

The Origins of Firm Heterogeneity: A Production Network Approach

*Andrew B. Bernard, Emmanuel Dhyne, Glenn Magerman, Kalina Manova,
Andreas Moxnes*

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

The Origins of Firm Heterogeneity: A Production Network Approach

Abstract

This paper quantifies the origins of firm size heterogeneity when firms are interconnected in a production network. Using the universe of buyer-supplier relationships in Belgium, the paper develops a set of stylized facts that motivate a model in which firms buy inputs from upstream suppliers and sell to downstream buyers and final demand. Larger firm size can come from high production capability, more or better buyers and suppliers, and/or better matches between buyers and suppliers. Downstream factors explain the vast majority of firm size heterogeneity. Firms with higher production capability have greater market shares among their customers, but also higher input costs and fewer customers. As a result, high production capability firms have lower sales unconditionally and higher sales conditional on their input prices. Counterfactual analysis suggests that the production network accounts for more than half of firm size dispersion. Taken together, our results suggest that multiple firm attributes underpin their success or failure, and that models with only one source of firm heterogeneity fail to capture the majority of firm size dispersion.

JEL-Codes: F100, F120, F160.

Keywords: production networks, productivity, firm size heterogeneity.

Andrew B. Bernard
Tuck School of Business at Dartmouth
USA - Hanover, NH 03755
andrew.b.bernard@tuck.dartmouth.edu

Emmanuel Dhyne
National Bank of Belgium
Brussels / Belgium
Emmanuel.Dhyne@nbb.be

Glenn Magerman
ECARES - Université Libre de Bruxelles
Brussels / Belgium
glenn.magerman@ulb.ac.be

Kalina Manova
University College London
London / United Kingdom
k.manova@ucl.ac.uk

Andreas Moxnes
University of Oslo / Norway
andreamo@econ.uio.no

October 2018

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreements No 715147 and No 724880). The views expressed here are those of the authors and do not necessarily reflect those of the National Bank of Belgium. We thank Julian di Giovanni, Isabelle Mejean, Jasmine Xiao, Thomas Åstebro and two anonymous referees at the National Bank of Belgium for helpful comments and valuable discussion.

1 Introduction

Why are firms large or small? Even within narrowly defined industries, there is evidence of massive dispersion in firm outcomes such as revenue, employment, labor productivity and measured total factor productivity (see Syverson, 2011 for a recent overview). In Belgium, a firm at the 90th percentile of the size distribution has turnover more than 34 times greater than a firm at the 10th percentile in the same industry.¹ Understanding the origins of firm size heterogeneity has important micro- and macro-economic implications. At the micro level, bigger firms perform systematically better along many dimensions, including survival, innovation, and participation in international trade (e.g., Bernard et al., 2012). At the macro level, the skewness and granularity of the firm size distribution affect aggregate productivity, the welfare gains from trade, and the impact of idiosyncratic and systemic shocks (e.g., Pavcnik, 2002, Gabaix, 2011, di Giovanni et al., 2014, Melitz and Redding, 2015 and Gaubert and Itskhoki, 2016).

While the literature has made progress in identifying underlying firm-specific supply- and demand-side factors driving firm size (e.g., Hottman et al., 2016), much less is known about the role of firm-to-firm linkages in production networks. In particular, the focus has been on one-sided heterogeneity in either firm productivity on the supply side (e.g., Jovanovic, 1982, Hopenhayn, 1992, Melitz, 2003, Luttmer, 2007) or final consumer preferences on the demand side (e.g., Foster et al., 2016, Fitzgerald et al., 2016). To the extent that the literature has considered firm-to-firm trade, it has typically remained anchored in one-sided heterogeneity by assuming that firms source inputs from anonymous upstream suppliers or sell to anonymous downstream buyers, without accounting for the heterogeneity of all trade partners in the production network.

This paper examines how buyer-supplier connections in a complete production network are related to the firm size distribution.² The basic premise of the analysis is intuitive: Firms can be large because they have inherently attractive capabilities such as productivity or product quality, because they interact with more or better buyers, and/or because they are particularly well matched to their buyers. Moreover, firms can have high product quality or low marginal costs if they have good inherent capabilities, or if they buy inputs from high-quality, efficient, and/or well-matched suppliers. There may be higher-order effects in a production network as well, because the customers of the customers (and so on) of any one firm may ultimately also matter for that firm's economic performance.

The paper makes four main contributions. First, we document new stylized facts about

¹ Averaged across all NACE 4-digit industries in 2014.

² Throughout the paper, firm size, sales, revenue and turnover are used interchangeably.

a complete production network using data on the universe of firm-to-firm domestic transactions in Belgium, and present the first extensive analysis of how upstream, downstream and final demand heterogeneity translate into firm size heterogeneity. Second, we provide a theoretical framework with minimal assumptions on production and demand that relates firm size to firm-specific characteristics, buyer and supplier characteristics, and buyer-supplier match characteristics. Third, the model allows us to develop a new methodology for inferring firm primitives from production network data, overcoming a reflection problem that is fundamental in all networked environments. Finally, we close the model in general equilibrium and simulate counterfactual shocks to firm capabilities and intermediate input shares to assess the role of the production network.

The paper leverages unique data on firm-to-firm sales between virtually all firms in a domestic production network. A key feature of production network data is that sales from firm i to j can be decomposed into seller-, buyer- and match specific components (fixed effects), similar to the analysis of employer-employee data (Abowd et al., 1999). High dispersion in seller effects means that firms vary in how much they sell to their customers, controlling for demand by those customers, i.e. firms differ in their average market share across customers. Conversely, high dispersion in buyer effects means that some firms match with large customers while others do not, leading to larger sales even as the average market share remains the same.

Given estimates of these fixed effects, the total sales of a firm can be decomposed into three distinct factors: (i) an upstream component that captures the firm's ability to obtain large market shares across its customers, (ii) a downstream component that captures the firm's ability to attract many and/or large customers, and (iii) a final demand component that captures the firm's ability to sell outside the domestic network, i.e., to final demand or to foreign customers.

The results are striking: 81 percent of the variation in firm sales within narrowly defined (4-digit NACE) industries is associated with the downstream component, while the upstream component contributes only 18 percent. The variation in firm size is largely unrelated to the final demand component. These findings imply that trade in intermediate goods and firm-to-firm connections are essential to understanding firm performance and, consequently, aggregate outcomes.

Motivated by these results and additional stylized facts on the Belgian production network, we develop a quantitative theoretical framework that features two-sided firm heterogeneity in an input-output production network. In the model, firms use a constant elasticity of substitution production technology that combines labor and inputs from upstream suppliers. Firms sell their output to final consumers and to domestic producers. Firms differ in

their *production capability* (a combination of efficiency and quality conditional on the firm’s connections), as well as in their network characteristics - their upstream and downstream connections and their match quality with each partner. Marginal costs, employment, prices, and sales are endogenous outcomes because they depend on the outcomes of all other firms in the economy. A link between two firms increases the total sales of both the seller and the buyer; for the seller this occurs mechanically because it gains a customer, while for the buyer this arises because a larger supplier base implies greater opportunities to source cheaper or higher-quality inputs.

The model yields three main insights. First, the estimated buyer and seller effects are a function of the fixed effects of all other firms - in other words, the fixed effects are contaminated by a reflection problem, in the spirit of Manski (1993). A low labor share exacerbates the reflection problem because purchased inputs then constitute a larger share of marginal costs, and suppliers’ costs consequently matter more for sales. Using the model we show how to overcome the reflection problem and isolate firms’ production capability which is independent of characteristics of other firms in the network (Proposition 1). Second, one can recover a firm’s input price index from the fixed effects, and this price index respects the general equilibrium constraints of the model (Proposition 2).³ Third, there is a unique mapping from the estimated buyer and seller effects to model parameters (Proposition 3). This powerful result implies that our methodology can be applied in a variety of settings to discipline and calibrate network models.

The theoretical and empirical framework is silent about the network formation process and instead conditions on the observed equilibrium network. This approach minimizes the assumptions required. However, there is a potential concern that the estimated fixed effects are a function of the network formation process itself, similar to the concern about conditional endogenous mobility in the analysis of employer-employee data. We provide evidence, and develop new statistical tools, to test the identifying assumptions. Overall, we find empirical support for our approach (Section 6).

These insights allow us to give the initial firm size decomposition a structural interpretation and further decompose the upstream and downstream components. We draw two main conclusions. First, larger firms have lower input prices, more customers, and higher market shares among their customers relative to smaller firms, consistent with the previous literature (e.g., Bøler et al., 2015). At the same time, firms with higher production capability have greater market shares among their customers, but also higher input costs and fewer customers. As a result, more productive firms have lower sales unconditionally, but nevertheless

³ In general equilibrium, the input price index is the solution to a fixed point problem. The input price index obtained from the fixed effects is proportional to the general equilibrium solution.

higher sales conditional on their input costs. These empirical results are difficult to reconcile with canonical models, and suggest that multiple dimensions of firm activity underpin their success or failure. For example, one interpretation of our findings is that firm attributes that matter for finding customers and suppliers (e.g., managerial talent and marketing capacity) are orthogonal, or negatively related, to firm attributes that determine sales conditional on a match (e.g., productivity or quality).

Second, most of the downstream variance in network sales is determined by the number of buyers and the allocation of activity towards well-matched partners of high quality, rather than by average partner capability. The main reason why the production network enables firms to sell more downstream is through the number of buyers, not because their buyers tend to purchase more intermediates. Conversely, nearly all the upstream variation is driven by own production capability rather than input purchases from the network.

Finally, we exploit the general equilibrium structure of the model and perform two counterfactual exercises to evaluate the contribution of the production network to firm size dispersion. In the first counterfactual, heterogeneity in production capability is shut down. This eliminates direct heterogeneity due to variation in a firm's production capability, but also removes heterogeneity in the capability of upstream suppliers and downstream customers. The remaining variation in firm size then comes from the network itself, via differences in the number of connections and their match quality across firms. In the second counterfactual, we focus on a key parameter in our model, the cost share of inputs purchased from the network (goods and services) in total production costs (the network input share). The results from the counterfactuals provide additional evidence for the importance of the production network for firm size variation. Even after eliminating all traditional sources of firm heterogeneity, the network explains over half (56%) the variance in firm sales. Firm size dispersion also widens when we alternatively increase the network input cost share.

This paper contributes to several strands of literature. Most directly, the paper adds to the large literature on the extent, causes and consequences of firm size heterogeneity. The vast dispersion in firm size has long been documented, with a recent emphasis on the skewness and granularity of firms at the top end of the size distribution (e.g., Gibrat, 1931, Syverson, 2011). This interest is motivated by the superior growth and profit performance of bigger firms at the micro level, as well as by the implications of firm heterogeneity and superstar firms for aggregate productivity, growth, international trade, and adjustment to various shocks (e.g., Gabaix, 2011, Bernard et al., 2012, Freund and Pierola, 2015, Gaubert and Itskhoki, 2016, Oberfield, 2018).

Traditionally, this literature has analyzed own-firm characteristics on the supply side as the driver of firm size heterogeneity. The evidence indicates an important role for firms'

production efficiency, management ability, and capacity for quality products (e.g., Jovanovic, 1982, Hopenhayn, 1992, Melitz, 2003, Sutton, 2007, Bender et al., 2016, Bloom et al., 2017). Recent work has built on this by also considering the role of either upstream suppliers or downstream demand heterogeneity, but not both. Results suggest that access to inputs from domestic and foreign suppliers matters for firms' marginal costs and product quality, and thereby performance (e.g., Goldberg et al., 2010, Manova et al., 2015, Fieler et al., 2018, Bernard et al., forthcominga, Antràs et al., 2017, Boehm and Oberfield, 2018), while final consumer preferences affect sales on the demand side (e.g., Foster et al., 2016, Fitzgerald et al., 2016).

By contrast, we provide a comprehensive treatment of both own firm characteristics and production network features, on both the upstream and the downstream sides. The paper is related to Hottman et al. (2016) who also find that demand-side factors such as variation in firm appeal and product scope rather than prices (marginal costs) drive firm size dispersion. However, as these authors do not observe the production network, they cannot distinguish between the impact of serving more customers, attracting better customers, and selling large amounts to (potentially few) customers. Since they have no information on the supplier margin, they also cannot compare own versus network supply factors. On the other hand, while rich in network features, our data do not provide information on prices or products and thus do not allow for a comparable decomposition into firm appeal and product scope.

The paper also adds to a growing literature on buyer-supplier production networks (see Bernard and Moxnes, forthcoming for a recent survey). Bernard et al. (forthcominga) study the impact of domestic supplier connections on firms' marginal costs and performance in Japan, whereas Bernard et al. (2018), Eaton et al. (2016) and Eaton et al. (2018) explore the matching of exporters and importers using data on firm-to-firm trade transactions for Norway, US-Colombia and France, respectively. While we confirm some of the findings in these papers about the distributions of buyers and suppliers, we examine transaction-level data on a complete domestic production network and focus on the implications of two-sided heterogeneity and production networks for the firm size distribution. Using the Belgian production network data, Magerman et al. (2016) analyze the contribution of the network structure of production to aggregate fluctuations, while Tintelnot et al. (2017) study the impact of trade on the domestic production network. In recent work, Baqaee and Farhi (2018a), Baqaee and Farhi (2018b) and Lim (2017) study the impact of microeconomic shocks on macroeconomic outcomes in networked environments.

Finally, the methodology in this paper is related to the econometrics of two-sided heterogeneity in other economic contexts (see Bonhomme et al., 2017 for a review). In particular,

we estimate seller and buyer fixed effects from production network data in a log linear model that is conditional on the observed network. A related recent contribution is Kramarz et al. (2016), who estimate buyer and seller effects in a bipartite trade network. Our work also builds on employer-employee econometric models in the labor literature (e.g., Abowd et al., 1999, Card et al., 2013). However, each economic agent plays a unique role in the labor market - either a firm or a worker - such that both panel data and worker transitions across firms are necessary to identify the employer, employee and match effects. We extend the existing empirical bipartite matching literature along important dimensions. Our setting pertains to a many-to-many non-bipartite network, as each firm is both a buyer and a supplier in a production network. This permits the identification of the fixed effects in the cross-section, such that they are not required to remain constant over time. Moreover, it attenuates the incidental parameter problem as the number of suppliers per customer and the number of customers per supplier is relatively large.

The rest of the paper is organized as follows. Section 2 introduces the data and presents stylized facts about the Belgian production network. Section 3 agnostically decomposes firm sales into upstream, downstream and final demand components. Section 4 develops a theoretical framework with heterogeneous firms in a production network which provides a structural interpretation of the decomposition in Section 3 and a finer model-based decomposition of the upstream and downstream components. Section 5 presents the results of this model-based decomposition, and Section 6 discusses potential issues with the empirical framework. Section 7 introduces a general equilibrium formulation of the model and performs counterfactual exercises. The last section concludes.

