Gravity and the Organization of Firm-to-Firm Production Networks *

Gustavo González, Luca Macedoni, Guzmán Ourens, and Brian Pustilnik

July 9, 2026

Abstract

We study the extent to which gravity explains the organization of firm-to-firm production networks. Using detailed transaction data from Chile, we estimate gravity equations separately for narrowly defined industries and show that gravity explains a large share of the observed variation in buyer-supplier relationships. Gravity coefficients differ substantially across industries but remain remarkably stable over time, even after the disruption caused by the COVID-19 pandemic. Links that conform more closely to gravity are more likely to survive, and newly formed links closely resemble those destroyed. We further show that industry-specific gravity coefficients estimated in Chile successfully predict production networks in Denmark, suggesting convergence in within-industry organization. Together, these findings suggest that gravity captures a persistent and transferable component of production-network organization. Recovering network topology and industry-specific heterogeneity proves sufficient to predict economically meaningful firm-level networks where transaction-level data are unavailable.

Keywords: Production Networks, Input-Output, Linkages, Gravity.
JEL Codes: F10, R10, R40.

*Acknowledgments: We are grateful for the constructive comments received from Gustavo de Souza, Yoto Yotov, Maurizio Zanardi and participants at LACEA and ETSG. This work greatly benefited from the amazing research assistance of Jose Elias Rishmawi. The views expressed are those of the authors and do not necessarily represent the views of the Central Bank of Chile (CBC) or its board members. All errors are our own. González: PUC Chile; gggonzal@uc.cl. Macedoni: Department of Economics, Management, and Quantitative Methods-University of Milan; luca.macedoni@unimi.it. Ourens: Department of Economics-Tilburg University; g.ourens@uvt.nl. Pustilnik: Banco Central de Chile; bpustilnik@bcentral.cl

1 Introduction

Economic production relies on a wide array of specialized firms connected by buyer-supplier relationships. These production networks play a central role in determining how productivity gains diffuse across firms, how disruptions propagate through the economy, and how policies affect aggregate outcomes. As a result, firm-level production networks have become an important object of study in international trade, macroeconomics, and industrial organization.[1] Despite this growing literature, relatively little is known about the empirical regularities governing how these networks are organized. While recent work has documented many of their consequences, much less is understood about the forces shaping the formation and persistence of the underlying firm-to-firm linkages.[2]

This paper studies the extent to which gravity can explain the organization of production networks. Gravity equations have long been one of the most successful empirical tools in economics, explaining trade flows between countries through simple measures of economic size and geographic distance. We ask whether similar forces also organize transactions between firms. Because firms are much more specialized than countries, the adequate unit of analysis must forcibly be much more disaggregated. Firms do not potentially trade with every other firm: a buyer can only source from firms producing the specific inputs it demands, making the set of feasible trading partners inherently product-specific. Consequently, gravity must be studied at the product level, where the set of feasible trading relationships is properly defined. Using detailed firm-to-firm transaction data from Chile, we estimate gravity equations separately for each product and evaluate how well they explain the structure of production networks.

Three findings emerge. First, gravity explains a large share of the observed variation in firm-to-firm transactions, highlighting gravity as a quantitatively important force in shaping the production network. At the same time, we document substantial heterogeneity in gravity coefficients across products, indicating that the importance of size and distance differs systematically across industries. We characterize this heterogeneity and show that the importance of distance is greater in industries with higher value added shares and lower capital intensity, suggesting that gravity reflects meaningful differences in how production is organized across sectors. These findings imply that production networks are fundamentally industry-specific: differences in production technologies, input requirements, and logistics generate distinct network structures, making the aggregate produc-

[1] Early examples of analyses that incorporate inter-sectoral linkages include Acemoglu et al. (2012), Jones (2013), and Caliendo and Parro (2015). More recent developments, considering firm-to-firm connections, can be found in di Giovanni et al. (2014), Bernard et al. (2019), Dhyne et al. (2021), Carvalho et al. (2021), and Tascherau-Dumouchel (2025).

[2] Much of the existing literature treats the network as exogenous. A notable exception is Arkolakis et al. (2023), who show that production networks adjust to shocks, with important consequences for welfare analysis.

tion network better understood as the superposition of industry-specific sub-networks rather than a single homogeneous network.

Second, gravity proves to be remarkably stable over time. Estimated coefficients change little before, during, and after the COVID-19 pandemic despite the large disruption experienced by production networks. Moreover, links that are more closely aligned with gravity are less likely to disappear, and when destroyed links are replaced following the shock, newly formed relationships exhibit gravity patterns very similar to those observed before the disruption. This suggests that large disruptions primarily reshape individual trading relationships rather than the underlying organization of production. Gravity therefore captures a persistent component of production networks, reflecting stable industry characteristics rather than transitory features of the observed network.

Third, we show that industry-specific gravity coefficients estimated in one country can successfully predict production networks in another. We combine gravity estimates obtained from Chile with administrative data from Denmark that identify firms' size, location, purchases, and sales but not their trading partners. We construct a predicted network of linkages for each industry in Denmark and assess the validity of our predicted network through several approaches. First, we evaluate how well the predicted network captures the correlation between changes in firms' purchases of specific products and the sales of their predicted suppliers of those products. Second, we analyze whether foreign shocks propagate through the Danish economy via our predicted network. For each firm, we calculate a direct exposure measure, capturing how global demand affects the firm through its exported products, and an indirect exposure measure, reflecting its vulnerability to global shocks through its sales to directly exposed firms. We demonstrate that the indirect exposure measure significantly explains the sales performance of non-directly exposed firms. In our last exercise, we construct measures of industry-level shocks as deviations from the long-term trend in total economic activity that each industry exhibits. We then calculate firm-specific shocks as the weighted average of these industry-level shocks, where the weights reflect the importance of each industry to a firm's output. We show that these firm-specific shocks also propagate to non-directly exposed firms through our predicted network. Collectively, these exercises provide robust validation of our out-of-sample predictions, and suggests that the relationship between gravity and production-network structure is not country specific. Instead, industries exhibit common organizational regularities that persist across economies despite substantial differences in geography, income levels, and industrial composition. Our robustness exercises further show that successful prediction depends primarily on recovering the topology of the network and the heterogeneity of gravity relationships across industries. Once these features are captured, moderate estimation error in the estimated gravity coefficients has only a limited impact on the economic content of the predicted network.

Taken together, these results indicate that production networks are considerably more predictable than their apparent complexity might suggest. While individual supplier relationships are shaped by many idiosyncratic considerations, a substantial fraction of network structure can be explained by a small set of gravity forces that are stable across time and remarkably similar across countries within narrowly defined industries. Rather than viewing observed production networks as entirely country-specific objects, our findings suggest that industries possess persistent organizational patterns that can be exploited to understand and predict firm-to-firm linkages.

Our paper contributes to three strands of the literature. First, we contribute to the growing literature studying firm-to-firm production networks and their endogenous formation. Recent work has emphasized that buyer-supplier relationships are the outcome of firms' endogenous sourcing decisions and that accounting for network adjustments is important for understanding the propagation of shocks, productivity, and welfare (Oberfield, 2018, Eaton et al., 2022, Arkolakis et al., 2023, Miyauchi, 2024, Alviarez et al., 2023, Boehm et al., 2024, Fontaine et al., 2024, and Kortum and Ribeiro, 2025). Our contribution is complementary. Rather than modelling the mechanisms underlying network formation, we document that gravity provides a remarkably stable empirical characterization of firm-to-firm linkages across industries, over time, and across countries. These regularities provide a simple and empirically disciplined way of predicting production networks in settings where transaction-level data are unavailable.

Second, we contribute to the literature on the dynamics of production networks. Previous work documents substantial persistence in buyer-supplier relationships despite continuous link creation and destruction (Lim, 2018, Dhyne et al., 2021, Huneeus, 2020, and Khanna et al., 2022), while other studies show that these linkages are central to the transmission of shocks across firms (Carvalho et al., 2021). Our results provide an empirical characterization of this persistence. We show that links more closely aligned with gravity survive longer and that network reorganization following a major disruption largely preserves the same gravity structure observed before the shock. This suggests that gravity is not merely a cross-sectional regularity but also an important feature of the dynamics of production networks.

Third, our paper relates to the literature on technological convergence across countries. Existing work has documented convergence in production technologies and productivity within industries (Comin and Hobijn, 2010, Madsen and Timol, 2011, or Rodrik, 2013). Our results suggest that convergence extends beyond production technologies inside firms to the organization of production around them. The success of Chilean gravity coefficients in predicting firm-to-firm linkages in Denmark indicates that firms operating in the same industries organize their supplier relationships according to similar principles despite substantial differences in national environments.

Finally, our findings have practical implications for empirical research. Detailed transaction-level data remain available for only a small number of countries. We make publicly available the full set of industry-specific gravity coefficients together with the methodology used to translate them into predicted firm-to-firm linkages, providing researchers with a straightforward framework for predicting production networks in any setting where information on firms' size, location, and product-level activity is available. This substantially expands the range of environments in which firm-level production networks can be incorporated into analyses of productivity, distortions, and shock propagation that increasingly rely on production-network information (e.g., Baqaee and Farhi, 2019; Bigio and La'o, 2020; Demir et al., 2022; Carvalho et al., 2021). The importance of being able to observe firm-to-firm input-output networks cannot be overstated, as it offers substantial advantages over analysis based on sector-level connections. Firm-level networks capture the considerable heterogeneity in firms' positions within the production network, allowing researchers to identify key suppliers, bottlenecks, and vulnerable nodes that would otherwise remain hidden. In addition, because firm-level networks incorporate firms' geographic locations, they also make it possible to study the spatial dimension of production networks and the role of geography in regional disparities and shock propagation. More generally, replacing sector-level relationships with firm-level linkages enables a much finer characterization of production networks, improving both positive analyses of productivity, distortions, and shock transmission, and the design of industrial and regional policies targeted at the firms where interventions are likely to be most effective.

The remainder of the paper is structured as follows. Section 2 outlines the methodology used to estimate gravity coefficients. Section 3 provides an overview of the data and highlights key summary statistics. Section 4 presents findings on the role of gravity in explaining input-output networks. We first document significant heterogeneity in the importance of each gravity dimension across sectors. Next, we assess the role of gravity in explaining networks in the cross-section. Finally, we examine how gravity influences network persistence and its response to major shocks. Section 5 details our proposed method for using industry-specific gravity coefficients to predict out-of-sample networks and presents results from various tests that validate the predicted network. Section 6 concludes.

2 Empirical Gravity Framework

The gravity equation was first introduced in the economic literature to explain trade flows between countries (Tinbergen, 1962), and the methods for estimating it have been perfected since then by multiple efforts. Compared to countries, the production network at the firm level is much sparser, as firms produce and demand a significantly smaller range of products. Most firm pairs are not feasible trading partners regardless of their size or proximity. Estimating gravity at the aggregate firm level would therefore mix technologically feasible and infeasible matches, obscuring the role of gravity. Moreover, firm-level output supply and input demand are influenced by technological constraints and other industry-specific factors (Eaton et al., 2022, Miyauchi, 2024, Fontaine et al., 2024). To address these issues and account for the sparsity of the network, we estimate gravity equations per product. This means that for each product p we consider the network conformed only by firms buying or selling p at each moment in time. This drastically reduces the number of firm pairs to be considered in each regression, allowing us to isolate the forces governing supplier selection among feasible partners.

Following the structural gravity literature (Head and Mayer, 2014), for each time period t and product p, we estimate the following equation:

$$f_{ijpt} = \beta_{pt} + \beta_{A,pt} log(A_{jpt}) + \beta_{S,pt} log(S_{ipt}) + \beta_{D,pt} log(D_{ijt}) + \epsilon_{ijpt} \tag{1}$$

where f_{ijpt} is a measure of the value sold by firm i to firm j in product p, A_{jpt} is the total value acquired of product p by firm j, S_{ipt} is the total value sold of product p by firm i, and D_{ijt} is the geographical distance between both firms at time t. Equation (1) should be interpreted as describing how transaction values are allocated across feasible buyer-supplier pairs within a product. Conditional on producing and demanding the same product, we expect larger buyers to purchase more, larger suppliers to sell more, and geographically closer firms to transact more. The estimated coefficients therefore summarize the relative importance of market size and spatial frictions in organizing firm-to-firm trade within each industry.

