Multinational Firms and International Business Cycle Transmission*

Javier Cravino
University of Michigan
and NBER

Andrei A. Levchenko
University of Michigan
NBER and CEPR

August 19, 2016

Abstract

We investigate how multinational firms contribute to the transmission of shocks across countries using a large multi-country firm-level dataset that contains cross-border ownership information. We use these data to document two novel empirical patterns. First, foreign affiliate and headquarter sales exhibit strong positive comovement: a 10% growth in the sales of the headquarter is associated with a 2% growth in the sales of the affiliate. Second, shocks to the source country account for a significant fraction of the variation in sales growth at the source-destination level. We propose a parsimonious quantitative model to interpret these findings and to evaluate the role of multinational firms for international business cycle transmission. For the typical country, the impact of foreign shocks transmitted by all foreign multinationals combined is non-negligible, accounting for about 10% of aggregate productivity shocks. On the other hand, since bilateral multinational production shares are small, interdependence between most individual country pairs is minimal. Our results do reveal substantial heterogeneity in the strength of this mechanism, with the most integrated countries significantly more affected by foreign shocks.

Keywords: international business cycle comovement, multinational firms

JEL Codes: F23, F44

*We are grateful to the editor (Pol Antràs), four anonymous referees, Ariel Burstein, Aaron Flaaen, Jim Hines, Sebastian Krautheim, Sebastian Sotelo, Linda Tesar, Carolina Villegas-Sanchez, Andrei Zlate, and seminar participants at various institutions for helpful suggestions and to Laurien Gilbert, Nitya Pandalai-Nayar, and Rishi Sharma for excellent research assistance. This research is supported by the Michigan Institute for Teaching and Research in Economics (MITRE). Email: jcravino@umich.edu, alev@umich.edu.

1 Introduction

Multinational companies are a first-order feature of the world economy, accounting for about one-third of gross output in many developed countries (see, e.g., Alviarez, 2013). Since multinationals encompass production facilities that are spread across different parts of the globe, a natural conjecture is that their rapid growth in recent decades has had an impact on how economic shocks are transmitted across countries. However, the relationship between multinational firms and transmission of shocks is not yet well-understood. At the micro level, there is limited empirical evidence on how the activities of the different parts of a multinational are interrelated at business cycle frequencies. At the macro level, it is yet to be established whether multinational production (MP) matters quantitatively for aggregate comovement.

This paper uses novel firm-level data and a quantitative multi-country model to examine the role of multinational firms in aggregate business cycle transmission. Our data come from ORBIS, a firm-level database that covers several million domestic and multinational firms operating in 34 countries over the period 2004-2012. The key feature of the dataset is that it contains information on domestic and foreign ownership. Hence, for the first time in this context, the operations of parents and affiliates are observed in the same dataset as well as through time and in a broad cross-section of countries. This information allows us to study micro-level cross-country comovement between the different parties of the multinational corporations. At the same time, the data cover the bulk of economic activity in our sample of countries, making it possible to aggregate the firm-level results and derive their implications for business cycle comovement.

Our analysis goes from micro patterns to macro implications in three stages. First, at the firm level, we document strong comovement between multinational affiliates and their parents: a 10% growth in the sales of the parent is associated with a 2% growth in the sales of the affiliate. This correlation is computed after controlling for sectoral and aggregate trends using source-sector-destination-sector-year fixed effects, so that it captures the role of linkages within the multinational firm. The correlation is pervasive across firms in different sectors, including services, which suggests that it is not driven solely by vertical production linkages. The strong correlation between the parent's and affiliate's growth is also present when we use value added or employment to measure firms' growth, and is highly significant and robust to different samples, time periods, fixed effects, and aggregation methods.

The firm-level estimates show that units of the same firm comove together at the business cycle frequency. However, precisely because they are obtained controlling for very detailed aggregate trends, they may not capture transmission of shocks that are common across parent firms in the source country. With this in mind, in our second step we aggregate multinational sales to the source-destination level (e.g., combined sales of all US multinational affiliates operating in the UK), and estimate whether the variation in source-destination growth rates is driven by source-specific or destination-specific factors. Source-specific factors account for about 10% of the variation in bilateral growth rates, compared to 18% accounted for by destination-specific factors. We interpret this result as evidence that shocks to the source country are important for the variation in total sales.

Our empirical results thus demonstrate strong interdependence between source countries and their foreign affiliates. This interdependence is detectable both at the firm and the source-destination level. The third step of our analysis assesses the quantitative importance of this phenomenon for aggregate business cycle transmission using a multi-country model that can be taken to the data. In the model, each country produces a final good by aggregating the output of intermediate producers. These intermediate producers may be local firms or foreign multinational affiliates. We introduce comovement between multinational firms and their foreign affiliates by assuming that the productivity of the affiliate is affected by the productivity of the parent.[1] In particular, the productivity of foreign affiliates is a combination of a source-specific and destination-specific component. The relative importance of the source vs. the destination component is governed by a crucial parameter that we discipline with the data.

In the model, the extent to which multinationals contribute to the transmission of shocks across countries is driven by: (i) what share of the firm's technology shock originates in the source vs. the destination country; (ii) the distribution of bilateral multinational shares in the economy; and (iii) general equilibrium effects. We use the model's structural equations to interpret our empirical results, and to calibrate the extent to which shocks in the source country are transmitted by multinationals. We estimate that between 20 and 40 percent of the foreign affiliates' shocks originate in the source country. The multinational production shares are taken directly from the data. Finally, the magnitude of the general equilibrium effects depends on a composite parameter that combines the elasticity of substitution across intermediates and the Frisch labor supply elasticity. We

[1] This is a common approach in the literature on multinational production, see, among many others, Helpman (1984); Markusen (1984); Helpman et al. (2004) and more recently McGrattan and Prescott (2009, 2010); Burstein and Monge-Naranjo (2009); Keller and Yeaple (2013); Ramondo and Rodríguez-Clare (2013); Ramondo (2014); Alviarez (2013). See Antràs and Yeaple (2014) for a recent overview of modeling approaches. Bilir and Morales (2016) show that US multinational parents' R&D expenditures affect the productivity of their foreign affiliates.

benchmark these parameters using micro estimates of these elasticities, and check the sensitivity of the results to alternative values.[2]

We use the calibrated model to conduct three quantitative exercises to measure the importance of multinational firms for the transmission of shocks across countries. First, we compute impulse responses to productivity shocks in each source country, and track the propagation of these shocks across countries. A 1% productivity shock in the rest of the world as a whole raises productivity by 0.12% in the average country, and by as much as 0.2-0.35% in the most integrated countries such as Ireland, the Netherlands, and Slovakia. Not surprisingly, the external impact of individual source country shocks is considerably smaller. A shock that increases GDP by 1% in one of the 4 most important source countries – US, Germany, UK, and France – raises output in the rest of the sample by between 0.01 and 0.02%. Shocks to other source countries have a negligible impact, since multinational affiliates from other source countries tend to have small output shares in a typical destination.

Second, we use the model to compute the business cycle correlation between each pair of countries assuming that the primitive productivity shocks are uncorrelated. This is an assessment of how much correlation can be generated purely by propagation of shocks through multinationals under the observed levels of multinational activity. The variation in model-implied correlations is driven entirely by the pattern of multinational output shares. On the one hand, in most country pairs bilateral multinational shares are small, and thus the model generates little business cycle comovement: the mean model-implied correlation is 0.01 in the full sample of country pairs. On the other hand, for country pairs involving either a major source or a major recipient of multinational firms, the model generates between one-tenth and one quarter of the correlation observed in the data. In addition, in the cross-section of country pairs the model-implied correlations have a positive and highly significant relationship to the GDP correlations in the data.

Third, we conduct two counterfactual exercises that evaluate how the cross-country dispersion of growth rates changes as we change the shares of multinational firms in the world economy. In the first counterfactual, we consider a world in which there are no multinational firms operating in foreign destinations. The counterfactual cross-country dispersion in growth rates is 7.5% larger in this scenario than in our benchmark calibration. In the second counterfactual, we simulate a "full integration" equilibrium, in which multinationals from any source country operate with the same intensity in all destina-

[2]Our analysis abstracts from the extensive margin adjustment of MP in response to business cycle shocks. While the business cycle frequency response of MP location to aggregate shocks remains a fruitful area for future research, the extensive margin cannot be disciplined with our data as it does not distinguish between firms entering/exiting production and firms entering/exiting the scope of ORBIS data collection.

tions. Under full integration the counterfactual cross-country dispersion in growth rates is over 30% smaller than in our benchmark calibration.

Our main takeaway from these exercises is that the combined impact of all foreign multinationals is small but significant, accounting for about 10% of the productivity shocks in a typical country and leading to a somewhat more synchronized international business cycle. The impact is highly heterogeneous across countries. The transmission of shocks and positive business cycle correlations induced by multinational presence are clearly detectable for the country pairs involving the most important source and destination countries. On the other hand, aggregate interdependence between most individual country pairs is minimal, since most bilateral shares are small.

We highlight three key advantages of our dataset relative to existing empirical analyses of multinationals and business cycle comovement. First, ORBIS provides information on the activities of both the multinational parents and affiliates at a yearly frequency, which allows us to estimate parent-affiliate sales correlations. Second, it includes the local firms in each country along with the domestic and foreign multinationals. This allows us to compute the importance of multinationals in each economy relative to the domestic firms, and also to better estimate the country components of business cycle shocks. Finally, ORBIS covers a broad cross-section of countries. This permits a decomposition of growth rates into source and destination components, an exercise requiring data from multiple sources and destinations. In addition, we can document the large heterogeneity in the impact of multinationals across country pairs.

While the driving mechanism in our model is that productivity shocks are directly transferred across countries within multinational firms, our model is isomorphic to a setup in which comovement arises from the transmission of demand shocks for the firms' product or from certain types of intermediate input linkages (see Online Appendix C.4). It has not (yet) been established empirically that the transmission of shocks through input trade by multinationals is a quantitatively important phenomenon. Ramondo et al. (2016) show that US multinational affiliates abroad sell mostly in the local market, with the median affiliate having no shipments to the parent. In a non-international context, Atalay et al. (2014) show that most vertical ownership links are not primarily motivated by input trade within the firm. In our own results, the correlation between affiliate and parent sales occurs even among service sector firms, for which input trade is likely to be much less relevant. While our model can accommodate the input linkage interpretation, our empirical results show that intermediate input linkages are unlikely to be the sole determinant of parent-affiliate comovement.

This paper contributes to three strands of the literature. The first is the research agenda on the role of multinational firms in the transmission of international business cycles (see, e.g., Burstein et al., 2008; Contessi, 2010; Menno, 2014; Zlate, 2016).[3] This literature has focused mainly on the role of within-multinational trade and vertical integration for business cycle synchronization, and has predominantly employed 2-country models. In contrast, we develop a parsimonious multi-country quantitative framework that can be directly taken to the firm-level data.[4]

Second, we contribute to the empirical literature on multinational firms and comovement. A number of papers (e.g., Budd et al., 2005; Boehm et al., 2014) explore whether parents and affiliates are correlated, including due to internal capital markets within multinationals (Desai and Foley, 2006; Desai et al., 2009).[5] Buch and Lipponer (2005) and Kleinert et al. (2015) use sectoral and regional data to study whether greater multinational presence is associated with greater comovement. All of these papers feature only one source, or only one destination country, and frequently the information on either the parent or the affiliate is limited. Our work is the first to study aggregate comovement with multi-country data in which parents and affiliates are observed within the same dataset. In addition, these papers by and large do not attempt to go from micro estimates to business cycle comovement between countries. We develop a quantitative framework to interpret the empirical findings and evaluate their implications for aggregate comovement.

Finally, a large theoretical literature studies multinationals and technology transfers (see, among many others, McGrattan and Prescott, 2009; Keller and Yeaple, 2013; Ramondo and Rodríguez-Clare, 2013), and an extensive body of empirical work investigates the effects of FDI on productivity.[6] Also related is the literature on the domestic employment impact of increased multinational activity.[7] Our empirical contribution is to use firm-level data to quantify the extent to which parents and affiliates are affected by common shocks at the business cycle frequency.