2 Data and Stylized Facts

2.1 Data

We exploit several comprehensive data sources on annual firm operations in Belgium: (i) the NBB B2B Transactions Dataset, containing the universe of domestic firm-to-firm sales relationships, (ii) annual accounts, with typical firm characteristics for firms above a minimum size threshold, (iii) VAT declarations, with more limited firm characteristics for small firms, and (iv) the Crossroads Bank of Enterprises dataset, containing firms' industry affiliation. Unique firm identification numbers allow us to unambiguously match these datasets. We can thus examine an entire economy in unprecedented detail observing the complete domestic production network with information on seller firm characteristics, buyer firm characteristics, and seller-buyer transaction values. We use the 2014 cross-section in the main analysis and data for the 2002-2014 period in robustness exercises and extensions.

The primary data source is the NBB B2B Transactions Dataset, administered by the National Bank of Belgium (NBB), which documents both the extensive and the intensive margins of domestic buyer-supplier relationships in Belgium. The dataset reports the sales relationships between any two VAT-liable enterprises across all economic activities within Belgium.⁴ In particular, an observation is the sum m_{ij} of sales invoices (in euros, excluding any value-added tax due) from enterprise i to enterprise j in a given calendar year. Observations are directed, i.e. $m_{ij} \neq m_{ji}$. Coverage is quasi universal, as all relationships with annual sales of at least 250 euros must be reported, and pecuniary sanctions on late and erroneous reporting ensure high data quality.

Data on total sales (turnover), total input purchases, employment and labor costs come from firm annual accounts maintained by the Central Balance Sheet Office (CBSO) at the NBB.⁵ Annual accounts are collected by fiscal year and have been annualized to match the calendar year in the NBB B2B data. Since there is a firm-size threshold for reporting turnover and input purchases to CBSO, data on these two variables for small firms below the threshold comes from firms' VAT declarations. We keep only firms with at least one full-time equivalent employee. The main economic activity of each enterprise is available at the NACE 4-digit level (harmonized over time to the NACE Rev. 2 (2008) version).


Firms' sales to final demand is the difference between their turnover and the sum of all their B2B sales to other enterprises in the domestic production network. Final demand thus contains sales to final consumers at home, potentially unobserved links in B2B with very small transaction values, and exports. Firms' purchases from outside the observed production network (including imports) is the difference between their total input costs and the sum of all their B2B purchases.⁶ The labor share in production at the NACE 4-digit level is computed as the sum of total employment expenses across all firms in an industry divided by total production costs in that industry. Similarly, average wages by industry are calculated as the sum of total labor costs divided by total employment. Further details on data coverage and preparation are in Appendix A.

⁴ See Dhyne et al. (2015) for details on the construction of this dataset. "Enterprise" and "firm" are used interchangeably in this paper. The unit of observation is the unique firm identification number, i.e. the legal entity of the enterprise.

⁵ Total input purchases are the sum of material and service inputs, and include both new inputs and net changes in input stocks. Employment is reported as average full-time equivalent employees. Total labor costs include wages, social security, and pension contributions.

⁶ The estimation procedure requires at least two customers or suppliers to identify the seller or buyer effect respectively. Some links (less than 1%) therefore drop out from the analysis (see Section 3).

Figure 1: Firm sales distribution (2014).



2.2 Stylized Facts

This section documents three stylized facts about firm size and firm linkages in the Belgian domestic production network.⁷ These facts provide evidence that buyer-supplier relationships are key to understanding firm size dispersion and motivate the subsequent theoretical and empirical analyses. We present cross-sectional evidence for 2014, but the patterns are stable over each year in 2002-2014.

Fact 1. *The distributions of firms' total sales, buyer-supplier connections, and buyer-supplier bilateral sales exhibit high dispersion and skewness.*

Firm size varies dramatically in Belgium, as in other countries. Table 1 provides summary statistics for firm sales in 2014, both overall and within six broad sectors (primary and extraction, manufacturing, utilities, construction, market services, and non-market services).⁸ Across the 109,739 firms with sales data that are active in the production network, average turnover is 6.8 million euros, with a standard deviation of 145 million euros. Similar patterns hold within each broad sector category.

The cross-sectional distribution is extremely skewed. Overall, firms at the 90th percentile generate turnover over 34 times higher than firms at the 10th percentile, while the top 10% of firms account for 84% of total sales. Although there is some variation in average firm size across sectors, the dispersion is similar, with large firms being up to four orders of

⁷ A subset of these stylized facts echo patterns established for the extensive margin of firm-to-firm linkages in the domestic production network in Japan (Bernard et al., forthcoming) and for both the extensive and the intensive margins of firm-to-firm export transactions in Norway (Bernard et al., 2018). The dispersion in transaction values in a buyer-supplier production network was first documented in the Belgian data by Dhyne et al. (2015).

⁸ See Table 14 in Appendix A for the classification of industry groups at the 2-digit NACE level.

Table 1: Firm sales (million euros, 2014).

Sector	NACE	N	Mean	St Dev	10th	50th	90th	95th	99th
Primary & Extraction	01-09	3,061	12.0	432.6	0.2	0.8	4.8	9.5	52.0
Manufacturing	10-33	18,077	14.4	250.8	0.2	1.1	13.8	34.6	201.8
Utilities	35-39	897	39.2	442.9	0.3	1.9	25.7	68.6	495.6
Construction	41-43	20,201	2.3	13.4	0.2	0.6	3.6	6.9	25.9
Market Services	45-82	65,175	5.5	79.9	0.2	0.8	6.3	13.4	63.9
Non-Market Services	84-99	2,328	2.2	26.3	0.1	0.3	2.6	5.5	24.9
All		109,739	6.8	145.1	0.2	0.8	6.6	14.3	78.4

Note: Summary statistics for the matched CBSO-B2B data. 10th, 50th, etc. refers to values at the 10th, 50th, etc. percentile of the distribution.

Table 2: Number of firm buyers and suppliers (2014).

	N	Mean	St Dev	10th	25th	50th	75th	90th	95th	99th
# of buyers	590,271	29.3	394.0	1	1	4	11	42	98	400
# of suppliers	840,607	20.6	49.5	1	3	9	22	46	71	177

Note: Summary statistics for the B2B data. 10th, 25th, etc. refers to values at the 10th, 25th, etc. percentile of the distribution.


magnitude bigger than their industry mean, as shown in Figure 1a. The histogram in Figure 1b illustrates the full firm size distribution, after demeaning at the NACE 4-digit industry level. Even within narrowly defined industries, these patterns remain.

Turning to firm-to-firm connections in the domestic production network, the number of downstream customers per seller (out-degree) and the number of upstream suppliers per buyer (in-degree) are also very skewed. In 2014, there are 17.3 million sales relationships among 859,733 firms within Belgium.⁹ Of these, 590,271 enterprises sell to other firms in the network, while 840,607 buy from other firms in the network. 31.5% of firms sell only to final demand, while a small minority of 2.2% do not purchase inputs from the domestic production network (or do so in an amount less than 250 euros).

Table 2 summarizes the overall distribution of buyer and supplier connections. Across all sellers, the average number of customers is 29.3, with a standard deviation of 394. Across all buyers, the average number of suppliers is 20.6, with a standard deviation of 49.5. The distribution of buyers per seller is more dispersed than that of suppliers per buyer. Firm-to-firm

⁹ The number of firms in the B2B production network is larger than the number of firms in the matched B2B-CBSO sample with turnover data, because B2B also contains small firms that do not submit annual accounts to CBSO.

Figure 2: Distribution of firm buyer and supplier connections (2014).



Note: The number of customers and suppliers is demeaned at the NACE 4-digit level.

links in the network are also highly concentrated among a few very connected participants: The median number of customers and suppliers is only 4 and 9 respectively, while the top 1 percent of firms transact with more than 400 buyers and 177 sellers. This dispersion and skewness across firms within NACE 4-digit industries is also evident in the histograms in Figure 2. Again, firms with the most customers or suppliers are several orders of magnitude more connected than the average firm in their industry.


The intensive margin of firm-to-firm bilateral sales is also very dispersed and skewed, with the vast share of economic activity concentrated in a small number of buyer-supplier transactions.¹⁰ The mean transaction amounts to 28,893 euros. At the same time, the median purchase totals only 1,392 euros, and the top 10% of relationships account for 92% of all domestic firm-to-firm sales by value.

Fact 2. *Bigger firms have more buyers and suppliers.*

A sharp pattern in the data is that bigger firms interact with more buyers and suppliers in the production network. Figure 3a plots the fitted line and 95% confidence band based on a local polynomial regression of firm turnover on the number of firm downstream customers, on a log-log scale. Both variables have been demeaned by their NACE 4-digit industry average, such that the latter corresponds to the point with coordinates (1,1) in the graph. The dotted line represents the 45 degree line that would obtain if the elasticity of turnover with respect to the number of customers were 1. Figure 3b repeats the exercise for the relationship

¹⁰ See Table 17 in Appendix B.

Figure 3: Firm size and number of buyers and suppliers (2014).



Note: Local polynomial regressions for firm turnover and number of customers and suppliers are demeaned at the NACE 4-digit level. Graphs are trimmed at the 0.1st and 99.9th percentiles of the number of customers and suppliers respectively. The dotted line represents the 45 degree line. Implied elasticities and R-squared from OLS regressions with NACE 4-digit industry fixed effects are reported in the lower left corner of each graph.

between firm sales and number of upstream suppliers. Both figures display tightly estimated upward-sloping lines.¹¹


Fact 3. *The distribution of sales across buyers does not vary with the number of buyers. The distribution of purchases across suppliers widens with the number of suppliers.*

Figure 4a illustrates the dispersion of downstream sales across buyers within a seller. For each firm with at least 10 customers, we take the 10th, 50th and 90th percentile values of its bilateral sales, demean by its NACE 4-digit industry, and plot the fitted lines from local polynomial regressions of these percentile values against firms' out-degree, including 95% confidence bands. The three lines are almost parallel, albeit slightly declining. The spread of sales to the 10th percentile, median and 90th percentile customers is essentially the same for firms with 100 customers and for firms with 10 customers. The slight decline is consistent with the out-degree elasticity of turnover being less than one in Figure 3a. Together with Fact 2, this suggests that larger sellers have higher sales primarily because they serve more customers, but they do not vary sales more across buyers.

Figure 4b shows the distribution of input purchases across upstream suppliers within a buyer. For each firm with at least 10 input providers, we obtain the 10th, 50th and

¹¹ Results are similar using downstream sales within the B2B domestic network instead of total turnover in Appendix B.

Figure 4: Sales distribution across buyers and suppliers within firms.



Note: Local polynomial regressions for the value of firm-to-firm transactions at the 10th, 50th and 90th percentile of the distribution. Firm-to-firm sales are demeaned by the NACE 4-digit industry of the seller and the customer in each figure respectively. The number of customers and suppliers respectively has been trimmed at the 0.1st and 99.9th percentiles.

90th percentile values of its bilateral purchases, demean by its NACE 4-digit industry, and graph the fitted lines from local polynomial regressions of these percentile values against firms' in-degree, with 95% confidence bands. While purchases from the median supplier are essentially unchanged across firms with broad and narrow supplier bases, firms that source inputs from more suppliers systematically buy more from their larger suppliers and less from their smallest.

These stylized facts signal an important role for (i) downstream input demand relative to final output demand, (ii) the number of buyers and suppliers of a firm, (iii) seller and buyer firm characteristics, and (iv) seller-buyer match characteristics. Motivated by these facts, we develop a theoretical framework that accommodates them by introducing two-sided firm heterogeneity in an input-output production network. Importantly, this model enables the decomposition of the variation in firm size into economically meaningful components related to both firms' own characteristics and their participation in the production network.

3 Initial Decomposition

In this section, we develop an exact decomposition of firm sales into upstream, downstream, and final demand margins. The downstream component reflects characteristics of a firm's customers (i.e., their number and size), while the upstream component captures firm characteristics that remain constant across customers (i.e., average sales to customers, controlling

for their size). Final demand includes factors unrelated to the domestic production network, such as sales to final consumers or foreign customers. This approach exploits the granularity of firm-to-firm transactions to inform the micro-foundations of firm size in a way that would be impossible without production network data.

In this part of the paper, the decomposition of firm sales does not rely on a specific model. Sections 4 and 7 develop a theoretical framework that delivers reduced-form expressions consistent with the decomposition approach here. The theory also gives the reduced-form parameters clear economic interpretation.

3.1 Methodology

We start by estimating buyer, seller and buyer-seller match effects using data on sales between firms in the production network. The specification is a two-way fixed effects regression for (log) firm-to-firm sales:

$$\ln m_{ij} = \ln \psi_i + \ln \theta_j + \ln \omega_{ij}, \quad (1)$$

where $\ln m_{ij}$ is log sales from i to j relative to the grand mean, $\ln m_{ij} \equiv \ln \ddot{m}_{ij} - G$, where $\ln \ddot{m}_{ij}$ is non-demeaned log sales and G is the log of the grand mean. In this OLS regression, the seller effect $\ln \psi_i$ is identified by the magnitude of sales by i to all customers j , controlling for total purchases by j . The seller effect is thus related to the average market share of i among her customers. Intuitively, attractive sellers account for a large share of input expenditures across all their customers and receive a high $\ln \psi_i$. Analogously, the buyer effect $\ln \theta_j$ is identified by the magnitude of purchases by j from all suppliers i , controlling for total sales of i . Intuitively, attractive buyers purchase a disproportionate share of suppliers' sales and receive a high $\ln \theta_j$. A positive residual $\ln \omega_{ij}$ reflects match-specific characteristics that induce a given firm pair to trade more with each other, even if they are not fundamentally attractive trade partners. We assign structural interpretations to the seller, buyer, and match effects in Section 4.

To interpret the variation in ψ_i and θ_j , consider the case where the variation in $\ln m_{ij}$ is only due to ψ_i . Sellers i and i' then differ because i sells more to every customer (while buyers j and j' purchase the same amount from i). Consider next the opposite case where the variance in $\ln m_{ij}$ is only due to θ_j . Sellers i and i' now differ because i happens to match with bigger customers than i' (while sales to a common customer j are identical). In the first case firm heterogeneity is driven by differences in sales ability, while in the second case it is driven by differences in matching ability.

To obtain unbiased OLS estimates, the assignment of suppliers to customers must be exogenous with respect to ω_{ij} , so-called conditional exogenous mobility (Abowd et al., 1999).

This identification assumption, as well as tests for exogenous mobility and functional form relevance, are discussed in Section 6. Overall, we find support for the log linear model and the conditional exogenous mobility assumption.

In order to estimate the two-way fixed effects model, firms must have multiple connections. Specifically, identifying a seller fixed effect requires a firm to have at least two customers, and identifying a buyer fixed effect requires a firm to have at least two suppliers. Therefore, single-customer or single-supplier links are dropped in the estimation procedure. Furthermore, dropping customer A might result in supplier B having only one customer left. Supplier B would then also be removed from the sample. This iterative process continues until a connected network component remains (i.e. a within-projection matrix of full rank), in which each seller has at least two customers and each customer has at least two suppliers. This component is known as a mobility group in the labor literature on firm-employee matches.