Following the literature, we estimate the above equation using different methods: OLS and PPML. The latter is the standard method for estimating gravity equations, as it controls for heteroscedasticity in trade data and better accounts for sparsity (Santos Silva and Tenreyro, 2006). These issues are common in cross-country trade data and even more pronounced in firm-to-firm trade data. When estimating with OLS, the dependent variable is the log of the transaction value. In contrast, when using PPML, the dependent variable is the transaction value in levels. Unlike the standard gravity literature, our objective is not to recover structural trade elasticities but to characterize the empirical organization of firm-to-firm transactions within products. Accordingly, we report both baseline specifications and versions including buyer and seller fixed effects.

Because production networks are inherently sparse, characterizing their organization requires understanding not only the intensity of existing buyer-supplier relationships but also the process by which links form. Following the distinction between the extensive and intensive margins in the gravity literature (e.g., Helpman et al., 2008), we therefore estimate an analogous specification in which the dependent variable equals one whenever a positive transaction is observed and zero otherwise. We estimate these extensive-margin specifications using both PPML and probit.

The resulting set of coefficients $[\hat{\beta}, \hat{\beta}_A, \hat{\beta}_S, \hat{\beta}_D,]_{pt}$, estimated for 2018 under each procedure, is made publicly available as supplementary material. As we show in Section 5, these estimates can be combined with information on firms' size and location to predict firm-to-firm production networks in settings where transaction-level data are unavailable.

3 Data

3.1 Chile

Widely regarded as one of the better performing economies in Latin-America, Chile is a middle-income economy that belongs to the OECD and is well integrated to international markets with good exports mainly concentrated in minerals and agriculture. Our primary data source is the Electronic Invoice from the Chilean Internal Revenue Service (SII, by its Spanish acronym), which records the monthly sum of all bilateral transactions between any pair of firms per product from 2014 onward. Records include all purchases and sales involving a manufacturing firm. We consider domestic transactions at the product level, defined in the domestic product classification (CUP-270). For confidentiality reasons we can only access data for transactions in product-codes where more than 24 firms interact, a threshold that allows us to observe transactions in about 244 products (out of 292 possible) of which 164 are manufacturing products and the reminder are services.[3]

SII records the exact address of each firm's headquarters, which we use to compute bilateral distances between all firms. Table A.1 presents the main statistics describing the data. It shows that multi-product firms are the norm, with the median firm selling five distinct products. Most firms operate from a single location (branch), with fewer than 10% of firms producing in multiple locations. Importantly, firms are widely interconnected: the median firm supplies to four buyers, while the median number of suppliers to a firm is six. The table also reports statistics on the distribution of the (log) values of connections. All variables exhibit a relatively balanced distribution, with trade, demand

[3]This study was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in economic and financial affairs of its competence. The CBC has access to anonymized information from various public and private entities, by virtue of collaboration agreements signed with these institutions. To secure the privacy of workers and firms, the CBC mandates that the development, extraction and publication of the results should not allow the identification, directly or indirectly, of natural or legal persons. Officials of the CBC processed the disaggregated data. All the analysis was implemented by the authors and did not involve nor compromise the SII, Aduanas and AFC. The information contained in the databases of the Chilean IRS is of a tax nature originating in self-declarations of taxpayers presented to the Service; therefore, the veracity of the data is not the responsibility of the Service.

size, and bilateral distances showing a slight right tail, while the size of selling nodes has a light left tail.

3.2 Denmark

The economy in Denmark is also well integrated in international markets and its size is comparable to that in Chile. Nevertheless, it features many important differences: the population is about 3 times lower, the average income is about 3 times higher and its specialization pattern is mainly dominated by Services and Chemicals. The intrinsic difference between the two economies are enough to make our network predictions in Denmark less than obviously successful.

We use three main datasets that include data on total purchases, production, and international trade flows at the firm-product level. This section describes these datasets in detail. In the *Manufacturers' Purchases of Goods survey* (VARK, henceforth referred to as the "purchases survey"), firms are asked to report both quantities and values for each product they procure. The survey includes all product purchases, covering both domestic and imported items, regardless of origin. The data on input values are classified at the national six-digit level, with the first four digits aligning with the four-digit CN classification. Purchases of machinery (capital goods) and energy (e.g., electricity and gas) are excluded. The survey encompasses all firms in the 'Mining and Quarrying' and 'Manufacturing' sectors with at least 50 employees. However, the data does not provide information on suppliers' identities or locations, so it does not detail the linkages between firms.

In the *Manufacturers' Sales of Goods survey* (VARS; "sales survey" henceforth), firms report the quantities and values of each product they produce. Sales are documented independently of the market in which the product is sold and, therefore, include both domestic and export sales. Firms with any manufacturing activity and a minimum of 10 employees are required to participate in this survey. We aggregate firm sales at the four-digit CN classification for consistency with the purchases survey.

In the *International Trade in Goods statistics* (UHDI; "trade statistics" henceforth), firms report their exports and imports by product and destination. We aggregate these flows at the four-digit level of the CN classification to merge the dataset with the previous two. This allows us to disentangle domestic from international transactions. The reporting unit is the firm (CVRNR).[4]

[4]The data is derived from two sources: Intrastat and Extrastat. Intrastat covers trade with other EU countries and is based on data reported by Danish enterprises with annual exports exceeding DKK 6 million (approximately EUR 0.8 million) in 2015. The reporting threshold varies by year and is set annually to capture 97% of total exports. Extrastat covers trade with non-EU countries and is based on customs and supply data collected by Danish tax authorities.

Using firm (and product) identifiers, we can link the three datasets. For any given year, we calculate domestic sales by taking the difference between the total sales of a firm-product in the sales survey and the total exports of the same firm-product in the trade statistics. Likewise, we determine domestic purchases by subtracting the total imports of a firm-product in the trade statistics from the total purchases of the same firm-product in the purchases survey.[5]

Using the firm identifiers, we incorporate additional information from the general firm statistics (FIRM), such as firm location. Specifically, these statistics provide information only on the *kommune* of each firm, which corresponds roughly to a municipal area, but do not include the firm's exact address. In Denmark, there are 98 municipal areas. We approximate the firm's location by using the center of the municipal area, allowing us to compute all possible bilateral distances between firms based on this central point. The coordinates of the municipal areas are kindly provided by Chan et al. (2024). As a result, for each product traded we can observe the topology of the network, i.e., the size and location for all receiver and sender nodes, but not the networks themselves as we do not observe the connections between the nodes.

To predict the network of firm-to-firm trade using Danish data, we retain, for each CN four-digit product, only those firms that either sell it domestically (i.e., have positive domestic sales) or purchase it domestically (i.e., have positive domestic purchases). We restrict our analysis to products with at least one buyer and one seller. For confidentiality, we exclude all products with fewer than five firms in the network.

Given the sample restrictions mentioned, we count on a panel of about three thousand distinct firms over the period 2000-2018. Firms are concentrated around a few important cities. On averages, there are 79 firms per municipality, with a standard deviation of 60. The largest municipalities are Copenhagen with 331 firms, Aarhus (312), and Odense (224). Table A.2 provides statistics on the observed nodes. Comparing these to the figures in Table A.1 reveals that our sample of Danish firms are somewhat smaller and more narrow in scope compared to Chilean firms. In fact, single-product firms appear as the norm in our sample. Table A.3 provides further details on the distribution of activity by broad sector of the economy. It shows that the Wood Products, Foodstuffs, and Chemical and Mineral Products sectors have the highest number of trading firms in our time period. Firms in these sectors are also larger, both in terms of sales and the number of distinct purchases, although the average value of purchases is not necessarily higher in these sectors. Finally, Figure A.2 shows that the number of manufacturing

[5] If the calculated domestic sales or purchases are negative, due to misreporting or other firm-level activities such as Carry-Along-Trade, we set them equal to zero.

firms in Denmark underwent a steady decline over the period 2000-2013, and has since recovered only partially.[6]

3.3 Matching process for codes

We develop a cross-walk from the product classification used in the Chilean data (CUP) to the Combined Nomenclature (CN) used in the Danish data, so that for each four-digit CN code there is a corresponding CUP code.[7] In our Danish data, we only keep products for which we have obtained gravity coefficients using the Chilean data.

The four-digit CN classification contains about 1.2 thousand codes, while the CUP contains 169 manufacturing codes (and we obtain gravity coefficients for 164 of them). This is a potential source of measurement error for our predicted links in cases where true gravity coefficients vary a lot for different CN-codes matched to the same CUP. Table A.4 provides a summary of the correspondence we use. The table shows that many CUP-codes are matched one to one with a CN-code, and in some cases where the match is one CUP to many CN's, the category suggests gravity coefficients should be similar. This is the case for example for CUP 6 *Other Cereals*, as logistics around cereals should result in homogeneous gravity coefficients for all of them. Of course, other CUP codes matched with several CN-codes are expected to be much more heterogeneous. For example, CUP 150 *Iron and steel products*, matched with 42 CN codes, contains a wide array of different goods for which gravity determinants may vary. On the aggregate, we find that the number of CN-codes matched to a CUP code presents a significantly negative correlation to the R^2 of the gravity regression in Chile for that CUP.

4 Gravity and the Organization of Production Networks

The empirical exercises in this section assess the extent to which production networks reflect systematic organizational forces rather than idiosyncratic buyer-supplier relationships. We begin by documenting substantial heterogeneity across industries, then examine the explanatory power of gravity in the cross section, and finally study whether gravity also explains the persistence and resilience of production networks over time.

[6] The decline of the number of firms in manufacturing is a well-studied phenomenon and linked to de-industrialization. See, for instance, Bernard et al. (2017).

[7] The full correspondence is provided as supplementary material to this paper.

4.1 Sectoral Heterogeneity

While gravity is expected to shape transactions in all industries, there is little reason to expect a common gravity relationship across them. Industries differ in the extent to which production relies on standardized or highly specialized inputs, in the importance of transportation costs, and in the geographic concentration of suppliers. If gravity captures economically meaningful aspects of production organization, its coefficients should therefore vary systematically across industries rather than representing a single economy-wide relationship.

To show this, we focus on the PPML estimations performed by product, excluding buyer and seller fixed effect. This provides an accurate measurement of the relative importance of each gravity dimension to explain transactions in each product. Figure 1 shows the distribution of each gravity coefficient across industries: $log(\widehat{S}_{jp})$ in panel (a), $log(\widehat{A}_{ip})$ in panel (b), and $log(\widehat{D}_{ij})$ in panel (c). Results show the heterogeneity of the estimated elasticity of the value traded to each gravity dimension. Additionally, we include a histogram showing the distribution of the R^2 values from the regression (panel d), which serves as a measure of the overall importance of gravity in explaining trade flows across industries.

As expected, the size of the nodes impact flows positively, while the distance between them affects flows negatively. This holds for all estimated coefficients, except in three industries, where the coefficients for distance are not significantly different from zero. Interestingly, the distribution of coefficients peaks at a value of one for both size of sales $log(\widehat{S}_{jp})$ and size of purchases $log(\widehat{A}_{ip})$. This result closely mirrors findings from the gravity equation literature in international trade, where the estimated coefficients on the GDP of the origin and destination countries typically cluster around one (Head and Mayer, 2014). Analogous to country-level gravity, our results suggest that trade flows between two firms are proportional to the product of their sizes.

The table reveals considerable heterogeneity in the coefficients. The dispersion in estimated coefficients indicates that industries differ not only in the intensity of trade but also in the mechanisms governing supplier selection. In some industries, supplier size dominates location, reflecting highly standardized markets with broad sourcing possibilities. In others, geographic proximity becomes considerably more important, suggesting stronger transport costs, coordination requirements, or localized production systems. Specifically, we find that the size of the supplier and the buyer are of particular importance in *Fish Oil* or *Liquid Natural Gas*, whereas these coefficient are notably lower in most services. We observe the negative impact of distance on transactions to be larger in industries where the size of buyer or seller is smaller and vice-versa. In general, we find that the importance of distance is negatively correlated to the capital intensity, and positively correlated to the labour intensity and the value added share in the industry.

