and $\bar{z}_{ijk}(Y_i)$ ($\underline{z}_{ijk}(Y_i)$), for $j \neq k$, are the minimum (maximum) productivity for firms producing in j to start (stop) exporting to k . In the next section, we use the sorting properties of the model to write the share of affiliates entering and exiting as a simple truncation of the underlying distribution of firm productivity $G(z)$, which we assume Pareto in our calibration.²⁵

Overall, using the myopic firm problem delivers a fast algorithm for quantitative analysis, while also providing intuition about the properties of the firm problem. While this solution mutes some dynamic responses, it allows us to keep general equilibrium feedbacks and more realistic features of both MNE and domestic firm problems, such as rich geography and aggregate productivity shocks.

3.5 Aggregation and equilibrium

We now turn to the equilibrium conditions of the model. We start by presenting the evolution of the mass of firms in each location and destination. The mass of firms from i located in j and serving k evolve as

$$M_{ijk}(t + \Delta t) = M_{ijk}(t) - M_{ijk}^X(t + \Delta t) + M_{ijk}^E(t + \Delta t), \quad (19)$$

²⁴Propositions 1 and 2 in Online Appendix O.6 present a formal proof of these results under the additional assumption that 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.

²⁵Proposition 3 in Online Appendix O.7 presents the formal argument.

where $M_{ijk}^E(t + \Delta t)$ is the mass of new affiliates from i located in j that start serving k , drawn from the pool of firms that are already active in j but not yet serving k . The variable $M_{ijk}^X(t + \Delta t)$ denotes the mass of affiliates that exit k . All affiliates in j serve market j , so that $M_{ijj}(t)$ is the total mass of affiliates from i in j , and $M_{ijj}^E(t + \Delta t)$ is the mass of new entrants into j drawn from the (fixed) mass of potential entrants in i , but not yet in j , $M_i - M_{ijj}(t)$. For $i = j = k$, the expression in (19) simply captures the evolution of the mass of domestic firms, while for $i = j \neq k$, (19) captures the evolution of domestic exporters.

It is convenient to define the aggregate productivity index in an origin-location-destination country triplet as

$$z_{ijk}(t) \equiv \left[\int_{z(v):v \in \Omega_{ijk}(t)} z^{\eta-1} dG(z) \right]^{\frac{1}{\eta-1}}. \quad (20)$$

Aggregate sales to destination k of firms from i producing in j are given by

$$X_{ijk}(t) = \frac{(\tau_{ijk} W_j(t))^{1-\eta} Y_i(t) z_{ijk}^{\eta-1}(t)}{\sum_{i',j'} (\tau_{i'j'k} W_{j'}(t))^{1-\eta} Y_{i'}(t) z_{i'j'k}^{\eta-1}(t)} X_k(t), \quad (21)$$

where $X_k(t) = \sum_{i,j} X_{ijk}(t)$, and the denominator is proportional to $P_k(t)^{1-\eta}$.

Lastly, we present the market clearing conditions for labor and goods markets. Labor market equilibrium entails that total payments to labor in country j at each time t add up to payments to production workers, marketing workers, and workers devoted to the sunk entry cost,

$$W_j(t) L_j = \frac{\eta-1}{\eta} \sum_{i,k} X_{ijk}(t) + W_j(t) \sum_{i,k} [f_{ijk} M_{ijk}(t) + F_{ijk} M_{ijk}^E(t)],$$

where $\sum_{i,k} X_{ijk}(t)$ denotes total output in j . Total absorption in country j has to be equal to national income, at each time t ,

$$X_j(t) = W_j(t) L_j + \sum_{l,k} \Pi_{jlk}(t),$$

where $\Pi_{jlk}(t)$ are the aggregate profits associated with sales X_{jlk} ,

$$\Pi_{jlk}(t) = \frac{1}{\eta} X_{jlk}(t) - W_l(t) [f_{jlk} M_{jlk}(t) + F_{jlk} M_{jlk}^E(t)].$$

At each time t , the equilibrium conditions of our model are the same as the conditions of static models of trade and multinational production (e.g. Arkolakis et al., 2018). Dynamics enter through

changes over time of the sets $\Omega_{ijk}(t)$. The strict sorting properties of the solution of the myopic firm problem imply that the sets $\Omega_{ijk}(t)$ are connected intervals.²⁶ Hence, the mass of new, incumbent, and exiting firms can be written as truncations of the distribution $G(z)$. In this way, we can keep track of the firm distribution by status, making it easy to compute $M_{ijk}^E(t + \Delta t)$ and $M_{ijk}^X(t + \Delta t)$ in (19), as well as the productivity index in (20). Together with the assumptions on the shock structure and the continuous-time formulation, this property ensures aggregation. In turn, the myopic-firm assumption reduces the computation of the model’s equilibrium aggregates to a period-by-period problem, and enables the use of fast computational algorithms such as the one in Alvarez and Lucas (2007).

Traditionally, general equilibrium models of trade dynamics include firm-level shocks but exclude sunk costs, as in Luttmer (2007) and Arkolakis (2016), among others. In contrast, existing dynamic models that include sunk costs characterize the discrete-choice problem and equilibrium dynamics for a single firm, as in Das et al. (2007) and Morales et al. (2019), focus on stationary equilibria where aggregate variables do not change over time, as in Alessandria and Choi (2007), or restrict the geography to a two-country world, as in Impullitti et al. (2013). Since these models are often formulated in discrete time, the firm’s value function itself needs to be solved numerically. The compound-option formulation, coupled with unit-root aggregate shocks, allow us to solve for the value functions in an efficient way. Relative to the literature on exporter dynamics, we lose firm-level time-varying shocks, but we gain the ability to solve for the general equilibrium with many countries, maintaining a realistic spatial MNE problem and a realistic MNE shock structure. By solving the model using a myopic firm, we lose the general equilibrium effects on the firm’s entry and exit decisions, but we are able to capture the responses of aggregate prices and quantities, as well as wages, to the firm’s dynamic decisions regarding which market to serve and from which location.

4 Calibration

We calibrate the model to match the expansion of US MNEs during the period 1987-2011. Since we only observe MNEs headquartered in the United States, we restrict the source of MNEs to that country. All the other countries host US MNEs and are populated by domestic firms that produce only locally, but can also export to foreign destinations.

We construct data moments using the sample of all manufacturing affiliates operating in the top-ten host countries for US FDI: Brazil, Canada, China, France, United Kingdom, Germany,

²⁶See Proposition 3 in Online Appendix O.7).

Ireland, Japan, Mexico, and Singapore. This sample includes 83,214 affiliate-year observations, which account for almost 70 percent of all sales by foreign affiliates of US MNEs.²⁷

4.1 Procedure

We set the parameters capturing preferences and technology using estimates from the literature and direct observations from the data. Next, 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 set the elasticity of substitution to $\eta = 5$, in line with estimates in the literature (e.g. Broda and Weinstein, 2006). We set the time preference rate to $\rho = 0.085$ so that it does not violate the technical condition that ensures that the present discounted value of profits does not diverge ($\rho > \mu_i$ for all i).²⁸ We assume that the distribution of firm productivities z is Pareto with location parameter one and shape parameter $\vartheta = 4.25$, estimated by Kondo et al. (2023) using Census data for US firms.²⁹

We calibrate the aggregate shock process for each country in our sample using data on manufacturing value added (in current US dollars) and employment, from the World Bank, together with data on manufacturing trade flows, from Feenstra and Romalis (2014), for the period 1987-2011. We construct absorption in manufacturing, in each country and year, as manufacturing output minus the value of net manufacturing exports.³⁰

We assume that the drift of the aggregate shock is common across countries, and equal to the average growth rate of manufacturing value added per worker among the countries in our sample, $\mu_i = 0.065$, for all i , over 1987-2011.³¹

We set the variance-covariance matrix of Y_i to match the variance-covariance matrix of the observed growth rates of manufacturing value added per worker across countries. To initialize the shock

²⁷Online Appendix O.5 shows that using this smaller sample does not change the facts documented in Section 2.