A production network is a directed graph because firms are simultaneously buyers and sellers. Production networks are richer and more complex than typical bipartite networks that have been studied extensively, such as the labor market for firms and workers, the marriage market for men and women, or the organ market for donors and recipients. For example, in employer-employee data, worker transitions across firms over time (panel data) are necessary to identify the employer, employee and match effects. An immediate implication is that cross-sectional production network data is sufficient to identify the parameters of interest in our setting. This has two advantages. First, it attenuates the incidental parameter problem as the number of suppliers per customer and the number of customers per supplier is relatively large (see Section 2) compared to e.g. the number of job switchers over time. Second, this setting does not require the typical assumption that the fixed effects are constant over time, which might be increasingly hard to justify as the time dimension grows.

Given estimates of $\Psi = \{\psi_i, \theta_j, \omega_{ij}\}$, firm sales can be exactly decomposed into upstream, downstream, and final demand factors. Total sales of firm i are by construction $S_i = \sum_{j \in \mathcal{C}_i} \ddot{m}_{ij} + \mathcal{F}_i$, where \mathcal{C}_i is the set of firm i 's customers and \mathcal{F}_i is final demand (i.e., sales outside of the domestic network). Combining this with equation (1) yields

$$\ln S_i = G + \ln \psi_i + \ln \xi_i + \ln \beta_i, \quad (2)$$

where $\xi_i \equiv \sum_{j \in \mathcal{C}_i} \theta_j \omega_{ij}$ and β_i is total sales relative to network sales, $\beta_i \equiv S_i / (S_i - \mathcal{F}_i) \geq 1$, i.e. an inverse measure of final demand. The components ψ_i and ξ_i represent upstream and downstream fundamentals that shape firm size, respectively. To fix ideas, consider the case where sales dispersion is only due to variance in ψ_i . Then, large firms have greater market shares among their customers than small firms. Next, consider the case where sales

Table 3: Full vs. Estimation Sample

Full Sample			Estimation Sample			
# Links	# Sellers	# Buyers	Links	Value	Sellers	Buyers
17,304,408	590,271	840,607	99%	95%	74%	88%

Note: Summary statistics for firm-to-firm transactions in the raw B2B data and in the estimation sample.

dispersion is only due to variance in ξ_i . Then, large firms transact with more, bigger, and/or better-matched customers than small firms (while market shares are the same).

Note that all components of equation (2) are known: S_i , β_i and G come directly from the data, while ψ_i and ξ_i are estimated from equation (1). In order to assess the role of each margin, we follow the literature (Eaton et al., 2004, Hottman et al., 2016) and regress each component ($\ln \psi_i$, $\ln \xi_i$, and $\ln \beta_i$) separately on log sales. By the properties of ordinary least squares, those three coefficients will sum to unity, and the coefficient magnitudes will represent the share of the overall variation in firm size explained by each margin (see Appendix D). All observed and constructed variables are first demeaned by their NACE 4-digit industry average, such that systematic variation across industries is differenced out.

3.2 Results

As mentioned above, the estimation sample is a subset of the full sample because firms are required to have at least two customers or suppliers. In practice, the estimation sample covers the vast majority of observations in the production network. This underlines the highly connected structure of the production network across all economic activities, even while it is relatively sparse. For the baseline year, 2014, the estimation spans 17,054,274 firm-to-firm transactions which represent 99% of all links in the data and 95% of their sales value. We thus obtain seller fixed effects for 436,715 firms and buyer fixed effects for 743,326 firms. The characteristics of the initial and estimation samples are given in Table 3.

The results from estimating equation (1) are reported in Table 4. Three patterns stand out. First, the adjusted R^2 from the regression is 0.39, showing that the buyer and seller fixed effects explain a large share of variation in the network data. Second, the variation in the seller effect $\ln \psi_i$ is larger than that in the buyer effect $\ln \theta_j$. Third, the correlation between the fixed effects is close to zero, suggesting that high- ψ_i sellers match with both high- θ_j and low- θ_j buyers (and vice versa).

The results from Table 4 inform us about the variation in transaction values, m_{ij} , but

Table 4: Buyer and Seller Effects

	N	$\frac{\text{var}(\ln\psi_i)}{\text{var}(\ln\psi_i+\ln\theta_j)}$	$\frac{\text{var}(\ln\theta_j)}{\text{var}(\ln\psi_i+\ln\theta_j)}$	$\frac{2\text{cov}(\ln\psi_i,\ln\theta_j)}{\text{var}(\ln\psi_i+\ln\theta_j)}$	Adjusted R^2
$\ln m_{ij}$	17,054,274	0.66	0.32	0.02	0.39

Note: The table reports the (co)variances of the estimated seller and buyer fixed effects from equation (1). Estimation is based on the high-dimensional fixed effects estimator from Correia (2016).


not about the variation in firm sales, S_i , which is given by the exact firm sales decomposition in equation (2). To operationalize (2), we use the estimates of $\Psi = \{\psi_i, \theta_j, \omega_{ij}\}$ and balance sheet data on total sales S_i .¹² We calculate the network sales ratio, β_i , based on total sales and the sum of sales to other firms in the domestic network from B2B ($\sum_{j \in \mathcal{C}_i} \ddot{m}_{ij}$). Since firm-to-firm links are only available within Belgium, sales to foreign firms (exports) are classified as part of final demand.

Table 5 reports the results from the decomposition of total firm sales. The downstream side accounts for 81% of the size dispersion across firms, upstream fundamentals account for 18%, and final demand explains only 1%. The upstream factor $\ln \psi_i$ represents, loosely speaking, the average market share of i among its customers. The relatively small role for upstream fundamentals means that average market share is not strongly correlated with total firm sales. In other words, being an important supplier to one's customers is only weakly related to overall firm success. This does not mean, however, that supply-side factors in general are unimportant in explaining firm size. Rather, the results suggest that supply-side factors that are orthogonal to our upstream component might be important. Examples of such factors are efficiency in marketing or skills in finding and attracting a customer base. On the other hand, relative differences in final demand across firms, as captured by the ratio of total sales to sales to final consumers, $\ln \beta_i$, account for an economically negligible 1% of the overall variation in firm size. Thus large firms are not systematically selling relatively more (or less) to final demand than small firms.

The importance of each component can be illustrated using a binned scatterplot. In Figure 5, we group log sales into 50 equal-sized bins, compute the mean of log sales and the components $\ln \psi_i$, $\ln \xi_i$ and $\ln \beta_i$ within each bin, and then create a scatterplot of these data points. The result is a non-parametric visualization of the conditional expectation function, where the sum of the three components on the vertical axis equals log sales on the horizontal axis. Again, the dominance of the downstream component is apparent, and the relationship

¹² The population of firms in the balance sheet data and production network data is partly non-overlapping. All firms with estimated fixed effects enter the calculation of $\xi_i \equiv \sum_{j \in \mathcal{C}_i} \theta_j \omega_{ij}$, even if they are not in the balance sheet data themselves.

Figure 5: Overall Decomposition.



Note: This binned scatterplot groups firms into 50 equal-sized bins by log sales, computes the mean of log sales and the components $\ln \psi_i$, $\ln \xi_i$ and $\ln \beta_i$ within each bin, and graphs these data points. The result is a non-parametric visualization of the conditional expectation function.

Table 5: Overall Decomposition.

	N	Upstream $\ln\psi_i$	Downstream $\ln\xi_i$	Final Demand $\ln\beta_i$
$\ln S_i$	94,330	.18*** (.00)	.81*** (.00)	.01** (.00)

Note: The table reports coefficient estimates from OLS regressions of a firm size margin (as indicated in the column heading) on total firm sales. All variables are first demeaned by their 4-digit NACE industry average. Standard errors in parentheses. Significance: * < 5%, ** < 1%, *** < 0.1%.

is close to linear across the entire distribution of firm sales.

These findings suggest that key to understanding the vast firm size heterogeneity observed in modern economies is how firms manage their sales activities, and specifically how they match and transact with buyers in the production network.

4 Theoretical Framework

This section develops a theoretical framework that serves several purposes. First, the model allows for various sources of firm heterogeneity both on the demand side (e.g., being connected to many or large customers) and on the supply side (e.g., having access to many or cheap intermediate inputs). Second, the framework gives a clear mapping between model parameters and firm-level estimated coefficients from Section 3. Third, the framework permits a finer model-based decomposition of firm sales into various upstream and downstream margins with clear economic interpretations (Section 5.2). And finally, the model can be used for counterfactual analyses (Section 7).

Our starting point is a model in which firms are heterogeneous in productivity or quality. Firms operate in a production network and sell to other firms and to final demand. In addition to productivity or quality, a firm's size can depend on its input prices and on how many and what type of buyers it is connected to. Input prices can be low (and sales high) if firms have many suppliers, low-price, and/or better-matched suppliers. We do not model the firm-to-firm matching decision itself, but rather condition on the equilibrium network structure that is observed in the data. We discuss the empirical implications of this approach in Section 6.

4.1 Technology

To implement our approach, we start with the following production function of firm i :

$$y_i = \kappa_i z_i l_i^{\alpha_i} (u_i^{1-\gamma_i} v_i^{\gamma_i})^{1-\alpha_i},$$

where y_i is output (in quantities), z_i is productivity, l_i is the amount of labor used by firm i , α_i is the labor share, u_i is intermediate inputs purchased from outside the domestic network (e.g., imported inputs), and $\kappa_i > 0$ is a normalization constant.¹³ v_i is a constant elasticity of substitution (CES) domestic network input bundle with associated cost share γ_i :

$$v_i = \left(\sum_{k \in \mathcal{S}_i} (\phi_{ki} \nu_{ki})^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)},$$

where ν_{ki} is the quantity purchased from firm k , \mathcal{S}_i is the set of suppliers to firm i , and $\sigma > 1$ is the elasticity of substitution across suppliers.¹⁴ ϕ_{ki} is a demand shifter that captures the idea that firms (and industries) may have very different production technologies, and that their purchases from a given supplier may vary. The corresponding inverse input price index is $P_i^{1-\sigma} = \sum_{k \in \mathcal{S}_i} (p_{ki}/\phi_{ki})^{1-\sigma}$, where p_{ki} is the price charged by supplier k to firm i . The marginal cost of the firm is then

$$c_i = \frac{w_i^{\alpha_i} (\tilde{w}_i^{1-\gamma_i} P_i^{\gamma_i})^{1-\alpha_i}}{z_i}, \quad (3)$$

where w_i and \tilde{w}_i are the wage and the price of non-network inputs, respectively. We allow for firm-level heterogeneity in most of the parameters of the model (e.g., wages and the labor share). For tractability the elasticity of substitution σ is identical across firms and inputs.¹⁵

4.2 Firm-to-Firm Sales, Total Sales and Total Purchases

Each firm faces demand from other firms as well as from final demand. Given the assumptions about technology, sales from firm i to j are

$$\dot{m}_{ij} = \left(\frac{\phi_{ij}}{p_{ij}} \right)^{\sigma-1} P_j^{\sigma-1} \gamma_j M_j, \quad (4)$$

¹³ In particular, $\kappa_i \equiv \alpha_i^{-\alpha_i} (1-\gamma_i)(1-\alpha_i)^{-(1-\gamma_i)(1-\alpha_i)} \gamma_i (1-\alpha_i)^{-\gamma_i(1-\alpha_i)}$. This normalization maps the production function to the cost function, and simplifies the expression for the cost function without any bearing on our results.

¹⁴ The input bundle v_i captures capital inputs acquired from the network (e.g. machinery, equipment). On the one hand, this may overstate the flow of new capital inputs used in current production, since new investment goods will be used over many periods. On the other hand, this may understate the flow of total capital inputs used in current production, since it ignores the stock of accumulated capital.

¹⁵ For the remainder of the paper, the elasticity of substitution will play no important role, and all results in the paper are independent of the exact value of σ .

where M_j are total intermediate purchases by firm j , and $\gamma_j M_j = \sum_{i \in \mathcal{S}_j} \ddot{m}_{ij}$ corresponds to the sum of domestic network purchases by firm j , where \mathcal{S}_j is the set of domestic network suppliers of firm j .

Match quality ϕ_{ij} can be written as $\phi_{ij} = \phi_i \tilde{\phi}_{ij}$, where ϕ_i captures the average quality of firm i and $\tilde{\phi}_{ij}$ is an idiosyncratic match term. Similarly, prices can be written as $p_{ij} = \tau_i \tilde{\tau}_{ij} c_i$, where c_i is marginal cost, τ_i reflects the average markup and trade cost of i across its customers j , and $\tilde{\tau}_{ij}$ is the match-specific trade cost/markup term. This assumption implies that one can separate the systematic variation across firms from the variation across matches in sales from i to j .

$\tilde{\tau}_{ij}$ can contain any type of price variation, such as heterogeneous trade costs or markups across a seller's customers. It will be useful to collapse parameters that are related to either the buyer, the seller, or the buyer-seller pair. Equation (4) can then be rewritten as:

$$\ddot{m}_{ij} = \ddot{\psi}_i \ddot{\theta}_j \ddot{\omega}_{ij}, \quad (5)$$

where

$$\begin{aligned} \ddot{\psi}_i &\equiv \left(\frac{\phi_i}{\tau_i c_i} \right)^{\sigma-1} \\ \ddot{\theta}_j &\equiv P_j^{\sigma-1} \gamma_j M_j \\ \ddot{\omega}_{ij} &\equiv \left(\frac{\tilde{\phi}_{ij}}{\tilde{\tau}_{ij}} \right)^{\sigma-1}. \end{aligned} \quad (6)$$

There is a well-defined mapping from the model in this section to the empirical model from Section 3, and so we normalize the logs of $\ddot{\psi}_i$, $\ddot{\theta}_j$, and $\ddot{\omega}_{ij}$ by their respective means (across all seller-buyer pairs), $\ln \psi_i \equiv \ln \ddot{\psi}_i - \overline{\ln \ddot{\psi}}$, $\ln \theta_j \equiv \ln \ddot{\theta}_j - \overline{\ln \ddot{\theta}}$ and $\ln \omega_{ij} \equiv \ln \ddot{\omega}_{ij} - \overline{\ln \ddot{\omega}}$. Firm-to-firm sales can then be written as

$$\ln m_{ij} = \ln \psi_i + \ln \theta_j + \ln \omega_{ij}. \quad (7)$$

We refer to ψ_i as a seller effect, θ_j as a buyer effect, and ω_{ij} as a match effect. In the model, the seller effect is decreasing in average quality adjusted prices, $\tau_i c_i / \phi_i$, while the buyer effect θ_j is increasing in total purchases M_j and the input price index P_j . The seller, buyer and match-specific components are by construction mean zero. The model thus delivers a simple log linear expression for firm-to-firm sales, just as in the reduced-form equation (1) in Section 3.

For the model-based results in Section 5, the only required assumptions are the production function, cost minimization and the functional forms of p_{ij} and ϕ_{ij} . In particular, there is no need to assume anything about market structure or firms' pricing behavior. However,

a few additional elements are necessary to solve the general equilibrium and to perform counterfactuals. These are introduced when needed in Section 7.