Figure 1: Distribution of gravity statistics (PPML estimations)

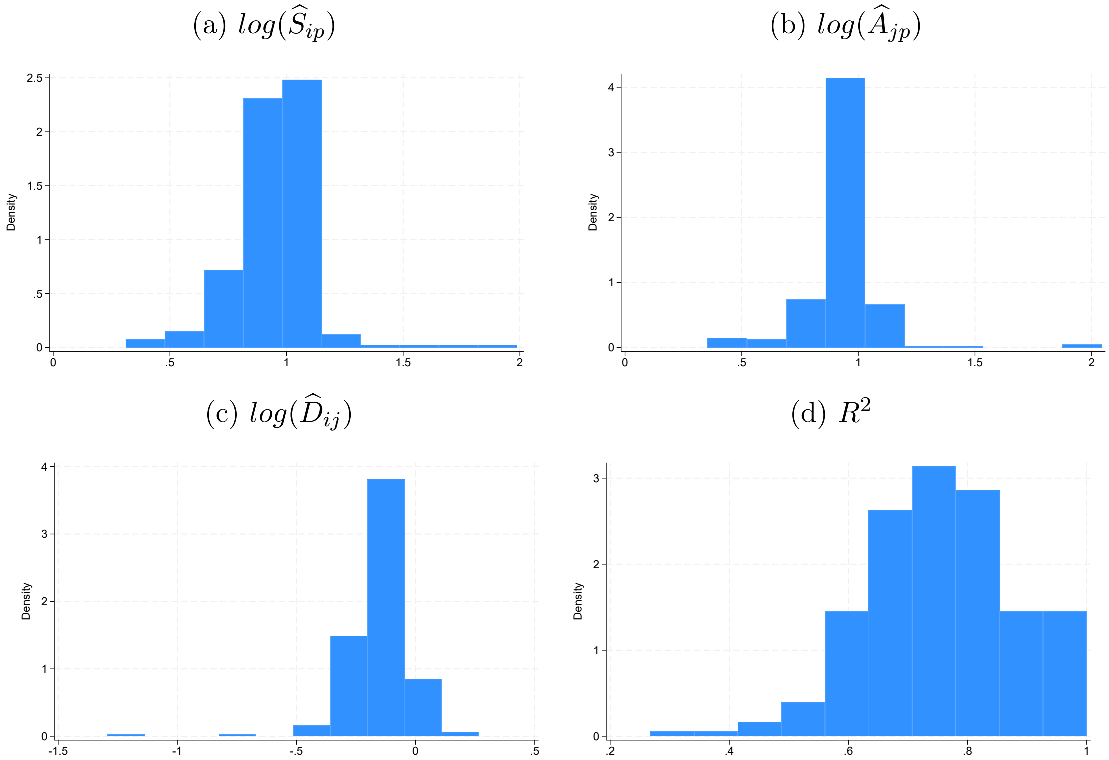

Notes: Histograms for the different statistics obtained by running PPML gravity estimations, without fixed effects, per product. A_{jp} is value acquired of product p by firm j, S_{ip} is value sold of product p by firm i, and D_{ij} is the geographical distance between both firms.

We also observe substantial heterogeneity in the overall explanatory power of gravity across industries. Even within narrowly defined product markets, gravity is not expected to fully determine buyer-supplier relationships. Firms may differ in quality, reliability, contractual arrangements, or other dimensions that generate matching frictions and idiosyncratic supplier choices beyond the basic gravity determinants.[8] As shown in panel (d), the explanatory power of gravity ranges from just above 30% for code 150 *Iron and Steel Products* to almost complete explanatory power for code 103 *Beer* and code 33 *Poultry*. A natural explanation for this variation is the degree of technological heterogeneity within product categories. Broad categories such as code 61 *Silver*, which includes both raw materials and refined jewelry, encompass production processes and supplier relationships that differ substantially, reducing the explanatory power of a single gravity equation. More narrowly defined products, by contrast, exhibit much more homogeneous trading patterns. Nevertheless, the average R^2 of 0.764 indicates that gravity explains most of the observed variation in firm-to-firm transactions for manufacturing industries, suggesting

[8]See, for example, Arkolakis et al. (2023), Eaton et al. (2022), Miyauchi (2024), and Fontaine et al. (2024).

that much of the apparent complexity of production networks reflects systematic rather than idiosyncratic forces.

Finally, we observe the coefficients to be largely persistent over time. For the PPML coefficient of distance, less than 7% of the variance comes from within-industry variation over time, with the rest being between-industry variation. Coefficients are also very much consistent across varieties and we typically observe coefficients to be similar across similar products. A good example of this is the great similarity between the coefficients for different types of non-exotic fruit, for example code 23 *Apples* and 24 *Pears, fresh quinces, and other pome fruit*).

4.2 Gravity in the Cross Section

While our analysis focuses on product-specific gravity relationships, pooled estimates provide a useful benchmark summarizing the average role of gravity across the economy and facilitating comparison with the international trade literature. Table 1 presents the results of intensive-margin gravity estimations for the entire pool of observations across multiple sectors in Chile from 2014 to 2023. To focus solely on the intensive margin, we run the regressions on a subsample of transactions with strictly positive trade values. These estimates provide insights into the economy-wide average importance of each gravity dimension for the intensity of transactions, conditional on the transaction occurring. Columns (1) and (2) report the results from OLS estimations, while columns (3) and (4) present the PPML estimates. The analysis focuses on results using monthly data from multi-establishment firms within the manufacturing sector. Results at the yearly frequency, using single-establishment firms and including non-manufacturing trading partners, are reported in the Appendix.

The results reaffirm the standard findings: the size of the nodes positively impacts flows, while the distance between them negatively affects flows (see columns 1 and 3). Specifically, we observe that flows are larger when buyers and suppliers are bigger, and when partners are geographically closer. Notably, the predictive power of these gravity estimations (as measured by R^2) is quite high: given a set of connections for the economy, the intensity of transactions predicted by the gravity model closely matches the actual intensity observed. It is important to emphasize that this conclusion holds even before introducing any additional controls into the specification. The limited increase in explanatory power after introducing extensive fixed effects (columns 2 and 4) suggests that a substantial fraction of the variation in observed transaction values is already captured by simple gravity fundamentals. In other words, gravity alone explains much of the organization of existing buyer-supplier relationships.

Table 1: Pooled gravity estimation, intensive margin (2014-2023)

	OLS		PPML	
	(1)	(2)	(3)	(4)
Dep, var.:	$\log(f_{ijp})$		f_{ijp}	
$\log(A_{jp})$	0.7178***	0.5835***	0.8607***	0.8534***
	(0.0001)	(0.0001)	(0.0000)	(0.0000)
$\log(S_{ip})$	0.0715***	0.2331***	0.0072***	0.0218***
	(0.0001)	(0.0001)	(0.0000)	(0.0000)
$\log(D_{ij})$	-0.0498***	-0.0200***	-0.0046***	-0.0018***
	(0.0001)	(0.0002)	(0.0000)	(0.0000)
Obs.	35169769	35126778	35169769	35126778
Product FE	-	Yes	-	Yes
Buyer FE	-	Yes	-	Yes
Seller FE	-	Yes	-	Yes
Month FE	-	Yes	-	Yes
R^2	0.7318	0.8264		
Pseudo R^2			0.0659	0.0750

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Regressions include all monthly flows among multi-establishment manufacturing firms, and exclude flows with zero values. f_{ijp} is value sold by firm i to firm j in product p, A_{jp} is value acquired of product p by firm j, S_{ip} is value sold of product p by firm i, and D_{ij} is the geographical distance between both firms.

A natural concern regarding our results is the potential measurement error in bilateral distances, as the Chilean data only provide information on the address of the firm headquarters, not each establishment. To assess the magnitude of this issue, we reproduce our results using a sample of single-establishment firms. The results, presented in the Appendix (Table A.9), are very similar to those reported here, indicating that this source of measurement error does not significantly affect our findings. The results also remain consistent when we include nodes (either buyers or sellers) outside the manufacturing industry.[9]

4.3 Gravity and Network Persistence

Production networks exhibit two seemingly contradictory features. On the one hand, buyer-supplier relationships display substantial turnover, with firms continuously creating and destroying links over time (Lim, 2018 or Huneeus, 2020). On the other hand, production networks exhibit remarkable aggregate persistence, with many of their key

[9]These results are omitted here to save space but are available upon request to the authors.

structural properties remaining stable despite ongoing link reallocation (Dhyne et al., 2021 or Khanna et al., 2022). An important question is therefore how these two observations can coexist. One possibility is that firms replace existing suppliers with new ones that satisfy similar economic criteria, preserving the overall organization of production despite substantial link turnover. If gravity captures these underlying organizational forces, its explanatory power should remain largely unchanged even after major disruptions that destroy a large number of buyer-supplier relationships.

The evidence presented so far establishes that gravity explains the cross-sectional organization of production networks. However, this evidence alone cannot distinguish between alternative interpretations of the observed regularities. Gravity may characterize a long-run equilibrium toward which production networks gradually converge after shocks, implying that its explanatory power should temporarily weaken following major disruptions. Alternatively, gravity may only describe a short-run configuration, with other forces gradually reshaping the network over time, in which case its explanatory power should strengthen after a shock. Finally, gravity may instead represent a persistent organizing force governing supplier selection throughout the adjustment process. Under this interpretation, gravity coefficients should remain broadly unchanged before, during, and after major disruptions.

We evaluate these competing hypotheses by exploiting the COVID-19 pandemic, which generated a large and sudden disruption to production networks in 2020. The pandemic severely affected the Chilean economy, leading to the worst economic downturn in over four decades, with a per capita GDP decline of 7.4% in 2020. However, the economy rebounded strongly in 2021, with a 10.6% increase in per capita GDP, and since 2022, is regarded as having returned to its pre-pandemic growth path. The pandemic destroyed many buyer-supplier relationships while leaving the underlying technologies and production requirements of industries largely unchanged, providing a natural setting to assess whether gravity governs not only the structure of production networks at a point in time but also their reorganization following a major shock.

Table 2: Pooled gravity OLS estimations, intensive margin, over time

	2018-2019		2020-2021		2022-2023	
	(1)	(2)	(3)	(4)	(5)	(6)
Dep. var.:	$\log(f_{ijp})$					
$\log(A_{jp})$	0.5481***	0.4653***	0.5586***	0.4729***	0.5598***	0.4649***
	(0.0004)	(0.0005)	(0.0004)	(0.0005)	(0.0004)	(0.0005)
$\log(S_{ip})$	0.1521***	0.2928***	0.1417***	0.3024***	0.1450***	0.3132***
	(0.0004)	(0.0005)	(0.0003)	(0.0005)	(0.0003)	(0.0005)
$\log(D_{ij})$	-0.0084***	-0.0257***	-0.0109***	-0.0226***	-0.0121***	-0.0278***
	(0.0006)	(0.0012)	(0.0006)	(0.0011)	(0.0006)	(0.0011)
Obs.	2359351	2332589	2449276	2418037	2572018	2539325
Product FE	-	Yes	-	Yes	-	Yes
Buyer FE	-	Yes	-	Yes	-	Yes
Seller FE	-	Yes	-	Yes	-	Yes
Month FE	-	Yes	-	Yes	-	Yes
R^2	0.5777	0.7245	0.5868	0.7319	0.5925	0.7421

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Regressions include all monthly flows among multi-establishment manufacturing firms, and exclude flows with zero values. f_{ijp} is value sold by firm i to firm j in product p, A_{jp} is value acquired of product p by firm j, S_{ip} is value sold of product p by firm i, and D_{ij} is the geographical distance between both firms.

Our analysis reveals that the influence of gravity on the production network is highly persistent. Both the overall explanatory power of the regressions and the relative importance of each gravity dimension remain stable across different time periods, even during the COVID-19 shock. Table 2 presents results from specifications similar to those in columns 1 and 2 of Table 1, covering three distinct periods: pre-shock (2018-2019, columns 1 and 2), the shock and recovery phase (2020-2021, columns 3 and 4), and the post-shock years (2022-2023, columns 5 and 6). A remarkable feature from comparing these periods is the consistency of the regression coefficients over time. Despite some coefficients being statistically different across periods (for instance, when comparing columns 1, 3, and 5), the variations are economically negligible. Similarly, the explanatory power of the models remains remarkably stable, even during the economic upheaval of 2020. This stability strongly supports the idea that gravity is a persistent force shaping the production network.[10] The influence of gravity materializes in the short term and is sustained over time.

These findings align with the view that gravity reflects how economic activity is organized within each sector, which in turn is shaped by the prevailing state of technology at a given time. Since changes in technology tend to occur gradually over long periods, we would not expect dramatic shifts in the short run. The COVID-19 pandemic, despite being a major shock that destroyed many nodes and links in the production network, likely had little impact on the technological state of the economy. As a result, when new links formed to replace those that were lost, they followed gravity-driven patterns with similar intensities as before, even as the network's node structure itself changed with the shock. More generally, our results highlight that network adjustment occurs primarily through the replacement of individual links rather than through changes in the broader organization of industries. Firms may change suppliers following disruptions, but they appear to search within a relatively stable gravity structure. This helps reconcile two stylized facts in the literature: production networks exhibit substantial link turnover while simultaneously displaying remarkable aggregate persistence.