²⁸This value might appear high, but it may include variables other than just the time preference rate. For example, if the model included an exogenous death rate, this parameter would be added to the time preference rate and the technical condition would allow for a lower time preference rate.

²⁹Since productivity is Pareto, the firms' sales distribution is also Pareto with shape $\tilde{\vartheta} \equiv \vartheta/(\eta - 1)$. Hence, using firm-level data, this parameter can be estimated by regressing the log-rank on the log-size of a firm. Kondo et al. (2023) estimate $\tilde{\vartheta}$ ranging from 0.75 to 1.25, depending on the cutoff used for employment on the right tail. We choose a value around the middle range of these estimates, which corresponds to larger firms (as US parent firms are), and that, given $\eta = 5$, satisfies the condition $\vartheta > \eta - 1$.

³⁰We use the share of value added in manufacturing, as reported by OECD STAN, to transform value added in manufacturing into gross output.

³¹Growth rates of manufacturing value added per worker across countries are highly heterogeneous, ranging from 0.025 for Japan to 0.163 for China. Using these country-specific values for the calibration of μ_i would force us to use an extremely high value of ρ to ensure that $\rho - \mu_i > 0$, and would imply unreasonably high effective discount rates for profits for the countries with lower μ_i .

process, we normalize the initial value in the United States to one. We then set $Y_i(0)$, for all $i \neq US$, to be equal to the average of country i 's manufacturing value added per worker relative to the United States over the period 1987-2011.

We set the potential mass of firms in each country to one, $M_i = 1$, for all i , and L_i to match the observed average manufacturing employment levels, for each country relative to the United States, over the period 1987-2011.

It remains to calibrate the parameters related to the costs of MNE expansion: the per-period fixed costs, f_{ijk} ; the sunk entry costs, F_{ijk} ; and the iceberg-type costs, τ_{ijk} . We impose the following restrictions.

First, since we do not observe entry and exit of US firms in the data, we impose that $F_{US,US,US} = f_{US,US,US} = 0$, which implies that the measure of US firms is fixed and all US firms operate in the US market. We allow for changes in the mass of domestic firms in the other countries in our sample, even if we do not observe those margins in the data. To such end, we assume that the sunk and fixed costs for firms entering and serving market j are the same for US MNE affiliates and domestic firms: $F_{US,jj} = F_{jjj} = F_j^h$ and $f_{US,jj} = f_{jjj} = f_j^h$. Furthermore, since we only observe MNEs headquartered in the United States, $F_{ijk}, f_{ijk} \rightarrow \infty$ for $i \neq US$ and $i \neq j$, for all k .

Second, we assume that the fixed and sunk costs of US MNE affiliate exports are symmetric across all destination countries, except for the United States: $f_{US,jk} = f_j^e$ and $F_{US,jk} = F_j^e$, for $k \neq j$, and $k \neq US$; and $f_{US,j,US} = f_{j,US}^e$ and $F_{US,j,US} = F_{j,US}^e$, for $j \neq US$. Since we do not observe export entry and number of exporters systematically for all the other countries, we assume that domestic firms in each country $j \neq US$ face the same sunk and fixed export costs to each destination, including the United States, as affiliates of US MNEs operating in j .

Lastly, we assume that the iceberg-type costs adopt the following separable form: $\tau_{ijk} = \tau_{ij}^M \tau_{jk}^T$ with $\tau_{ij}^M = \tau_j^M > 1$ for $i = US \neq j$, and $\tau_{ij}^M \rightarrow \infty$, for $i \neq US \neq j$. This cost can be interpreted as the efficiency loss of using the home technology abroad and it is not incurred by local firms, $\tau_{jj}^M = 1$ for all j .

The second term τ_{jk}^T , when $j \neq k$, denotes the standard iceberg trade costs of exporting from j to k . Since we observe export sales of US MNE affiliates from market j to Canada, Japan, the United Kingdom, and the United States, we use those sales as targets to calibrate τ_{jk}^T for $k \neq j$ and $k = CAN, JPN, UK, US$. For the remaining destinations, we assume that these costs are proportional to the geographic distance between j and k , an exporter-specific constant, and an error term,

$$\log \tau_{jk}^T = D_j + \beta^d \log Dist_{jk} + u_{jk}. \quad (22)$$

Under the assumption that productivity z is Pareto and given the properties of the sets $\Omega_{ijk}(t)$, export sales in (21), at each time t , have a gravity structure as defined by Arkolakis et al. (2012). Leveraging this structure, we calculate the distance elasticity β^d by estimating a standard gravity equation for the log of bilateral trade flows in manufacturing, using data from Feenstra and Romalis (2014), including importer and exporter fixed effects, and assuming a trade elasticity of $\eta - 1 = 4$.³² The variable D_j is chosen to match the export share of country j , an average across the years in our sample period. Notice that domestic exporters in each host country j , including exporters located in the United States, also face these iceberg trade costs.

We are left with 123 parameters to calibrate, for which we target 123 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, it is relatively easy to isolate the moment that drives the identification of a given parameter. Specifically, affiliate entry rates and the share of US MNE affiliates in each country help identify the sunk and fixed MNE costs, F_j^h and f_j^h . 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 . Export sales shares drive the identification of the iceberg trade cost τ_{jk}^T , while sales to country j of US MNE affiliates located in j , relative to US parent sales, help identify the efficiency loss τ_j^M .

We choose the values of the parameters to best fit the data moments, for each country. To this end, we simulate the model 100 times, each time for a different realization of the vector of aggregate shocks. Each simulation solves for the entry and exit problems for 500 firms for $t = 1, \dots, 30$. Computationally, this entails solving $2N^2 + 3N$ systems of four equations (given by value-matching and smooth-pasting conditions) in four unknowns (given by option value terms and thresholds in the realization of the shocks), for each firm, time period, and simulation. For each time period and simulation we condition on a set of aggregate variables $\{P_j(t), X_j(t), W_j(t)\}$, for $j = 1, \dots, N$, and $t = 1, \dots, 30$, which are then solved using the iterative algorithm in Alvarez and Lucas (2007). Appendix B describes the algorithm and updating rules in detail.

4.2 Results

Table 3 reports simulated and data moments, averaged across the top-ten host countries for US FDI and across years. The model matches extremely well static moments such as the affiliate sales shares to the different destinations and the share of affiliates serving those destinations. The calibrated model underestimates the affiliate entry margin, while it slightly overestimates the export entry

³² For this estimation, we use a large cross-section of countries, much larger than the ten countries used in the other parts of the calibration.

margin.³³

The table also includes three sets of non-targeted moments: affiliate size advantage, MNE sorting patterns, and affiliate exit.

First, panel 4 of Table 3 shows moments related to the affiliate size advantage.³⁴ While in the data the average sales to the host market of an affiliate that exports from birth are 6.3 times larger than those of an affiliate that never exports, our calibrated model generates a premium of about 26 times, much larger than the one observed in the data. Similarly, the observed average host-market sales of an affiliate that exports from birth are 3.7 times larger than those of an affiliate that starts exporting later in life, but in the calibrated model this early-exporter premium reaches 9.8. In contrast, the model is able to reproduce quite accurately the fact that the sales to the host market of a first affiliate of an MNE are larger than those of subsequent affiliates. As discussed in various papers in the literature (e.g. Armenter and Koren, 2015), trade models with sorting properties and Pareto tails fail to match the exporter premium observed in the data, conditional on matching the share of exporters. Our model suffers from a similar shortcoming: it creates large firms that are too large, a feature evident when comparing average and median moments.

Second, in panel 5 of Table 3, we show moments for MNEs analogous to the ones in Eaton et al. (2011) for exporters. Our model predicts that the MNEs with the largest US parent sales should enter more markets, including less popular ones. The calibrated model captures fairly accurately the magnitude of the elasticity of US parent sales with respect to the number of countries where US MNE affiliates operate, but it delivers an elasticity of average parent size to market popularity more than twice as large as the one observed in the data.