4.3 A Reflection Problem and Solution

While one can estimate the parameters $\Psi = \{\psi_i, \theta_i, \omega_{ij}\}$ from production network data, their interpretation is not straightforward because they will embody information about the firm itself as well as information about the firm's suppliers. This can be seen directly from equation (6) as both the seller and the buyer effects for firm i , ψ_i and θ_i , depend on its suppliers' prices via the input price index P_i .

This reflection property means that firm-level fundamentals, such as productivity or quality, cannot be isolated directly from the seller and buyer effects. However, a manipulation of the equations in (6) results in

$$Z_i \equiv k_i \left(\frac{\phi_i z_i}{\tau_i} \right)^{\sigma-1} = \psi_i \left(\frac{\theta_i}{\gamma_i M_i} \right)^{\gamma_i(1-\alpha_i)}, \quad (8)$$

where Z_i captures a cluster of parameters only related to the firm itself (productivity/quality, markups/trade costs).¹⁶ Henceforth, we refer to Z_i as *production capability* and to θ_i as *sourcing capability*. Production capability Z_i is a combination of efficiency and quality conditional on the firm's connections and can be isolated simply by multiplying the seller effect with a transformation of the buyer effect, $[\theta_i / (\gamma_i M_i)]^{\gamma_i(1-\alpha_i)}$, as summarized in Proposition 1:

Proposition 1. *The seller and buyer effect of firm i are both functions of the prices charged by firm i 's suppliers. The transformation $Z_i \equiv \psi_i [\theta_i / (\gamma_i M_i)]^{\gamma_i(1-\alpha_i)}$ isolates the production capability of firm i , which is independent of the characteristics of firm i 's suppliers.*

By exploiting the production network data, this methodology overcomes the reflection problem. What is the economic intuition for this result? Recall that $\theta_i / (\gamma_i M_i)$ is proportional to the input price index $P_i^{\sigma-1}$, e.g. if θ_i is small and total purchases M_i are large, input purchases must be spread out over many suppliers, leading to a low input price index P_i . A firm with low input prices will sell more to each customer, leading to a high seller effect ψ_i . The seller effect for this particular firm therefore overstates its inherent production capability. By multiplying the seller effect with $\theta_i / (\gamma_i M_i)$, we penalize firms with low input prices and therefore correct for the reflection bias. Note that the reflection problem is weaker the higher the labor share, α_i . For example, if the labor share is close to one, then input costs will be negligible and the seller effect ψ_i will be highly correlated with production capability. In

¹⁶ $k_i \equiv \bar{\psi}^{-1} \bar{\theta}^{-\gamma_i(1-\alpha_i)} \left(w_i^{\alpha_i} \tilde{w}_i^{(1-\gamma_i)(1-\alpha_i)} \right)^{1-\sigma}$, where $\bar{\psi}$ and $\bar{\theta}$ are the geometric means of ψ_i and θ_i , respectively.

the limit as $\alpha_i \rightarrow 1$, i.e. when input-output linkages are absent, the bias stemming from the reflection problem becomes zero.

Proposition 1 will be useful in Section 5 where we back out Z_i and perform additional decompositions, as well as in Section 7 where we calibrate the model.

4.4 The Input Price Index

We next discuss the relationship between firms' input price index and buyer effect θ_j . Using the expressions for c_i , Z_i and ω_{ij} , the inverse input price index, $P_i^{1-\sigma}$, can be written as

$$\begin{aligned} P_i^{1-\sigma} &= \sum_{k \in \mathcal{S}_i} \left(\frac{p_{ki}}{\phi_{ki}} \right)^{1-\sigma} \\ &= \sum_{k \in \mathcal{S}_i} P_k^{(1-\sigma)\gamma_k(1-\alpha_k)} \tilde{g}_k Z_k \omega_{ki} \quad \forall i, \end{aligned} \quad (9)$$

where \tilde{g}_k is a parameter.¹⁷ The input price index of firm i depends on the input price indices of i 's suppliers, and so on. The input price index is therefore a fixed point of the function in equation (9). Appendix C proves that $P_i^{1-\sigma}$ retrieved from the buyer effect in equation (6), $P_i^{1-\sigma} \propto \gamma_i M_i / \theta_i$, obeys the equilibrium constraints in equation (9) as summarized in Proposition 2:

Proposition 2. *The inverse input price index is a fixed point of the function*

$$P_i^{1-\sigma} = \sum_{k \in \mathcal{S}_i} P_k^{(1-\sigma)\gamma_k(1-\alpha_k)} \tilde{g}_k Z_k \omega_{ki}.$$

$\gamma_i M_i / \theta_i$ is proportional to $P_i^{1-\sigma}$ from the fixed point of the function above.

Proof. See Appendix C. □

An immediate implication of this result is that the buyer fixed effect θ_i estimated in Section 3 can be used to calculate the equilibrium $P_i^{1-\sigma}$ (up to a constant) simply by using the formula $\gamma_i M_i / \theta_i$. Figure 10 in Section 7 confirms that Proposition 2 holds perfectly in our quantitative application.

This concludes the first part of the model. The next section presents the first part of the quantitative application using the insights from Propositions 1 and 2. Section 7 returns to the theory to characterize the general equilibrium and perform counterfactuals.

¹⁷ $\tilde{g}_k \equiv \bar{\psi} \bar{\theta}^{\gamma_k(1-\alpha_k)} \bar{\omega}$.

5 Model-Based Results

This section will use the insights from the theoretical model to explore in more detail the origins of firm heterogeneity. Armed with Propositions 1 and 2, we document the correlation between the inverse input price index, $\ln P_i^{1-\sigma}$, production capability, $\ln Z_i$, and other firm outcomes. We then decompose the upstream and downstream margins of firm size to provide additional evidence on its sources of variation.

The ingredients for the analysis are the estimates from the two-way fixed effects model in Section 3, $\Psi = \{\ln \psi_i, \ln \theta_j, \ln \omega_{ij}\}$, along with balance sheet data on total sales, S_i , total purchases, M_i , the share of inputs sourced from the domestic network, γ_i , and the labor share, α_i .¹⁸ These enable the calculation of production capability from Proposition 1, $Z_i = \psi_i [\theta_i / (\gamma_i M_i)]^{\gamma_i(1-\alpha_i)}$, and the inverse input price index from Proposition 2, $P_i^{1-\sigma} \propto \gamma_i M_i / \theta_i$.

5.1 Correlations

We start by considering the correlations between various firm characteristics in our data. Column 1 in Table 6 shows that firm sales are strongly positively correlated with both the upstream, $\ln \psi_i$, and the downstream, $\ln \xi_i$, components. The correlation coefficient with the downstream component is more than twice as large, mirroring the earlier decomposition. The downstream and upstream components are themselves negatively correlated, implying that firms with many customers or with particularly good customers (leading to a high ξ_i) tend to have smaller average market shares among those customers (leading to a low $\ln \psi_i$).¹⁹ Our interpretation of this finding is that firms are unlikely to succeed on both the intensive and the extensive margins: some firms become large by accumulating a customer base while other firms become large by being important suppliers to a smaller number of firms, and few firms manage to do both.

Firm sales are strongly positively correlated with the inverse input price index, $P_i^{1-\sigma}$, such that larger firms tend to benefit from cheaper input prices. In the model, low input prices can arise because a firm has many suppliers or matches with particularly good suppliers. The inverse price index, $P_i^{1-\sigma}$, is positively correlated with ψ_i , showing that firms with low input prices tend to have higher average market shares among their customers. Recall from Proposition 2 that $P_i^{1-\sigma}$ is calculated from $\gamma_i M_i / \theta_i$; seller effects ψ_i are therefore negatively correlated with $\theta_i / (\gamma_i M_i)$. This result confirms a theoretical prediction from Proposition

¹⁸ The baseline case uses 4-digit industry averages of α_i and γ_i . Section 6 shows results with firm-level α_i and γ_i . Note that inputs purchased from outside the network can be either from foreign suppliers or from domestic suppliers outside the network estimation sample (because they supply less than 250 euro worth of inputs or have too few network connections).

¹⁹ Recall from Section 3 that the downstream component is large when firms transact with more, bigger and/or better-matched customers.

Table 6: Correlation Matrix.

Firm Size Component	$\ln S_i$	$\ln \psi_i$	$\ln \xi_i$	$\ln P_i^{1-\sigma}$	$\ln n_i^c$	$\ln n_i^s$	$\ln Z_i$
Total Sales, $\ln S_i$	1						
Upstream, $\ln \psi_i$.23	1					
Downstream, $\ln \xi_i$.66	-.16	1				
Inverse input price index, $\ln P_i^{1-\sigma}$.91	.14	.61	1			
# Customers, $\ln n_i^c$.49	-.33	.85	.50	1		
# Suppliers, $\ln n_i^s$.76	-.02	.63	.76	.57	1	
Production Capability, $\ln Z_i$	-.31	.79	-.51	-.40	-.59	-.51	1

Note: All correlations are significant at 5%. All variables are demeaned at the NACE 4-digit level.

1, namely that the seller and buyer effects are related because low input prices (low buyer effect) feed into high market shares (high seller effect).

Recall that production capability Z_i is a combination of efficiency and quality conditional on the firm's connections. Production capability is strongly positively correlated with the upstream component and strongly negatively correlated with the downstream component and with the inverse input price index. Thus firms with higher production capability have bigger market shares among their customers (high ψ_i), but their input costs are higher (low $P_i^{1-\sigma}$) and they have fewer or smaller downstream customers (low n_i^c and low ξ_i). As a result, high production capability firms have lower sales unconditionally, but nevertheless higher sales conditional on their input prices: When we regress $\ln S_i$ on $\ln Z_i$ controlling for $\ln P_i^{1-\sigma}$, the coefficient estimates on production capability (0.061***) and the inverse input price index (1.01***) are both positive and significant at 0.1%.

These results, coupled with the initial decomposition findings in Section 3, are difficult to reconcile with standard heterogenous firm models. They suggest that multiple upstream and downstream dimensions of firm activity underpin sales dispersion when firms interact in production networks. One interpretation of our findings is that firm attributes that matter for finding customers and suppliers (e.g., managerial talent and marketing capacity) are orthogonal, or negatively related, to firm attributes that determine sales conditional on a match (e.g., productivity or quality).

5.2 Downstream and Upstream Decompositions

The initial decomposition results in Section 3 provide evidence for the importance of upstream, downstream, and final demand margins in the variance of firm sales. Downstream

factors contribute over 80 percent of total variance and upstream components 18 percent. This section uses the theoretical framework developed above to further decompose the upstream and downstream margins, and thereby shed more light on the sources of firm size heterogeneity.

Starting with the downstream side, which accounts for the majority of firm sales dispersion, the parameter $\xi_i \equiv \sum_{j \in \mathcal{C}_i} \theta_j \omega_{ij}$ can be expressed as

$$\ln \xi_i = \ln n_i^c + \ln \bar{\theta}_i + \ln \Omega_i^c, \quad (10)$$

where n_i^c is the number of customers and $\bar{\theta}_i \equiv \left(\prod_{j \in \mathcal{C}_i} \theta_j \right)^{1/n_i^c}$ is the average customer capability.²⁰ The covariance term Ω_i^c is defined as

$$\Omega_i^c \equiv \frac{1}{n_i^c} \sum_{j \in \mathcal{C}_i} \omega_{ij} \frac{\theta_j}{\bar{\theta}_i}.$$

Each of these components has an intuitive economic interpretation. First, firms face high network demand if they are linked to many customers (high n_i^c). Second, they face high network demand if their average customer has high sourcing capability (high $\bar{\theta}_i$). Third, they face high network demand if the covariance term Ω_i^c is large, i.e. if large customers (high θ_j) also happen to be good matches (high ω_{ij}). These components are directly available in the data (n_i^c) or can be calculated from $\Psi = \{\ln \psi_i, \ln \theta_j, \ln \omega_{ij}\}$ ($\bar{\theta}_i$ and Ω_i^c).

Turning to the upstream decomposition, a firm may be large because it has high production capability (high Z_i), or because it benefits from cheap or high-quality inputs (low P_i). In turn, the input price index can be decomposed into the number of suppliers, average supplier capability, and a covariance term. This can be shown in three steps. First, from equation (8), the production capability of a firm, Z_i , is a function of its estimated buyer and seller effects. Second, from equation (5), log total network purchases are

$$\ln(\gamma_i M_i) = G + \ln \theta_i + \ln \sum_{k \in \mathcal{S}_i} \psi_k \omega_{ki}. \quad (11)$$

Third, solving equation (8) for $\ln \psi_i$ and substituting for $\ln(\gamma_i M_i/\theta_i)$ using equation (11) yields

$$\ln \psi_i = \ln Z_i + \gamma_i (1 - \alpha_i) [G + \ln n_i^s + \ln \bar{\psi}_i + \ln \Omega_i^s], \quad (12)$$

where n_i^s is the number of suppliers, $\bar{\psi}_i \equiv \left(\prod_{k \in \mathcal{S}_i} \psi_k \right)^{1/n_i^s}$ is average supplier capability, and the covariance term Ω_i^s is

$$\Omega_i^s \equiv \frac{1}{n_i^s} \sum_{k \in \mathcal{S}_i} \omega_{ki} \frac{\psi_k}{\bar{\psi}_i}.$$

²⁰ By the properties of ordinary least squares, the average term $(1/n_i^c) \sum_{j \in \mathcal{C}_i} \ln \omega_{ij} = (1/n_i^c) \sum_{k \in \mathcal{S}_i} \ln \omega_{ki} = 0$ and is therefore omitted from the expression.

Table 7: Downstream Decomposition.

	# Customers $\ln n_i^c$	Avg Customer Capability $\ln \bar{\theta}_i$	Customer Covariance $\ln \Omega_i^c$
$\ln \xi_i$.71*** (.00)	.03*** (.00)	.26*** (.00)

Note: The table reports coefficient estimates from OLS regressions of a firm size margin (as indicated in the column heading) on the downstream factor, $\ln \xi_i$. All variables are demeaned by 4-digit industry averages. Standard errors in parentheses. Significance: * < 5%, ** < 1%, *** < 0.1%.

Detailed derivations are found in Appendix C.2. Again, each component of this expression is either observed directly (α_i , γ_i and n_i^s) or can be calculated from $\Psi = \{\ln \psi_i, \ln \theta_j, \ln \omega_{ij}\}$ (Z_i , $\bar{\psi}_i$ and Ω_i^s).

The interpretation of each element is as follows. A firm has a large market share among its customers (high ψ_i) because it is inherently productive or high-quality (high Z_i), because it has many suppliers (high n_i^s), because those suppliers are on average attractive suppliers (high $\bar{\psi}_i$), or because attractive suppliers also happen to be a good match (high Ω_i^s).

As with the overall decomposition, we regress each component in equations (10) and (12) on $\ln \xi_i$ and $\ln \psi_i$, respectively, to evaluate its contribution to the variation in $\ln \xi_i$ or $\ln \psi_i$. The coefficient estimates across components will mechanically sum to one because the left and right hand side of equations (10) and (12) are by construction identical. As above, all components are demeaned by their 4-digit industry average, so that variation across industries is differenced out.²¹


The limited assumptions placed on the economic environment imply that this is an agnostic firm size decomposition that enables an assessment of the contribution of different margins to the overall variation in firm size. Our approach imposes no restrictions on the absolute or relative contribution of these margins.