4.4 Gravity and Link Survival

The previous sections show that gravity explains the cross-sectional organization of production networks and that this relationship remains remarkably stable over time. A stronger implication follows from these findings. If gravity captures economically meaningful determinants of buyer-supplier relationships rather than merely fitting observed transactions, then links that conform more closely to gravity should also be more resilient. Conversely, relationships sustained primarily by idiosyncratic factors should be less stable and therefore more likely to disappear. In this section, we use this implication to assess whether gravity contains information about the stability of production networks beyond its ability to explain the cross section.

To do so, we use the residuals from the product-level PPML gravity estimations without buyer or seller fixed effects. Residuals are then divided by the total transaction value to re-scale them relative to trade size. These residuals summarize the component of each transaction that is not explained by basic gravity fundamentals. Large residuals therefore identify relationships whose observed transaction values differ substantially from those predicted solely by the size of buyers, the size of suppliers, and the distance between

[10] As in the previous section, results look very similar to the reported above when the sample is restricted to single-establishment firms (see Tables A.10-A.12 in the Appendix), or expanded to include nodes (either buyers or sellers) that do not belong to the manufacturing industry. The latter results are omitted here to save space, but available upon request to the authors.

them. If gravity indeed reflects persistent organizational forces, links displaying larger deviations from these fundamentals should exhibit lower survival probabilities.

We estimate a probit model in which the dependent variable equals one if a buyer-supplier relationship observed in period t disappears in period $t+1$. The main explanatory variable is the absolute value of the gravity residual, interpreted as a measure of the extent to which a relationship departs from gravity fundamentals. Product fixed effects are included so identification comes entirely from comparisons among links within the same product market. A positive relationship between the absolute residual and the probability of link destruction would indicate that relationships deviating further from gravity are less stable. The results show that, out of the 233 industries for which we can report results 175 give a significant and positive coefficient. From the rest, 48 coefficients are not significantly different to zero and 10 are significantly negative. Figure A.3 in the appendix shows all coefficients and their confidence intervals.

We also estimate a complementary specification using the signed residual rather than its absolute value. Whereas the previous regression evaluates whether departures from gravity reduce link stability irrespective of their direction, the signed residual allows us to assess whether relationships trading above or below the gravity prediction exhibit systematically different survival patterns. Results are very similar to those obtained using the absolute residual (see Figure A.4), indicating that what matters is not whether a relationship trades above or below its gravity prediction, but rather whether it departs substantially from gravity at all.

Overall, these results reinforce the interpretation of gravity as capturing the persistent component of buyer-supplier relationships. Links that conform more closely to gravity are systematically more likely to survive, while those that deviate from gravity are disproportionately short-lived. This suggests that the systematic forces summarized by gravity do not merely explain the contemporaneous organization of production networks; they also identify the relationships that are more likely to persist as networks evolve over time.

5 Predicting Production Networks

The previous sections establish that gravity explains a large share of the observed variation in firm-to-firm transactions and that these relationships are remarkably stable across industries and over time. A natural question is whether this information can be used to recover production networks in settings where firm-to-firm transactions are unobserved. This question is particularly relevant because detailed transaction data are available for only a handful of countries, whereas information on firms' size, location, and product-level activity is considerably more widespread.

If gravity captures persistent organizational regularities rather than country-specific features of observed networks, industry-specific gravity coefficients estimated in one economy should be informative about the organization of production networks in another. We evaluate this hypothesis using administrative data from Denmark, where we observe firms' size, location, purchases, and sales, but not their trading partners. We therefore use the gravity coefficients estimated for each product in Chile to predict the complete network of firm-to-firm linkages in Denmark. This exercise provides a stringent out-of-sample test of whether industry-specific gravity relationships generalize across economies. We then assess the validity of the predicted network by asking whether it successfully captures economically meaningful patterns of shock propagation.

5.1 Predicting Firm-to-Firm Linkages

The prediction procedure combines two ingredients: the industry-specific gravity relationships estimated from Chile and information on the size and location of firms in Denmark. The underlying assumption is that firms operating in the same narrowly defined industry organize their supplier relationships according to similar gravity forces across economies. Conditional on observing the set of potential buyers and suppliers, gravity determines both the likelihood that a link forms and, conditional on formation, the expected intensity of trade.

Because production networks are sparse, predicting transaction values alone would incorrectly assign positive trade to every feasible buyer-supplier pair. We therefore predict network formation and transaction intensity separately, combining an extensive-margin prediction for link existence with an intensive-margin prediction for trade volumes.

The main inputs for the predictions are the coefficients estimated for product p on data from Chile (CL) $[\hat{\beta}, \hat{\beta}_A, \hat{\beta}_S, \hat{\beta}_D,]_{pt}^{m,CL}$. These coefficients are obtained estimating (1) through different methods m: OLS, PPML, probit and extensive-margin PPML (Ext-PPML). The latter two are obtained using a dummy variable for the transaction being positive as a dependent variable. Armed with these coefficients, we produce predictions for the trade flows between any pair ij of Danish firms (DK) in industry p by following a two-step procedure.

First, we obtain predicted values by means of the following equation:

$$\exp[\hat{\beta}_{pt}^{m,CL} + \hat{\beta}_{p,A}^{m,CL} log(A_{jpt}^{DK}) + \hat{\beta}_{p,S}^{m,CL} log(S_{jpt}^{DK}) + \hat{\beta}_{p,D}^{m,CL} log(D_{ijt}^{DK})] = \\ = \begin{cases} \widehat{int}_{ijpt}^{m,DK} & \text{if } m \in (OLS,\ PPML) \\ \widehat{ext}_{ijpt}^{m,DK} & \text{if } m \in (probit,\ Ext\text{-}PPML) \end{cases} \tag{2}$$

where p denotes a CN four-digit code.

When $m \in$ (*OLS,PPML*), the outcome $\widehat{int}_{ijpt}^{m,DK}$ is interpreted as the predicted value of the trade flow assuming the flow exists. We generate these including all sellers of p

($S_{ipt} > 0$) and all buyers of p ($A_{jpt} > 0$). In other words, this step yields an intensity level for each seller-buyer-product-year tuple. When $m \in$ (*probit, Ext-PPML*), the predicted value $\widehat{ext}_{ijpt}^{m,DK}$ is interpreted as measuring the probability of positive trade flows for each specific tuple.

Finally, we combine our intensive and extensive margin results to obtain predicted trade flows as follows:

$$\widehat{f}_{ijpt}^{m,DK} = \begin{cases} \widehat{int}_{ijpt}^{m,DK} & \text{if } \widehat{ext}_{ijpt}^{m,DK} \geq t_{pt}^{m,DK}, \\ 0 & \text{if } \widehat{ext}_{ijpt}^{m,DK} < t_{pt}^{m,DK} \end{cases} \tag{3}$$

In words, we let our predicted intensities constitute the value of all links if the probability of the link existing exceeds certain threshold $t_{pt}^{m,DK}$, and set the intensity to zero when the threshold is not reached. Threshold $t_{pt}^{m,DK}$ is set at 80% of predicted probabilities within a given year. Although the precise threshold is ultimately a modelling choice, our objective is not to optimize classification performance but to evaluate whether a simple gravity-based rule is sufficient to recover economically meaningful production networks.[11] In what follows we restrict our attention to two alternative predictions. The first one, denoted *OLS* prediction results from combining the *OLS* intensity estimations with *probit* estimations for the extensive margin. The second prediction is denoted *PPML* as it combines estimations using that method for both margins. The coefficients $\hat{\beta}$ are derived from the gravity estimation based on Chilean data at the CUP level.

In predicting flows, we incorporate the estimated coefficients on the constant derived from the Chilean data. These coefficients partially reflect characteristics unique to Chile, such as trade values being expressed in Pesos. Consequently, our predictions of trade flows in Denmark in levels may include measurement error. To address this, our validation exercise will rely on share-based measures, which help mitigate measurement errors that are common across products. Nevertheless, we retain the constant term to account for sectoral differences in firms' propensity to engage in trade.

Table A.13 presents descriptive statistics on the predicted networks that obtain from our procedure. Comparing the numbers with those from observed networks in Chile (Table A.1), we can see that magnitudes are comparable in order of magnitude. The average predicted size of trade values are somewhat smaller for firms in Denmark, while the average number of buyers per sellers we obtain are slightly smaller in the European economy, while the average number of sellers per buyer much larger, though with a large standard deviation.

[11] As highlighted by the literature, predicting the extensive margin of trade flows in sparse networks can be a complex task, and more sophisticated approaches could have been chosen for it. We decided to adopt the approach above to underscore that the resulting predictions can find validation even when a simple and straightforward rule for the extensive margin is followed.

It must be noted that our prediction procedure deliberately abstracts from several features that may also influence network formation, including final demand, cross-product complementarities, and input substitution across products. Incorporating these factors would likely contribute to the accuracy of the predictions, and could be explored in future research. The fact that the predicted networks nevertheless perform well in the validation exercises suggests that industry-specific gravity relationships capture an important component of production-network organization even in the absence of these additional mechanisms.

5.2 Validating the Predicted Network

The objective of this section is to assess whether the gravity-based predicted networks capture economically meaningful buyer-supplier relationships. Because the true Danish network is unobserved, validation must rely on implications of the predicted network. Specifically, if the predicted links approximate the true production network, shocks affecting buyers should propagate to predicted suppliers in ways consistent with observed firm outcomes. We assess this hypothesis through three complementary exercises. First, we examine whether changes in buyers' activity are associated with changes in the sales of their predicted suppliers. Second, we analyze whether foreign demand shocks propagate through the predicted network. Finally, we evaluate whether domestic industry-specific shocks are likewise transmitted through the predicted buyer-supplier relationships.

Co-movement Between Predicted Buyers and Suppliers. The first implication of a correctly predicted production network is that firms identified as buyers and suppliers should co-move over time. When buyers expand or contract, demand should propagate to their suppliers through the predicted network. We therefore begin by asking whether firms identified as suppliers through the predicted network experience changes in sales when their predicted buyers expand or contract. To achieve this, we construct a measure of suppliers' exposure to buyers' scale and test whether changes in this exposure measure predict changes in the sales of suppliers.

To compute the measure of indirect exposure, we use our predicted firm-to-firm trade values. Let $\widehat{f}_{ijpt}$ denote the predicted sales of product p from firm i to firm j in year t. From here on we do everything either using the OLS or PPML predictions separately, so we omit the method m subscript. Define $\widehat{s}_{ijt}$ as the predicted share of firm i's sales to firm j relative to firm i's total sales, given by:

$$\widehat{s}_{ijt} = \sum_{p} \frac{\widehat{f}_{ijpt}}{S_{ipt}} \tag{4}$$

We compute a measure of firm i's predicted exposure to the scale of its buyers (ESB) as:

$$ESB_{it} = 1 + \sum_j \widehat{s}_{ijt} x_{jt} \tag{5}$$

where x_{jt} is a measure of performance of firm j, which in this case can amount to total sales in year t, or alternatively, total exports in year t.

We then consider the following regression:

$$\Delta log(S_{it}) = \theta \Delta log(ESB_{it}) + \Gamma_i + \Gamma_t + \epsilon_{it} \tag{6}$$

where $\Delta H_t = H_t - H_{t-1} \forall H_t$, and we include firm and year fixed effects (Γ_i and Γ_t respectively). A positive θ indicates that increases in the sales of buyers are associated with increased domestic sales of suppliers, providing validation to our predicted network.

Given previous findings in the literature, we expect shock pass-through to be imperfect. This is because the structure of the network is expected to react to the shock, so the importance of connections involving firms negatively (positively) affected should fall (increase) preventing the shock to spread to some degree (Huneeus, 2020). However, a mechanical relationship may exist between these two variables, since the aggregate domestic purchases of a product must equal the aggregate domestic supply of the same product. As a result, any observed correlation between buyer-side variables and supplier outcomes could simply reflect this identity, rather than capturing the role of network structure in shock transmission. It is important for our exercise to distinguish genuine network effects from this mechanical relationship. To address this, we control for the log change in the aggregate size of buyers by including the yearly log change in total buyers' sales. Specifically, for each supplier i, we compute the sum of the sales of all its potential buyers, defined as all firms for which we have estimated trade flows—i.e., firms that are either buyers, sellers, or both for a given product that i sells:

$$BS_{it} = \sum_{i'} \mathbb{I}_{ii't} \sum_p S_{i'pt} \ , \forall i' \neq i \tag{7}$$

where $\mathbb{I}_{ii't}$ is an indicator of a potential link from seller i to buyer i', so it takes value of one if at moment t, i sells any of the products that i' buys, and zero otherwise.