Third, as shown in the last panel of Table 3, the model underestimates exit rates of affiliates from their host market, but less so from their export markets. This result is not surprising given the absence of an exogenous death rate and of time-varying firm-level shocks in the model; exit rates are driven only by firms' endogenous responses to aggregate shocks.

Finally, our quantitative model reproduces the observation, first documented by Brainard (1997), that bilateral exports relative to bilateral affiliate sales to the host market decline with higher trade costs and increase with higher fixed costs of MNE activity. In a regression using model-generated data, the export-to-affiliate sales ratio, $X_{US,US,j}(t)/X_{US,j,j}(t)$, falls by 0.35 percent for every one percent increase in (relative) trade costs, $\tau_{US,j}^T/\tau_j^M$, while it rises by 4.19 percent with a one percent

³³Online Appendix O.9 reports the full set of simulated and data moments, as well as the calibrated parameters, for each host country.

³⁴See Online Appendix Figures O.2 and O.3 for more details.

Table 3: Moments: model versus data, averages.

	Data	Model
Targeted Moments		
1. Intensive margin		
1.1 Affiliate sales share to host country	0.026	0.026
1.2 Affiliate export share to the US	0.139	0.139
1.3 Affiliate export share to third countries	0.276	0.249
1.4 Affiliate export share to Canada	0.015	0.010
1.5 Affiliate export share to the U.K.	0.069	0.061
1.6 Affiliate export share to Japan	0.033	0.033
2. 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	0.563
2.3 Share of affiliates in j exporting to third countries	0.671	0.658
3. Entry		
3.1 Share of MNEs opening affiliates in j	0.035	0.024
3.2 Share of affiliates in j that start exporting to the US	0.030	0.032
3.3 Share of affiliates in j that start exporting to third countries	0.030	0.033
Untargeted Moments		
4. Affiliate size advantage		
4.1 Exporter size advantage	6.27	26.7 (12.0 [†])
4.2 Early exporter size advantage	3.68	9.83 (6.66 [†])
4.3 First affiliate size advantage	2.57	2.30 (1.46 [†])
5. MNE sorting		
5.1 Elasticity of average US parent sales w.r.t. # of markets entered	0.736	0.920
5.2 Elasticity of average US parent sales w.r.t # of firms selling to at least n markets	-0.424	-0.980
6. Exit		
6.1 Share of MNEs shutting down affiliates in j	0.113	0.065
6.2 Share of affiliates in j that stop exporting to the US	0.025	0.015
6.3 Share of affiliates in j that stop exporting to third countries	0.027	0.021

Notes: 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 sales to j of affiliates located 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 affiliate sales to their host market. The elasticities in 5. are computed by OLS aggregating the firm-level observations of MNEs that enter the same number of countries (5.1) and that sell to at least n markets. ([†]): Median values.

Table 4: Calibrated MNE costs: shares of sales and monetary values. Average across host countries.

	As % of sales			In thousands \$
	5th	50th	95th	
Sunk MNE cost $\tilde{F}_j^h(t)$, as % of US parent sales	0.61	0.32	0.06	2,389
Fixed MNE cost $f_j^h(t)$, as % of sales to j	51.3	28.5	4.15	3,565
Sunk export cost $\tilde{F}_{jk}^e(t)$, as % of sales to j				
To United States	28.4	19.8	10.6	1,951
To other destinations	46.4	29.0	14.6	2,865
Fixed export cost $\tilde{f}_{jk}^e(t)$, as % of sales to k				
To United States	10.7	5.74	0.90	55.9
To other destinations	9.43	5.21	1.98	315.3

Notes: We define $\tilde{x}_j(t) \equiv W_j(t)x_j$. US parent sales are evaluated in the year of affiliate entry. Sales to host market j are evaluated in the year the affiliate first exports to the destination, for the calculation of $\tilde{F}_{jk}^e(t)$, and averaged across years, for the calculation of $\tilde{f}_{jk}^e(t)$. Export sales to destination $k \neq j$ are averaged across years for the calculation of $\tilde{f}_{jk}^e(t)$. 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, reported median sales are an average of the 9 observations around the median.

increase in fixed costs, $W_j(t)f_j^h$. Larger markets in terms of expenditure $X_j(t)$ are also associated with a higher ratio, with a coefficient of 0.15.³⁵ These findings highlight that, as a result of general equilibrium effects on aggregate prices and wages, our model delivers substitution effects between aggregate trade and MNE activities.

The costs of MNE expansion. We now evaluate the magnitude of the costs to MNE expansion in time and space. Table 4 shows the calibrated MNE costs as shares of firm revenues and in monetary values.

On average, opening an affiliate requires only a small share of U.S. parent revenues—0.32 percent for the median case. Fixed operating costs vary widely, from about four percent of host-market sales for the largest affiliates to more than 50 percent for the smallest, with a median of 28.5 percent. Starting to export from an affiliate is much more costly, amounting to 29 percent of host-country revenues for the median affiliate, for sales to non-US destinations, and 19.8 percent when exporting back to the United States. Once established, however, affiliate export operations involve relatively low fixed per-period costs, about five percent of export sales for the median affiliate.

³⁵The unit of observation in our regression is at the time-simulation-market level, so that we have 29,966 observations. We include time and simulation fixed effects and robust standard errors, and the coefficients are all statistically significant at one percent level. Results are virtually unchanged if we include annualized sunk MNE costs, $W_j(t)(f_j^h + \rho F_j^h)$. We do not include the fixed cost of exporting from the United States to country j because it is common across destinations ($f_{US,j}^e = f_{US}^e$).

The right panel of Table 4 shows the costs for the median affiliate expressed in thousands of current US dollars. Starting MNE export operations to non-US destinations is more costly in monetary terms (around 2.8 millions US dollars) than starting affiliate operations in a host market (around 2.3 millions US dollars). However, the fixed costs of affiliate exports are much smaller than the fixed costs of host-market operations, more so if exports are directed to the United States.

Online Tables O.26 and O.27 show calibrated costs by country, highlighting the heterogeneity across countries, across cost type, and across firms of different size.

5 Quantitative Analysis

Armed with the calibrated model, we evaluate quantitatively a key feature of our theory: the role played by the extensive margin of MNE expansion, both over time and across space. To such end, we evaluate the impact on the extensive margin of: export platforms (i.e. the compound-option structure); sunk costs (i.e. dynamics); various types of trade-cost shocks (i.e. variable, fixed, sunk); and general equilibrium effects (i.e. wage adjustments).

Since our sample includes the United Kingdom, Ireland, Germany, and France, and since about half of US MNEs have UK-based affiliates, we perform exercises that change costs related to operating in—or exporting from and to—the United Kingdom, loosely mimicking the United Kingdom abandoning the European Union (EU), “Brexit.”³⁶ Additionally, given that our sample also includes China, Mexico, and Canada, among other countries that were impacted by the US import tariffs announced in 2025, we perform an exercise in which we increase the trade costs of shipping goods to the United States.

In all the exercises, we simulate the model for 30 periods and impose a permanent change in parameters related to the operations of MNEs in a host country or, alternately, to export costs, at $t = 15$.

5.1 The role of export platforms

We analyze the effects of an increase in the barriers to MNE activities in a model with and without export platforms —i.e. we allow for direct exports by domestic firms in each country, including US MNE parents, but not by foreign US MNE affiliates.

Our first exercise uses the baseline calibrated model and removes the option of doing export

³⁶Online Appendix Table O.22 reports the share of US affiliates in each of the host countries in our sample.

Table 5: The role of export platforms in MNE affiliate entry.