5.2.1 Downstream Decomposition

Table 7 reports the results for the downstream decomposition. Much of the variation in the downstream component across firms (71 percent) can be attributed to the extensive margin, i.e. the number of (domestic) buyers, $\ln n_i^c$. On the other hand, the average sourcing capability across a firm's customers, $\ln \bar{\theta}_i$, and the customer covariance term, $\ln \Omega_i^c$, contribute

²¹ The term in G , the grand mean, drops out from the upstream decomposition since α and γ only vary at the NACE 4-digit level and we demean all components by NACE 4-digit sector.

Figure 6: Downstream decomposition.



Note: This binned scatterplot groups firms into 50 equal-sized bins by downstream sales component $\ln \xi_i$, computes the mean of $\ln \xi_i$ and its sub-components $\ln n_i^c$, $\ln \bar{\theta}_i$ and $\ln \Omega_i^c$ within each bin, and graphs these data points. The result is a non-parametric visualization of the conditional expectation function.

a much more modest 3% and 26%, respectively. As above, the results are displayed in a binned scatterplot in Figure 6, which reveals that the patterns are stable across 50 bins by firms' downstream sales component.

On the downstream side, the single most important advantage of large firms is that they successfully match with many buyers. The covariance term is also substantial, suggesting that relative to smaller firms, bigger firms more effectively concentrate sales among large buyers with high sourcing capability that are very good bilateral matches. On the other hand, the negligible role for average customer capability shows that large firms do not match with more capable buyers on average.

5.2.2 Upstream Decomposition

Table 8 reports the results for the upstream decomposition. As above, the results are also shown using a binned scatterplot in Figure 7. The seller-specific production capability, $\ln Z_i$, drives practically the entire upstream factor (93%). The remaining variation comes from the

Table 8: Upstream Decomposition.

	Own Prod Capability $\ln Z_i$	# Suppliers $\ln n_i^s$	Avg Suppl. Capability $\ln \bar{\psi}_i$	Suppl. Cov. $\ln \Omega_i^s$
$\ln \psi_i$.93*** (.00)	-.01*** (.00)	.03*** (.00)	.06*** (.00)

Note: The table reports coefficient estimates from OLS regressions of a firm size margin (as indicated in the column heading) on the upstream factor, $\ln \psi_i$. All variables are demeaned by 4-digit industry averages. Standard errors in parentheses. Significance: * < 5%, ** < 1%, *** < 0.1%.

covariance term (6%), average supplier capability (3%) and the number of suppliers (-1%).

These results reveal that, first and foremost, inherent firm characteristics such as productivity or quality ($\ln Z_i$) explain differences in average market shares among customers, ψ_i . In addition, firms that have good suppliers (high $\ln \bar{\psi}_i$), or that source relatively more from good suppliers that they are well-suited to (high $\ln \Omega_i^s$), are also more successful in terms of sales, but the economic magnitude of these effects is small.

6 Discussion


This section first discusses the empirical relevance of the functional form chosen for firm-to-firm sales in equation (1). It then examines the necessary assumptions on the assignment process of buyers and sellers for OLS to identify the underlying parameters of interest, and develops a test for conditional exogenous mobility in the context of a production network. Finally, it presents more evidence on the robustness of the results. Appendix D.5 provides additional extensions and sensitivity checks.

6.1 Functional form

The log-linear relationship in equation (1) predicts the following: (i) expected sales from seller i to customer j are increasing in the average sales of i to other customers k ; (ii) expected purchases by buyer j from seller i are increasing in the average purchases by j from other suppliers k .

Properties (i)-(ii) can be tested non-parametrically as follows. For each seller i and buyer j , calculate the leave-out mean of log sales (\bar{s}_i^{-l}) and purchases (\bar{m}_j^{-l}) across its buyers and

Figure 7: Upstream Decomposition.



Note: This binned scatterplot groups firms into 50 equal-sized bins by upstream sales component $\ln \psi_i$, computes the mean of $\ln \psi_i$ and its sub-components $\ln Z_i$, $\ln n_i^s$, $\ln \bar{\psi}_i$ and $\ln \Omega_i^s$, and graphs these data points. The result is a non-parametric visualization of the conditional expectation function.

suppliers, excluding customer/supplier l , respectively:²²

$$\bar{s}_i^{-l} = \frac{\sum_{j \in \mathcal{C}_i \setminus l} \ln m_{ij}}{n_i^c - 1}$$


$$\bar{m}_j^{-l} = \frac{\sum_{k \in \mathcal{S}_j \setminus l} \ln m_{kj}}{n_j^s - 1}.$$

Then sort firms into decile groups based on \bar{s}_i^{-l} and \bar{m}_j^{-l} , denoting the decile group the firm belongs to as $q_{\bar{s}} = 1, \dots, 10$ and $q_{\bar{m}} = 1, \dots, 10$, respectively. Finally, calculate the mean of $\ln m_{ij}$ for every decile group pair, $\overline{\ln m_{q_{\bar{s}}, q_{\bar{m}}}}$, e.g. the average $\ln m_{ij}$ for the seller-buyer pairs in $(q_{\bar{s}}, q_{\bar{m}}) = (1, 1)$, and so on.

Figure 8 illustrates the results using a heatmap. The decile groups $q_{\bar{s}}$ and $q_{\bar{m}}$ are plotted on the horizontal and vertical axes, respectively. $\ln m_{ij}$ is increasing in the average sales from

²² Using the overall mean generates a mechanical relationship between e.g. seller size and sales between i and j . We calculate \bar{s}_i^{-l} and \bar{m}_j^{-l} for all (i, l) and (j, l) pairs respectively. Firms with only one customer or supplier are by construction omitted from the sample.

Figure 8: Average log sales across seller and buyer decile groups.



Note: The figure shows the average of $\ln m_{ij}$ in all decile group pairs $(q_{\bar{s}}, q_{\bar{m}})$.

i to other customers k (moving from left to right in the diagram), and $\ln m_{ij}$ is increasing in the average purchases of j from other suppliers k (moving from bottom to top in the diagram).

6.2 Assumptions on the Assignment Process

Equation (1) is a two-way fixed effects model similar to the models that are used in the employer-employee literature (Abowd et al., 1999; Card et al., 2013).²³ OLS estimates of $\ln \psi_i$ and $\ln \theta_j$ will identify the effect of seller and buyer characteristics if the following moment conditions are satisfied:

$$\begin{cases} E[s'_i r] = 0 & \forall i \\ E[b'_j r] = 0 & \forall j. \end{cases} \quad (13)$$

Here $S = [s_1, \dots, s_N]$ is the $N^* \times N_s$ seller fixed effects design matrix, $B = [b_1, \dots, b_N]$ is the $N^* \times N_b$ buyer fixed effects design matrix, r is the $N^* \times 1$ vector of residual match effects, and N^* , N_s and N_b are the number of matches, sellers and buyers, respectively. The first condition states that for each seller i , the average $\ln \omega_{ij}$ across buyers j is zero, while the second condition states that for each buyer j , the average $\ln \omega_{ij}$ across sellers i is zero. Intuitively, a high $\ln \omega_{ij}$ that is common across customers j of i will be automatically loaded onto i 's seller effect (and similarly for suppliers i of j). In other words, these moment conditions require that the assignment of suppliers to customers is exogenous with respect to ω_{ij} , so-called conditional exogenous mobility in the labor literature.

It is instructive to review four cases when these moment conditions hold. First, they hold if firms match based on their seller and buyer effects, e.g., highly productive firms match with more and/or different customers/suppliers than less productive ones. Second, the assumption holds if firms match based on idiosyncratic pair-wise shocks that are unrelated to $\ln \omega_{ij}$. One example of this is idiosyncratic fixed costs, such as costs related to search and matching, which affect profits for a potential match but not the value of bilateral sales.²⁴ Third, the moment conditions are consistent with our theory. In the model, a buyer j that receives a favorable shock from i will get a lower CES input price index P_j .²⁵ The estimation allows for

²³ The linear fixed-effects approach imposes no restrictions on the seller and buyer effects, unlike random or mixed effects models. With random effects, one also needs to model the network formation game to assess the plausibility of the required distributional assumptions for unobserved heterogeneity (see Bonhomme (Forthcoming)).

²⁴ Bernard et al. (forthcominga) and Lim (2017) develop models where idiosyncratic fixed costs determine matching. Eaton et al. (2018) develop a quasi-random matching model where matching shocks are unrelated to ω_{ij} .

²⁵ The buyer effect θ_j is proportional to $\gamma_j M_j / P_j^{\sigma-1}$ and the input price index P_j is determined by ω_{ij} along with other variables, see Proposition 2.

arbitrary seller-specific correlations between ω_{ij} and θ_j , i.e. $\text{corr}(\ln \omega_{ij}, \ln \theta_j) \neq 0 \forall i$. Finally, and along the same lines, the moment conditions also allow for a seller i to charge higher markups (low ω_{ij}) to certain buyers (e.g. high θ_j buyers), and for a buyer j to be charged higher markups by certain sellers (e.g., high ψ_i sellers). Formally, $\text{corr}(\ln \omega_{ij}, \ln \theta_j) \neq 0 \forall i$ and $\text{corr}(\ln \omega_{ij}, \ln \psi_i) \neq 0 \forall j$. The key requirement is that the average of $\ln \omega_{ij}$ for every buyer and seller is zero.

Now consider the case of endogenous mobility. To fix ideas, assume that matching is based on the idiosyncratic match component of sales, ω_{ij} , together with the seller effect ψ_i . In that case, only high ψ_i sellers would want to match with low ω_{ij} buyers. OLS would then give a downward bias in the estimated ψ_i , because OLS imposes that the average $\ln \omega_{ij}$ across customers is zero.

To explore the possibility that matching shocks are correlated with sales shocks, we test conditional exogenous mobility as follows. Consider firm i selling to customers 1 and 2. The expected difference in bilateral sales is


$$\Delta \ln m_i \equiv E[\ln m_{i2} - \ln m_{i1} \mid (i, 1), (i, 2)] = \ln \theta_2 - \ln \theta_1 + E[\ln \omega_{i2} - \ln \omega_{i1} \mid (i, 1), (i, 2)].$$

Consider the case $\theta_2 > \theta_1$. Under exogenous mobility, the last expectation term is zero, and $\Delta \ln m_i$ is unrelated to firm i characteristics. Under endogenous mobility, the last expectation term is non-zero, and $\Delta \ln m_i$ is potentially a function of firm i characteristics. Now seller i will only want to match with customer 1 if ω_{i1} is sufficiently large. The expectation $E[\ln \omega_{i2} - \ln \omega_{i1} \mid (i, 1), (i, 2)]$ is then negative. Moreover, for small sellers (low ψ_i), the size of ω_{i1} is important for whether a match occurs or not, while for large sellers (high ψ_i), the size of ω_{i1} is less important (since matching is determined by both ψ_i and ω_{ij}). Under endogenous mobility, the expectation is therefore less negative for high- ψ_i than low- ψ_i firms, so that $\Delta \ln m_i$ is greater for high- ψ_i than low- ψ_i firms. Under exogenous mobility, by contrast, $\Delta \ln m_j$ should be unrelated to ψ_i .²⁶

Going back to the seller and buyer decile groups constructed above, these predictions can be tested by looking at $\overline{\ln m_{q_s, q_m}}$ when moving from a small to a big customer, for different groups of sellers. Figure 9 shows the results. Each line represents the mean of log sales for a given seller decile group (1,...,10). Within a seller group, we calculate $\overline{\ln m_{q_s, q_m}}$ to small customers (buyer decile group 1) and to big customers (buyer decile group 10). Under exogenous mobility, those lines should be parallel, i.e. for buyer bins q_m and q'_m , $\overline{\ln m_{q_s, q_m}} - \overline{\ln m_{q_s, q'_m}}$ does not depend on the seller decile group. The lines are, to a large degree, parallel, in particular for the seller decile groups 2 to 9. Parallel lines are a sufficient

²⁶ Card et al. (2013) test for endogenous mobility for employer-employee matches using a related, but different, test.

Figure 9: Average log sales across seller and buyer decile groups.



Note: The figure shows $\ln \bar{m}_{ij}$ across buyer decile groups $q_{\bar{m}} = 1, \dots, 10$. Each line represents a seller decile group, $q_{\bar{s}} = 1, \dots, 10$.

but not necessary condition for exogenous mobility: If the data generating process is not linear in logs, then one could find non-parallel lines even under exogenous mobility.

One can test for this non-parametrically as follows. Using the buyer and seller bins defined above, exogenous mobility implies that

$$\overline{\ln m}_{q'_{\bar{s}}, q'_{\bar{m}}} - \overline{\ln m}_{q_{\bar{s}}, q_{\bar{m}}} - (\overline{\ln m}_{q_{\bar{s}}, q'_{\bar{m}}} - \overline{\ln m}_{q_{\bar{s}}, q_{\bar{m}}}) = 0, \quad (14)$$

for any bins $q_{\bar{s}}$, $q'_{\bar{s}}$, $q_{\bar{m}}$ and $q'_{\bar{m}}$. We form these averages for $q'_{\bar{s}} = q_{\bar{s}} + 1$ and $q'_{\bar{m}} = q_{\bar{m}} + 1$ and test the null hypothesis that the double difference equals zero. This yields 81 separate hypothesis tests across all buyer-seller pair bins.²⁷ Overall, the results mirror those in Figure 9: the double differences are not significantly different from zero in the middle of the distribution, whereas we find significant deviations in the tails. Significant deviations are typically relatively small: e.g. moving from a 6th to 7th decile buyer yields 12% more sales for seller decile 9 and 14% more sales for seller decile 10. All 81 hypothesis tests are reported in Appendix D.1.

Separately, Appendix D.3 presents a Monte Carlo simulation to examine how well the empirical model recovers the structural parameters.

²⁷ t-values are calculated using Welch's t-test.

Table 9: Business Groups: Overall Decomposition.

	N	Upstream $\ln\psi_i$	Downstream $\ln\xi_i$	Final Demand $\ln\beta_i$
$\ln S_i$	86,485	.16*** (.00)	.83*** (.00)	.01** (.00)

Note: The table reports coefficient estimates from OLS regressions of a firm size margin (as indicated in the column heading) on total firm sales. All variables are first demeaned by their 4-digit NACE industry average. Standard errors in parentheses. Significance: * < 5%, ** < 1%, *** < 0.1%.

6.3 Business Groups

There remains the possibility that intra-firm trading and ownership structure across VAT enterprises is affecting the results. In particular, while the VAT ID is the legal entity of a firm in Belgium, some firms might be owned by other firms, generating intra-firm trade between parents and affiliates, which might not be subject to typical market forces. While the VAT ID is typically used in firm-level analysis of Belgian data (see Amiti et al., 2014, Magerman et al., 2016 and Bernard et al., forthcomingb), we follow the procedure in Tintelnot et al. (2017) and aggregate variables across multiple VAT IDs owned by the same firm as a robustness check. VAT IDs are grouped into a single firm if the same parent company owns at least 50% of their shares. Turnover, inputs, employment and labor costs are summed across subsidiaries to the group level, after subtracting within group transactions from turnover and inputs to avoid double counting. The NACE code of the firm with the largest turnover is assigned to the group. There are 11,737 groups with multiple VAT IDs in the raw data in 2014, but they account for a sizable fraction of output. The resulting decomposition is almost identical across all components, e.g. the overall decomposition in Table 9.