Tables A.20 and 4 show the results using respectively buyers' domestic sales and export sales in the measures of exposure. Results are similar for both exposure measures and show that higher exposure to positive changes in predicted buyers' sales and exports correspond to larger domestic sales for suppliers. Results are robust to controlling for total buyers' size. Results are also robust to controlling for time fixed and firm fixed effects, restricting results to within-firm variation that is unaffected by aggregate shocks.

Table 3: Domestic Sales and Exposure to Buyer's Sales

	Dependent Variable: Log Change in Supplier's Domestic Sales					
	(1)	(2)	(3)	(4)	(5)	(6)
Log Change in Indirect Exposure	0.019***	0.018***	0.018***	0.017***	0.017***	0.017***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Log Change in Buyers Size					0.051***	0.053***
					(0.005)	(0.005)
Grav. Est.	OLS	PPML	OLS	PPML	OLS	PPML
Time FE	No	No	Yes	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes	Yes	Yes
# Obs.	43545	43545	42698	42698	42698	42698
R^2	0.01	0.01	0.13	0.13	0.13	0.13

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses.

Table 4: Domestic Sales and Exposure to Buyer's Exports

	Dependent Variable: Log Change in Supplier's Domestic Sales					
	(1)	(2)	(3)	(4)	(5)	(6)
Log Change in Indirect Exposure	0.013***	0.012***	0.013***	0.012***	0.012***	0.011***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Log Change in Buyers Size					0.052***	0.054***
					(0.005)	(0.005)
Grav. Est.	OLS	PPML	OLS	PPML	OLS	PPML
Time FE	No	No	Yes	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes	Yes	Yes
# Obs.	43545	43545	42698	42698	42698	42698
R^2	0.01	0.01	0.13	0.13	0.13	0.13

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses.

Export Demand Shocks. The second implication concerns the transmission of exogenous shocks. Export demand shocks provide a source of plausibly exogenous variation in buyers' activity. If the predicted network identifies true buyer-supplier relationships, these shocks should propagate indirectly to predicted suppliers. We follow a standard procedure in the literature by constructing export demand shocks following a shift-share structure.[12]

[12] See for example Hummels et al. (2014) or Bernard et al. (2020). Huneeus (2020) constitutes a recent application of this procedure to observed firm-to-firm connections. To the best of our knowledge our exercise constitute the first in the literature using this approach on predicted firm-to-firm networks.

We define the direct export exposure of a firm i in year t, denoted as dir_{it}, as:

$$dir_{it} = \sum_{d} s_{id0} \Delta \log \left(\sum_{o} X_{odt} \right) \text{ , } \forall \text{ o and d} \neq \text{Denmark} \tag{8}$$

Here, X_{odt} represents the total exports from origin o to destination d, s_{id0} is the observed share of firm i's exports to destination d in the initial year the firm appears in our dataset. To construct this variable we utilize data from BACI, as reported by Gaulier and Zignago (2010), and exclude Denmark from both the pool of origin and destinations countries.

We compute a measure of firm i's indirect export exposure to foreign demand, ind_{it}, as:

$$ind_{it} = \sum_{j} \widehat{s}_{ijt} dir_{jt} \tag{9}$$

where $\widehat{s}_{ijt}$ represents the predicted trade shares, defined in (4) on the basis of our predicted domestic trade flows.

We then consider the following regression:

$$\Delta log(S_{it}) = \theta^{dir} dir_{it} + \theta^{ind} ind_{it} + \Gamma_i + \Gamma_t + \epsilon_{it} \tag{10}$$

where $\Delta log(S_{it})$ is the growth (measured as the log difference) of domestic sales and export by firm i in year t. We include firm and year fixed effects, to clean for year shocks and focus on within-firm changes.

Table 5 presents the results. The measure of direct export exposure is positively associated with a firm's exports. This finding is intuitive, as it suggests that positive foreign demand shocks lead to increased exports for firms exposed to these shocks through their export destinations. We also observe that positive foreign demand shocks do not significantly affect domestic sales. This indicates that the increased sales abroad do not boost or crowd-out the domestic activity of the firm.

In contrast, indirect export exposure—measured through sales to buyers—boosts domestic sales. This result provides supports to our predicted network, as it indicates that positive foreign demand shocks affecting buyers lead to increased domestic sales for their (estimated) suppliers. It is also worth noting that exports do not increase for firms indirectly exposed to foreign shocks. This contributes in ruling out potential direct effects to our indirectly exposed firms. Finally, Table A.14 in the appendix shows that the results are robust to including the log change in buyers size as a control variable.

Table 5: Export Demand Shocks, Domestic Sales, and Exports

	Dependent Variable: Growth Rate of Sales (Domestic or Exports)					
	(Dom.)	(Dom.)	(Dom.)	(Exp.)	(Exp.)	(Exp.)
Export Exposure	-0.004	-0.004	-0.004	0.095***	0.095***	0.095***
	(0.005)	(0.005)	(0.005)	(0.010)	(0.010)	(0.010)
Indirect Export Exposure		0.023***	0.023***		0.006	-0.001
		(0.005)	(0.005)		(0.007)	(0.007)
Grav. Est.	-	OLS	PPML	-	OLS	PPML
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	31565	31565	31565	27005	27005	27005
R^2	0.11	0.11	0.11	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Columns denoted with (Dom.) use the growth rate of domestic sales as the dependent variable and columns denoted with (Exp.) use the the growth rate of exports.

Domestic Demand Shocks. The final implication is that the predicted network should also capture the transmission of domestic shocks. While export shocks provide external variation, they affect directly only exporting firms. Industry-specific demand fluctuations allow us to assess whether the predicted network captures propagation more generally within the domestic economy in Denmark. This exercise builds on similar principles of firm exposure to industry-level shocks but adapts the analysis to the context of domestic markets and firm-to-firm linkages.

First, we identify industry-level shocks by computing deviations from the long-term production trends of each CN four-digit good p. Specifically, we regress the log of total sales for each product p on a time trend and use the residuals as measures of the deviation from the trend. Let these residuals be denoted as $log(Shock_{pt})$. To ensure robust estimates, we restrict the analysis to products with at least five yearly observations.

Next, we calculate firm-specific exposure to shocks. For each firm i, we compute a weighted sum of the industry-level shocks, where the weight corresponds to the importance of good p in the firm's sales portfolio. In particular, let:

$$w_{ipt} = \frac{S_{ipt}}{\sum_p S_{ipt}} \tag{11}$$

represent the share of sales of a firm i on product p in year t. The exposure to shocks of firm i to shocks equals:

$$exp\_shock_{it} = \sum_p w_{ipt} log(Shock_{pt}) \tag{12}$$

The exposure of a firm reflects its sensitivity to fluctuations—both positive and negative—in the goods it sells. A firm whose sales are concentrated in goods experiencing downturns will have a high negative exposure. Conversely, if the goods the firm sells are undergoing an upswing, the firm will experience a positive exposure.

Finally, we examine how these firm-specific shocks affect firms upstream. Specifically, we compute a measure of firm i's predicted indirect exposure to industry shocks, as:

$$ind\_shock_{it} = \sum_j \widehat{s}_{ijt} exp\_shock_{jt} \tag{13}$$

where $\widehat{s}_{ijt}$ is the share of firm i's sales to firm j, as defined in (4), on the basis of our predicted domestic trade flows.

We then consider the following regression:

$$\Delta log(S_{it}) = \vartheta \Delta ind\_shock_{it} + \Gamma_i + \Gamma_t + \epsilon_{it} \tag{14}$$

where ϑ would capture the elasticity of sales by firm i to shocks received through its customers, according to our predicted network. Again, we consider domestic sales and exports, as two distinct variants for the outcome variable of interest.

The results in Table 6 show that domestic sales are positively correlated with exposure to positive shocks to predicted domestic buyers. As expected, effects are not so evident for export sales: significance is lower on the PPML estimation and we find no significant results on the OLS estimation. Table A.15 in the appendix shows that the results are robust to including the log change in buyers size as a control variable.

Table 6: Domestic Demand Shocks, Domestic Sales, and Exports

	Dependent Variable: Growth of Sales			
	(Dom.)	(Dom.)	(Exp.)	(Exp.)
Exposure to Buyer's Industry Shocks	0.115***	0.126***	0.037	0.088*
	(0.025)	(0.024)	(0.050)	(0.048)
Grav. Est.	OLS	PPML	OLS	PPML
Time FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
# Obs.	42698	42698	27370	27370
R^2	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Columns denoted with (Dom.) use the growth rate of domestic sales as the dependent variable and columns denoted with (Exp.) use the the growth rate of exports.

Together, the three exercises progressively strengthen the evidence. The first shows that predicted buyers and suppliers co-move over time. The second demonstrates that exogenous foreign demand shocks propagate through the predicted network. The third shows that the same network also captures the transmission of domestic industry shocks. Collectively, these findings suggest that the gravity-based predicted network recovers economically meaningful buyer-supplier relationships. This conforms strong evidence in support of the idea that the way industries organize production locally shares important similarities across economies, and therefore, it is possible to use information from one location to predict the shape of the network elsewhere.

5.3 Decomposing the Sources of Predictive Performance

The previous section shows that gravity-based predicted networks reproduce economically meaningful patterns of co-movement and shock propagation. We now ask which elements of the prediction procedure are responsible for this performance. In particular, we evaluate the contribution of three ingredients: the treatment of the extensive margin, the use of industry-specific gravity coefficients, and the precise values of those coefficients. Results are reported in Tables A.16 to A.24 of the Appendix.

Ignoring the Extensive Margin Threshold. The first exercise evaluates the importance of predicting the topology of the network. If correctly identifying which buyer-supplier relationships exist is important, ignoring the extensive margin should weaken the ability of the predicted network to capture shock transmission.

In our baseline specification, predicted firm-to-firm trade flows combine intensive-margin predictions with an extensive-margin correction. Specifically, predicted intensities are set to zero whenever the estimated probability of a link is below 80 percent. We now construct an alternative predicted network in which no such threshold is imposed. That is, we retain the intensive-margin predictions for all firm pairs without censoring any flows based on predicted probabilities.

Using this alternative network, we recompute the indirect exposure measures and replicate the validation regressions described above. Results are reported in the tables of Appendix A.1, in the columns labeled "Ext". The results indicate that removing the extensive-margin threshold reduces the predictive power of the network. In particular, the estimated coefficients on indirect exposure are generally smaller in magnitude relative to the baseline specification, and in some cases—especially when using export-based exposure measures—lose statistical significance. This suggests that imposing the threshold to determine the existence of links plays an important role in disciplining the predicted network and improving its ability to capture economically meaningful transmission channels.

Minimum-Coefficient Network. We next evaluate whether the predictive content of the network comes from gravity itself or from allowing gravity relationships to differ across industries. For this we propose an exercise that alters the informational content of the gravity coefficients used to generate predicted trade flows. For each gravity dimension (seller size, buyer size, and distance), we identify the minimum estimated coefficient in the cross-sector distribution (i.e., the left tail of the estimated coefficient distribution). We then assign this minimum coefficient to all sectors and use this common coefficient vector to predict firm-to-firm trade flows. As in the baseline, we apply the 80 percent threshold rule to determine the existence of links.

This procedure generates a predicted network that preserves the general gravity structure but eliminates cross-sector heterogeneity in estimated coefficients and imposes the lowest observed elasticities uniformly across industries. Using this alternative network, we recompute the indirect exposure measures and replicate the validation regressions. Results are reported in the tables of Appendix A.1, in the columns labeled "Min".

Using a single minimum coefficient across all sectors eliminates the predictive power of the network. Across specifications, the indirect exposure measures are not statistically significant and their estimated coefficients are close to zero. This indicates that sector-specific heterogeneity in gravity coefficients is crucial for generating informative predictions of firm-to-firm linkages. Imposing a common, low elasticity across industries produces a network that fails to capture meaningful patterns of shock propagation.

Randomized-Coefficient Networks. Finally, we evaluate the sensitivity of our predictions to estimation error in the gravity coefficients. We do this by implementing a family of placebo exercises in which gravity coefficients are randomly perturbed before constructing predicted trade flows. For each sector and each gravity coefficient, we randomly add or subtract N times the cross-sector standard deviation of that coefficient, where $N \in \{1, 2, 3, 4\}$. This generates alternative coefficient vectors that deviate from the estimated values while preserving the overall scale of the coefficient distribution.