[3] Also related is the literature that explores the role of cross-border vertical production linkages in the international business cycle transmission (see, e.g., Kose and Yi, 2001; Arkolakis and Ramanarayanan, 2009; Johnson, 2014), though this line of research is not explicit on whether the production linkages take place within firms.

[4] Our analysis complements several contributions that address the opposite question: how the correlation between domestic and foreign business cycles affects MP location decisions (Ramondo and Rappoport, 2010; Ramondo et al., 2013), or multinationals' risk premia (Fillat and Garetto, 2015; Fillat et al., 2015). Online Appendix C.5 presents an extended model in which firm location decisions are made subject to uncertainty and sunk costs, and relates it to our baseline framework.

[5] Alfaro and Chen (2012) investigate whether the affiliates of multinational firms responded to the recent financial crisis differently than local establishments. Their focus is not, however, on parent and affiliate comovement. A number of recent studies examine how liquidity shocks are transmitted through international banks, see for example Acharya and Schnabl (2010), Cetorelli and Goldberg (2011), and Schnabl (2012).

[6] See for example Javorcik (2004), Guadalupe et al. (2012), and Fons-Rosen et al. (2013).

[7] See Becker et al. (2013); Boehm et al. (2015); Kovak et al. (2015) for some recent examples.

The rest of this paper is organized as follows. Section 2 describes the data and presents the basic summary statistics on multinationals' presence. Section 3 documents bilateral firm-level and source-destination-level comovement between multinational firms. Section 4 develops a structural framework to interpret our empirical results and to study the aggregate implications of multinationals for business cycle comovement and for the transmission of shocks. Section 5 describes the quantitative results from the model and counterfactuals, and Section 6 concludes. Detailed descriptions of data, robustness checks on the empirical results, and additional theoretical results are collected in the Online Appendix.

2 Data and summary statistics

The data come from ORBIS, a large cross-country database maintained by Bureau van Dijk. The ORBIS database includes information on both listed and unlisted firms collected from various country-specific sources, such as national registries and annual reports. Importantly, it contains information on the "global ultimate owner" of each firm in the database. This information enables us to build links between affiliates of the same firm, including cases in which the affiliates and the parent are in different countries. We specify that a parent should own at least 50% of an affiliate to identify an ownership link between the two firms. The time period is 2004-2012. The main variable used in the analysis is the total sales (turnover) of each firm.

ORBIS contains data on more than 100 countries, but coverage is extremely uneven, with most of those countries reporting information on very few firms. In addition, in order to analyze multinationals we must use the "unconsolidated" accounts of each firm, since the "consolidated" accounts may include operating revenue of the foreign affiliates. After extensive checking of the data, we retain a sample of 34 countries with sufficiently good coverage and data quality. In particular, the country sample satisfies the following criteria. First, we keep countries with data on more than 750 firms in the average year (as noted below, most countries in our sample are well above this threshold, the median country has data on 73,000 firms in the average year). Second, we keep countries for which the aggregate revenues in ORBIS are at least 40% of aggregate output as reported in standard sources. Third, we keep countries for which the correlation between the growth rate of aggregate revenues in ORBIS and of GDP as reported in the World Bank's World Development Indicators exceeds 0.50. Online Appendix A describes the data assembly, cleaning steps, and additional data validation exercises in detail.

Appendix Table A1 presents the resulting sample of countries along with some summary statistics and checks on the quality of the data. The sample is dominated by European countries, but includes both developed and developing countries, as well as countries outside of Europe. Column 1 reports the total number of firms in the average year for each country. The mean number of firms is about 140,000, and the median is about 73,000. There is a wide range of coverage even in our restricted sample of countries: the country with the smallest number of firms, Australia, has only 766 in an average year. Column 2 reports the number of foreign multinational affiliates in each country. In the median country there are about 2,250 foreign multinational firms in the average year.

Column 3 presents the correlation between the country's GDP growth rate and the growth rate of aggregate sales of all the continuing firms in ORBIS. The aggregate growth rate implied by ORBIS mimics the GDP growth quite well: the mean correlation between aggregate growth in ORBIS and GDP growth from the national accounts is 0.81, and the median is 0.83. This suggests that business cycle features are well captured in the ORBIS data. Column 4 reports the ratio of the total sales of firms in ORBIS to the gross output as reported in other sources. We use two data sources for this consistency check. For EU countries, the best source of gross output data is EUROSTAT. For countries outside of the EU, we take gross output data from the UN System of National Accounts. In this sample of countries, the ORBIS data captures the bulk of aggregate output as reported by national statistical agencies.

Figure 1 shows the relative importance of foreign multinational affiliates in the countries in our sample for the average year. In the average country, about 7.5% of all firms are affiliates of foreign multinationals, ranging from 0.1% in Japan to 29% in Australia. Multinational affiliates tend to be larger than domestically-owned firms, so they comprise higher shares of total revenue, 29% on average. Once again there is a wide range, from 1.5% for Japan to 84% for Ireland. Indeed, in a number of countries – Netherlands, Slovakia, Singapore, Belgium, Czech Republic, Austria, and the UK – multinational affiliates account for 40% or more of total sales in our data.

3 Empirical results

This section estimates how the growth rates of affiliates are related to the growth rates of parents, both at the firm level and at the source-destination level. Throughout the analysis below, we use growth rates and shares in the form suggested by Davis et al. (1996): for any variable x_j and time periods t and $t-1$, the growth rate is defined as $\gamma_{j,t} \equiv 2\left(\frac{x_{j,t}-x_{j,t-1}}{x_{j,t}+x_{j,t-1}}\right)$. That is, the denominator is the average of the beginning and end period levels, rather than the beginning period level. Davis et al. (1996) recommend using this growth rate because

Figure 1: The importance of foreign multinationals

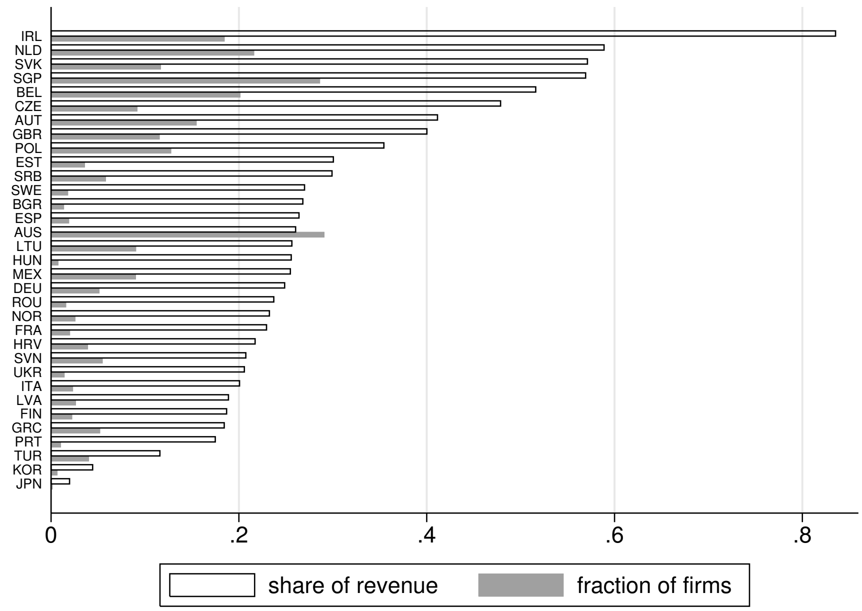

Notes: This figure reports, for each country, the share of foreign multinational affiliates in total revenue (light bars) and the total number of firms (dark bars).

it has a number of attractive properties: it is bounded between -2 and 2, is symmetric around zero, and lends itself to aggregation. If $x_t = \sum_j x_{j,t}$, the aggregate growth of x_t, γ_t, can be written as the weighted sum of the disaggregated growth rates, $\gamma_t = \sum \omega_{j,t}\gamma_{j,t}$, with weights that are defined as $\omega_{j,t} = \frac{x_{j,t}+x_{j,t-1}}{\sum_j (x_{j,t}+x_{j,t-1})}$. All of the firm-level and aggregate growth rates between years $t-1$ and t are computed using only firms present in ORBIS in both $t-1$ and t, and thus capture intensive margin growth rates.[8]

3.1 Firm-level comovement

We begin by documenting comovement at the firm level between parents and affiliates. In particular, we estimate the following specification:

$$\gamma_{in,t}(f) = \phi\gamma_{ii,t}(f) + \bar{a}_{inss',t} + \epsilon_{in,t}(f). \quad (1)$$

[8]Because ORBIS does not cover the universe of firms in each country, it cannot be used to measure entry and exit, since for newly observed firms we cannot distinguish between genuine entry and entry into the ORBIS data collection. Using a Census of French firms in which entry and exit can be measured relatively more accurately, di Giovanni et al. (2014) show that the extensive margin of entry and exit of firms is not important in accounting for aggregate fluctuations.

Table 1: Affiliate-parent comovement

	(1)	(2)	(3)	(4)	(5)	(6)
	All		Manufacturing		Services	
ϕ	0.278***	0.229***	0.402***	0.299***	0.233***	0.213***
	(0.005)	(0.012)	(0.014)	(0.039)	(0.006)	(0.013)
Obs.	182029	182029	19809	19809	105704	105704
N. mult.	18886	18886	2476	2476	12413	12413
R^2	0.047	0.724	0.102	0.789	0.032	0.674
FE	No	Yes	No	Yes	No	Yes

Notes: Standard errors clustered at the parent level in parentheses. ***: significant at the 0.1% level. This table presents the results of estimating equation (1). "FE" refers to source × destination × affiliate sector × parent sector × year fixed effects. Sectors are defined at the 2 digit level of the NACE classification.

Here $\gamma_{in,t}(f)$ is the sales growth rate of the firms in multinational group f from source country i, operating in destination country n, and $\gamma_{ii,t}(f)$ is the growth rate of multinational group f's parent firm in the source country i.[9] The specifications include source × destination × affiliate sector × parent sector × year fixed effects $\bar{a}_{inss',t}$ that control for comovement arising from country-specific sectoral and aggregate trends. We run equation (1) on the sample of firms that are foreign affiliates (so that the growth rate of the parent $\gamma_{ii,t}(f)$ exists), pooling observations across years. Standard errors are clustered at the parent level.

Table 1 presents the results. It reports estimates of a simple bivariate regression with no fixed effects, as well as with the fixed effects. The first panel of the table shows the results for a sample consisting of all firms, while the next two panels focus on a sample of firms in which both the parent and the affiliate are either in the manufacturing or in the service sector. There is a strong positive and highly significant correlation between affiliates and parents across all the specifications. Our benchmark estimate of ϕ using the full sample and controlling for fixed effects is 0.229. The estimated correlation is larger for firms in the manufacturing sector, although the last panel shows there is a strong positive correlation for service sector firms as well.

[9] To compute the growth rate of the multinational group in a (source or destination) country, $\gamma_{in,t}(f)$, we aggregate the sales of all the firms belonging to multinational f that operate in the country in the two consecutive years on which the growth rate is computed. This ensures that changes in the composition of the multinational group (i.e. by the acquisition of a new firm in a particular destination) are not reflected in the growth rate.

Robustness Online Appendix B and Appendix Tables A4-A8 explore the firm-level results in detail and present a large battery of robustness checks. First, Section B.1 follows several strategies to show that the results are not driven solely by vertical input linkages between parents and affiliates. We show that parent-affiliate comovement is present among (i) firms that operate only in the service sector; (ii) dropping wholesale and retail firms; and (iii) using value added or employment instead of sales. In addition, we use information on sectoral cross-border input linkages from the World Input-Output Database to show that parent-affiliate comovement is not differentially stronger in sector pairs that exhibit greater input trade.

Second, Section B.2 evaluates whether the parent-affiliate comovement is driven by tax shifting purposes or by internal capital markets within multinational firms. The results are unchanged when we drop Ireland and the Netherlands, countries most associated with tax shifting behavior. Comovement is not differentially stronger among country pairs with larger corporate tax rate differences. Parent-affiliate comovement is also not differentially stronger in destinations with less developed capital markets, and is equally evident in the years prior to the 2008 financial crisis.