6.4 Firm-specific input and labor shares

The baseline upstream decomposition uses labor (α_i) and network purchases (γ_i) shares calculated at the 4-digit industry level.²⁸ As all firm-level outcomes are demeaned at the 4-digit level, variation in α_i and γ_i is differenced out. Table 10 presents results for the upstream decomposition when using firm-specific α_i and γ_i instead. Note that, from equation (12), this introduces an additional margin, $(1 - \alpha_i)\gamma_i G$. This term represents the variation in the

²⁸ Note that α_i and γ_i are not needed for the overall and downstream decompositions.

Table 10: Upstream Decomposition.

	Own Prod Capability $\ln Z_i$	# Suppliers $\ln n_i^s$	Avg Suppl. Capability $\ln \bar{\psi}_i$	Suppl. Cov. $\ln \Omega_i^s$	Cluster term $(1 - \alpha_i) \gamma_i G$
$\ln \psi_i$	1.18*** (.00)	-.08*** (.00)	.03*** (.00)	.03*** (.00)	-.16*** (.00)

Note: The table reports coefficient estimates from OLS regressions of a firm size margin (as indicated in the column heading) on the upstream factor, $\ln \psi_i$. All variables are demeaned by 4-digit industry averages. Standard errors in parentheses. Significance: * < 5%, ** < 1%, *** < 0.1%.

seller effect ψ_i associated with variation in $(1 - \alpha_i) \gamma_i$. As in the baseline, own production capability is the dominant source of heterogeneity. Firms with high labor shares and low network input shares have lower seller effects.

7 General Equilibrium and Counterfactuals

The estimation and decomposition presented in Sections 4 and 5 provide parameter values for firm-level fundamentals. This section closes the model and solves for the general equilibrium to perform counterfactual analyses.

7.1 General Equilibrium

Final Demand. To close the model, three additional assumptions are required on final demand, markups and factor shares. For final demand, we choose the simplest possible case and assume CES utility with the same elasticity of substitution σ across firms:

$$U = \left(\sum_i (\phi_i \nu_i)^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}.$$

Using the same functional form for final demand and firm demand means that the estimates for production can also be used for final demand. The final consumer is an average input consumer, so that the terms $\tilde{\phi}_{ki}$ and $\tilde{\tau}_{ki}$ do not appear in final demand.²⁹ The value of final demand is then $\mathcal{F}_i = \left(\frac{\phi_i}{\tau_i c_i} \right)^{\sigma-1} \mathcal{P}^{\sigma-1} X$, where X is total income, and \mathcal{P} is the CES consumer price index:

$$\mathcal{P}^{1-\sigma} = \sum_i \left(\frac{p_i}{\phi_i} \right)^{1-\sigma} = \sum_i P_i^{(1-\sigma)\gamma_i(1-\alpha_i)} \tilde{g}_i Z_i. \quad (15)$$

²⁹ Since final demand is modeled as a representative consumer, there is by construction no match-specific component ϕ_{ki} . Since $\phi_{ki} = \phi_k \tilde{\phi}_{ki}$ and $(1/n_i^c) \sum_i \tilde{\phi}_{ki} = 1$, this implies that the perceived quality of firm k is identical for the final consumer and the average downstream firm i .

Labor shares. Factor shares are assumed to be constant across firms, $\alpha_i = \alpha$ and $\gamma_i = \gamma$, i.e. identical labor and import shares. This is needed to ensure a well-defined mapping from our firm-level moments to the equilibrium objects in the model (see Appendix C.5).

Markups and profits. Thus far the model has been completely agnostic about market structure and price determination. Define $\bar{p}_i \equiv \sum p_{ij}y_{ij}/y_i$ as the average price of firm i , i.e. total sales relative to total quantity sold, and define $\mu_i \equiv \bar{p}_i/c_i$ as its average markup. For tractability, we assume that the average markup is identical across firms and constant across equilibria, i.e. $\mu_i = \mu$. This restriction is needed in order to ensure that total income is proportional to labor income (see below). The set of firms is fixed and there is no free entry. The final consumer is the shareholder of all firms, so that aggregate profits Π become part of consumer income. Income X is therefore the sum of labor income and aggregate profits, $X = wL + \Pi$, where w is the wage and L is inelastically supplied labor. Appendix C.3 shows that in equilibrium, $X = (1 + (\mu - 1)/\alpha)wL$. Wages are chosen as the numeraire.

Backward fixed point. The general equilibrium can be found by solving two fixed points sequentially. The input costs of firm i depend on the input costs of the suppliers of i . The equilibrium input price index can be solved by iterating on a backward fixed point problem from equation (9) (Appendix C.5 and Proposition 2):

$$\tilde{P}_i = \sum_{k \in \mathcal{S}_i} \tilde{P}_k^{\gamma(1-\alpha)} Z_k \omega_{ki}, \quad (16)$$

where \tilde{P}_i is proportional to $P_i^{1-\sigma}$, $\tilde{P}_i \propto P_i^{1-\sigma}$.³⁰ Firm i 's input costs depend on the production capability of its suppliers, Z_k , the suppliers' inverse input costs, \tilde{P}_k , and the match terms ω_{ki} .

Forward fixed point. Sales of firm i relate to the sales of the customers of i . Total firm sales are $S_i = \mathcal{F}_i + \sum_{j \in \mathcal{C}_i} m_{ij}$. Using equations (3), (4) and (8), equilibrium sales can be solved by iterating on a forward fixed point:

$$S_i = Z_i \tilde{P}_i^{\gamma(1-\alpha)} \left(\frac{X}{\sum_j \tilde{P}_j^{\gamma(1-\alpha)} Z_j} + \frac{\gamma(1-\alpha)}{\mu} \sum_{j \in \mathcal{C}_i} \frac{S_j}{\tilde{P}_j} \omega_{ij} \right). \quad (17)$$

A detailed derivation is found in Appendix C.4 and C.5. Firm i 's sales depend on final demand, X , the production and sourcing capability of the firm itself, Z_i and \tilde{P}_i , as well as the sales, sourcing capabilities and match effects of its customers, S_j , \tilde{P}_j and ω_{ij} . Appendix C.6 proves the existence and uniqueness of the equilibrium.

Inspecting equations (16) and (17) above, there is a unique mapping from data and the estimates $\Psi = \{\psi_i, \theta_i, \omega_{ij}\}$ to the equilibrium objects \tilde{P}_i and S_i , summarized in the following proposition:

³⁰ $P_i^{1-\sigma} = \tilde{g}^{1/[1-\gamma(1-\alpha)]} \tilde{P}_i$ where $\tilde{g} \equiv \bar{\psi} \bar{\theta}^{\gamma(1-\alpha)} \bar{\omega}$.

Proposition 3. Define a variable \tilde{P}_i which is proportional to $P_i^{1-\sigma}$, where P_i is the input price index. The equilibrium \tilde{P}_i is a function of the parameters $\Psi = \{\psi_i, \theta_i, \omega_{ij}\}$ and data $\{M_i, \gamma, \alpha\}$. Equilibrium sales S_i are a function of the parameters Ψ and data $\{M_i, \gamma, \alpha, \mu, X\}$.

This result implies that our methodology can be applied in a variety of settings to discipline and calibrate network models. Note that the equilibrium distribution of sales can be solved for without imposing any assumption on the elasticity of substitution σ .

Welfare. Indirect utility equals the inverse of the final demand price index \mathcal{P} . Welfare can be evaluated with equation (15), using estimates of production capability Z_i and the solution for \tilde{P}_i from the backward fixed point.

7.2 Counterfactuals

The general equilibrium structure allows a return to the main research question of this paper, exploring the role of the network in explaining firm size dispersion. To do so, we conduct two counterfactuals. In the first, we turn off the traditional source of firm heterogeneity by eliminating dispersion in $\ln Z_i$. This shuts down direct heterogeneity due to variation in a firm’s production capability, but it also eliminates heterogeneity in the $\ln Z_i$ of a firm’s suppliers and customers. Intuitively, the remaining dispersion is then associated only with the network itself: the variation across firms stems solely from their number of connections and the match quality ω_{ij} . In the second counterfactual, we focus on a key parameter in our model, namely the purchased input cost share, $1 - \alpha$. A higher input share means that purchases from the network constitute a bigger share of production costs, such that the network becomes more important.

Applying Proposition 3, a baseline equilibrium is calculated using the estimated parameters $\Psi = \{\psi_i, \theta_i, \omega_{ij}\}$ and data $\{M_i, \gamma, \alpha, \mu, X\}$ and iterating on the two fixed points in equations (16) and (17). Next, the counterfactual equilibria are constructed holding everything else fixed. For the first counterfactual, all $\ln Z_i$ are set equal to their sample average. For the second counterfactual, the purchased input cost share, $1 - \alpha$, is increased by 10 percent. While the fixed point problem is conceptually straightforward, it is numerically difficult because of the large size of the observed network. Since standard matrix operations are not feasible in this setting, we have developed a custom algorithm using dictionary data types in Python to solve the problem.³¹

Parameterization. In addition to the estimated firm-level parameters Ψ , information is needed on γ (share of network purchases in total purchases), α (labor cost share), μ

³¹ Using the custom algorithm, the fixed points converge after 20 minutes on a standard computer. The algorithm is available upon request.

Table 11: Summary of Parameters.

Parameter	Definition	Value	Source
γ	Network purchases/total purchases	.58	Mean of $\frac{\sum_j \sum_{i \in S_j} m_{ij}}{\sum_j M_j}$ (j is NACE4)
α	Labor cost share	.20	Mean of $\frac{\sum_j w_j L_j}{\sum_j (w_j L_j + M_j)}$ (j is NACE4)
μ	Markup	1.23	Mean of $\sum_i \frac{S_i}{w_i L_i + M_i}$ (i is firm)
X	Aggregate final demand	470 bln euros	$\sum_i S_i (\beta_i - 1) / \beta_i$ (i is firm)

(markup) and X (aggregate income). γ and α are constructed by taking the simple average of the industry-specific γ and α used in Section 5. μ is calculated as the unweighted average of sales to total costs. X is inferred from sales going out of the network, i.e. $X = \sum_i \mathcal{F}_i = \sum_i S_i (\beta_i - 1) / \beta_i$ (or, equivalently, $\sum_i S_i - \sum_i \sum_{j \in \mathcal{C}_i} \ddot{m}_{ij}$). Table 11 summarizes the parameters of the model, their definitions, and the values assigned to them.

Model Fit. The model fit is shown in Figures 10 and 11. According to Proposition 2, the equilibrium \tilde{P}_i from equation (16) is proportional to the inverse input price index backed out from the network data ($P_i^{1-\sigma} \propto \gamma_i M_i / \theta_i$). Figure 10 confirms this. There is also a strong correlation between observed sales S_i and equilibrium sales from equation (17), reported in Figure 11. The correlation coefficient is 0.50. The primary reason why the correlation is less than 1 is that firm-level sales to final demand are not a targeted moment in the parameterization. This implies that prediction errors in sales to final demand propagate to prediction errors in network sales through firm-to-firm links in the network. Appendix D.4 shows the fit between observed sales S_i and equilibrium sales if sales to final demand are also a targeted moment. In this case, the correlation coefficient is 0.96.³²


Results. The baseline and counterfactual equilibria are reported in Table 12. The first counterfactual eliminates dispersion in firm production capabilities, $\ln Z_i$, which reduces the standard deviation of log sales by 43%. When heterogeneity in $\ln Z_i$ is eliminated, firms become identical except for their connections to other firms. All remaining heterogeneity in sales is entirely driven by the network, because some firms have more customers than others, which raises their sales, or because some firms have more suppliers than others, which lowers their costs. In the baseline calibration, the standard deviation in log sales is 1.41. The network itself therefore explains over half (56%) of the dispersion of firm size (0.80/1.41).³³

In the second counterfactual, a 10% increase in the purchased input share increases the

³² In addition, the assumption of homogeneous values of γ , α and μ generates prediction errors in sales, although this channel is quantitatively less important.

³³ The 90/10 percentile ratio of sales is roughly 6 ($e^{1.78}$) when dispersion in $\ln Z_i$ is eliminated versus 36 ($e^{3.59}$) in the baseline, implying that heterogeneity in network connections and in $\ln Z_i$ both raise sales dispersion by a factor of six.

Figure 10: Model Fit: $\ln \tilde{P}_i$.



Note: The figure shows the equilibrium $\ln \tilde{P}_i$ calculated from equation (16) on the horizontal axis and the inverse price index calculated from $\gamma_i M_i / \theta_i$ on the vertical axis.

standard deviation of log sales by 2%. Thus increasing the role of the network generates more inequality across firms. Firm size dispersion can widen for two reasons. First, large firms may have initially lower input prices than small firms, i.e. high \tilde{P}_i 's. A lower labor share will then benefit big relative to small firms. This can be seen by differentiating sales S_i with respect to $\gamma(1-\alpha)$, from equation (17), holding all other \tilde{P}_j 's and S_j 's constant. This yields the elasticity

$$\varepsilon_{\gamma(1-\alpha)} S_i = \gamma(1-\alpha) \ln \tilde{P}_i.$$

Second, large firms may experience a bigger decline in input costs relative to small firms, i.e. the change in \tilde{P}_i may be systematically related to firm size.

Table 13 sheds light on both mechanisms operating in the second counterfactual. Column 1 first confirms that initially larger firms grow faster with a convergence regression of the counterfactual change in firm size, $\Delta \ln S_i$, on the baseline equilibrium firm size, $\ln S_i$. The next two columns provide evidence consistent with the first transmission channel: Bigger firms tend to have lower input prices and thus a higher $\ln \tilde{P}_i$ (Column 2),³⁴ and producers that start out with lower input prices expand their sales more (Column 3). The last two columns indicate that the second channel also plays an important role: Initially larger firms experience a bigger decline in their input price index, i.e. a higher $\Delta \ln \tilde{P}_i$ (Column 4), and greater reductions in input prices are associated with higher sales growth (Column 5). In sum, the results reveal that reducing the labor cost share benefits especially large firms, and this effect is quantitatively important. In other words, the production network amplifies firm size heterogeneity. This is driven by the facts that larger firms are better at obtaining lower input costs and that equilibrium input prices fall more for larger firms.


Other outcomes. One might be tempted to use the last counterfactual to analyze other outcomes such as changes in input/output prices or consumer welfare. Unfortunately, these outcomes are not identified using our methodology. From Section 7, input prices are $P_i^{1-\sigma} = \tilde{g}^{1/[1-\gamma(1-\alpha)]} \tilde{P}_i$. While \tilde{P}_i is identified according to Proposition 3, $P_i^{1-\sigma}$ and the consumer price index, \mathcal{P} , are not, because they require information about the unobserved variable $\tilde{g} \equiv \bar{\psi} \bar{\theta}^{\gamma(1-\alpha)} \bar{\omega}$. Intuitively, only relative input prices $(P_i/P_j)^{1-\sigma}$ are identified, but their absolute level is not. However, one would need information about the absolute price level (relative to wages) to infer the impact of a greater labor share α on consumer welfare.