For each value of N, we use the perturbed coefficients to predict trade flows and apply the 80 percent threshold rule to determine link existence. We then recompute the indirect exposure measures and replicate the validation regressions. Results are reported in the tables of Appendix A.1, in the columns labeled "1 SD", "2 SD", "3 SD", and "4 SD", corresponding to the magnitude of the perturbation applied to the gravity coefficients.

The results from the randomized-coefficient exercises show that the validation findings are relatively robust to sizeable perturbations of the gravity parameters. When using OLS-based predictions, the estimated indirect exposure effects remain broadly similar to the baseline even when coefficients are shifted by multiple standard deviations. In contrast, PPML-based predictions exhibit a more pronounced decline in coefficient magnitudes. For the domestic industry shock exercise, both OLS and PPML specifications display a reduction in precision relative to the baseline, with PPML estimates in particular losing statistical significance. These patterns suggest that while the exact values of the gravity coefficients matter for precision, the overall structure of the predicted network remains informative under substantial parameter perturbations.

Taken together, these exercises highlight the ingredients that make gravity-based prediction successful. Correctly predicting which links exist is quantitatively important, and allowing gravity relationships to vary across industries is essential. By contrast, the exact values of the estimated coefficients matter considerably less: moderate perturbations leave the economic content of the predicted network largely intact. One possible explanation is that aggregation at the firm level attenuates product-level idiosyncratic variation, so that moderate deviations in product-level gravity coefficients do not substantially alter firm-level exposure measures. Overall, these findings suggest that successful prediction depends primarily on recovering the topology of the network and the industry-specific heterogeneity in gravity relationships. Once these features are captured, moderate estimation error in the gravity coefficients has only a limited impact on the economic content of the predicted network.

6 Concluding remarks

This paper studies the forces that shape firm-to-firm production networks. Using detailed transaction data from Chile, we show that gravity explains a large share of the observed variation in firm-to-firm transactions within narrowly defined industries, yet its importance differs substantially across products. These differences are systematic and reflect heterogeneity in the technological and logistical organization of production across industries rather than idiosyncratic variation across firms.

A central finding of the paper is that gravity characterizes not only the cross-sectional structure of production networks but also their evolution over time. We show that gravity coefficients are remarkably stable despite large disruptions such as the COVID-19 pandemic. Links that are more consistent with gravity are more likely to survive, while new supplier relationships formed after major disruptions closely resemble the gravity structure observed before the shock. Taken together, these results suggest that the observed production network should be viewed as the realization of a relatively stable underlying organizational structure rather than as an arbitrary collection of bilateral relationships. In other words, gravity appears to describe the persistent matching technology that governs how firms organize production within industries.

The paper also shows that these organizing principles extend beyond the economy in which they are estimated. Industry-specific gravity coefficients obtained from Chile successfully predict firm-to-firm linkages in Denmark, even though the two economies differ substantially in their industrial composition, income levels, and geography. The success of these out-of-sample predictions suggests that firms operating within narrowly defined industries rely on similar spatial and organizational principles when selecting suppliers and buyers. This result complements the literature documenting technological convergence across countries by showing that convergence extends beyond production technologies inside firms to the organization of production networks connecting them.

More broadly, our findings suggest that production networks exhibit considerably more regularity than is often assumed. While individual supplier relationships inevitably reflect idiosyncratic factors, the aggregate organization of production within narrowly defined industries follows stable gravity forces that persist across time, survive major disruptions, and generalize across countries. Our results further suggest that recovering economically meaningful production networks depends primarily on correctly identifying their topology and the industry-specific heterogeneity in gravity relationships. Once these features are captured, moderate estimation error in the estimated coefficients has only a limited effect on the network's predictive content. This perspective opens new opportunities for empirical and quantitative research. In settings where firm-to-firm transaction data are unavailable, industry-specific gravity parameters combined with commonly observed information on firms' size and location provide a tractable way to recover economically meaningful production networks, thereby extending the scope for studying productivity, shock propagation, industrial policy, and regional development in environments where direct observations of production linkages do not exist.

References

D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016, 2012.

V. I. Alviarez, M. Fioretti, K. Kikkawa, and M. Morlacco. Two-sided market power in firm-to-firm trade. Technical Report w31253, National Bureau of Economic Research, 2023.

C. Arkolakis, F. Huneeus, and Y. Miyauchi. Spatial Production Networks. *Working Paper*, 2023.

D. R. Baqaee and E. Farhi. The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. *Econometrica*, 87(4):1155–1203, 2019.

A. Bernard, V. Smeets, T. Fort, and F. Warzynski. Heterogeneous globalization: Offshoring and reorganization. Technical report, 2020.

A. B. Bernard, V. Smeets, and F. Warzynski. Rethinking deindustrialization. *Economic Policy*, 32(89):5–38, 2017.

A. B. Bernard, A. Moxnes, and Y. U. Saito. Production networks, geography, and firm performance. *Journal of Political Economy*, 127(2):639–688, 2019. doi: 10.1086/700764.

S. Bigio and J. La'o. Distortions in production networks. *Quarterly Journal of Economics*, 135(4):2187–2253, 2020.

J. Boehm, R. South, E. Oberfield, and M. Waseem. The network origins of firm dynamics: Contracting frictions and dynamism with long-term relationships. Technical report, Geneva Graduate Institute, September 2024.

L. Caliendo and F. Parro. Estimates of the Trade and Welfare Effects of NAFTA. *The Review of Economic Studies*, 82(1):1–44, Jan. 2015.

V. M. Carvalho, M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi. Supply Chain Disruptions: Evidence from the Great East Japan Earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321, Mar. 2021.

J. M. Chan, M. Koch, and L. Ma. Globalization, migration, and the spatial organization of production. Technical report, 2024.

D. Comin and B. Hobijn. An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059, 2010.

B. Demir, B. Javorcik, T. K. Michalski, and E. Ors. Financial Constraints and Propagation of Shocks in Production Networks. *Review of Economics and Statistics*, 2022.

E. Dhyne, A. K. Kikkawa, M. Mogstad, and F. Tintelnot. Trade and Domestic Production Networks. *Review of Economic Studies*, 88(2):643–668, 2021.

J. di Giovanni, A. A. Levchenko, and I. Mejean. Firms, Destinations, and Aggregate Fluctuations. *Econometrica*, 82(4):1303–1340, 2014.

J. Eaton, S. S. Kortum, and F. Kramarz. Firm-to-firm trade: Imports, exports, and the labor market. Technical Report w29685, National Bureau of Economic Research, 2022.

F. Fontaine, J. Martin, and I. Mejean. Frictions and adjustments in firm-to-firm trade. Technical report, 2024. Working paper, November 4, 2024.

G. Gaulier and S. Zignago. BACI: International Trade Database at the Product-Level. The 1994-2007 Version. Technical report, CEPII, 2010.

K. Head and T. Mayer. Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of international economics*, volume 4, pages 131–195. Elsevier, 2014.

E. Helpman, M. J. Melitz, and Y. Rubinstein. Estimating Trade Flows: Trading Partners and Trading Volumes. *Quarterly Journal of Economics*, CXXIII(2):441–487, 2008.

D. Hummels, R. Jørgensen, J. Munch, and X. C. The wage effects of offshoring: Evidence from danish matched worker-firm data. *American Economic Review*, 04(6):1597—1629, 2014.

F. Huneeus. Production Network Dynamics and the Propagation of Shocks. *Working Paper*, 2020.

C. I. Jones. Misallocation, Economic Growth, and Input-Output Economics. In D. Acemoglu, M. Arellano, and E. Deckel, editors, *Advances in Economics and Econometrics, Tenth World Congress, Volume II: Applied Economics*, pages 419–456. Cambridge University Press, 2013.

G. Khanna, N. Morales, and N. Pandalai-Nayar. Supply chain resilience: Evidence from indian firms. Technical report, 2022.

S. Kortum and B. S. C. Ribeiro. Growth with firm-to-firm trade. Technical Report 2459, Cowles Foundation for Research in Economics, Yale University, September 2025.

K. Lim. Firm-to-firm trade in sticky production networks. *American Economic Review*, 108(9):2427–2467, 2018. doi: 10.1257/aer.20161970.

J. B. Madsen and I. Timol. Long-run convergence in manufacturing and innovation-based models. *The Review of Economics and Statistics*, 93(4):1155–1171, 2011.

Y. Miyauchi. Matching and agglomeration: Theory and evidence from japanese firm-to-firm trade. *Econometrica*, 92(6):2413–2444, 2024. doi: 10.3982/ECTA19697.

E. Oberfield. A Theory of Input-Output Architecture. *Econometrica*, 86(2):559–589, 2018.

D. Rodrik. Unconditional Convergence in Manufacturing. *Quarterly Journal of Economics*, 128(1):165–204, 2013.

J. Santos Silva and S. Tenreyro. The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658, 2006.

M. Tascherau-Dumouchel. Cascades and fluctuations in an economy with an endogenous production network. *Review of Economic Studies*, 2025.

J. Tinbergen. Shaping the world economy: Suggestions for an international economic policy. *The Twentieth Century Fund*, 1962.

A Appendix

Table A.1: Manufacturing firms and observed trade flows in Chile

	Mean	SD	p50	p99	Max	Obs
Extensive margin (yearly)						
Branches by Seller	1.10	0.79	1	3	160	89,083
Products by Seller	9.00	11.30	5	54	160	89,083
Products by Seller-Branch	8.78	11.00	5	53	160	97,932
Buyers by Seller	26.83	223.41	4	385	25,587	89,083
Buyers by Seller-Branch	25.73	171.78	4	395	14,319	97,932
Branches by Buyer	1.55	1.29	1	6	131	171,486
Products by Buyer	16.36	16.56	11	74	185	171,486
Products by Buyer-Branch	12.16	14.59	6	66	179	265,413
Sellers by Buyer	13.94	31.14	6	135	1,681	171,486
Sellers by Buyer-Branch	9.43	23.59	3	99	1,555	265,413
Intensive margin (monthly)						
$\log(f_{ijp})$	11.27	2.23	11.09	17.24	50.93	43,390,737
$\log(A_{jp})$	11.96	2.57	11.73	19.01	50.93	43,390,737
$\log(S_{ip})$	15.88	3.14	15.97	21.76	51.36	43,390,737
$\log(D_{ij})$	10.89	1.74	10.56	14.32	15.94	35,169,769

Notes: Transactions data between firms belonging to the manufacturing sector, over the period 2014-2023. f_{ijp} is value sold by firm i to firm j in product p, A_{jp} is value acquired of product p by firm j, S_{ip} is value sold of product p by firm i, and D_{ij} is the geographical distance between both firms.

Table A.2: Manufacturing firms in Denmark

	Mean	SD	p50	p99	p100	Obs
Extensive margin (yearly)						
Products by Seller	1,9	2,3	1,0	9,2	18,7	51,004
Products by Buyer	5,1	6,8	3,0	30,6	50,1	14,038
Branches by Seller	2,2	5,4	1,4	11,8	37,8	51,004
Branches by Buyer	3,1	5,8	1,4	19,0	45,4	14,038
Intensive margin (yearly)						
$\log(A_{jp})$	7,2	2,4	7,3	11,9	12,9	70,565
$\log(S_{ip})$	8,4	2,6	8,9	13,0	14,1	94,495

Notes: All moments calculated in two steps. First, we compute each metric across firms within each year. Next, we average the yearly metric across all years. For confidentiality reasons, for p50, p99, and p100, firms are first divided into 100 bins per year, and the average is calculated within each bin. Finally, the bin averages are averaged across the 19-year period. A_{jp} is value acquired of product p by firm j. S_{ip} is value sold of product p by firm i.

Table A.3: Summary Statistics - Denmark

Sector	# Firms	# Prod. Sold	$\log(S)$	# Prod. Acq.	$\log(A)$
Chemicals and Mineral Products	151	124	18	134	16
Foodstufss	182	68	17	98	16
Minerals, Machinery, and Electrical	1307	198	18	208	17
Miscellaneous	333	97	17	125	15
Plastics and Rubbers	213	83	16	77	14
Textiles, Leather, and Fur	51	38	14	37	13
Transportation	45	20	14	29	13
Vegetable Products	20	29	15	21	13
Wood Products	382	52	17	67	15

Notes: All variables are averaged within broad sector and then averaged across years. A is value acquired and S is value sold.