Model with:	Share of US MNE affiliates in host country j		% change in sales to host country j
	export platforms	no export platforms	no export- relative to export-platform
Brazil	0.20	0.12	-24.6
Canada	0.54	0.24	-3.79
China	0.18	0.11	2.48
France	0.31	0.14	4.43
Great Britain	0.55	0.22	10.2
Germany	0.37	0.16	3.33
Ireland	0.12	0.04	15.4
Japan	0.16	0.09	-2.23
Mexico	0.30	0.16	-6.63
Singapore	0.13	0.05	4.21
Average	0.29	0.13	0.28

Notes: The model with and without export platforms uses the baseline calibration. Averages over 30 periods.

platforms. Table 5 compares the share of US MNE affiliates in host country j with and without that option. The results indicate that the share of US MNE affiliates in the average host country is twice as high when export platforms are allowed. US MNE presence in countries such as Canada and Ireland would decrease by more than half if affiliates there were not allowed to export to other markets; the incentives to open an MNE affiliate in a country coming from the possibility of exporting are absent. Removing the possibility of export platforms decreases host-market sales in some host countries, like Canada, while it increases them in others, like Ireland. These substitution effects operate through the general-equilibrium forces of the model, a channel that we further explore below in Section 5.4.

Our second exercise re-calibrates the model assuming no export platforms (i.e. no compound option). We target the same moments as in our baseline calibration that do not involve US MNE affiliate exports (moments 1.1, 2.1, and 3.1 in Table 3).³⁷ The calibrated model without export platforms delivers estimates of sunk MNE costs that are virtually zero for all sales percentiles, while estimates of fixed costs are around half the magnitude of the estimates for the baseline model.³⁸ Not surprisingly, when we include the option of exporting from the host market, the value of opening MNE affiliates abroad increases, and hence, to match the shares of US MNE affiliates observed in the data, sunk (fixed) costs must be higher.


³⁷For domestic exporters, we keep the baseline calibration values of τ_{jk}^T , F_j^e , and f_j^e . See Online Appendix Table O.19.

³⁸Results are shown in Online Appendix Table O.27.

With both calibrations in hand, our next exercise increases the per-period fixed cost and the one-time sunk cost of horizontal MNE activities in the United Kingdom, f_{UK}^h and F_{UK}^h , by an amount such that, in either case, $f_{UK}^h + \rho F_{UK}^h$ increases by 20 percent. Figure 5 compares the dynamics of the share of US MNE affiliates in the United Kingdom after each of those changes, using the calibrated models with and without export platforms.

Following an increase in f_{UK}^h , in Figure 5a, the share of U.S. MNE affiliates in the United Kingdom drops immediately by about 15 percentage points in both models. Around eight years after the cost increase, the share of affiliates is 1.5 percentage points lower in the model without export platforms than in the baseline model (84 vs 82.5 percent). In the baseline case, incentives to use affiliates as export platforms gradually induce additional entry despite higher operating costs. Turning to the effects of an increase in sunk costs in Figure 5b, the higher F_{UK}^h widens the affiliates' band of inaction, reducing both new entry and exit. Even though in the longer run the aggregate share of affiliates declines in both models, the dynamics and magnitudes are very different. Because the export option sustains more entry in the model with export platforms than in the model without them, the share of affiliates increases at impact by ten percent in that model, and then, slowly decreases, reaching a mere two percent decline after 15 periods. In contrast, in the model without export platform, the decline is steady and large, reaching an almost 15 percent reduction in the share of US affiliates in the United Kingdom after 15 periods.

Figure 5: The role of export platforms: share of US MNE affiliates in the United Kingdom.



Notes: Increase in fixed cost f_{UK}^h and sunk cost F_{UK}^h . Results are shown as deviations from the calibrated models with and without export-platforms.

Overall, the results highlight the importance of incorporating both dynamic and spatial dimensions when evaluating MNE expansion strategies. They also point to the need to distinguish between recurring and one-time costs—something static models cannot capture, a point that we explore

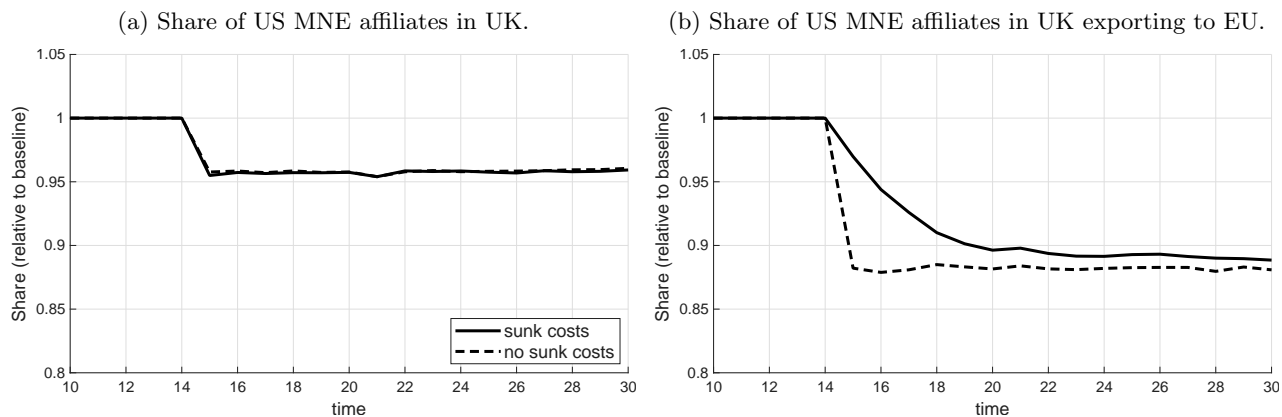
next.

5.2 The role of sunk costs

We now explore the role of sunk costs in shaping the dynamics of MNE expansion. To such end, we perform two exercises. First, we show the effects of including dynamics in the model by comparing the evolution of affiliate activities in a model with and without sunk costs (i.e. a repeated static model). Second, we show the effects of changing sunk and fixed costs on firm dynamics.

We start by comparing the evolution of the extensive margin of US MNE affiliates in the United Kingdom after a 20 percent increase in variable trade costs between the United Kingdom and EU countries, using our calibrated model and a model without any sunk costs.³⁹ Figure 6a shows that the dynamics of the share of US MNE affiliates operating in the United Kingdom are extremely similar between the two models. This is not surprising since barriers to host-market sales are unchanged, and hence, the effects are only indirect. In contrast, the dynamics of exporters are *quantitatively* very different between the two models. Figure 6b shows that an increase in variable trade costs decreases the shares of active US MNE exporting affiliates. However, in the static model, this decrease happens immediately, while our dynamic model delivers a slow fifteen-year decline.

Figure 6: The role of sunk costs.



Notes: Changes in $\tau_{UK,k}^T$ and $\tau_{k,UK}^T$ where k refers to Ireland, Germany, and France.

Next, we evaluate the effects of changing the frictions driving firm dynamics. Figure 7 shows the effects of increasing the sunk and fixed costs of exports from the United Kingdom for all firms located in that country, as well as the sunk entry costs and fixed operating costs in that country,

³⁹Results for British firms are very similar (not shown).

for US MNEs, one, five, ten, and fifteen years after the shock. Naturally, for small changes, dynamic effects are smaller than for large changes in the shock. Long-run effects (i.e. after 15 periods) are larger than short-run effects (i.e. year 1) for shocks to the sunk costs; the time horizon has a minimal role in shocks to the fixed costs.⁴⁰ This comparison between differences in the dynamic responses of MNE activities when fixed or sunk costs change are of course not possible to make in static setups.

5.3 The effects of trade-cost shocks

Our next set of counterfactual exercises compares the effects of different trade-cost shocks on firm dynamics. To this end, we increase, one at a time and permanently, the barriers to export between the United Kingdom and the EU countries in our sample: $\tau_{UK,k}^T$ and $\tau_{k,UK}^T$, $f_{UK,k}^e$, and $F_{UK,k}^e$, for $k = \text{Ireland, Germany, and France}$. We increase each friction by an amount equivalent to a 20-percent increase in the total per-period cost of exporting, $\left(f_{j,k}^e + \rho F_{j,k}^e\right) \tau_{j,k}^{\eta-1}$. This index of fixed-plus-annuitized-sunk export costs, expressed relative to variable profits, is used to scale the different shocks in a comparable way.

Increasing trade barriers has three main effects. First, when exporting from the United Kingdom to the EU becomes more costly, the incentive to open affiliates (and domestic firms) in the United Kingdom decreases 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, decreasing the incentive to open affiliates in those countries. Second, increases in trade costs affect the firms' export band of inaction, which affects firms' export entry and exit decisions—an effect only present in dynamic models. Finally, the increase in trade frictions changes aggregate prices and wages through general equilibrium effects.