8 Conclusions

This paper quantifies the origins of firm size heterogeneity when firms are interconnected in a production network. We first document new stylized facts about a complete produc-

³⁴ This is also seen from the correlation in Table 6.

Figure 11: Model Fit: $\ln S_i$.



Note: The figure shows the equilibrium $\ln S_i$ calculated from equation (17) on the horizontal axis and $\ln S_i$ from balance-sheet data on the vertical axis.

Table 12: Firm Size Dispersion. Baseline and Counterfactual.

$\ln S_i$ dispersion	St. Dev.	P90-P10
<i>Baseline</i>	1.41	3.59
<i>Counterfactuals</i>		
1: $\text{var}(\ln Z_i) = 0$	0.80	1.78
2: 10% higher $(1 - \alpha)$	1.44	3.68
N	83,741	

Note: The table shows the standard deviation and the difference between the 90th and 10th percentile of log sales in the baseline and counterfactual equilibria.

Table 13: Firm Size Dispersion in Counterfactual 2. Mechanisms.

	Dependent variable:				
	$\Delta \ln S_i$ (1)	$\ln \tilde{P}_i$ (2)	$\Delta \ln S_i$ (3)	$\Delta \ln \tilde{P}_i$ (4)	$\Delta \ln S_i$ (5)
$\ln S_i$.02 (.00)	.33 (.00)		.003 (.000)	
$\ln \tilde{P}_i$.06 (.00)		
$\Delta \ln \tilde{P}_i$.78 (.00)
N			83,741		

Note: The table shows regression results based on the counterfactual exercise of raising the purchased input share by 10%. Standard errors in parentheses. Significance: * < 5%, ** < 1%, *** < 0.1%.

tion network using data on the universe of buyer-supplier relationships among all firms in Belgium. These stylized facts suggest that the network of buyer-supplier links is key to understanding the firm size distribution. Specifically, they signal the important roles played by downstream input demand as distinct from final demand, by both seller- and buyer-specific firm characteristics, and by seller-buyer match characteristics.

Motivated by these facts, we outline a model in which firms buy inputs from upstream suppliers and sell to downstream buyers and final demand. In the model, firms can be large for the standard reason that they have high production capability (i.e. productivity or product quality). However, firms can also be large because they interact with more, better and larger buyers and suppliers and because they are better matched to their buyers and suppliers. This framework delivers an exact decomposition of firm size into upstream and downstream margins with firm, buyer/supplier and match components. We design an estimation methodology that makes it possible to back out these firm size components from data on firm-level balance sheets and firm-to-firm transactions in a production network. We implement the methodology using detailed network data, and quantify the contribution of each component to the overall dispersion in firm size in the economy.

We establish three empirical results for the origins of firm size heterogeneity. First, downstream factors explain 82 percent of firm size heterogeneity, while upstream (supply) factors only 18 percent. Second, nearly all the variation on the demand side is driven by network sales to other firms rather than by final demand. At the same time, nearly all

the variation on the supply side is driven by own production capability rather than input purchases from other firms in the network. Third, most of the variance in network sales is determined by the number of buyers and the allocation of sales towards well-matched buyers of high quality, rather than by average buyer capability. Conversely, most of the variance in network purchases comes from average supplier capability and the allocation of purchases towards well-matched suppliers of high quality, rather than from the number of suppliers. Counterfactual analyses also establish the importance of firm-to-firm connections: Even in the absence of heterogeneity in production capability, dispersion in firm size is still substantial.

These theoretical, methodological and empirical contributions open interesting avenues for future research. We have taken the production network as given in order to assess its role in shaping the firm size distribution. Our results nevertheless shed light on the various challenges and opportunities that firms face in the presence of input-output linkages in the economy. Future work can examine how firm-specific characteristics determine the matching of buyers and suppliers in the production network in light of our findings. Separately, we have dissected the origins of firm size heterogeneity, but not explored its implications for the aggregate economy. Future studies can analyze whether different sources of the dispersion in firm size have different implications for aggregate outcomes such as growth or income inequality. Finally, we have focused on the relationship between the production network and firm size heterogeneity in steady state. Future research can explore how this relationship affects the propagation and aggregate welfare impact of firm-specific and macroeconomic shocks.

References

- Abowd, J.M., F. Kramarz, and D.N. Margolis**, “High Wage Workers and High Wage Firms,” *Econometrica*, 1999, *67* (2), 251–333. 1, 3.1, 6.2
- Allen, Treb, Costas Arkolakis, and Xiangliang Li**, “On the Existence and Uniqueness of Trade Equilibria,” Technical Report, Working paper 2016. C.6
- Amiti, Mary, Oleg Itskhoki, and Jozef Konings**, “Importers, Exporters, and Exchange Rate Disconnect,” *American Economic Review*, July 2014, *104* (7), 1942–1978. 6.3
- Antràs, P., T.C. Fort, and F. Tintelnot**, “The Margins of Global Sourcing: Theory and Evidence from US Firms,” *American Economic Review*, September 2017, *107* (9), 2514–64. 1
- Baqae, David Rezza and Emmanuel Farhi**, “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” Technical Report, Working paper 2018. 1
- and –, “Productivity and Misallocation in General Equilibrium,” Technical Report, Working paper 2018. 1
- Bender, S., D. Card, N. Bloom, J. Van Reenen, and S. Wolter**, “Management Practices, Workforce Selection and Productivity,” *NBER Working Paper*, 2016, (22101). 1
- Bernard, A., A. Moxnes, and Y.U. Saito**, “Production Networks, Geography and Firm Performance,” *Journal of Political Economy*, forthcoming. 1, 7, 24
- Bernard, A.B., J.B. Jensen, S.J. Redding, and P.K. Schott**, “The Empirics of Firm Heterogeneity and International Trade,” *Annual Review of Economics*, 2012, *4* (1), 283–313. 1
- Bernard, Andrew B. and Andreas Moxnes**, “Networks and Trade,” *Annual Review of Economics*, forthcoming. 1
- , –, and **Karen Helene Ulltveit-Moe**, “Two-sided Heterogeneity and Trade,” *Review of Economics and Statistics*, 2018, *100* (3). 1, 7
- , **Emily J. Blanchard, Ilke Van Beveren, and Hylke Y. Vandenbussche**, “Carry-Along Trade,” *Review of Economic Studies*, forthcoming. 6.3

- Bloom, Nick, Kalina Manova, John Van Reenen, Stephen Sun, and Zhihong Yu**, “Managing Trade: Evidence from China and the US,” *NBER Working Paper 24718.*, 2017. 1
- Boehm, Johannes and Ezra Oberfield**, “Misallocation in the Market for Inputs: Enforcement and the Organization of Production,” Technical Report, Working paper 2018. 1
- Bøler, Esther Ann, Andreas Moxnes, and Karen Helene Ulltveit-Moe**, “R&D, International Sourcing, and the Joint Impact on Firm Performance,” *American Economic Review*, 2015, 105 (12). 1
- Bonhomme, Stéphane**, “Econometric Analysis of Bipartite Networks,” Forthcoming. 23
- , **Thibaut Lamadon, and Elena Manresa**, “A Distributional Framework for A Distributional Framework for Matched Employer Employee Data,” 2017. 1
- Card, D., J. Heining, and P. Kline**, “Workplace Heterogeneity and the Rise of West German Wage Inequality*,” *The Quarterly Journal of Economics*, 2013, 128 (3), 967–1015. 1, 6.2, 26
- Correia, Sergio**, “Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator,” Technical Report, Working Paper 2016. 4
- Dhyne, E., G. Magerman, and S. Rubinova**, “The Belgian Production Network 2002–2012,” *National Bank of Belgium Working Paper*, 2015, 288. 4, 7, A.1
- di Giovanni, J., A.A. Levchenko, and I. Mejean**, “Firms, Destinations, and Aggregate Fluctuations,” *Econometrica*, 07 2014, 82 (4), 1303–1340. 1
- Eaton, J., D. Jinkins, J. Tybout, and D.Y. Xu**, “Two-sided Search in International Markets,” 2016. NBER. 1
- , **S. Kortum, and F. Kramarz**, “Dissecting Trade: Firms, Industries, and Export Destinations,” *American Economic Review*, May 2004, 94 (2), 150–154. 3.1
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz**, “Firm-to-Firm Trade: Imports, Exports, and the Labor Market1,” Technical Report, Working paper 2018. 1, 24
- Fieler, Ana Cecília, Marcela Eslava, and Daniel Yi Xu**, “Trade, Quality Upgrading, and Input Linkages: A Theory with Evidence from Colombia,” *American Economic Review*, 2018, 108. 1

- Fitzgerald, D., S. Haller, and Y. Yedid-Levi**, “How Exporters Grow,” Staff Report 524, Federal Reserve Bank of Minneapolis January 2016. 1
- Foster, L., J. Haltiwanger, and C. Syverson**, “The Slow Growth of New Plants: Learning about Demand?,” *Economica*, 2016, 83 (329), 91–129. 1
- Freund, C. and M.D. Pierola**, “Export Superstars,” *The Review of Economics and Statistics*, 2015, 97 (5), 1023–1032. 1
- Gabaix, Xavier**, “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 2011, 79, 733–772. 1
- Gaubert, Cecile and Oleg Itskhoki**, “Granular Comparative Advantage,” 2016. Princeton. 1
- Gibrat, Robert**, *Les inègalités économiques*, Librairie du Recueil Sirey, Paris, 1931. 1
- Goldberg, P., A.K. Khandelwal, N. Pavcnik, and P. Topalova**, “Imported Intermediate Inputs and Domestic Product Growth: Evidence from India,” *The Quarterly Journal of Economics*, 2010, 125 (4), 1727–1767. 1
- Hopenhayn, H.A.**, “Entry, Exit and Firm Dynamics in Long Run Equilibrium,” *Econometrica*, 1992, 60, 1127–1150. 1
- Hottman, C. J., S.J. Redding, and D. E. Weinstein**, “Quantifying the Sources of Firm Heterogeneity,” *Quarterly Journal of Economics*, 2016, 131 (3), 1291–1364. 1, 3.1
- Jovanovic, Boyan**, “Selection and the Evolution of Industry,” *Econometrica*, May 1982, 50 (3), 649–70. 1
- Kramarz, F., J. Martin, and I. Mejean**, “Volatility in the Small and in the Large: The Lack of Diversification in International Trade,” CEPR Discussion Papers 11534, C.E.P.R. Discussion Papers September 2016. 1
- Lim, Kevin**, “Firm-to-firm Trade in Sticky Production Networks,” 2017 Meeting Papers 280, Society for Economic Dynamics 2017. 1, 24
- Luttmer, Erzo G. J.**, “Selection, growth and the size distribution of firms,” *The Quarterly Journal of Economics*, 2007, 122 (3), 1103–1144. 1
- Magerman, Glenn, Karolien De Bruyne, Emmanuel Dhyne, and Jan Van Hove**, “Heterogeneous Firms and the Micro Origins of Aggregate Fluctuations,” Technical Report, KU Leuven 2016. 1, 6.3

- Manova, Kalina, Shang-Jin Wei, and Zhiwei Zhang**, “Firm Exports and Multinational Activity Under Credit Constraints,” *The Review of Economics and Statistics*, July 2015, 97 (3), 574–588. 1
- Manski, Charles F.**, “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies*, 1993, 60 (3). 1
- Melitz, Marc.**, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry,” *Econometrica*, 2003, 71 (6), 1695–1725. 1
- Melitz, Marc and Stephen Redding**, “New Trade Models, New Welfare Implications,” *American Economic Review*, 2015, 105 (3), 1105–1146. 1
- Oberfield, Ezra**, “A Theory of Input-Output Architecture,” *Econometrica*, 2018, 86 (2), 559–589. 1
- Pavcnik, Nina**, “Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants,” *The Review of Economic Studies*, 2002, 69, 245–276. 1
- Sutton, John**, “Market Share Dynamics and the "Persistence of Leadership" Debate,” *American Economic Review*, March 2007, 97 (1), 222–241. 1
- Syverson, C.**, “What Determines Productivity?,” *Journal of Economic Literature*, 2011, 49 (2), 326–365. 1
- Tintelnot, Felix, Ken Kikkawa, Magne Mogstad, and Emmanuel Dhyne**, “Trade and Domestic Production Networks,” Technical Report, Working paper 2017. 1, 6.3

Appendix

A Data Sources and Data Construction

A.1 Data sources

The empirical analysis draws on three main data sources administered by the National Bank of Belgium (NBB): (i) the NBB B2B Transactions Dataset, (ii) annual accounts from the Central Balance Sheet Office at the NBB supplemented by VAT declarations, and (iii) the Crossroads Bank at the NBB. Firms are identified by a unique enterprise number, which is common across all databases and allows for unambiguous merging.

Firm-to-firm relationships The confidential NBB B2B Transactions Dataset contains the value of yearly sales relationships among all VAT-liable Belgian enterprises for the years 2002 to 2014, and is based on the VAT listings collected by the tax authorities. The Belgian value-added tax (VAT) system requires that the vast majority of enterprises located in Belgium across all economic activities charge VAT on top of the delivery of their goods and services. This includes foreign companies with a branch in Belgium and firms whose securities are officially listed in Belgium. Enterprises that only perform financial transactions, medical or socio-cultural activities such as education are exempt. The standard VAT rate in Belgium is 21%, but for some goods a reduced rate of 12% or 6% applies.³⁵

At the end of every calendar year, all VAT-liable enterprises have to file a complete listing of their Belgian VAT-liable customers over that year.³⁶ An observation in this dataset refers to the value of sales in euros by enterprise i to enterprise j within Belgium, excluding the VAT due on these sales. The reported value is the sum of invoices from i to j in a given calendar year. Whenever this aggregated value is 250 euros or greater, the relationship has to be reported.³⁷ Fines for late or incomplete reporting ensure a very high quality of the data. Note that each relationship is directed, as the observation from i to j is different from the observation from j to i ; i.e. firm i might be both a supplier to and a customer of j . The dataset thus covers both the extensive and the intensive margins of the Belgian production network. A detailed description of the collection and cleaning of this dataset is given in Dhyne et al. (2015).

Firm-level characteristics We extract information on enterprises' annual accounts from

³⁵ See ec.europa.eu/taxation_customs for a complete list of rates. These rates did not change over our sample period.

³⁶ Sample VAT listings forms can be found at [here](#) (French) and [here](#) (Dutch).