Figure A.1: Distribution of Firms across Municipalities - Denmark

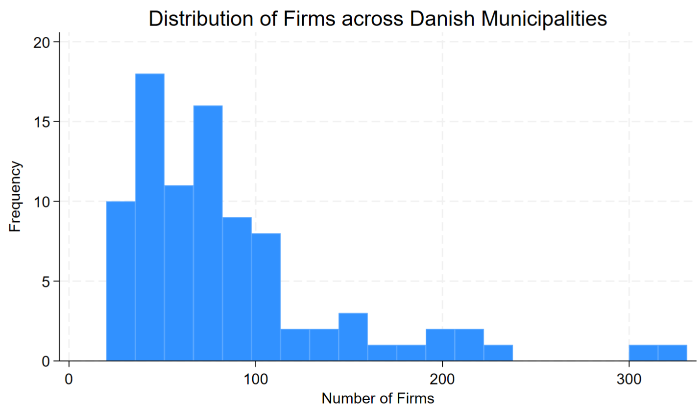

The vertical axis represents the number of municipalities and the horizontal axis the number of firms pooled across all years. The data is truncated so as to include only municipalities with at least 20 firms.

Figure A.2: Number of Firms per Year - Denmark

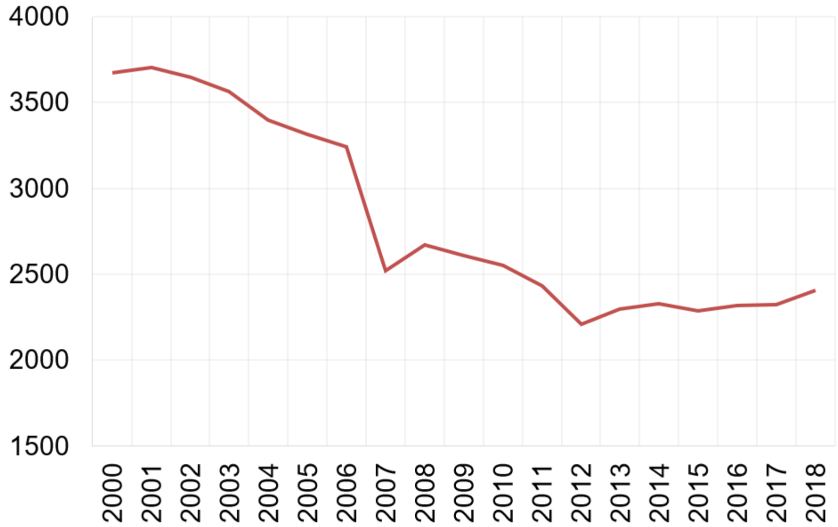

Table A.9: Pooled gravity estimation (2014-2023), single establishment firms only

	OLS		PPML	
	(1)	(2)	(3)	(4)
	$\log(Trade_{ijp})$	$\log(Trade_{ijp})$	$Trade_{ijp}$	$Trade_{ijp}$
$\log(Demand_{jp})$	0.6897***	0.6076***	0.0608***	0.0551***
	(0.0004)	(0.0005)	(0.0001)	(0.0001)
$\log(Supply_{ip})$	0.0418***	0.1734***	0.0040***	0.0161***
	(0.0003)	(0.0005)	(0.0000)	(0.0001)
$\log(Distance_{ij})$	-0.0292***	0.0031**	-0.0026***	0.0002*
	(0.0005)	(0.0014)	(0.0000)	(0.0001)
Observations	2025204	1999063	2025204	1999061
Product FE	-	Yes	-	Yes
Buyer FE	-	Yes	-	Yes
Seller FE	-	Yes	-	Yes
Month FE	-	Yes	-	Yes
R^2	0.6372	0.7474		
Pseudo R^2			0.0536	0.0633

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Regressions include all monthly flows among single-establishment manufacturing firms, and exclude flows with zero values.

Table A.4: CUP categories (1/5)

CUP	CUP Description	# CN
1	Wheat	1
2	Maize	1
3	Rice	1
4	Barley	1
5	Oats	1
6	Other cereals	12
7	Tubers	2
8	Legumes	3
9	Roots and bulbs; leafy, stem, or flowering vegetables that produce fruits; mushrooms and truffles	1
10	Oilseeds, beverage plants, and spices	12
11	Tobacco	1
12	Plants used in sugar and sweetener production	1
13	Other agricultural products for industrial use	7
14	Forage crops	1
15	Nursery products, live plants, flowers, and seeds	3
16	Raw plant materials for textiles, perfumery, and medicine; other raw vegetable materials n.e.c.	1
19	Table grapes	1
20	Avocados	1
21	Kiwis and other subtropical and tropical fruits	1
22	Citrus fruits, fresh or dried	2
23	Apples	1
24	Pears, fresh quinces, and other pome fruits	1
26	Other stone fruits, fresh	1
27	Berries	1
28	Nuts, fresh or dried, of temperate climates, even shelled or peeled	9
30	Cattle	1
31	Sheep	1
32	Pigs	1
33	Poultry	1
34	Other live animals not considered above	3
35	Milk	1
36	Eggs and other animal products for consumption	2
37	Raw animal materials for textiles and other non-edible animal products	2
43	Firewood and other forestry products	2
49	Other fish	1

Table A.5: CUP categories (2/5)

CUP	CUP Description	# CN
50	Crustaceans, mollusks, and other aquatic invertebrates	4
51	Seaweed and other aquatic products	1
52	Copper ores and concentrates	2
53	Unrefined copper; ashes, residues, and waste from smelting and refining	17
54	Refined copper	9
55	Coal	3
56	Crude petroleum	1
57	Natural gas (gaseous state)	1
59	Iron ore and its unagglomerated concentrates	1
60	Iron pellets and other agglomerated iron ore concentrates	9
61	Silver	3
62	Gold	16
63	Molybdenum ores and concentrates	2
64	Other metallic mining products	13
66	Stone, sand, clay, and other industrial minerals	10
67	Other non-metallic minerals	9
69	Beef and edible offal	2
70	Pork and edible offal	1
71	Poultry meat and edible offal	1
72	Other meats and meat products	8
73	Preserved and prepared meat, offal, or blood products	5
75	Other fish, fillets, and fish offal, fresh, chilled, or frozen	3
76	Preserved or prepared fish; caviar	2
77	Frozen or prepared crustaceans, mollusks, and other aquatic invertebrates	1
78	Fishmeal and other aquatic animal meals	1
79	Fish oil	1
80	Prepared or preserved fruits and vegetables	10
81	Animal fats and oils, except fish oil	7
82	Vegetable oils and fats	15
83	Margarine and similar products; residues of oil and fat production	4
84	Processed liquid milk and cream	2
85	Solid milk and cream; concentrated or sweetened milk and cream; modified milk	2
86	Other dairy products	1
87	Wheat flour and by-products	2
88	Inulin	1
89	Other milling products	10

Table A.6: CUP categories (3/5)

CUP	CUP Description	# CN
90	Bread	1
92	Sugar	2
93	Molasses, beet pulp, bagasse, and other sugar industry waste	2
94	Cocoa, chocolate, and confectionery products with sugar	3
95	Noodles and pasta	1
96	Miscellaneous foods	12
99	Feed for pigs, cattle, and other animals, except pets	1
100	Pet food	1
101	Pisco and other spirits	4
102	Wines	1
103	Beer	1
104	Non-alcoholic beverages	2
105	Tobacco products	3
106	Fibers, yarns, threads, fabrics, textile articles	73
107	Clothing, accessories, and fur articles	48
108	Leather and leather products; travel goods and saddlery	21
109	Footwear	6
110	Sawn, planed, and treated wood; impregnated or preserved logs	4
111	Wood chips and sawdust	1
112	Wood panels, sheets, veneers, and carpentry works	10
113	Wooden containers and packaging; cork, straw, and wicker products	11
114	Cellulose pulp (from wood or other fibrous materials)	7
115	Newsprint	1
116	Paper and cardboard containers; recyclable paper or cardboard	3
117	Tissue products	2
118	Other paper and cardboard products; waste and scrap	13
119	Printed materials and recordings	19
120	Diesel oil	1
123	Fuel oil	1
125	Petroleum gas and other gaseous hydrocarbons, except natural gas	2
128	Sulfuric acid	2
129	Other basic chemical products	64

Table A.7: CUP categories (4/5)

CUP	CUP Description	# CN
130	Fertilizers and pesticides	5
131	Ammonium nitrate	2
132	Molybdenum oxides	15
133	Paints, varnishes, inks, and related products	15
134	Soap, detergents, cleaning and polishing preparations; cosmetics and toiletries	17
135	Explosives and detonators	6
136	Other chemical products n.e.c.	46
137	Manufactured textile fibers	46
138	Pharmaceutical products	6
139	Tires	3
140	Other rubber products	12
141	Plastic rods, profiles, tubes, hoses, sheets, and films	2
142	Plastic packaging items, caps, stoppers, and closures	10
143	Glass bottles	2
144	Other glass and glass products	16
145	Articles of cement, concrete, asphalt, ceramics, and stone	27
146	Cement, lime, and gypsum	6
147	Ready-mixed concrete	1
148	Asphalt	5
149	Ferromolybdenum and other basic steel products	30
150	Iron and steel products	42
151	Other non-ferrous metal products	18
152	Structural metal products; prefabricated buildings; tanks, boilers, and containers	21
153	Other fabricated metal products	18
154	Computer equipment	1
155	Office, accounting, and IT machinery	2
156	Televisions	4
157	Mobile phones and other wireless network phones	1
158	Other electronic products	9
159	Electrical machinery and apparatus	24
160	Domestic appliances, except electronic items such as TVs and radios; parts and accessories	25
161	Industrial machinery	56
162	Parts and accessories for industrial machinery	21
163	Buses	1
164	Automobiles and SUVs	1

Table A.8: CUP categories (5/5)

CUP	CUP Description	# CN
165	Trucks, vans, tractors, and other motor vehicles (except buses, cars, and SUVs)	5
166	Other land transport equipment; parts and accessories	19
167	Ships, boats, aircraft, and their parts and accessories	14
168	Furniture	4
169	Other manufactured products	84
174	Electric power and energy	1

Notes: Full list of product-codes (CUP) in manufacturing activities with their description. The last column is the count of CN four-digit codes that are matched with the corresponding CUP.

Table A.10: Pooled gravity estimation (2018-2019), single-establishment firms only

	OLS		PPML	
	(1)	(2)	(3)	(4)
	$\log(Trade_{ijp})$	$\log(Trade_{ijp})$	$Trade_{ijp}$	$Trade_{ijp}$
$\log(Demand_{jp})$	0.6551***	0.5551***	0.0578***	0.0506***
	(0.0007)	(0.0009)	(0.0001)	(0.0001)
$\log(Supply_{ip})$	0.0608***	0.2145***	0.0057***	0.0200***
	(0.0005)	(0.0009)	(0.0001)	(0.0001)
$\log(Distance_{ij})$	-0.0138***	0.0122***	-0.0012***	0.0011***
	(0.0009)	(0.0025)	(0.0001)	(0.0002)
Observations	723313	706989	723313	706989
Product FE	-	Yes	-	Yes
Buyer FE	-	Yes	-	Yes
Seller FE	-	Yes	-	Yes
Month FE	-	Yes	-	Yes
R^2	0.6065	0.7371		
Pseudo R^2			0.0531	0.0653

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Regressions include all monthly flows among single-establishment manufacturing firms, and exclude flows with zero values.

Table A.11: Pooled gravity estimation (2020-2021), single-establishment firms only

	OLS		PPML	
	(1)	(2)	(3)	(4)
	$\log(Trade_{ijp})$	$\log(Trade_{ijp})$	$Trade_{ijp}$	$Trade_{ijp}$
$\log(Demand_{jp})$	0.6556***	0.5497***	0.0568***	0.0490***
	(0.0006)	(0.0009)	(0.0001)	(0.0001)
$\log(Supply_{ip})$	0.0569***	0.2160***	0.0053***	0.0198***
	(0.0005)	(0.0008)	(0.0000)	(0.0001)
$\log(Distance_{ij})$	-0.0127***	0.0097***	-0.0011***	0.0008***
	(0.0008)	(0.0022)	(0.0001)	(0.0002)
Observations	829936	810692	829936	810692
Product FE	-	Yes	-	Yes
Buyer FE	-	Yes	-	Yes
Seller FE	-	Yes	-	Yes
Month FE	-	Yes	-	Yes
R^2	0.6091	0.7430		
Pseudo R^2			0.0518	0.0639

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Regressions include all monthly flows among single-establishment manufacturing firms, and exclude flows with zero values.