Third, Section B.3 checks whether parent-affiliate comovement is conditioned by country characteristics that affect the ease of technology transfer within the firm. We show that parent-affiliate comovement is not significantly greater in destinations with better intellectual property rights (IPR) protection, though the absence of significant results may be due to limited variation in the strength of IPR protection among the destination countries in our sample. Similarly, comovement is not more pronounced for country pairs with lower geographical distance. Parent-affiliate comovement does appear to be modestly higher for country pairs with larger bilateral MP shares and somewhat higher in high-income destinations.[10,11]

Finally, the Online Appendix B also presents a large battery of additional robustness checks on data construction, aggregation, and subsamples of firms and countries.

[10] The coefficient on parent growth is 0.191 for the bottom quartile of MP shares, and 0.252 for the top quartile. High-income countries have a 0.091 greater coefficient on parent growth than lower-income countries.

[11] All of the Online Appendix Tables control for source-sector-destination-sector-year fixed effects. With source-sector-year and destination-sector-year effects, we do find stronger parent-affiliate comovement in sectors with greater input linkages. However, without the most stringent fixed effects that we use in the baseline, we cannot separate this particular mechanism from other omitted variables affecting comovement at the country-pair-sector level.

3.2 Bilateral comovement

The estimates from the previous section show that units of the same firm comove together at the business cycle frequency. However, precisely because they are obtained after partialling out detailed aggregate trends, they may not capture transmission of shocks that are common across parent firms in the source country. We would like to establish that there is a common component to the combined overall sales of multinationals from a particular country. We thus estimate the contribution of source- and destination-specific shocks to the variation in the bilateral growth rates:

$$\gamma_{in,t} = \quad s_{i,t} + d_{n,t} + a_{in,t}. \tag{2}$$

Equation (2) writes the growth rate $\gamma_{in,t}$ of total sales of firms owned by country i operating in country n (e.g., the growth rate of the total sales of all i =US multinationals operating in n =UK) as a sum of the source effect $s_{i,t}$ common to all firms owned by i worldwide, the destination effect $d_{n,t}$ common to all firms from all countries selling in market n, and an idiosyncratic term $a_{in,t}$. This decomposition of a cross-section of data into different types of shocks draws on a standard approach in macroeconomics (see, e.g., Stockman, 1988, and the literature that followed), but to our knowledge has never been applied to foreign multinational operations to establish the existence of a source country shock.

The empirical model (2) is estimated by regressing observed growth rates $\gamma_{in,t}$ on source and destination fixed effects (when carried out year-by-year), or source-year and destination-year effects (when carried out in a pooled sample of years). The regression for the pooled sample of years also includes non-time-varying source-destination fixed effects. There is a large amount of variation in the size of source-destination pairs. Smaller *in* pairs tend to have fewer firms and thus tend to be more volatile. To account for this fact, we employ a Generalized Least Squares estimation in which the observations are weighted by the inverse of the Herfindahl index of firm-level sales shares in an *in* pair.[12] This approach underscores the usefulness of firm-level data even for the estimation of source- or destination-level outcomes, as firm-level information helps capture the heteroskedasticity in the source-destination data.

[12] Let the variance of the residual of an individual firm's growth rate be $\sigma^2(f)$, and let $\tilde{\omega}_{in,t}(f)$ be the share of firm f in the total sales of firms from source i in destination n. Assuming that $\sigma^2(f)$ does not differ by firm, the variance of the residual of the source-destination level observation is equal to $Var(a_{in,t}) = \sigma^2(f)\sum_{f\in\Omega_{in}}\tilde{\omega}^2_{in,t}(f) \equiv \sigma^2(f)Herf_{in,t}$, where Ω_{in} is the set of firms from i selling in n. The GLS estimator weights the observations by the inverse of the variance of the error term, which in this case is proportional to the Herfindahl index of firm sales shares.

Table 2: Importance of source and destination effects

	(1)	(2)	(3)	(4)	(5)	(6)
	Source			Destination		
	Part. R^2	F-stat.	p-val.	Part. R^2	F-stat	p-val.
2005	0.11	2.65	0.00	0.13	4.53	0.00
2006	0.09	2.22	0.00	0.14	5.21	0.00
2007	0.06	1.49	0.02	0.12	4.60	0.00
2008	0.17	5.01	0.00	0.24	10.03	0.00
2009	0.08	2.08	0.00	0.19	7.72	0.00
2010	0.13	3.60	0.00	0.22	9.16	0.00
2011	0.09	2.44	0.00	0.17	6.89	0.00
2012	0.11	2.74	0.00	0.22	8.89	0.00
Mean	**0.10**	**2.78**	**0.00**	**0.18**	**7.13**	**0.00**
Median	**0.10**	**2.55**	**0.00**	**0.18**	**7.31**	**0.00**
Pooled	0.11	6.57	0.00	0.16	7.91	0.00

Notes: This table reports the results of estimating equation (2). The first column reports the partial R^2 associated with the source and destination effects. The second column reports the F-statistic associated with the hypothesis that all of the source/destination effects are zero, and the third column reports the p-value associated with that hypothesis test. The results are reported year-by-year as well as pooled across years. The pooled estimation uses source-year and destination-year effects.

Table 2 reports the results. Source effects account for about 10% of the variation in the cross-section of source-destination growth rates, compared to 18% for the destination shocks. The table reports the F-statistics and p-values associated with the hypothesis that the source effects as a group are zero. The source effects are jointly highly significant in accounting for the variation in the data.

4 A structural framework for interpreting the data

The preceding empirical results underscore two key features of the data. First, there is significant comovement between multinational parents and their foreign affiliates. This comovement is detectable in overall source-destination sales. Second, there is a large amount of heterogeneity across sources, destinations, and country pairs in the extent of multinational presence. This suggests that the impact of multinational firms on business cycle comovement may differ significantly across country pairs. These two features of the data inform the design of the quantitative multi-country model that we use to study the implications of the empirical findings for aggregate cross-country comovement. After setting up the theoretical framework, we circle back to the empirical results in Section 3 and interpret them through the lens of the model.

4.1 Model

Preliminaries The world economy consists of multiple countries indexed by i and n. Each country is populated by differentiated intermediate good producers potentially owned by firms from different countries. The output of the intermediate producers cannot be traded internationally. In each country, intermediates are aggregated into a final good by competitive final goods producers. We assume that the final good is homogeneous across countries and can be freely traded.[13] The final good is the numeraire of the world economy and its price is set to one. We focus on the model's predictions for productivity and aggregate output. As discussed below, these assumptions coupled with a standard functional form for agent preferences imply that production allocations are independent of the international asset market structure.

Technologies and market structure The production function of the final good in each country n is given by:

$$Q_{n,t} = \left[\sum_i \widetilde{A}_{in,t}^{\frac{1}{\rho}} Q_{in,t}^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}}, \tag{3}$$

where $Q_{in,t}$ is a bundle of the output produced by firms from source country i that operate in country n, and ρ denotes the elasticity of substitution across goods produced by firms from different source countries. $\widetilde{A}_{in,t}$ is a source-destination specific productivity parameter, normalized such that $\sum_i \widetilde{A}_{in,t} = 1$ for each n. Thus, the production function is an Armington aggregator of goods produced by firms owned by various countries, including domestically owned and operated firms. As will become clear below, the $\widetilde{A}_{in,t}$ parameters allow us to match the full distribution of multinational production shares in the cross-section of source and destination countries. In addition, changes in $\widetilde{A}_{in,t}$ at the business cycle frequency allow us to accommodate source-destination specific idiosyncratic shocks.

In turn, the intermediate output bundle $Q_{in,t}$ aggregates the output of all the firms

[13] The assumption that the final good is homogeneous is not crucial for the results that follow. Online Appendix C.2 derives the equations under the assumption that country-specific goods are imperfect substitutes, so that there are terms of trade movements in response to productivity shocks.

from source country i operating in n:

$$Q_{in,t} = \left[\sum_{f \in \Omega_i} Q_{in,t}(f)^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}}, \quad (4)$$

where Ω_i is the set of firms from country i and $Q_{in,t}(f)$ is the output of firm f from country i in the destination country n.

As in Melitz (2003), firms are monopolistically competitive and differ in productivity. Each firm operates a linear technology that uses labor in the destination country as the only input in production. Following the literature on multinational production and technology transfers, we assume that the multinational's technology can be partially shared across all destination countries.[14] In particular, the output of the firm is given by:

$$Q_{in,t}(f) = Z_{in,t}(f) L_{in,t}(f) = \widetilde{Z}^{\phi}_{i,t}(f) \widetilde{Z}^{1-\phi}_{n,t}(f) L_{in,t}(f), \quad (5)$$

where $L_{in,t}(f)$ is the firm-specific labor input, and $Z_{in,t}(f)$ is a firm-destination specific productivity component.

The second equality in (5) states the key assumption in our framework. Productivity of an affiliate of firm f selling in n and whose parent is from i is a Cobb-Douglas aggregate of the parent's productivity $\widetilde{Z}_{i,t}(f)$ and a local productivity component $\widetilde{Z}_{n,t}(f)$.[15] The parent thus transfers productivity to the affiliate, with the share ϕ of the affiliate's productivity coming from the parent.[16] This is the only potential endogenous source of aggregate comovement in the model. It is worth noting that, while formally technology transfers are the drivers of comovement between parent and affiliate firms, our model is isomorphic to a setup in which comovement is driven by shocks to demand for the firms' product. In this alternative setup, the term $Z_{in,t}(f)$ would come out of equation (5), and demand

[14] See, among others, McGrattan and Prescott (2009, 2010), Ramondo and Rappoport (2010), Keller and Yeaple (2013), Ramondo (2014), Antràs and Yeaple (2014), and Tintelnot (2016) for theoretical treatments that adopt this assumption, and Bilir and Morales (2016) for empirical evidence that R&D spending by the parent is the key determinant of foreign affiliate performance.

[15] The assumption that the productivity in the source and destination are combined by a Cobb-Douglas aggregator is not crucial. Online Appendix C.3 derives the equations under CES aggregation of productivities, and shows they are the same to a first-order approximation.

[16] We assume that the parameter ϕ is the same for all country pairs. The model can be easily extended to let ϕ vary across i and n. Our firm-level empirical results evaluate many types of heterogeneities in the parent-affiliate comovement across countries, including, but not limited to, by quality of IPR protection, geographical distance, and directly by MP shares themselves. There is some heterogeneity, for instance country-pairs with higher multinational presence do experience slightly higher parent-affiliate comovement. However, quantitatively the differences are minor, and thus we favor a more parsimonious specification and do not make explicit use of this variation in the analysis below.

shifters for the firm's product would enter as $Z_{in,t}(f)^{\frac{1}{\rho}}$ in equation (4).[17] In addition, our setup can be reinterpreted in a version of the model in which the transmission of shocks is driven by cross-border intermediate input linkages, as in e.g. Irarrazabal et al. (2013). Online Appendix C.4 presents this alternative model.[18]

Finally, we assume that the source component of the firm productivity is given by $\widetilde{Z}_{i,t}(f) = Z_{i,t}Z_{i,t}(f)$, where $Z_{i,t}$ is common to all firms from source i and $Z_{i,t}(f)$ is idiosyncratic to firm f. We make the same assumption for the destination component $\widetilde{Z}_{n,t}(f) = Z_{n,t}Z_{n,t}(f)$.

Preferences Consumers in country n experience utility from consumption of the final good and disutility from supplying labor according to the GHH preferences (Greenwood et al., 1988):

$$u(C_{n,t}, L_{n,t}) = \sum_t \delta^t \nu \left(C_{n,t} - \frac{\psi_0}{\bar{\psi}} L_{n,t}^{\bar{\psi}} \right),$$

where $C_{n,t}$ is consumption, $L_{n,t}$ the labor supply, and the function ν is increasing and concave.