Results in Figure 8 combine these three effects. Since trade frictions affect activities in the host market only indirectly through the compound option, Figure 8a shows that effects on the share of UK-based affiliates of US MNEs are small: Increasing sunk or fixed export costs causes a permanent decrease in this share of around two percent, while the increase of variable trade frictions drives a decrease of almost five percent.⁴¹

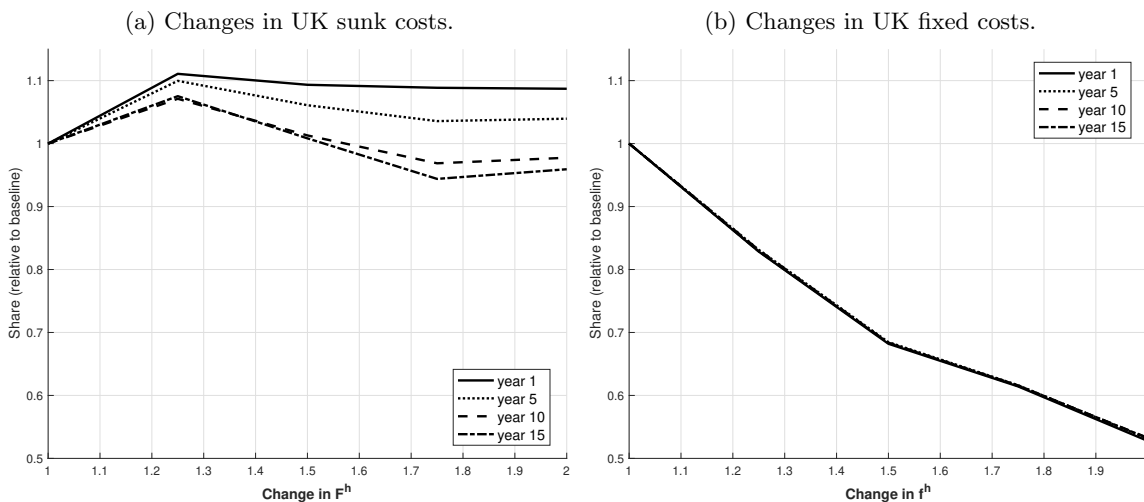
Figures 8b and 8c illustrate that different trade frictions generate distinct quantitative effects on

⁴⁰Online Appendix Figure O.7 shows the changes in the share of US MNE affiliates operating in each host market in our sample.

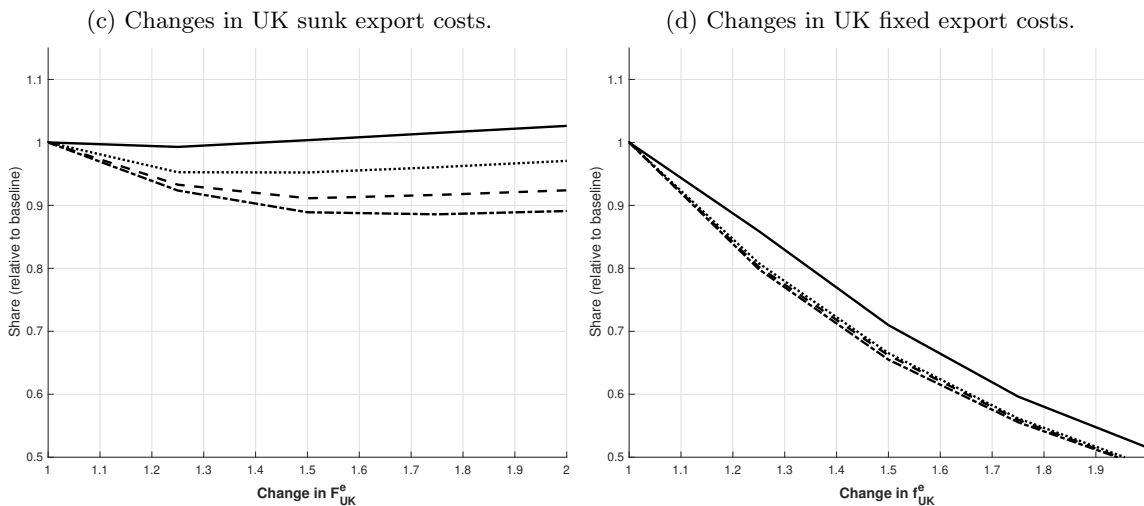
⁴¹Online Appendix Figure O.8 shows the effects of trade costs changes on sales.

Figure 7: The dynamic effects of sunk and fixed costs.

Share of US MNE affiliates in UK.



Share of US MNE affiliates in UK exporting to EU.




Notes: Changes in $f_{UK,k}^e$ and $F_{k,UK}^e$ where k refers to Ireland, Germany, and France, and changes in f_{UK}^h and F_{UK}^h . Years 1, 5, 10, and 15, after the shock.


export participation, even when the shock results in the same increase in the per-period cost of exporting. In general, higher export costs significantly reduce the share of British and U.S. MNE firms that export. Increases in fixed export costs lead to the largest decline, since these costs are intimately tied to a firm’s exit decision. By contrast, increases in sunk export costs—which decreases both entry and exit—have a less pronounced effect and, in the case of British firms exporting to the EU, actually slightly raise participation, as the reduction in exit dominates.

Figure 8: UK trade shock: Firms in the United Kingdom and France.

Firms in the United Kingdom.



Firms in France.



Notes: Increase in F_{UK}^e , f_{UK}^e , and $\tau_{UK,k}^T$ and $\tau_{k,UK}^T$, where k refers to Ireland, Germany, and France.

Results for firms located in France are quantitatively different from the results for the United Kingdom.⁴² The lower panel of Figure 8 shows that higher trade costs between France and the United Kingdom reduce US MNEs' incentives to locate in France and to export to the UK from there. While shocks to export fixed or sunk costs deliver almost identical dynamics, shocks to variable trade costs are smoother and produce a smaller decrease in the share of US MNE affiliates in France exporting to UK.

Summing up, once again, these results show that shocks to variable, fixed, or sunk trade costs have different *quantitative* effects on MNE activity. For this reason, it is particularly important to include dynamics in evaluating the response of MNE expansion to trade shocks.

⁴²Online Appendix Figures O.9 and O.10 show results for Ireland and Germany, respectively.

5.4 General equilibrium effects and welfare

To conclude the portion of our quantitative analysis that is concerned with frictions to and from the UK market, we quantify the model’s general equilibrium effects on firms’ reallocation over time and space, and present welfare calculations. Table 6 reports the dynamics of MNE sales and export participation after an increase in variable trade costs (equivalent to a 20-percent increase in total per-period cost of exporting) between the United Kingdom and EU countries, both under fixed and flexible price indices and wages.

The price and wage adjustments act as buffers to the decrease in sales of UK-based affiliates of US MNEs, both to the United Kingdom and export markets. The effect on UK sales is strong enough to reverse the partial-equilibrium pattern both in the short and long run (-0.4 vs 4.7 percent). In the long run, decreases in the share of US MNE affiliates operating in the United Kingdom and the share of them exporting to EU are lower when prices and wages are allowed to adjust.

Having a general equilibrium model allows us to compute the welfare consequences of the increase in trade costs for the United Kingdom, calculated as the change in real income per capita. Both in the short and long run, the United Kingdom experiences a decline in welfare of 0.9 percent.

Finally, general equilibrium responses create substitution effects between aggregate export and affiliate sales. As the cost of intra-Europe exports increases, the lower panel of Table 6 shows that US parent exports increase, both on the extensive and intensive margins—that is, US exports increase when the cost of multinational activity (here, the cost of export platforms) increases. General equilibrium effects change the relative costs of production across countries, making US domestic production and exports a cheaper way to serve the British and EU markets. The effects differ in the short and long run.