Table A.12: Pooled gravity estimation (2022-2023), single-establishment firms only

	OLS		PPML	
	(1)	(2)	(3)	(4)
	$\log(Trade_{ijp})$	$\log(Trade_{ijp})$	$Trade_{ijp}$	$Trade_{ijp}$
$\log(Demand_{jp})$	0.6567***	0.5402***	0.0564***	0.0476***
	(0.0006)	(0.0009)	(0.0001)	(0.0001)
$\log(Supply_{ip})$	0.0601***	0.2344***	0.0056***	0.0211***
	(0.0005)	(0.0008)	(0.0000)	(0.0001)
$\log(Distance_{ij})$	-0.0272***	-0.0037*	-0.0024***	-0.0004**
	(0.0008)	(0.0021)	(0.0001)	(0.0002)
Observations	843287	824200	843287	824200
Product FE	-	Yes	-	Yes
Buyer FE	-	Yes	-	Yes
Seller FE	-	Yes	-	Yes
Month FE	-	Yes	-	Yes
R^2	0.6092	0.7482		
Pseudo R^2			0.0512	0.0634

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Regressions include all monthly flows among single-establishment manufacturing firms, and exclude flows with zero values.

Figure A.3: Gravity misalignment and link survival (per industry)

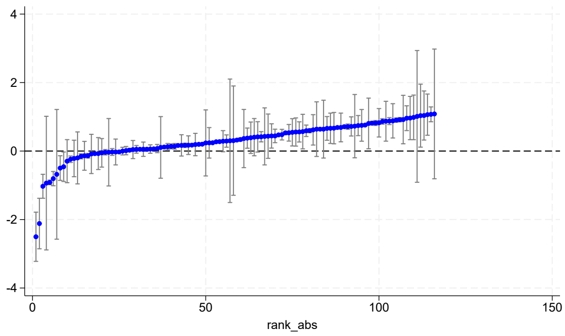

(a) Lowest half of coefficients

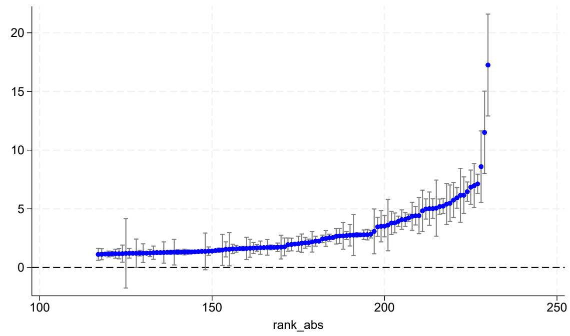

(b) Highest half of coefficients

Notes: Coefficient and 95% confidence intervals from a probit regression on the probability that a link disappears next period, where the main regressor is departure from gravity. The latter is the absolute value of the error term of the industry estimated from equation (1) by PPML in 2018, rescaled by trade value. For exposition purposes panel (a) shows the lower half of coefficients and panel (b) the higher half. The horizontal axis is sorted by the value of the coefficient.

Figure A.4: Gravity misalignment and link survival (per industry)

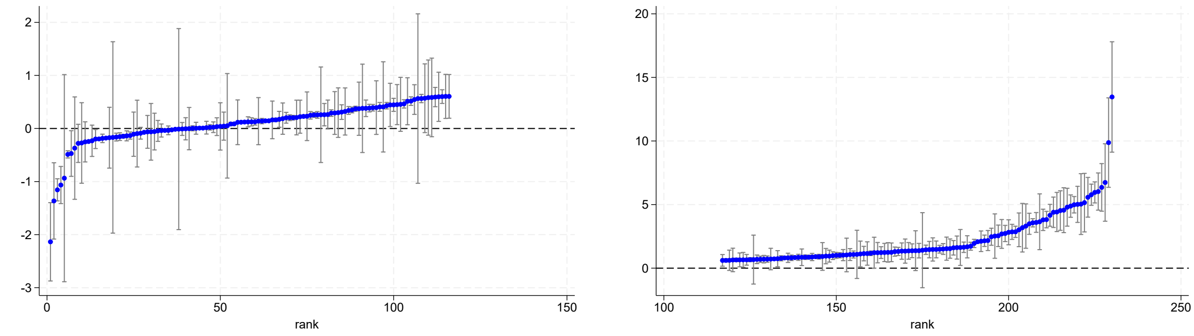

(a) Lowest half of coefficients

(b) Highest half of coefficients

Notes: Coefficient and 95% confidence intervals from a probit regression on the probability that a link disappears next period, where the main regressor is departure from gravity. The latter is the signed value of the error term of the industry estimated from equation (1) by PPML in 2018, rescaled by trade value. For exposition purposes panel (a) shows the lower half of coefficients and panel (b) the higher half. The horizontal axis is sorted by the value of the coefficient.

Table A.13: Predicted trade flows in Denmark

	Mean	SD	Obs
OLS			
Log f_{ijp}	7,4	1,8	413,914
Buyer per Seller	14,2	14,0	381,136
Seller per Buyer	45,2	63,8	381,136
PPML			
Log f_{ijp}	-0,8	4,2	335,540
Buyer per Seller	12,5	12,0	315,979
Seller per Buyer	38,8	57,1	315,979
Log D_{ij}	4,5	1,0	413,914

Notes: Averages and standard deviations are calculated in two steps. First, we compute these metrics across firms within each year. Next, we average the yearly metrics across all years. Under OLS gravity estimates, there are 413,914 firm-pair-product-year non-zero predicted trade flows and 381,136 firm-pair-year non-zero predicted trade flows. Under PPML gravity estimates, there are 335,540 firm-pair-product-year non-zero predicted trade flows and 315,979 firm-pair-year non-zero predicted trade flows. f_{ijp} is value sold by firm i to firm j in product p and D_{ij} is the geographical distance between both firms at time t.

Table A.14: Export Demand Shocks, Domestic Sales, and Exports

	Dependent Variable: Growth Rate of Sales (Domestic or Exports)					
	(Dom.)	(Dom.)	(Dom.)	(Exp.)	(Exp.)	(Exp.)
Export Exposure	-0.004	-0.004	-0.004	0.095***	0.095***	0.095***
	(0.005)	(0.005)	(0.005)	(0.010)	(0.010)	(0.010)
Log Change in Buyers Size	0.064***	0.063***	0.063***	0.012	0.012	0.012
	(0.005)	(0.005)	(0.005)	(0.009)	(0.009)	(0.009)
Indirect Export Exposure		0.021***	0.021***		0.006	-0.001
		(0.005)	(0.005)		(0.007)	(0.007)
Grav. Est.	-	OLS	PPML	-	OLS	PPML
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	31565	31565	31565	27005	27005	27005
R^2	0.11	0.11	0.11	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Columns denoted with (Dom.) use the growth rate of domestic sales as the dependent variable and columns denoted with (Exp.) use the the growth rate of exports.

Table A.15: Domestic Demand Shocks, Domestic Sales, and Exports

	Dependent Variable: Growth Rate of Sales (Domestic or Exports)							
	(Dom.)	(Dom.)	(Dom.)	(Dom.)	(Exp.)	(Exp.)	(Exp.)	(Exp.)
Exposure to Buyer's Industry Shocks	0.115***	0.126***	0.102***	0.115***	0.037	0.088*	0.034	0.085*
	(0.025)	(0.024)	(0.025)	(0.024)	(0.050)	(0.048)	(0.050)	(0.048)
Log Change in Buyers Size			0.058***	0.058***			0.013	0.013
			(0.005)	(0.005)			(0.009)	(0.009)
Grav. Est.	OLS	PPML	OLS	PPML	OLS	PPML	OLS	PPML
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	42698	42698	42698	42698	27370	27370	27370	27370
R^2	0.12	0.12	0.13	0.13	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Columns denoted with (Dom.) use the growth rate of domestic sales as the dependent variable and columns denoted with (Exp.) use the the growth rate of exports.

A.1 Robustness and Placebo

Table A.16: Domestic Sales and Exposure to Buyer's Sales (OLS)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Log Change in Indirect Exposure	0.014***	-0.001	0.021***	0.021***	0.021***	0.022***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Grav. Est.	OLS	OLS	OLS	OLS	OLS	OLS
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	42698	42698	42698	42698	42698	42698
R^2	0.12	0.12	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (3) of Table A.20.

Table A.17: Domestic Sales and Exposure to Buyer's Sales (PPML)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Log Change in Indirect Exposure	0.012***	-0.001	0.007***	0.005***	0.008***	0.019**
	(0.002)	(0.036)	(0.001)	(0.001)	(0.001)	(0.009)
Grav. Est.	PPML	PPML	PPML	PPML	PPML	PPML
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	42698	42698	42698	42698	42698	42698
R^2	0.12	0.12	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (4) of Table A.20.

Table A.18: Domestic Sales and Exposure to Buyer's Exports (OLS)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Log Change in Indirect Exposure	0.006***	-0.000	0.015***	0.015***	0.015***	0.015***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Grav. Est.	OLS	OLS	OLS	OLS	OLS	OLS
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	42698	42698	42698	42698	42698	42698
R^2	0.12	0.12	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (3) of Table 4.

Table A.19: Domestic Sales and Exposure to Buyer's Exports (PPML)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Log Change in Indirect Exposure	0.005***	-0.001	0.005***	0.003***	0.005***	0.015**
	(0.001)	(0.023)	(0.001)	(0.001)	(0.001)	(0.006)
Grav. Est.	PPML	PPML	PPML	PPML	PPML	PPML
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	42698	42698	42698	42698	42698	42698
R^2	0.12	0.12	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (4) of Table 4.

Table A.20: Domestic Sales and Exposure to Buyer's Exports (PPML)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Log Change in Indirect Exposure	0.005***	-0.001	0.005***	0.003***	0.005***	0.015**
	(0.001)	(0.023)	(0.001)	(0.001)	(0.001)	(0.006)
Grav. Est.	PPML	PPML	PPML	PPML	PPML	PPML
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	42698	42698	42698	42698	42698	42698
R^2	0.12	0.12	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (3) of Table **??**.

Table A.21: Export Demand Shocks and Domestic Sales (OLS)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Export Exposure	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Indirect Export Exposure	0.008	-0.004	0.024***	0.022***	0.020**	0.027***
	(0.005)	(0.009)	(0.009)	(0.009)	(0.008)	(0.009)
Grav. Est.	OLS	OLS	OLS	OLS	OLS	OLS
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	31565	31565	31565	31565	31565	31565
R^2	0.11	0.11	0.11	0.11	0.11	0.11

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (2) of Table 5.

Table A.22: Export Demand Shocks and Domestic Sales (PPML)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Export Exposure	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Indirect Export Exposure	0.008*	0.018	0.015***	0.013**	0.029**	-0.011
	(0.005)	(0.111)	(0.004)	(0.006)	(0.014)	(0.039)
Grav. Est.	PPML	PPML	PPML	PPML	PPML	PPML
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	31565	31565	31565	31565	31565	31565
R^2	0.11	0.11	0.11	0.11	0.11	0.11

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (3) of Table 5.

Table A.23: Domestic Demand Shocks and Domestic Sales (OLS)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Exposure to Buyer's Industry Shocks	0.104***	-0.155**	0.106*	0.150**	0.122*	0.107*
	(0.023)	(0.065)	(0.064)	(0.063)	(0.064)	(0.058)
Grav. Est.	OLS	OLS	OLS	OLS	OLS	OLS
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	42698	42698	42698	42698	42698	42698
R^2	0.12	0.12	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (1) of Table 6.

Table A.24: Domestic Demand Shocks and Domestic Sales (PPML)

	(Ext)	(Min)	(1 SD)	(2 SD)	(3 SD)	(4 SD)
Exposure to Buyer's Industry Shocks	0.072***	-0.026	0.064***	0.056	0.031	0.207
	(0.019)	(0.641)	(0.021)	(0.034)	(0.033)	(0.458)
Grav. Est.	PPML	PPML	PPML	PPML	PPML	PPML
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
# Obs.	42698	42698	42698	42698	42698	42698
R^2	0.12	0.12	0.12	0.12	0.12	0.12

Notes: ***, **, * denote significance at the 1%, 5% and 10% level respectively. Standard errors in parentheses. Each column corresponds to a different predicted network used to construct the indirect exposure measures. Column "(Ext)" reports results using the baseline gravity predictions, where intensive-margin predicted trade flows are set to zero whenever the predicted probability of a link is below 80 percent. Column "(Min)" reports results using a predicted network constructed by assigning the minimum estimated gravity coefficients (across sectors) to all sectors, while maintaining the 80 percent threshold rule. Columns "(1 SD)"–"(4 SD)" report results using predicted networks constructed from gravity coefficients that are randomly perturbed by adding or subtracting N times the cross-sector standard deviation of each coefficient, with $N \in \{1, 2, 3, 4\}$, and applying the same 80 percent threshold rule. The specification used corresponds to column (2) of Table 6.