Equilibrium Let $W_{n,t}$ denote the wage earned by labor in n, and $P_{in,t}(f)$ denote the price charged by firm f from country i operating in n. A *monopolistically competitive equilibrium* at time t is a set of prices $\left\{W_{n,t}, \{P_{in,t}(f)\}_{i,f}\right\}_n$ and resource allocations $\left\{C_{n,t}, L_{n,t}, \{L_{in,t}(f)\}_{i,f}\right\}_n$ such that (i) consumers maximize utility; (ii) firms maximize profits, and (iii) all goods and factor markets clear. Online Appendix C.1 characterizes the equilibrium.

Let lower-case variables denote growth rates of the corresponding upper-case variables. Online Appendix C.1 shows that aggregate growth in country n is approximated by:

$$\gamma_{n,t} = \psi \sum_i \sum_{f \in \Omega_i} \omega_{in,t}(f) \left[\frac{\tilde{a}_{in,t}}{\rho - 1} + z_{in,t}(f) \right], \tag{6}$$

[17] Of course, this isomorphism claim relies on the assumption that firm demand shocks across locations have the same posited structure as productivity shocks in the baseline, that is, a demand shock in a location is a combination of a global demand shock for the multinational's product and a destination-specific component.

[18] The production function can be easily generalized to a case in which firms use local intermediate inputs to produce. In particular, if the firms' production function is given by $Q_{in,t}(f) = [Z_{in,t}(f) L_{in,t}(f)]^{\alpha} X_{in,t}^{1-\alpha}(f)$, where $X_{in,t}(f)$ is an input produced with the final good (3), the results below go through with $\rho - 1$ replaced by $\alpha(\rho - 1)$ in equations (6), (7), (8) and (10).

where $\omega_{in,t}(f) \equiv \frac{P_{in,t}(f) Q_{in,t}(f)}{P_{n,t} Q_{n,t}}$ denotes the share of country $n's$ revenues generated by firm f from source country i, and $\psi \equiv \frac{\bar{\psi}}{\bar{\psi}-1} > 1$. Using the functional form for $z_{in,t}(f)$, equation (6) becomes

$$\gamma_{n,t} = \frac{\psi}{\rho - 1} \sum_i \omega_{in,t} \left[a_{in,t} + \phi (\rho - 1) z_{i,t} \right] + \psi (1 - \phi) z_{n,t}, \tag{7}$$

where $\omega_{in,t} \equiv \frac{P_{in,t} Q_{in,t}}{P_{n,t} Q_{n,t}}$ denotes the share of country n's revenues generated by firms from source country i, and the bilateral term encompasses the idiosyncratic terms specific to the country pair: $a_{in,t} = (\rho - 1) \sum_{f \in \Omega_i} \frac{\omega_{in,t}(f)}{\omega_{in,t}} \left[\frac{\tilde{a}_{in,t}}{\rho - 1} + \phi z_{i,t}(f) + (1 - \phi) z_{n,t}(f) \right]$.

Discussion Equation (7) encapsulates the role of multinationals in business cycle comovement. It states that growth in country n depends on its own productivity shock, $z_{n,t}$, and a weighted average of the productivity shocks $z_{i,t}$ to all countries that have firms operating in country n. Because foreign multinational affiliates inherit part of the shock to the parent $z_{i,t}$, their presence implies that productivity and output of countries will be positively correlated even if the primitive productivity shocks $z_{n,t}$ are not. This equation connects our framework to the international business cycle literature in the tradition of Backus et al. (1995, henceforth BKK). The canonical BKK model has no multinationals, but it typically assumes that TFP shocks across countries are correlated. Equation (7) provides a possible micro foundation for this correlation.

The equation illuminates the key parameters and quantities that determine the strength of the shock transmission through multinationals. The first is the share of the affiliate productivity shock that originates in the source country, ϕ. The more foreign affiliates inherit the source country productivity, the more comovement there will be in the aggregate. The second is the multinational shares, $\omega_{in,t}$. Larger shares will imply more comovement, since more of the shocks are shared. Finally, the combination of parameters $\frac{\psi}{\rho-1}$ captures the strength of general equilibrium effects that occur in response to a particular productivity shock $z_{i,t}$. It regulates how the rest of the economy responds to a shock in a particular country.[19]

Note that output in the model is determined independently of the structure of international asset markets and of how multinational firms' profits are distributed across

[19] As a side note, if the $z_{i,t}$'s are not perfectly correlated, increased multinational presence reduces the volatility of output growth by diversifying shocks, a mechanism reminiscent of Caselli et al. (2015). In particular, when $z_{i,t}$'s are uncorrelated and the variance of $z_{i,t}$ is equal to σ_z^2 for all i, and ignoring idiosyncratic shocks $a_{in,t}$, the variance of output growth is given by $\sigma_n^2 = \psi^2 \left[2\phi (1 - \phi) \omega_{nn} + \phi^2 \sum_i \omega_{in}^2 + (1 - \phi)^2 \right]^2 \sigma_z^2$. Output growth variance is highest in autarky ($\omega_{nn} = 1$).

countries. While these factors will determine how countries split the consumption of the final good, under GHH preferences the labor supply is independent of the level of consumption. As long as firms maximize profits, output growth is solely determined by productivity growth in each country. In this sense, our model is closely related to a standard international business cycle model with one good and no capital. The assumption of a homogeneous final good thus allows us to isolate the comovement arising from the transmission of shocks within multinational firms, while abstracting from the transmission arising from factor supply and relative price movements that are emphasized in the international business cycle literature. For the same reason, for the purposes of the analysis below we do not need to take a stand on how multinationals' profits are allocated (the allocation of profits is important in other contexts, see, e.g., Arkolakis et al., 2014). Allocation of profits across countries will affect consumption, but not output in our framework.

In the counterfactual exercises that follow, the MP shares along with the elasticities ϕ and $\frac{\psi}{\rho-1}$ are sufficient statistics to evaluate the propagation of country level shocks across countries. An advantage of this approach is that it avoids the need to take an explicit stand on a large number of parameters of the model, such as the levels of productivities and the demand shifters $\widetilde{A}_{in,t}$ in every country. Note that in our framework MP shares are endogenous, as they arise from firm optimization decisions given productivities. Online Appendix C.5 lays out a model in which firms' decisions to open foreign affiliates are made subject to sunk costs and uncertainty, as in Ramondo and Rappoport (2010) and Ramondo et al. (2013), and can potentially depend on the covariance of shocks across countries. The appendix shows that the key equations, the parameter estimation approaches, and the counterfactual analyses in the main text are still valid in this extended model. It also discusses the model's relationship to the frameworks of Ramondo and Rappoport (2010) and Ramondo et al. (2013), and provides parameterizations under which the covariance of shocks does not affect the equilibrium MP shares, despite the sunk cost.[20]

Finally, we note that the counterfactual analysis below will evaluate the impact of multinationals on business cycle comovement through the mechanism posited by the model, namely the within-firm productivity transmission coupled with general-equilibrium effects. We acknowledge that multinationals may have other effects on international comovement that we leave unmodeled, for instance through engagement in arms-length

[20] The Appendix model assumes that MP entry only happens in period 0. This assumption implies that we do not need to keep track of each firm's entry and exit decisions in each period, and can solve for one non-time-varying productivity cutoff for MP for each source-destination pair. Hence, the model abstracts from the extensive margin adjustment of MP in response to shocks, highlighted by Fillat and Garetto (2015) and Fillat et al. (2015). With time-varying entry subject to sunk costs, a country experiencing a (persistent) positive productivity shock would generate more entry from all source countries. These extensive margin responses may generate correlation over and above what is implied by our quantitative exercises.

trade, or technology spillovers outside the firm.

We now interpret the empirical results from Section 3 in light of the conceptual framework, and use these results to disentangle the different shocks and discipline the model.

4.2 Interpreting affiliate-parent comovement

The empirical results in Section 3.1 can be given a structural interpretation and used to estimate the share of the firm's technology that gets transferred across destinations, ϕ. Online Appendix C.1.1 shows that the affiliate growth rate can be expressed in the model as a function of parent growth:

$$\gamma_{in,t}(f) \quad = \quad \bar{\bar{a}}_{in,t} + \phi\gamma_{ii,t}(f) + \epsilon_{in,t}(f), \tag{8}$$

where $\bar{\bar{a}}_{in,t} \equiv \bar{a}_{in,t} - \phi\bar{a}_{ii,t}$, with $\bar{a}_{in,t} \equiv (1-\rho)[w_{n,t} - \phi z_{i,t} - (1-\phi) z_{n,t}] + \rho p_{in,t} + q_{in,t}$, and $\epsilon_{in,t}(f) \equiv (\rho - 1)(1-\phi) z_{n,t}(f)$.

Equation (8) states that, after controlling for source-destination-year effects, the coefficient on the parent's growth rate can be interpreted as ϕ. Hence, the empirical results in Table 1 imply that the share of a firm's productivity that is transferred across countries is approximately 20% ($\phi \approx 0.2$).

Econometrically, an identifying assumption required for the structural interpretation of the regression coefficients is that the idiosyncratic component of the destination-specific shock $z_{n,t}(f)$ is orthogonal to the idiosyncratic component of the source shock $z_{i,t}(f)$. Substantively, however, this orthogonality assumption is without loss of generality, as the technology transfer coefficient can simply be reinterpreted as technology transfer-cum-shock correlation, in which case the destination-specific shock becomes orthogonalized with respect to the source-specific shock (see Online Appendix C.1.1 for details).

Note that (8) can be thought of as a purely cross-sectional specification, in spite of the variables being indexed by t. Indeed, the empirical results are unchanged when we estimate the specification separately for each year. Under the assumption that affiliate productivity is a function of only the contemporaneous (and not lagged) parent's productivity, (8) is the correct specification even if there is time dependence in the underlying productivity shocks $z_{n,t}(f)$ and $z_{i,t}(f)$. Indeed, since the idiosyncratic component of the destination-specific productivity $z_{n,t}(f)$ is part of the error term, time dependence in it does not bias coefficient estimates, and clustering at the parent level adjusts for the autocorrelated error structure.[21]

[21] Irarrazabal et al. (2013) estimate the share of inputs imported from the parent in total costs of foreign affiliates of Norwegian multinationals, and find that it is quite high (0.9). Though not the same object, it

4.3 Interpreting source- and destination-specific shocks

We now use the model's implications for aggregate source-destination growth rates $\gamma_{in,t}$ to interpret the empirical results in Section 3.2. Online Appendix C.1.1 shows that the total revenues by multinationals from source country i operating in country n can be expressed as:

$$P_{in,t}Q_{in,t} \quad = \quad A_{in,t}S_{i,t}D_{n,t}, \tag{9}$$

where $S_{i,t} = Z_{i,t}^{\phi(\rho-1)}$ is a term common to all firms from source country i, $D_{n,t} = \left[\sum_i A_{in,t}Z_{i,t}^{\phi(\rho-1)}\right]^{\frac{\psi-\rho+1}{\rho-1}} Z_{n,t}^{\psi(1-\phi)}$ is a term common to all firms operating in destination country n, and $A_{in,t} = \widetilde{A}_{in,t}\sum_{f\in\Omega_i}\left[Z_{i,t}^{\phi}(f)Z_{n,t}^{(1-\phi)}(f)\right]^{\rho-1}$. Expressed in growth rates, this becomes:

$$\gamma_{in,t} \quad = \quad s_{i,t} + d_{n,t} + a_{in,t}, \tag{10}$$

which is identical to the decomposition (2) estimated in Section 3.2, and
$a_{in,t} = (\rho-1)\sum_{f\in\Omega_i}\frac{\omega_{in,t}(f)}{\omega_{in,t}}\left[\frac{\tilde{a}_{in,t}}{\rho-1} + \phi z_{i,t}(f) + (1-\phi)\, z_{n,t}(f)\right]$.

Equation (10) provides a structural interpretation for the source and destination dummies estimated in Section 3.2. The fact that a significant fraction of the variation of the bilateral growth rates is accounted for by the source dummies, as reported in Table 2, implies a role for the transmission of technology from the source country, $\phi > 0$.