5.5 The effects of the 2025 trade war

Having some of the most relevant affected countries included in our sample, we can evaluate the effects of the tariffs announced by the Trump Administration between April and June 2025 on the activities of US MNEs. We compute the applied US tariffs to each country in our sample, other than China, Mexico, and Canada, before the trade war, using WTO MFN tariffs and import value data in 2022. For China, we use a tariff rate of 20.8 percent, computed by Chad Bown for December 2024.⁴³ For Mexico and Canada, we set tariffs to zero given their USMCA membership.

⁴³Trump’s trade war timeline 2.0: An up-to-date guide (Peterson Institute for International Economics). <https://www.piie.com/blogs/realtime-economics/2025/trumps-trade-war-timeline-20-date-guide>.

Table 6: The role of general equilibrium effects and welfare.

% change relative to baseline:	One period after shock		Fifteen periods after shock	
	General eq	Partial eq	General eq	Partial eq
Sales to UK of US MNE affiliates in UK	4.76	-0.43	4.77	-0.41
Share of US MNE affiliates in UK	-4.27	-7.25	-4.07	-7.04
Sales to EU of US MNE affiliates in UK	-14.8	-17.3	-15.1	-17.6
Share of US MNE affiliates in UK exporting to EU	-5.61	-5.02	-11.2	-11.6
Real income in UK	-0.91	–	-0.90	–
US parent sales to UK	3.42	0.00	3.52	0.00
Share of US parents exporting to UK	1.35	0.00	0.92	0.00
US parent sales to EU	1.10	0.00	1.21	0.00
Share of US parents exporting to EU	1.00	0.00	0.88	0.00

Notes: The shock refers to increasing the variable trade costs $\tau_{UK,k}^T$ and $\tau_{k,UK}^T$, where k refers to Ireland, Germany, and France.

We calculate new tariffs (including threats) using the April 2nd announcement as well as news sources (up to August 1st, 2025). We then compute the implied change in iceberg-trade costs as $(1 + t'_{j,US}/1 + t_{j,US})$. Table 7 shows the results.

Not surprisingly, total exports from foreign affiliates of US MNEs to the United States decrease by almost 23 percent, with a smaller effect in the long run. But the two host countries that experience the lowest increase in tariffs, Great Britain and Singapore, increase exports towards the United States as well as the share of US affiliates located in those countries. After 15 periods, the share of US affiliates located in Great Britain would increase further (0.76 vs 1.00 percent relative to baseline), while in Singapore, this change would decrease (0.72 vs 0.44 percent). While there is “re-shoring” (i.e. US parents located in the United States increase their US sales by around seven percent), the effects on real income are negative and large, of around 2.5 percent relative to baseline, both in the short and long run.

6 Conclusions

This paper studies the expansion patterns of multinational enterprises (MNEs) over time and across space. Using a long panel of US MNEs, we document a set of facts that guide the construction of a dynamic model of MNEs, which is tractable yet rich enough to capture the spatial complexity of MNE activities observed in the data. The model includes heterogeneity in firm productivity, persistent aggregate shocks, a realistic structure of MNE costs, and general equilibrium responses.

Table 7: The effects of the 2025 trade war.

Affiliate Ctry	Tariff rate (%)		% Change in MNE sales to US		% Change in MNE affiliate share		% Change in real income	
	initial	new	one period	15 periods	one period	15 periods	one period	15 periods
Brazil	2.2	50	-65.3	-64.6	-2.86	-3.39	-0.66	-0.76
Canada	0.0	35	-46.2	-44.9	-7.54	-7.67	-1.39	-1.48
China	20.8	30	-34.6	-32.9	-3.38	-3.26	-0.82	-0.85
France	1.4	15	-11.3	-10.0	-0.25	-0.18	-0.09	-0.08
Great Britain	1.4	10	4.92	6.27	0.76	1.00	0.24	0.30
Germany	1.4	15	-11.2	-9.82	-0.68	-1.46	-0.18	-0.18
Ireland	0.4	15	-11.6	-10.4	-0.25	-0.07	-0.24	-0.21
Japan	1.5	15	-14.7	-12.8	-0.57	-0.24	-0.38	-0.38
Mexico	0.0	30	-41.4	-40.4	-3.72	-3.76	-1.13	-1.24
Singapore	1.9	10	3.77	5.01	0.72	0.44	0.49	0.59
Average	3.1	23	-22.8 [†]	-21.5 [†]	-1.78	-1.86	-0.42	-0.43
United States	–	–	7.65*	6.90*	–	–	-2.57	-2.73

Notes: Tariff rates in 2025 include effective changes as well as announced threats. MNE sales to US refers to imports into the United States from country j by affiliates of US MNEs. Change in MNE affiliate share refers to changes in the share of US MNE affiliates located in country j . Results are shown one period and 15 periods after the increase in tariffs. (*) refers to sales to the US by US MNEs located in the United States (i.e. US parents). (†): Change in total imports from affiliates of US MNEs located in the countries in our sample.

We introduce a compound-option formulation for the dynamic problem of the MNE, which is novel to the literature, and allows us to simultaneously decouple the firm’s locations of production and sales, as well as to incorporate entry and exit decisions both into and from host and export markets. Our quantitative exercises reveal that the compound option structure is important for understanding the reallocation of MNE activity in time and space after a globalization shock. These exercises also reveal that the nature of the frictions to MNE activities (variable, fixed, or sunk), as well as general equilibrium responses, 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 decisions happen over time and they are likely to be affected by uncertainty in demand and other market characteristics. This structure, however, can be useful in other contexts. Problems related to global sourcing decisions, which are likely to occur over time and under uncertainty, are good applications for future research.

References

- Alessandria, G., C. Arkolakis, and K. J. Ruhl (2021). Firm dynamics and trade. *Annual Review of Economics* 13(1), 253–280.
- 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 (2021). Trade adjustment dynamics and the welfare gains from trade. *Journal of International Economics* 131, 1034–58.
- Alfaro-Ureña, A., J. Castro-Vincenzi, S. Fanelli, and E. Morales (2024). Firm export dynamics in interdependent markets. *NBER Working Paper* (31610).
- Alvarez, F. and R. E. Lucas (2007). General equilibrium analysis of the Eaton–Kortum model of international trade. *Journal of Monetary Economics* 54(6), 1726–1768.
- Alviarez, V. (2019). Multinational production and comparative advantage. *Journal of International Economics* 119, 1–54.
- Antràs, P. and S. R. Yeaple (2014). Multinational firms and the structure of international trade. *Handbook of International Economics* 4, 55–130.
- Antràs, P., E. Fadeev, T. C. Fort, and F. Tintelnot (2024). Exporting, global sourcing, and multinational activity: Theory and evidence from the United States. *The Review of Economics and Statistics*, 1–48.
- Arkolakis, C. (2016). A unified theory of firm selection and growth. *Quarterly Journal of Economics* 131(1), 89–155.
- Arkolakis, C., A. Costinot, and A. Rodríguez-clare (2012). New trade models, same old gains? *American Economic Review* 102(1), 94–130.
- Arkolakis, C., F. Eckert, and R. Shi (2025). Combinatorial discrete choice: Theory and application to multinational production. *NBER WP 31877*.
- 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.
- Armenter, R. and M. Koren (2015). Economies of scale and the size of exporters. *Journal of the European Economic Association* 13(3), 482–511.

- 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.
- Bernard, A. B. and J. B. Jensen (1999). Exceptional exporter performance: cause, effect, or both? *Journal of International Economics* 47, 1–25.
- Brainard, S. L. (1997). An empirical assessment of the proximity-concentration trade-off between multinational sales and trade. *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.
- Castro-Vincenzi, J. (2024). Climate hazards and resilience in the global car industry. *Manuscript, Harvard University*.
- Castro-Vincenzi, J., E. Menaguale, E. Morales, and A. Sabal (2025). Market entry and plant location in multi-product firms. *Manuscript, Harvard University*.
- 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. (2024). Talent, geography, and offshore r&d. *Review of Economic Studies* 92(2), 1022–1060.
- Feenstra, R. C. and J. Romalis (2014, 02). International prices and endogenous quality. *The Quarterly Journal of Economics* 129(2), 477–527.
- 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 (2024). How exporters grow. *Review of Economic Studies* 91(4), 2276–2306.
- Foster, L., J. C. Haltiwanger, and C. Syverson (2016). The slow growth of new plants: Learning about demand? *Economica* 83(329), 91–129.
- 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.
- Hsieh, C.-T. and P. J. Klenow (2014). The life cycle of plants in India and Mexico. *The Quarterly Journal of Economics* 129(3).
- 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.