4.4 Calibrating the comovement parameter with source-destination data

These structural equations together with the estimates for the source- and destination-specific shocks provide a means to pin down the technology transfer parameter ϕ. In particular, the model structure implies the destination components have the form:

$$d_{n,t} \quad = \quad \left[\frac{\psi}{\rho-1} - 1\right]\sum_i \omega_{in,t}\left[a_{in,t} + s_{i,t}\right] + \frac{\psi}{\rho-1}\frac{1-\phi}{\phi}s_{n,t}. \tag{11}$$

Foreign productivity shocks $z_{i,t}$ affect the destination effect in country n through two different channels. On the one hand, these changes affect competitiveness in country n

can be interpreted as related to our ϕ parameter. Our approach identifies ϕ from the comovement in sales between parents and affiliates, while Irarrazabal et al. (2013) rely instead on the elasticity of affiliate sales to distance. A higher value of ϕ would lead to an even greater importance of multinationals in business cycle transmission.

through $\left[\sum_i A_{in,t} Z_{i,t}^{\phi(\rho-1)}\right]^{-1}$ (i.e. in response to an increase in $Z_{i,t}$, firms from all other source countries i' will sell less in country n due to increased competition). On the other hand, these shocks affect the real wage (and real aggregate output) in country i through $\left[\sum_i A_{in,t} Z_{i,t}^{\phi(\rho-1)}\right]^{\frac{\psi}{\rho-1}}$ (i.e. in response to an increase in $Z_{i,t}$, aggregate demand in country n will increase, increasing the sales of all firms operating in country n). In the case of $\rho - 1 = \psi$ these two effects exactly offset each other, and the destination effect is independent of changes in foreign technologies.

Rearranging (11), taking variances of both sides, and solving for ϕ yields an estimate of ϕ based on observed variabilities of the source and destination effects:

$$\phi \quad = \quad \frac{\sigma_{s,t}}{\sigma_{s,t} + \sigma_{\Phi,t}}, \tag{12}$$

where $\sigma_{s,t}^2 \equiv \frac{1}{N-1} \sum_n \left(s_{n,t} - \frac{1}{N} \sum_m s_{m,t}\right)^2$ is the cross-sectional variance of $s_{n,t}$ at time t, and $\sigma_{\Phi,t}^2 \equiv \frac{1}{N-1} \sum_n \left[\frac{\rho-1}{\psi}\left(d_{,t} - \frac{\psi+1-\rho}{\rho-1} \sum_i \omega_{in,t}\left[a_{in,t} + s_{i,t}\right]\right) - \frac{\rho-1}{\psi} \frac{1}{N} \sum_m \left(d_{m,t} - \frac{\psi+1-\rho}{\rho-1} \sum_i \omega_{im,t}\left[a_{im,t} + s_{i,t}\right]\right)\right]^2$ is the variance of the destination effect adjusted for the general equilibrium impact of foreign shocks. Note that in the special case of $\rho - 1 = \psi$, $\sigma_{\Phi,t}^2 = \sigma_{d,t}^2 = \frac{1}{N-1} \sum_n \left(d_{n,t} - \frac{1}{N} \sum_m d_{m,t}\right)^2$ is simply the cross-sectional variance of the destination effects at time t.

Equations (11) and (12) use the model structure to connect observables – $s_{i,t}$, $d_{n,t}$, and $a_{in,t}$ estimated in Section 3.2 – to the two key model parameters, $\psi/(\rho-1)$ and ϕ. For each value of $\psi/(\rho-1)$ we can thus use (11) and (12) and the estimated $s_{i,t}$, $d_{n,t}$, and $a_{in,t}$ to pin down ϕ.

The basic intuition for this approach can be gleaned from (12) and the fact that the source effect is a scaled productivity shock: $s_{n,t} = \phi(\rho-1) z_{n,t}$. Ignoring the general equilibrium effects, (12) says that ϕ determines the relative variances of the estimated source and destination effects. In the world of no spillovers from source countries ($\phi = 0$), shocks to the source country do not affect bilateral growth rates, so that the variance of the source effects is zero. By contrast, high ϕ would manifest itself in a high variability of the source effects. The variance of the source effects is benchmarked by the variance of the destination effects, since those are driven by the same productivity shock process as the source effects, but affect all the firms operating in each market.

Table 3 presents the implied ϕ for different values of $\frac{\psi}{\rho-1}$. We focus on the special case of $\frac{\psi}{\rho-1} = 1$, in which the general equilibrium effects cancel out, and the alternative cases of $\frac{\psi}{\rho-1} = 2$ (the effect of a positive foreign shock on domestic income dominates the effects on increased competition) and $\frac{\psi}{\rho-1} = 2/3$ (the increase in competition dominates

Table 3: Estimated ϕ based on source-destination data

Year	$\frac{\psi}{\rho-1}=1$	$\frac{\psi}{\rho-1}=2$	$\frac{\psi}{\rho-1}=\frac{2}{3}$
2005	0.517	0.579	0.425
2006	0.441	0.518	0.370
2007	0.379	0.455	0.311
2008	0.417	0.521	0.326
2009	0.380	0.507	0.285
2010	0.415	0.531	0.322
2011	0.365	0.477	0.276
2012	0.401	0.485	0.330
Mean	**0.414**	**0.509**	**0.331**
Median	**0.408**	**0.512**	**0.324**

Notes: This table reports estimates of ϕ using bilateral data following equation (12). Each column represents the estimate under an alternative value of the GE parameter $\frac{\psi}{\rho-1}$.

the effect of increased income).[22] The estimates of ϕ range from 0.3 to 0.5, with a central tendency of about 0.4. This is higher than, but not too dissimilar from, the firm-level estimates in Section 4.2.

What are the relative merits of the firm-level based estimates of ϕ from Section 4.2 compared to the source-destination level estimates in this section? The firm-level estimates use stringent fixed effects, and thus represent the most convincing evidence that the correlation between parents and affiliates captures within-firm transmission of shocks rather than simply common shocks across countries and/or sectors. On the other hand, precisely because it nets out common shocks at the source-sector-destination-sector-year level, the firm-level estimation will omit the within-firm transmission of aggregate shocks. A shock that hits all the firms in the Chemicals sector in France may be transmitted from the French parent operating in the Chemicals sector to its subsidiaries in Spain. But the fixed effects in the firm-level specification net out the aggregate/sectoral shocks, and thus identify only the transmission of the idiosyncratic shock hitting the French Chemicals parent. The firm-level estimate will shed light on this channel to the extent that common shocks are transmitted with the same intensity as purely idiosyncratic shocks. Al-

[22]The special case of $\psi = \rho - 1$ is consistent with empirical estimates of the Frisch elasticity of labor supply and the elasticities of substitution across intermediate varieties used in the trade literature. In particular, estimates of the aggregate labor supply elasticity put it at about 0.5 (see Chetty et al., 2013), which implies a $\bar{\psi} = 3$ and $\psi = 1.5$. This implies that $\rho = 2.5$ – well within the range of estimates in Broda and Weinstein (2006) – is consistent with $\frac{\psi}{\rho-1} = 1$. Under an aggregate labor supply elasticity of about 0.5, $\frac{\psi}{\rho-1} = 2$ (resp., $\frac{2}{3}$) implies an elasticity of substitution of $\rho = \frac{7}{4}$ (resp., 3.25).

ternatively, one can focus on the source-destination level estimates, since the source and destination effects will capture not only the transmission of firm-level, but also of aggregate shocks in the parent country to the foreign destinations. An additional benefit of the source-destination estimates is the much lower data requirements, as they can be implemented without firm-level data, and thus potentially used even in contexts where only aggregate MP data are available.

5 Quantitative results

We now have the theoretical structure, the estimates of the key parameter, and the data to carry out a quantitative assessment of multinationals' role in the international business cycle transmission. This section performs three exercises. The first is an "impulse response" exercise designed to answer the question, how much does a productivity shock in one country affect output in another? The second is a counterfactual correlation exercise, that answers the question, if all the countries' productivity shocks were uncorrelated, how much correlation would the business cycles exhibit across countries under the current levels of multinational activity? And third, how much do multinationals contribute to observed dispersion in cross-country growth rates, and how much would that dispersion fall if integration increased further? The exercises in the next two subsections do not require time subscripts, and thus we suppress them to streamline notation.

5.1 Transmission of shocks across countries

We start by assessing the total impact of all foreign productivity shocks on a country's productivity. One way to gauge the importance of all foreign shocks combined is to consider the impact of a 1% change in all foreign productivities simultaneously. The change in destination n's productivity in this experiment is given by

$$\phi(1 - \omega_{nn}). \tag{13}$$

This expression has a clear and intuitive interpretation. The combined importance of foreign shocks in country n is the product of the total presence of multinationals, $1 - \omega_{nn}$, and the strength of the productivity transmission from foreign parents to the local affiliates, ϕ.

The top row of Table 4 and Figure 2 report the results, under the assumption that $\phi = 0.4$ and ω_{nn}'s for 2011. In our sample of 34 countries, the mean value of this combination

Table 4: Impulse responses for top source countries and the world

	Destination				
Source	All Countries	High-Income Europe	Emerging Europe	High-Income ROW	Emerging ROW
World	0.121	0.139	0.126	0.075	0.077
United States	0.023	0.038	0.009	0.019	0.019
Germany	0.017	0.017	0.023	0.003	0.010
United Kingdom	0.015	0.022	0.008	0.020	0.005
France	0.010	0.015	0.010	0.002	0.004

Notes: This table reports averages of the impulse responses (13) in the entire sample of countries and 4 regions for the top 4 most important source countries, and the average impulse responses to a world shock.

of parameters is 0.12, with the median of 0.11. This suggests that loosely speaking, foreign shocks can account for 12% of productivity shocks in the average country, or alternatively, foreign shocks are about one-ninth as important as domestic productivity shocks. Foreign shocks are most important on average in Europe (13.9% and 12.6% in the high-income and emerging groups, respectively), and less important in the other countries in our cample. In some countries foreign shocks are more significant. At the extreme, the value of this combination of parameters is 0.35 in Ireland, 0.25 in the Netherlands, and 0.23 in Slovakia.

We next evaluate how productivity shocks to any individual source country spread internationally. From (7), the response of output in n to a productivity shock in any source country i is given by:

$$\frac{\partial \gamma_n}{\partial z_i} = \psi \left[\omega_{in} \phi + (1-\phi) \mathbb{I}_{i=n} \right], \quad (14)$$

where $\mathbb{I}_{i=n}$ is an indicator function that equals 1 if $i = n$ and 0 otherwise. We can express the response in country n as a fraction of the effect of the shock in the source country i as:

$$\frac{\partial \gamma_n}{\partial z_i} / \frac{\partial \gamma_i}{\partial z_i} = \frac{\omega_{in} \phi}{\omega_{ii} \phi + (1-\phi)} \quad n \neq i. \quad (15)$$

Equation (15) answers the question, how much does aggregate output growth in country n change when output in country i goes up by 1? It is immediate that the answer depends on two key quantities: the magnitude of the spillover ϕ, and the extent of country i's multinational presence in n, ω_{in}. If either of these is large, there will be more interdependence between i and n. In contrast, note that given these parameters, the impulse

Figure 2: Response (in %) to a 1% shock in all foreign countries simultaneously

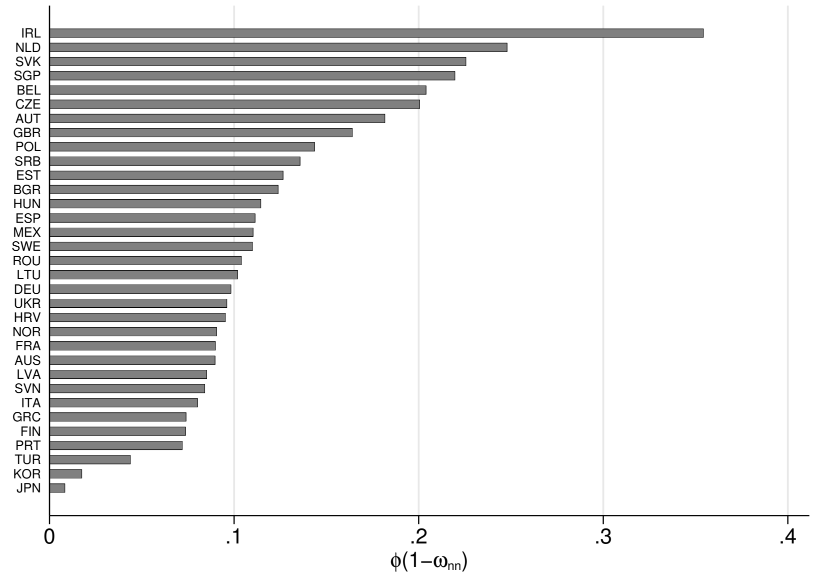

Notes: This figure displays the change in productivity in each destination that accompanies a change in productivity in every foreign source country (i.e. $i \neq n$) equal to 1.

response does not depend on the value of the general equilibrium parameter $\frac{\psi}{\rho-1}$. There is no simulation required to compute these impulse responses. Instead, they are computed directly from the data on ω_{in} and estimated ϕ. Since there are 34 countries in the sample, there are 34×33 cross-border impulse responses.

Figure 3 shows the impulse responses to source-specific shocks for all possible country pairs. We use $\phi = 0.4$ and ω_{in}'s for 2011 to construct the figure. Each square in the figure represents the impulse response in destination n to a productivity shock in source country i, z_i, relative to the response in the source country, as in (15). We can interpret each square of the figure as the percent change in country n GDP in response to a shock that increases GDP in country i by one percent. We rank countries on the x- and y-axes according to their importance as a source or as a destination, respectively. We omit the ii entries (they are all tautologically 1) to facilitate the presentation.

The figure shows that shocks to the productivity of most source countries do not have big aggregate consequence in most destinations. This reflects the fact that the bilateral shares ω_{in} in equation (15) are small for most country pairs. In about half of all source-destination pairs, the impact is less than 0.0001, reflecting the (near) absence of multinationals from most sources in most destinations. Even restricting attention to the top half of impulse responses, the mean impact is about 0.006, that is, an increase in a source country output of 1% changes foreign output by less than one-hundredth of that amount.

Figure 3: Response (in %) to a source shock that raises source country output by 1%

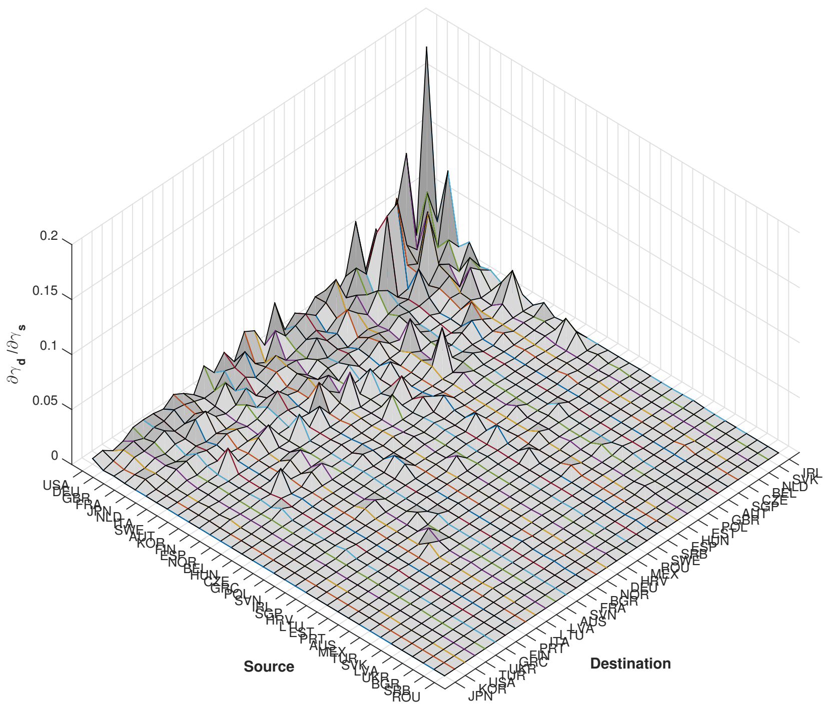

Notes: This figure displays the change in aggregate output of each destination that accompanies a change in source output equal to 1.

However, this low amount of transmission is in part a consequence of the fact that most countries are not quantitatively important sources of multinationals. Table 4 reports the average impulse responses to shocks in the top 4 most important source countries: the US, Germany, the UK, and France. In the entire sample, the average outward impact of these 4 countries ranges from 1 to 2.3 percent. The next four columns report averages by destination country regions. There is some heterogeneity in the regional impact: the shock to the US affects most strongly high-income Europe, a 3.8% average impulse response. By contrast, a shock to Germany has the largest effect in emerging Europe, 2.3%. In total, 18 country pairs have impulse response coefficients of above 0.03, with the maximum coefficient of 0.17 between US and Ireland.

5.2 Country-pair growth correlations and multinational shares

This section derives how much comovement in aggregate output would be generated by the presence of multinationals in a world where the only shocks are shocks to country-level productivities z_i. Consider a setting in which country productivity shocks have variance σ_z^2 and correlation $\rho_{z,z'}$, common across z and z'. Under these conditions, the output growth correlation between any pair of countries is:

$$\rho_{n,n'} = \frac{\left[\phi\left(1-\phi\right)\left[\omega_{n'n}+\omega_{nn'}\right]+\phi^2\sum_i \omega_{in}\omega_{in'}\right]\left[1-\rho_{z,z'}\right]}{\bar{\Theta}_n\bar{\Theta}_{n'}}+\frac{\rho_{z,z'}}{\bar{\Theta}_n\bar{\Theta}_{n'}}, \qquad (16)$$

where $\bar{\Theta}_n^2 \equiv \left[2\phi\left(1-\phi\right)\omega_{nn}+\phi^2\sum_i \omega_{in}^2+\left(1-\phi\right)^2\right]\left[1-\rho_{z,z'}\right]+\rho_{z,z'}$. Note that the correlation $\rho_{n,n'}$ is a function only of the correlation in firm-level growth ϕ, the multinational shares ω_{in}, and the correlation of the shocks $\rho_{z,z'}$. Given a value of ϕ, the size of the general equilibrium effects does not affect the results in this section.

Equation (16) illustrates how the parameters affect the correlation of growth rates across countries. First, if the primitive shocks are uncorrelated ($\rho_{z,z'} = 0$), and there are no multinational firms ($\omega_{in} = 0$ for $i \neq n$), then countries growth rates are uncorrelated, $\rho_{n,n'} = 0$. Other things equal, the correlation increases for country pairs that share more multinational links, as captured by the terms $\omega_{n'n} + \omega_{nn'}$ and $\sum_i \omega_{in}\omega_{in'}$. Second, the scope for multinational firms to induce cross-country correlations falls as the correlation of the primitive shocks increases. In the limit, if $\rho_{z,z'} = 1$, output is perfectly correlated across countries, $\rho_{n,n'} = 1$, irrespective of the multinational shares ω_{in}.

With this in mind, Table 5 evaluates the model's ability to generate positive cross-country growth correlations. The row labeled "Data" presents the summary statistics for the correlations of GDP growth over the period 1994-2007. The row labeled "Model" presents the correlations implied by the model when $\rho_{z,z'} = 0$. Consistent with our results from the previous section, on average in the whole sample the predicted correlations tend so be small. The mean is only 0.01, and 95% of all the bilateral correlations are below 0.03. However, this is partly a consequence of small multinational shares for most pairs of countries. The bottom two panels of the Table report the correlations for country pairs in which one country is a large source of multinational firms (such as the US, the UK, or Germany), and for country pairs in which one country is an important destination for multinationals (such as Ireland, Netherlands, or Slovakia). Not surprisingly, the predictions of the model come much closer to the data for these country pairs. At the extreme, the model generates about a quarter of the observed correlation for country pairs in which one of the countries is either the US or Ireland.

Table 5: Predicted and actual correlations

		Across all country pairs			
		Mean	St.Dev.	Min	Max
Data		0.132	0.352	-0.684	0.872
Model		0.009	0.017	0.000	0.255

	Pairs involving large sources					Pairs involving large destinations				
		Mean	St.Dev.	Min	Max		Mean	St.Dev.	Min	Max
Data	US	0.131	0.404	-0.623	0.752	Ireland	0.097	0.413	-0.684	0.868
Model		0.029	0.047	0.002	0.255		0.024	0.047	0.002	0.255
Data	UK	0.127	0.303	-0.547	0.588	Netherlands	0.232	0.429	-0.479	0.848
Model		0.024	0.029	0.002	0.119		0.020	0.023	0.004	0.119
Data	Germany	0.254	0.303	-0.243	0.752	Slovakia	0.049	0.326	-0.512	0.617
Model		0.023	0.018	0.005	0.080		0.011	0.014	0.001	0.067

Notes: This table reports summary statistics for aggregate correlations. The row labeled "Data" reports the actual correlations of aggregate GDP growth sourced from World Development Indicators over the period 1994-2007. The row labeled "Model" reports the results for correlations computed using equation (16) under the assumption that $\sigma_{z,z'} = 0$ for all country pairs.

Second, in the spirit of the literature on international trade and comovement (see, e.g., Johnson, 2014), we compare the correlations predicted by the model to those observed in the data.[23] Figure 4 plots the partial correlation between the GDP correlation in the data (y-axis) against the correlation implied by the model under uncorrelated shocks (x-axis), after controlling for source and destination country effects. There is a positive and highly significant conditional correlation between the model-implied correlations and the correlations in the data.[24] This is remarkable given that the model correlations are computed under the assumption of uncorrelated shocks, and the only source of variation across country pairs in the model is due to differences in multinational shares. The model also cannot yield negative correlation coefficients, which are observed in the data. As a result, relative to the data, the model generates substantially less variation in the cross-country correlations than observed in the data.

To underscore the way the model generates correlations, the figure labels the "inte-

[23] A challenge in empirically demonstrating a causal link between multinational presence and business cycle comovement is that multinational firms may locate in countries that for other reasons have more synchronized shocks. By imposing that the primitive shocks are uncorrelated and the same across country pairs, we can isolate the role of multinationals as a source of transmission from other factors that may induce comovement.

[24] The relationship is equally pronounced and significant unconditionally, without controlling for any fixed effects.

Figure 4: Actual and predicted correlations

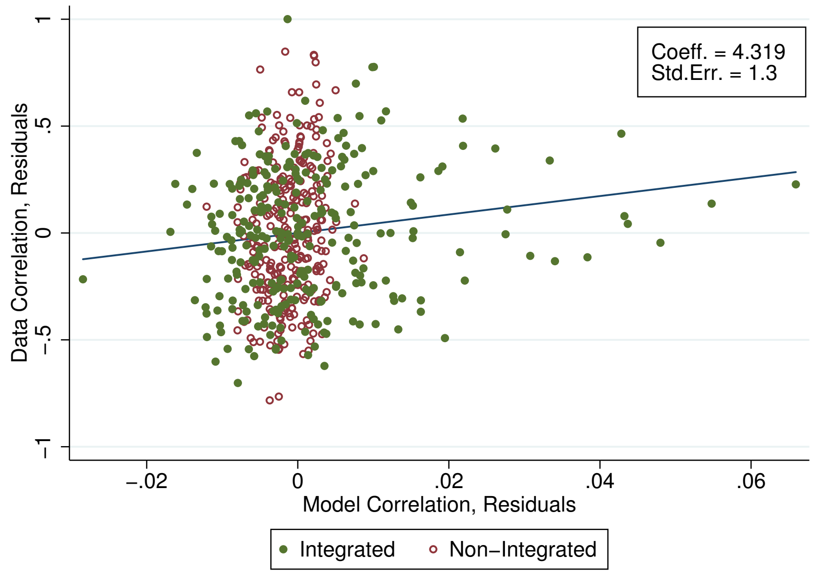

Notes: This figure plots the partial correlation between the GDP correlation in the data and the aggregate correlation implied by the model, after controlling for source and destination country effects. Dots labeled "Integrated" depict the country pairs with higher than the median combined bilateral multinational shares $(\omega_{n'n} + \omega_{nn'})$. Dots labeled "Non-Integrated" depict country pairs with the combined bilateral multinational shares below the median.

grated" and "non-integrated" country pairs differently (though the regression is run on all the data). We label a country pair integrated if its combined bilateral multinational shares $\omega_{n'n} + \omega_{nn'}$ are above the median, and non-integrated if they are below the median. The model generates little to no dispersion in predicted correlation among the non-integrated pairs (hollow dots), and as a result in this subsample the model has no predictive power over the data correlations, which range (in deviations from the country means) from large positive to large negative. On the other hand, for integrated pairs (solid dots) there is a clear positive relationship between the model-implied correlations and those in the data. Indeed, in the subsample of integrated pairs, the model correlations have a greater explanatory power than in the full sample.