- Kondo, I. O., L. T. Lewis, and A. Stella (2023). Heavy tailed but not Zipf: Firm and establishment size in the United States. *Journal of Applied Econometrics* 38(5), 767–785.
- Kovak, B., L. Oldenski, and N. Sly (2021). The labor market effects of offshoring by U.S. multinational firms. *Review of Economics and Statistics* 103(2).
- Kueng, L., M.-J. Yang, and B. Hong (2014). Sources of firm life-cycle dynamics: Differentiating size vs. age effects. NBER WP 20621.
- Leahy, J. V. (1993). Investment in competitive equilibrium: the optimality of myopic behavior. *Quarterly Journal of Economics* 108(4), 1105–1133.
- Luttmer, E. G. J. (2007). Selection, growth, and the size distribution of firms. *Quarterly Journal of Economics* 122(3), 1103–1144.
- McGrattan, E. R. and E. C. Prescott (2010). Technology capital and the us current account. *American Economic Review* 100(4), 1493–1522.
- McGrattan, E. R. and A. Waddle (2020, January). The impact of Brexit on foreign investment and production. *American Economic Journal: Macroeconomics* 12(1), 76–103.
- 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.
- Oberfield, E., E. Rossi-Hansberg, P.-D. Sarte, and N. Trachter (2024). Plants in space. *Journal of Political Economy* 132(3).
- 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.

- Saborowski, C. and F. Misch (2019). Firm size, life cycle dynamics and growth constraints: Evidence from Mexico. *IMF Working Papers 087*.
- Tintelnot, F. (2017). Global production with export platforms. *Quarterly Journal of Economics 132*(1), 157–209.
- Wang, X. (2021). Headquarters gravity: How multinationals shape international trade. *Journal of International Economics 131*.
- Yeaple, S. R. (2013). The multinational firm. *Annual Review of Economics 5*, 193–217.

Appendix

A Model Solution for the Myopic Firm: Derivations

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

Export entry and exit. To solve for the affiliate export problem, we write the Bellman equations (10) and (11) in the continuation region and take the limit as $\Delta t \rightarrow 0$,

$$\rho V_{ijk}^o(z; Y_i, W_j, \Psi_k) = E[dV_{ijk}^o(z; Y_i', W_j, \Psi_k) | Y_i, W_j, \Psi_k] \quad (\text{A.1})$$

$$\rho V_{ijk}^e(z; Y_i, W_j, \Psi_k) = \tilde{\pi}_{ijk}(z; W_j, \Psi_k) Y_i - W_j f_{ijk} + E[dV_{ijk}^e(z; Y_i', W_j, \Psi_k) | Y_i, W_j, \Psi_k]. \quad (\text{A.2})$$

where $\tilde{\pi}_{ijk}(z; W_j, \Psi_k)$ is defined in (18), $\Psi_k \equiv P_k^{\eta-1} X_k$, and Y_i is the aggregate productivity shock. Applying Ito's Lemma yields the no-arbitrage conditions given by

$$\rho V_{ijk}^o = \mu_i Y_i \frac{\partial V_{ijk}^o}{\partial Y_i} + \frac{1}{2} \sigma_i^2 Y_i^2 \frac{\partial^2 V_{ijk}^o}{\partial Y_i^2} \quad (\text{A.3})$$

$$\rho V_{ijk}^e = \tilde{\pi}_{ijk}(z; W_j, \Psi_k) Y_i - W_j f_{ijk} + \mu_i Y_i \frac{\partial V_{ijk}^e}{\partial Y_i} + \frac{1}{2} \sigma_i^2 Y_i^2 \frac{\partial^2 V_{ijk}^e}{\partial Y_i^2}, \quad (\text{A.4})$$

where we omit the dependence of the value functions on $\{z; W_j, \Psi_k\}$ to ease the notation.

We guess a solution for the value functions and apply the method of undetermined coefficients. Our guess is given by

$$\begin{aligned} V_{ijk}^o(z; Y_i, W_j, \Psi_k) &= B_{ijk}^o(z; W_j, \Psi_k) Y_i^{\beta_i} \\ V_{ijk}^e(z; Y_i, W_j, \Psi_k) &= \kappa_0^e + \kappa_1^e Y_i + A_{ijk}^e(z; W_j, \Psi_k) Y_i^{\alpha_i}. \end{aligned}$$

Here, α_i and β_i are parameters, and $A_{ijk}^e(z; W_j, \Psi_k)$ and $B_{ijk}^o(z; W_j, \Psi_k)$ are firm-specific terms to be determined. Substituting our guess for V_{ijk}^o into (A.3) yields

$$\rho B_{ijk}^o(z; W_j, \Psi_k) Y_i^{\beta_i} = \mu_i Y_i \beta_i B_{ijk}^o(z; W_j, \Psi_k) Y_i^{\beta_i-1} + \frac{1}{2} \sigma_i^2 Y_i^2 \beta_i (\beta_i - 1) B_{ijk}^o(z; W_j, \Psi_k) Y_i^{\beta_i-2}$$

which simplifies to

$$\frac{1}{2}\sigma_i^2\beta_i^2 + \left(\mu_i - \frac{1}{2}\sigma_i^2\right)\beta_i - \rho = 0. \quad (\text{A.5})$$

The positive solution to this quadratic equation is $\beta_i > 1$ as a function of μ_i , σ_i and ρ .

Substituting the guess for V_{ijk}^e into (A.4) yields

$$\begin{aligned} \rho\kappa_0^e + \rho\kappa_1^e Y_i + \rho A_{ijk}^e(z; W_j, \Psi_k) Y_i^{\alpha_i} = & \quad (\text{A.6}) \\ \tilde{\pi}_{ijk}(z; W_j, \Psi_k) Y_i - W_j f_{ijk} + \mu_i Y_i \kappa_1^e + \mu_i Y_i \alpha_i A_{ijk}^e(z; W_j, \Psi_k) Y_i^{\alpha_i - 1} \\ + \frac{1}{2}\sigma_i^2 Y_i^2 \alpha_i (\alpha_i - 1) A_{ijk}^e(z; W_j, \Psi_k) Y_i^{\alpha_i - 2}. \end{aligned}$$

Collecting the terms in $A_{ijk}^e(z; W_j, \Psi_k) Y_i^{\alpha_i}$ delivers again the fundamental quadratic equation in (A.5), with the negative solution α_i a function of μ_i , σ_i and ρ . Collecting the constant term and the term linear in Y_i delivers, respectively,

$$\kappa_0^e = -\frac{W_j f_{ijk}}{\rho} \quad \text{and} \quad \kappa_1^e = \frac{\tilde{\pi}_{ijk}(z; W_j, \Psi_k)}{\rho - \mu_i}. \quad (\text{A.7})$$

In this way, we obtain the expressions in (14) and (15), where the expected present discounted value of profits $\{\int_{\tau}^{\infty} e^{-\rho dt} \pi_{ijk}(z; Y_i(t), W_j, \Psi_k) dt\}$ is $\kappa_0^e + \kappa_1^e Y_i$.