Our main takeaway from these exercises is that in most country pairs, transmission of shocks through multinationals in and of itself cannot generate anything close to observed output correlations. This is unsurprising since in most country pairs bilateral multinational shares are small. However, among the more closely integrated country pairs, the model generates both non-negligible correlations, and a significantly positive relationship between model-implied and observed correlations.

5.3 Predicted and counterfactual comovement

This section studies how business cycle synchronization would change under different scenarios for multinational presence. Rather than assuming an exogenous parsimonious shock correlation structure as in the previous section, here we use the estimated $a_{in,t}$ and $s_{i,t}$ from Section 3.2 to compute aggregate growth rates using model-implied relationships (7) and (11). We then conduct two sets of counterfactual exercises to investigate how multinationals contribute to business cycle synchronization. Our metric of synchronization is the cross-sectional dispersion in country-level growth rates (see Kalemli-Ozcan et al., 2013, for a closely related metric of comovement). We exploit the cross-sectional dimension as the available time series is quite short.

Notation The aggregate growth rate in country n is:

$$\gamma_{n,t} \quad = \textstyle\sum_i \omega_{in,t}\gamma_{in,t} = \displaystyle\sum_i \omega_{in,t}\left[a_{in,t} + s_{i,t}\right] + d_{n,t}, \tag{17}$$

We can express country n's growth rate relative to the cross-sectional average growth rate at time t as:

$$\gamma_{n,t} - \bar{\gamma}_t \quad = \quad \mathcal{A}_{n,t} + \mathcal{S}_{n,t} + \mathcal{D}_{n,t} \tag{18}$$

where $\mathcal{A}_{n,t} \equiv \sum_i \omega_{in,t}a_{in,t} - \frac{1}{N}\sum_n \sum_i \omega_{in,t}a_{in,t}$ is the aggregation of all the idiosyncratic shocks; $\mathcal{S}_{n,t} \equiv \sum_i \omega_{in,t}\left[s_{i,t} - \bar{s}_{i,t}\right] - \frac{1}{N_n}\sum_n \sum_i \omega_{in,t}\left[s_{i,t} - \bar{s}_{i,t}\right]$ is the aggregation of all the source shocks, and $\mathcal{D}_{n,t} \equiv d_{n,t} - \bar{d}_t$ is the demeaned destination effect. In these expressions $\bar{x}_t \equiv \frac{1}{N}\sum x_{n,t}$ denotes the average of a variable across all destinations.

Changing multinational shares In the first set of counterfactuals, we ask what the cross-country dispersion in growth rates would look like if multinational shares were different. We focus on two polar opposite counterfactuals: (i) "No multinationals" and (ii) "Full Integration." Under "No Multinationals," we change the values of the ω_{in}'s so that $\omega^{NM}_{in,t} = 1$ if $i = n$, $\omega^{NM}_{in,t} = 0$ if $i \neq n$. That is, the only firms producing in country i are country i firms. Under "Full Integration" we change the ω_{in}'s so that $\omega^{FI}_{in,t} = \bar{\omega}^{FI}_{i,t} = \frac{1}{N}\sum_n^N \omega_{in,t}$.[25] That is, the production shares of firms of all source countries is the same in every country, and equal to the average share of each country i across destinations observed in the data.

[25]Note that $\sum_i \bar{\omega}_{i,t} = \frac{1}{N}\sum_n \sum_i \omega_{in,t} = 1$.

In each of the counterfactual exercises indexed by $c = \{NM, FI\}$, we compute the counterfactual components $\mathcal{S}^c_{n,t}$, $\mathcal{A}^c_{n,t}$, $\mathcal{D}^c_{n,t}$ using estimated $s_{i,t}$ and $a_{in,t}$ as:

$$\mathcal{S}^c_{n,t} = \sum_i \omega^c_{in,t} s_{i,t} - \frac{1}{N}\sum_n \sum_i \omega^c_{in,t} s_{i,t} \tag{19}$$

$$\mathcal{A}^c_{n,t} = \sum_i \omega^c_{in,t} a_{in,t} - \frac{1}{N}\sum_n \sum_i \omega^c_{in,t} a_{in,t} \tag{20}$$

$$\mathcal{D}^c_{n,t} = \frac{\psi + 1 - \rho}{\rho - 1}\left[\mathcal{A}^c_{n,t} + \mathcal{S}^c_{n,t}\right] + \frac{\psi}{\rho - 1}\frac{1-\phi}{\phi}\mathcal{S}^{own}_{n,t}, \tag{21}$$

where $\mathcal{S}^{own}_{n,t} \equiv \left[s_{n,t} - \frac{1}{N}\sum_n s_{n,t}\right]$ captures the deviation of country n's productivity shock from the world average. We use (19)-(21) to compute the counterfactual growth rates $\gamma^c_{n,t}$ as in equation (18) and then report the standard deviations $\sigma^c_{\gamma_{n,t}} = \sqrt{\frac{1}{N-1}\sum_n \left(\gamma^c_{n,t} - \bar{\gamma}^c_t\right)^2}$. Our baseline results adopt the assumption that $\frac{\psi}{\rho-1} = 1$ (the destination shocks are independent of the general equilibrium effects).

Table 6: Cross-sectional standard deviation of $\gamma_{n,t}$

	Mean	Median
$\sigma_{\gamma_{n,t}}$:		
Baseline	0.062	0.064
NM: No Multinationals	0.067	0.066
FI: Full Integration	0.043	0.042
Ratio of $\sigma^c_{\gamma_{n,t}}$ to baseline:		
NM: No Multinationals	1.074	1.075
FI: Full Integration	0.687	0.684
$\frac{\psi}{\rho-1} = 2; \phi = 0.5$		
Ratio of $\sigma^c_{\gamma_{n,t}}$ to baseline:		
NM: No Multinationals	1.091	1.095
FI: Full Integration	0.633	0.631
$\frac{\psi}{\rho-1} = 2/3; \phi = 0.3$		
Ratio of $\sigma^c_{\gamma_{n,t}}$ to baseline:		
NM: No Multinationals	1.056	1.055
FI: Full Integration	0.757	0.749

Notes: This table reports the mean and median cross-sectional standard deviations in aggregate growth rates over all years, and the mean and median ratios of the standard deviations in the two counterfactuals relative to the baseline. The bottom two panels summarize the ratios of standard deviations under alternative parameterizations of $\psi/(\rho-1)$.

The top panel of Table 6 reports the average baseline and counterfactual dispersions of growth rates across the years in our sample. To facilitate comparison, the second panel of the table reports the average ratios of the counterfactual $\sigma^c_{\gamma_{n,t}}$'s relative to the baseline. The standard deviation of growth rates under the "No multinationals" counterfactual is 7.5% higher compared to the baseline. Note from equation (19) that the dispersion of $\mathcal{S}^c_{n,t}$ is higher under this scenario, since multinationals are not there to spread the source shocks across countries.

The table also reports the average standard deviations of growth rates under the "Full Integration" counterfactual. Note from equation (19) that in this case $\mathcal{S}^c_{n,t} = 0$ (since $\omega_{in,t}$ is constant across destinations). Source shocks are completely shared across destinations under full integration, hence differences do not contribute to the dispersion in growth rates. As a consequence, the dispersion in growth rates is significantly smaller under this scenario. For the median year, the standard deviation of growth rates would decrease by over 30% if all barriers to multinationals are eliminated.

The bottom two panels of Table 6 report a sensitivity analysis to alternative values of the general equilibrium parameter $\frac{\psi}{\rho-1}$. We focus on the cases of $\frac{\psi}{\rho-1} = 2$ and $\frac{\psi}{\rho-1} = 2/3$ discussed in Section 4.4. Under each alternative parameterization, we re-calibrate the parameter ϕ according to equation (12). The table shows that the case $\frac{\psi}{\rho-1} = 2$ is associated with slightly larger counterfactual changes in the cross-sectional variance of growth rates, while the opposite is true for the case of $\frac{\psi}{\rho-1} = 2/3$. However, the alternative parameterizations do not change the order of magnitude of the results.

Changing the correlation in firm-level growth In the second set of counterfactuals, we maintain the observed multinational shares and change the correlation between parents and affiliates ϕ^c. In this case we can compute the counterfactual components as:

$$\begin{aligned}
\mathcal{S}^{c\phi}_{n,t} &= \frac{\phi^c}{\phi}\mathcal{S}_{n,t}, \\
\mathcal{A}^{c\phi}_{n,t} &= \mathcal{A}_{n,t} \\
\mathcal{D}^{c\phi}_{n,t} &= \left[\frac{\psi}{\rho-1} - 1\right]\left[\mathcal{A}_{n,t} + \mathcal{S}^{c\phi}_{n,t}\right] + \frac{\psi}{\rho-1}\frac{1-\phi^c}{\phi}\mathcal{S}^{own}_{n,t}.
\end{aligned}$$

Figure 5 shows the resulting standard deviation in growth rates for alternative values for counterfactual ϕ. As ϕ get closer to zero, there is no transmission of shocks between multinational firms and their foreign affiliates, and the standard deviation in growth rates increases and gets closer to that in the counterfactual of "No Multinationals." As ϕ gets closer to one, the correlation between multinationals and their foreign affiliates becomes

Figure 5: Correlation between multinationals and their foreign affiliates (ϕ) and the cross-sectional dispersion of aggregate growth rates

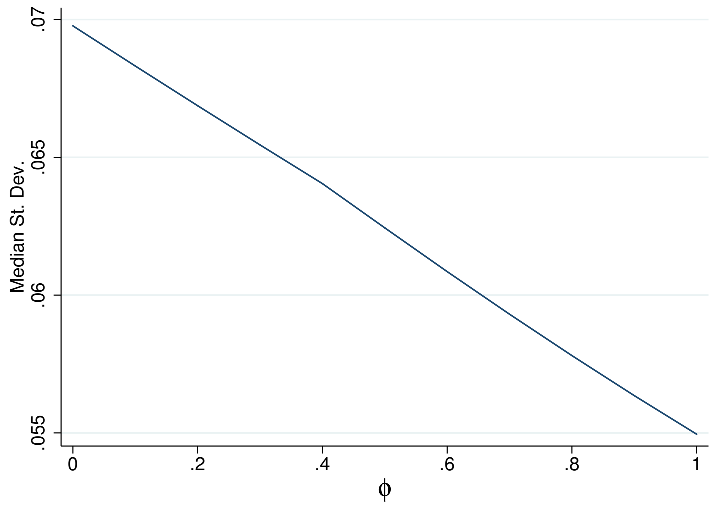

Notes: This figure plots the median standard deviation of aggregate growth rates on the y-axis against the share of source shocks in the affiliates' technology shocks (ϕ) on the x-axis.

stronger, and the dispersion in growth rates decreases. Yet, this effect is limited by the fact that the share of multinationals in the economy is small.

6 Conclusion

Understanding business cycle transmission across countries is one of the central questions in international macroeconomics. In this paper, we used new data and a quantitative model to assess how shocks are transmitted internationally through firms that operate in multiple countries. Our empirical results demonstrate important interdependence between source countries and their foreign affiliates. This interdependence is detectable both at the firm and the source-destination level. We use a quantitative model to interpret these findings and to evaluate the role of multinationals for international business cycle comovement.

All foreign multinationals together account for a large share of total output, and thus the rest of the world is responsible for about 10% of the productivity shocks in an average country. On the other hand, bilateral multinational production shares tend to be small, limiting the contribution of multinationals for observed comovement between individual country pairs. In the benchmark parameterization, eliminating barriers to multinational production decreases the cross-country standard deviation in growth rates by over 30 percent, indicating that international comovement may become significantly stronger as the share of multinationals in the world economy increases.

References

Acharya, Viral V and Philipp Schnabl, "Do Global Banks Spread Global Imbalances? Asset-Backed Commercial Paper during the Financial Crisis of 2007-09," *IMF Economic Review*, August 2010, *58* (1), 37–73.

Alfaro, Laura and Maggie Xiaoyang Chen, "Surviving the Global Financial Crisis: Foreign Ownership and Establishment Performance," *American Economic Journal: Economic Policy*, August 2012, *4* (3), 30–55.

Alviarez, Vanessa, "Multinational Production and Comparative Advantage," December 2013. mimeo, University of Michigan.

Antràs, Pol and Stephen R. Yeaple, "Multinational Firms and the Structure of International Trade," in Elhanan Helpman, Kenneth Rogoff, and Gita Gopinath, eds., *Handbook of International Economics*, Vol. 4, Elsevier, 2014, chapter 2, pp. 55–130.

Arkolakis, Costas and Ananth Ramanarayanan, "Vertical Specialization and International Business Cycle Synchronization," *Scandinavian Journal of Economics*, December 2009, *111* (4), 655–680.

_ , Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen R. Yeaple, "Innovation and Production in the Global Economy," December 2014. NBER Working Paper No. 18792.

Atalay, Enghin, Ali Hortasu, and Chad Syverson, "Vertical Integration and Input Flows," *American Economic Review*, April 2014, *104* (4), 1120–48.

Backus, David K., Patrick J. Kehoe, and Finn E. Kydland, "International Business Cycles: Theory and Evidence," in Thomas Cooley, ed., *Frontiers of business cycle research*, Princeton: Princeton University Press, 1995, pp. 331–356.

Becker, Sascha O., Karolina Ekholm, and Marc-Andreas Muendler, "Offshoring and the onshore composition of tasks and skills," *Journal of International Economics*, 2013, *90* (1), 91 – 106.

Bilir, L. Kamran and Eduardo Morales, "The Impact of Innovation in the Multinational Firm," April 2016. NBER Working Paper No. 22160.

Boehm, Christoph, Aaron Flaaen, and Nitya Pandalai-Nayar, "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tohoku Earthquake," October 2014. mimeo, University of Michigan.

_ , _ , and _ , "Multinationals, Offshoring and the Decline of U.S. Manufacturing," December 2015. mimeo, University of Michigan.

Broda, Christian and David Weinstein, "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, May 2006, *121* (2), 541–85.

Buch, Claudia M. and Alexander Lipponer, "Business Cycles and FDI: Evidence from German Sectoral Data," *Review of World Economics/Weltwirtschaftliches Archiv*, December 2005, *141* (4), 732–759.

Budd, John W., Jozef Konings, and Matthew J. Slaughter, "Wages and International Rent Sharing in Multinational Firms," *Review of Economics and Statistics*, February 2005, *87* (1), 73–84.

Burstein, Ariel T. and Alexander Monge-Naranjo, "Foreign Know-How, Firm Control, and the Income of Developing Countries," *Quarterly Journal of Economics*, February 2009, *124* (1), 149–195.

_ , Christopher Kurz, and Linda L. Tesar, "Trade, Production Sharing, and the International Transmission of Business Cycles," *Journal of Monetary Economics*, 2008, *55*, 775–795.

Caselli, Francesco, Miklós Koren, Milan Lisicky, and Silvana Tenreyro, "Diversification through Trade," August 2015. NBER Working Paper No. 21498.

Cetorelli, Nicola and Linda S Goldberg, "Global Banks and International Shock Transmission: Evidence from the Crisis," *IMF Economic Review*, April 2011, *59* (1), 41–76.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual*, 2013, *27* (1), 1 – 56.

Contessi, Silvio, "How Does Multinational Production Change International Comovement?," October 2010. mimeo, Federal Reserve Bank of St. Louis.

Davis, Steven J., John C. Haltiwanger, and Scott Schuh, *Job creation and destruction*, Cambridge and London: MIT Press, 1996.

Desai, Mihir A. and C. Fritz Foley, "The Comovement of Returns and Investment within the Multinational Firm," *NBER International Seminar on Macroeconomics*, 2006, pp. 197–237.

— , — , and James R. Hines, "Domestic Effects of the Foreign Activities of US Multinationals," *American Economic Journal: Economic Policy*, February 2009, *1* (1), 181–203.

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Méjean, "Firms, Destinations, and Aggregate Fluctuations," *Econometrica*, July 2014, *82* (4), 1303–1340.

Fillat, José L. and Stefania Garetto, "Risk, Returns, and Multinational Production," *The Quarterly Journal of Economics*, November 2015, *130* (4), 2027–2073.

Fillat, José L., Stefania Garetto, and Lindsay Oldenski, "Diversification, cost structure, and the risk premium of multinational corporations," *Journal of International Economics*, May 2015, *96* (1), 37–54.

Fons-Rosen, Christian, Sebnem Kalemli-Ozcan, Bent E. Sørensen, Carolina Villegas-Sanchez, and Vadym Volosovych, "Quantifying Productivity Gains from Foreign Investment," Mar 2013. NBER Working Paper No. 18920.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman, "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, June 1988, *78* (3), 402–17.

Guadalupe, Maria, Olga Kuzmina, and Catherine Thomas, "Innovation and Foreign Ownership," *American Economic Review*, December 2012, *102* (7), 3594–3627.

Helpman, Elhanan, "A Simple Theory of International Trade with Multinational Corporations," *Journal of Political Economy*, June 1984, *92* (3), 451–71.

— , Marc J. Melitz, and Stephen R. Yeaple, "Export versus FDI with Heterogeneous Firms," *American Economic Review*, March 2004, *94* (1), 300–316.

Irarrazabal, Alfonso, Andreas Moxnes, and Luca David Opromolla, "The Margins of Multinational Production and the Role of Intrafirm Trade," *Journal of Political Economy*, February 2013, *121* (1), 74–126.

Javorcik, Beata Smarzynska, "Does Foreign Direct Investment Increase the Productivity of Domestic Firms? In Search of Spillovers Through Backward Linkages," *American Economic Review*, June 2004, *94* (3), 605–627.

Johnson, Robert C., "Trade in Intermediate Inputs and Business Cycle Comovement," *American Economic Journal: Macroeconomics*, October 2014, *6* (4), 39–83.

Kalemli-Ozcan, Sebnem, Elias Papaioannou, and José-Luis Peydró, "Financial Regulation, Globalization and Synchronization of Economic Activity," *Journal of Finance*, June 2013, *68* (3), 1179–1228.

Keller, Wolfgang and Stephen Ross Yeaple, "The Gravity of Knowledge," *American Economic Review*, June 2013, *103* (4), 1414–44.

Kleinert, Jörn, Julien Martin, and Farid Toubal, "The Few Leading the Many: Foreign Affiliates and Business Cycle Comovement," *American Economic Journal: Macroeconomics*, October 2015, *7* (4), 134–159.

Kose, M. Ayhan and Kei-Mu Yi, "International Trade and Business Cycles: Is Vertical Specialization the Missing Link?," *American Economic Review Papers and Proceedings*, 2001, *91* (2), 371–75.

Kovak, Brian, Lindsay Oldenski, and Nicholas Sly, "The Labor Market Effects of Offshoring by U.S. Multinational Firms: Evidence from Changes in Global Tax Policies," *mimeo, Carnegie-Mellon, Georgetown, and Federal Reserve Bank of Kansas City*, June 2015.

Markusen, James R., "Multinationals, multi-plant economies, and the gains from trade," *Journal of International Economics*, May 1984, *16* (3-4), 205–226.

McGrattan, Ellen R. and Edward C. Prescott, "Openness, technology capital, and development," *Journal of Economic Theory*, November 2009, *144* (6), 2454–2476.

_ and _ , "Technology Capital and the US Current Account," *American Economic Review*, September 2010, *100* (4), 1493–1522.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, November 2003, *71* (6), 1695–1725.

Menno, Dominik, "Multinational Firms and Business Cycle Transmission," November 2014. mimeo, EUI.

Ramondo, Natalia, "A quantitative approach to multinational production," *Journal of International Economics*, May 2014, *93* (1), 108 – 122.

_ and Andrés Rodríguez-Clare, "Trade, Multinational Production, and the Gains from Openness," *Journal of Political Economy*, 2013, *121* (2), 273 – 322.

_ and Veronica Rappoport, "The Role of Multinational Production in a Risky Environment," *Journal of International Economics*, July 2010, *81* (2), 240–252.

_ , _ , and Kim J. Ruhl, "The Proximity-Concentration Tradeoff under Uncertainty," *Review of Economic Studies*, 2013, *80* (4), 1582–1621.

_ , _ , and _ , "Intrafirm Trade and Vertical Fragmentation in U.S. Multinational Corporations," *Journal of International Economics*, January 2016, *98* (1), 51–59.

Schnabl, Philipp, "The International Transmission of Bank Liquidity Shocks: Evidence from an Emerging Market," *Journal of Finance*, June 2012, *67* (3), 897–932.

Stockman, Alan C., "Sectoral and National Aggregate Disturbances to Industrial Output in Seven European Countries," *Journal of Monetary Economics*, 1988, *21*, 387–409.

Tintelnot, Felix, "Global Production with Export Platforms," April 2016. Forthcoming, *Quarterly Journal Economics*.

Zlate, Andrei, "Offshore Production and Business Cycle Dynamics with Heterogeneous Firms," *Journal of International Economics*, May 2016, *100*, 34–49.

Table A1: Sample and summary statistics

Country	Number of Firms	Number of Foreign Multinationals	Correlation between ORBIS growth and GDP growth	Ratio of ORBIS revenue to total revenue	Country	Number of Firms	Number of Foreign Multination-als	Correlation between ORBIS growth and GDP growth	Ratio of ORBIS revenue to total revenue
Austria	15299	2202	0.83	0.63	Lithuania	7473	631	0.96	0.53
Australia	766	208	0.60		Latvia	43887	1093	0.91	0.59
Belgium	18360	3603	0.91	0.70	Mexico	6102	485	0.49	0.93
Bulgaria	120520	1444	0.92	0.71	Netherlands	10061	2163	0.81	0.40
Czech Republic	85422	7007	0.86	0.81	Norway	148599	3708	0.80	0.81
Germany	224394	10010	0.89	0.69	Poland	56414	6780	0.82	0.68
Estonia	47131	1537	0.96	0.71	Portugal	212761	2047	0.89	0.93
Spain	519129	9034	0.82	1.07	Romania	319347	4700	0.86	0.55
Finland	106222	2301	0.93	0.93	Serbia	48083	2428	0.62	0.74
France	751858	14581	0.96	0.81	Sweden	222882	3942	0.79	0.93
United Kingdom	194711	22459	0.59	0.69	Singapore	1249	351	0.64	
Greece	24639	1262	0.74	0.54	Slovenia	29868	559	0.90	0.77
Croatia	60527	2292	0.96	0.75	Slovak Rep.	30377	3004	0.75	0.88
Hungary	174795	822	0.99	0.76	Turkey	7975	286	0.77	
Ireland	14131	2579	0.56	1.03	Ukraine	218489	2489	0.79	0.80
Italy	556756	12630	0.96	0.79	United States	97378	605	0.84	0.09
Japan	217024	282	0.81	0.84	**Mean**	137879	3828	0.81	0.74
Korea, Rep.	95111	598	0.68	0.78	**Median**	72974	2247	0.83	0.76

Notes: This table reports the sample of countries used in the analysis. It reports the total number of firms and total number of foreign multinational affiliates in the average year in each country, the correlation between the growth rates of aggregate sales in ORBIS and GDP growth over the period for which ORBIS data are available (2004-2012), and the ratio of combined sales in ORBIS to total gross output reported in EUROSTAT (for EU countries) or UN SNA data (for non-EU countries).