To fully solve the affiliate export problem, we are left to solve, for each triplet $\{i, j, k\}$, with $j \neq k$, and productivity z , for the thresholds of realizations of the shock that induce the affiliate to start and stop exporting, denoted by $\bar{Y}_{ijk}^e(z; W_j, \Psi_k)$ and $\underline{Y}_{ijk}^e(z; W_j, \Psi_k)$, and for $B_{ijk}^o(z; W_j, \Psi_k) > 0$ and $A_{ijk}^e(z; W_j, \Psi_k) > 0$. We use the system of value-matching conditions,

$$V_{ijk}^o(z; \bar{Y}_{ijk}, W_j, \Psi_k) = V_{ijk}^e(z; \bar{Y}_{ijk}, W_j, \Psi_k) - W_j F_{ijk} \quad V_{ijk}^o(z; \underline{Y}_{ijk}, W_j, \Psi_k) = V_{ijk}^e(z; \underline{Y}_{ijk}, W_j, \Psi_k) \quad (\text{A.8})$$

and smooth-pasting conditions,

$$\frac{\partial V_{ijk}^o(z; \bar{Y}_{ijk}, W_j, \Psi_k)}{\partial Y_i} = \frac{\partial V_{ijk}^e(z; \bar{Y}_{ijk}, W_j, \Psi_k)}{\partial Y_i} \quad \frac{\partial V_{ijk}^o(z; \underline{Y}_{ijk}, W_j, \Psi_k)}{\partial Y_i} = \frac{\partial V_{ijk}^e(z; \underline{Y}_{ijk}, W_j, \Psi_k)}{\partial Y_i} \quad (\text{A.9})$$

Host-country entry and exit. To determine the value of a firm from i with affiliates in j , we still need to solve for the value of horizontal sales. Writing the Bellman equations (12) and (13) in the continuation region yields

$$\rho V_{ij}^o(z; Y_i, W_j, \Psi) = E[dV_{ij}^o(z; Y_i', W_j, \Psi) | Y_i, W_j, \Psi] \quad (\text{A.10})$$

$$\rho V_{ij}^h(z; Y_i, W_j, \Psi) = \tilde{\pi}_{ijj}^h(z; W_j, \Psi_j) Y_i - W_j f_{ijj} + E[dV_{ij}^h(z; Y_i', W_j, \Psi) | Y_i, W_j, \Psi]. \quad (\text{A.11})$$

Taking the limit as $\Delta t \rightarrow 0$, and applying Ito's Lemma, yield

$$\rho V_{ij}^o = \mu_i Y_i \frac{\partial V_{ij}^o}{\partial Y_i} + \frac{1}{2} \sigma_i^2 Y_i^2 \frac{\partial^2 V_{ij}^o}{\partial Y_i^2} \quad (\text{A.12})$$

$$\rho V_{ij}^h = \tilde{\pi}_{ijj} Y_i - W_j f_{ijj} + \mu_i Y_i \frac{\partial V_{ij}^h}{\partial Y_i} + \frac{1}{2} \sigma_i^2 Y_i^2 \frac{\partial^2 V_{ij}^h}{\partial Y_i^2} \quad (\text{A.13})$$

where we omit the dependence of the value functions on $\{z; W_j, \Psi\}$ to ease the notation.

Using the same guess-and-verify procedure as for the export problem, the option value of opening an affiliate in country j for a firm from country i and the value of horizontal sales take the form

$$V_{ij}^o(z; Y_i, W_j, \Psi) = B_{ij}^o(z; W_j, \Psi) Y_i^{\beta_i} \quad (\text{A.14})$$

$$V_{ij}^h(z; Y_i, W_j, \Psi) = \kappa_0^h + \kappa_1^h Y_i + A_{ij}^h(z; W_j, \Psi) Y_i^{\alpha_i}, \quad (\text{A.15})$$

where $A_{ij}^h(z; W_j, \Psi)$ and $B_{ij}^o(z; W_j, \Psi)$ are firm-specific terms to be determined, and $\beta_i > 1$ and $\alpha_i < 0$ are the solutions of the fundamental quadratic equation (A.5).

Applying the method of undetermined coefficients yields

$$\kappa_0^h = -\frac{W_j f_{ijj}}{\rho} \quad \text{and} \quad \kappa_1^h = \frac{\tilde{\pi}_{ijj}(z)}{\rho - \mu_i}.$$

The entry and exit thresholds, $\bar{Y}_{ij}(z; W_j, \Psi)$ and $\underline{Y}_{ij}(z; W_j, \Psi)$, and the terms $A_{ij}^h(z; W_j, \Psi)$ and $B_{ij}^o(z; W_j, \Psi)$, can be solved from the system of value-matching and smooth-pasting conditions,

$$V_{ij}^a(z; \bar{Y}_{ij}, W_j, \Psi) - W_j F_{ijj} = V_{ij}^o(z; \bar{Y}_{ij}, W_j, \Psi) \quad V_{ij}^h(z; \underline{Y}_{ij}, W_j, \Psi) = V_{ij}^o(z; \underline{Y}_{ij}, W_j, \Psi), \quad (\text{A.16})$$

and

$$\frac{\partial V_{ij}^a(z; \bar{Y}_{ij}, W_j, \Psi)}{\partial Y_i} = \frac{\partial V_{ij}^o(z; \bar{Y}_{ij}, W_j, \Psi)}{\partial Y_i} \quad \frac{\partial V_{ij}^h(z; \underline{Y}_{ij}, W_j, \Psi)}{\partial Y_i} = \frac{\partial V_{ij}^o(z; \underline{Y}_{ij}, W_j, \Psi)}{\partial Y_i}. \quad (\text{A.17})$$

B Algorithm

Our algorithm proceeds as follows. We draw 500 firms with productivity z from a Pareto distribution with $\vartheta = 4.25$ and scale 1, in each country. We draw 100 sequences of origin-specific aggregate Brownian shocks, for $T = 30$, from a distribution with drift, variance, and covariance specified in Tables O.16 and O.17 in the Online Appendix.

1. For each sequence of aggregate shocks and a given guess for the set of aggregate variables $\{W_j(t), P_j(t), X_j(t)\}_{t=1}^{30}$, we solve for the firms' entry and exit problems as follows. In terms of notation, we remove dependence on $\{W_j(t), \Psi_k(t)\}$ from the firm's variables, and add instead their time dependence (t) .

(a) For each triplet $\{i, j, k\}$, and each z , we solve for $\{\bar{Y}_{ijk}^e(z, t), \underline{Y}_{ijk}^e(z, t), A_{ijk}^e(z, t), B_{ijk}^o(z, t)\}$ using the procedure described in Online Appendix O.8.

(b) We use the solutions to the export and affiliate entry and exit problems to compute the sets of active firms $\Omega_{ijk}(t)$, characterized in Proposition 3 in Online Appendix O.7. We set $\Omega_{iii}(0) = [1, \infty)$ for $i = US$, $\Omega_{iii}(0) = \emptyset$ for $i \neq US$, and $\Omega_{ijk}(0) = \emptyset, \forall j \neq i$. Using the solution for $\Omega_{ijk}(t)$ and the properties of the Pareto distribution, we compute aggregate (firm) productivity as

$$z_{ijk}(t)^{\eta-1} = \int_{\tilde{z}_{ijk}(t)} z^{\eta-1} dG(z) = \frac{\vartheta}{\vartheta + 1 - \eta} \tilde{z}_{ijk}(t)^{\eta-\vartheta-1}. \quad (\text{A.18})$$

Similarly, we compute the mass of firms, $M_{ijk}(t)$, $M_{ijk}^E(t)$, and $M_{ij}^E(t)$, using the properties of the Pareto distribution together with the definitions of the sets $\Omega_{ijk}(t)$.

(c) Given the export values implied by the affiliate export decisions, we solve the affiliate entry and exit problems using an analogous procedure as in (a), and compute aggregate productivity and the mass of firms in each host country as in (b).

2. Given the implied sequence of $\{M_{ijk}(t)\}_{t=1}^{30}$, we update the equilibrium sequence $\{W_j(t), P_j(t), X_j(t)\}_{t=1}^{30}$ using the iterative algorithm proposed by Alvarez and Lucas (2007). We normalize the wage in the United States in each period to one, $W_{US}(t) = 1$.

3. With the new equilibrium sequence $\{W_j'(t), P_j'(t), X_j'(t)\}_{t=1}^{30}$ we update the guess for the values of aggregate variables in Step 1 using a weighted average of the old and new values. We iterate till convergence.