³⁷ Pecuniary sanctions are given to firms for late or erroneous reporting.

the Central Balance Sheet Office at the NBB for the years 2002 to 2014. Enterprises above a certain size threshold have to file annual accounts at the end of their fiscal year.³⁸ We retain information on the enterprise identifier (VAT ID), turnover (total sales in euros, code 70 in the annual accounts), input purchases (total material and services inputs in euros and net changes in input stocks, codes 60+61), labor cost (total cost of wages, social securities and pensions in euros, code 62), and employment (average number of full-time equivalent (FTE) employees, code 9087). We annualize all flow variables from fiscal years to calendar years by pro-rating the variables on a monthly basis.³⁹

Enterprises below a size threshold can report abbreviated annual accounts. These firms report labor cost and employment, but are not required to report turnover or input purchases. For these small enterprises, we supplement information on turnover and inputs from their VAT declarations. All VAT-liable enterprises have to file periodic VAT declarations with the tax administration.⁴⁰ The VAT declaration contains the total sales value (including domestic sales and exports), the VAT amount charged on those sales (both to other enterprises and to final consumers), the total amount paid for inputs sourced (including domestic and imported inputs), and the VAT paid on those input purchases. This declaration is due monthly or quarterly depending on firm size, and it is the basis for the VAT due to the tax authorities every period. We aggregate the VAT declarations to the annual frequency.

We obtain information on the main economic activity of each enterprise at the NACE 4-digit level from the Crossroads Bank of Belgium for the years 2002 to 2014. We concord NACE codes over time to the NACE Rev. 2 version to deal with changes in the NACE classification over our panel from Rev. 1.1 to Rev. 2. Table 14 lists industry groups at the NACE 2-digit level.

A.2 Data construction and cleaning

We calculate the final demand for enterprise i in year t as i 's turnover minus the value of all of its B2B sales. The B2B Transactions dataset contains all seller-buyer relationships, including both intermediate and investment goods. This implies that final demand contains final domestic consumption, exports, and tiny business transactions below 250 euros that are not observed in the B2B dataset. Similarly, we calculate i 's total input purchases from the network as the value of all of its B2B purchases. We infer i 's input purchases from outside the network as its total input purchases minus its total B2B network purchases.

³⁸ See here for filing requirements and exceptions. See here for the size criteria and filing requirements for either full-format or abridged annual accounts.

³⁹ In our data, 78% of firms have annual accounts that coincide with calendar years, while 98% of firms have fiscal years of 12 months.

⁴⁰ Sample VAT declaration forms can be found at here (French) and here (Dutch).

Table 14: NACE classification of industry groups (NACE Rev 2).

NACE Section	NACE Division	Description	Industry
A	NACE 01-03	Agriculture, forestry and fishing	Primary and Extraction
B	NACE 05-09	Mining and quarrying	Primary and Extraction
C	NACE 10-33	Manufacturing	Manufacturing
D	NACE 35	Electricity, gas, steam and air conditioning supply	Utilities
E	NACE 36-39	Water supply; sewage, waste management and remediation activities	Utilities
F	NACE 41-43	Construction	Construction
G	NACE 45-47	Wholesale and retail trade; repair of motor vehicles and motorcycles	Market Services
H	NACE 49-53	Transportation and storage	Market Services
I	NACE 55-56	Accommodation and food service activities	Market Services
J	NACE 58-63	Information and communication	Market Services
K	NACE 64-66	Financial and insurance activities	Market Services
L	NACE 68	Real estate activities	Market Services
M	NACE 69-75	Professional, scientific and technical activities	Market Services
N	NACE 77-82	Administrative and support service activities	Market Services
O	NACE 84	Public administration and defence; compulsory social security	Non-Market Services
P	NACE 85	Education	Non-Market Services
Q	NACE 86-88	Human health and social work activities	Non-Market Services
R	NACE 90-93	Arts, entertainment and recreation	Non-Market Services
S	NACE 94-96	Other service activities	Non-Market Services
T	NACE 97-98	Activities of households as employers; undifferentiated goods- and services-producing activities of households for own use	–
U	NACE 99	Activities of extraterritorial organizations and bodies	–

This residual, unobserved part of input expenditures, contains imports and unobserved B2B input purchases under 250 euro.

We drop firms that have missing employment information or less than one FTE employee, as these may be shell companies or management companies. We also drop firms with less than 1,000 euro sales in a given year, which only amounts to 27 firms in 2014 after applying the other selection criteria. This is to avoid long left tails in the graphs in Section 2 of the paper, but has no bearing on any of the results.

Labor shares are calculated at the 4-digit NACE level, as the sum of labor cost over labor cost plus total input usage across all firms. $\alpha_S = \frac{\sum_j wL_j}{\sum_j (wL_j + M_j)}$, where wL_j is firm j 's wage bill and M_j its total expenditure on intermediate inputs (both from the annual accounts) for all j in sector S . For the general equilibrium calculations, we use the simple average of sector labor shares to obtain the common labor share α . We obtain the sector-level share of inputs sourced from the observed domestic network as $\gamma_S = \frac{\sum_j \sum_i m_{ij}}{\sum_j M_j}$, where $\frac{\sum_i m_{ij}}{M_j}$ is j 's share of inputs sourced from the domestic network. Again, we use the simple average across industries to obtain γ for the general equilibrium analysis.

Throughout the paper, we report statistics on both the full sample in the raw data and the estimation sample used in the firm size decomposition. For the full sample, we keep all B2B relationships in the NBB B2B dataset, even if there is missing firm-level information, as these contribute to the decomposition exercise. We thus keep all enterprises that show up in the network as either a buyer or a seller. For the estimation sample, in *Step One* we first estimate the two-way fixed effects regression on the full sample. Note that if a buyer or seller has only one business relationship, the fixed effect is not identified. This enterprise, together with its connections, is then dropped from the sample. This is done iteratively, until only enterprises that have at least two sellers or two buyers remain. Finally, for the decomposition exercise to contain the same number of observations across all (sub-)components, in *Step Two* and *Step Three* we keep only enterprises that have information on all the (sub-)components of the decomposition. In the general equilibrium case, we need to further restrict the sample to have a square matrix of buyers and sellers. We then iteratively drop firms and linkages until the set of sellers is the same as the set of buyers, and all have identified fixed effects.


Finally, for the counterfactual exercise, we obtain firm-level markups $\mu_i = \frac{S_i}{M_i}$, or the ratio of i 's sales revenue to its input expenditure, and calculate aggregate final demand wL by summing over final demand for all enterprises that are part of the fixed point algorithm. Note that this obtained value is very close to observed GDP in the National Accounts (420 billion euros in 2014).

B Additional Descriptive Statistics

The distributions of the number of suppliers and customers have similar features within different sectors, but they also display some heterogeneity in line with priors. For example, the number of buyers and suppliers is highest for firms in utilities, which are followed closely by manufacturing firms. These numbers are intermediate for producers in primary materials and extraction, and lowest among service providers.

Figure 12 replicates Figure 12 in the main text using firms' total B2B sales in the domestic network instead of total turnover. Recall that in addition to downstream domestic network sales, total turnover also includes domestic final demand and exports. Figure 12 plots the fitted line and 95% confidence band from a local polynomial regression of domestic network sales on the number of downstream customers or upstream suppliers, on a log-log scale. The pattern is very similar to the baseline in the main text: a strong monotonic relationship with implied elasticities of 0.78 and 1.24, respectively.

Figure 12: Firm size and number of buyers and suppliers (2014).



Note: The number of customers and suppliers is demeaned at the NACE 4-digit level. Graphs are trimmed at the 0.1st and 99.9th percentiles of the number of customers and suppliers respectively.

C The Model

C.1 Proposition 2

Proposition. *The input price index is a fixed point of the function $P_i^{1-\sigma} = \sum_{k \in S_i} P_k^{(1-\sigma)\gamma_k(1-\alpha_k)} \tilde{g}_k Z_k \omega_{ki}$. $\gamma_i M_i / \theta_i$ is proportional to the price index $P_i^{1-\sigma}$ from the fixed point of the function above.*

Table 15: Firm sales (million euros, 2014).

Sector	NACE	N	Mean	St Dev	10th	25th	50th	75th	90th	95th	99th
Primary & Extraction	01-09	3,061	12.0	432.6	0.2	0.4	0.8	1.9	4.8	9.5	52.0
Manufacturing	10-33	18,077	14.4	250.8	0.2	0.5	1.1	3.8	13.8	34.6	201.8
Utilities	35-39	897	39.2	442.9	0.3	0.7	1.9	6.9	25.7	68.6	495.6
Construction	41-43	20,201	2.3	13.4	0.2	0.3	0.6	1.4	3.6	6.9	25.9
Market Services	45-82	65,175	5.5	79.9	0.2	0.3	0.8	2.1	6.3	13.4	63.9
Non-Market Services	84-99	2,328	2.2	26.3	0.1	0.2	0.3	0.8	2.6	5.5	24.9
All		109,739	6.8	145.1	0.2	0.3	0.8	2.1	6.6	14.3	78.4

Note: Summary statistics for the matched CBSO-B2B data. 10th, 25th, etc. refers to values at the 10th, 25th, etc. percentile of the distribution.

Table 16: Number of firm buyers and suppliers (2014).

(a) Number of downstream buyers.

Sector	NACE	N	Mean	St Dev	10th	25th	50th	75th	90th	95th	99th
Primary & Extraction	01-09	50,706	12.1	60.1	1	2	4	8	18	40	154
Manufacturing	10-33	57,976	47.5	284.9	1	2	7	26	98	192	603
Utilities	35-39	2,734	192.7	3,304.8	1	2	7	36	154	336	1,514
Construction	41-43	104,566	14.6	107.9	1	2	4	10	24	45	174
Market Services	45-82	351,773	32.9	394.6	1	1	3	11	48	112	453
Non-Market Services	84-99	22,352	14.1	183.7	1	1	2	6	19	38	154
All		590,271	29.3	394.0	1	1	4	11	42	98	400

(b) Number of upstream suppliers.

Sector	NACE	N	Mean	St Dev	10th	25th	50th	75th	90th	95th	99th
Primary & Extraction	01-09	60,508	20.5	29.6	2	5	13	27	44	57	117
Manufacturing	10-33	72,698	38.0	89.5	2	5	15	38	89	148	348
Utilities	35-39	3,401	62.8	180.7	2	4	14	55	146	235	757
Construction	41-43	130,358	24.5	48.3	2	5	13	29	52	77	178
Market Services	45-82	506,145	18.3	41.5	1	3	8	19	42	64	150
Non-Market Services	84-99	60,223	10.6	39.2	1	2	4	10	20	33	97
All		840,607	20.6	49.5	1	3	9	22	46	71	177

Note: Summary statistics for the B2B data. 10th, 25th, etc. refers to values at the 10th, 25th, etc. percentile of the distribution.

Table 17: Firm-to-firm transaction values (euros, 2014).

Sector	N	Mean	St Dev	10th	25th	50th	75th	90th	95th	99th
Primary & Extraction	613,868	39,898	5,409,863	419	840	2,490	9,150	33,789	81,626	387,573
Manufacturing	2,755,457	44,303	2,007,421	359	613	1,661	6,185	25,436	63,467	411,379
Utilities	526,932	59,953	7,410,682	366	615	1,388	3,744	11,560	28,382	281,181
Construction	1,529,078	24,500	386,201	375	676	1,926	7,000	27,186	64,585	339,523
Market Services	11,562,445	24,373	2,886,213	341	546	1,266	4,060	15,579	37,960	224,363
Non-Market Services	315,529	8,044	319,407	315	472	998	2,739	8,395	18,908	92,879
All	17,304,408	28,893	2,988,881	348	571	1,392	4,669	18,280	44,770	269,153

Note: Summary statistics for the B2B data. 10th, 25th, etc. refers to values at the 10th, 25th, etc. percentile of the distribution. Industry refers to the main industry of activity of the seller.

Proof. Using the expressions for c_i , Z_i and ω_{ij} , the input price index can be written as

$$\begin{aligned}
P_i^{1-\sigma} &= \sum_{k \in \mathcal{S}_i} \left(\frac{p_{ki}}{\phi_{ki}} \right)^{1-\sigma} \\
&= \sum_{k \in \mathcal{S}_i} \left(\frac{\tau_k \tilde{\tau}_{ki} c_k}{\phi_k \tilde{\phi}_{ki}} \right)^{1-\sigma} \\
&= \sum_{k \in \mathcal{S}_i} P_k^{(1-\sigma)\gamma_k(1-\alpha_k)} \tilde{g}_k Z_k \omega_{ki},
\end{aligned} \tag{18}$$

where $\tilde{g}_k \equiv \bar{\psi} \bar{\theta}^{\gamma_k(1-\alpha_k)} \bar{\omega}$. The equilibrium input price index is a fixed point of the function in equation (18) above. Using the expression for Z_i , one can alternatively write the input price index as

$$P_i^{1-\sigma} = \bar{\psi} \bar{\omega} \sum_{k \in \mathcal{S}_i} \psi_k \omega_{ki}.$$

Sales from i to j are, according to equation (1), $\ddot{m}_{ki} = e^G \psi_k \theta_i \omega_{ki}$. Summing across suppliers, we get

$$\begin{aligned}
\sum_{k \in \mathcal{S}_i} \ddot{m}_{ki} &= \sum_{k \in \mathcal{S}_i} \psi_k \theta_i \omega_{ki} e^G = \gamma_i M_i \iff \\
\frac{\gamma_i M_i}{\theta_i} &= e^G \sum_{k \in \mathcal{S}_i} \psi_k \omega_{ki}.
\end{aligned} \tag{19}$$

Therefore, the assumption of log-linear sales from equation (5) alone guarantees that $\gamma_i M_i / \theta_i$ equals a weighted average of the suppliers' seller effects. The input price index calculated from $\gamma_i M_i / \theta_i$ is proportional to the equilibrium fixed point in equation (18). \square

C.2 The supply side decomposition

From equation (8), we get

$$\ln \psi_i = \ln Z_i + \gamma_i (1 - \alpha_i) (\ln (\gamma_i M_i) - \ln \theta_i).$$

Substituting for $\ln (\gamma_i M_i)$ from equation (11) yields

$$\ln \psi_i = \ln Z_i + \gamma_i (1 - \alpha_i) \left(G + \ln \sum_{k \in \mathcal{S}_i} \psi_k \omega_{ki} \right).$$

The term $\sum_{k \in \mathcal{S}_i} \psi_k \omega_{ki}$ can be further decomposed into

$$\ln \sum_{k \in \mathcal{S}_i} \psi_k \omega_{ki} = \ln n_i^s + \ln \bar{\psi}_i + \ln \left(\frac{1}{n_i^s} \sum_{k \in \mathcal{S}_i} \omega_{ki} \frac{\psi_k}{\bar{\psi}_i} \right),$$

where $\bar{\psi}_i = \left(\prod_{k \in \mathcal{S}_i} \psi_k \right)^{1/n_i^s}$. Combining the last two equations yields equation (12) in the main text.

C.3 Aggregate Profits

In this section, we show that aggregate income is proportional to labor income wL .

Define $\bar{p}_i \equiv \sum p_{ij}y_{ij}/y_i$ as the average price of firm i , i.e. total sales relative to total quantity sold. Define $\mu_i \equiv \bar{p}_i/c_i$ as the average markup of firm i . Profits can then be written

$$\begin{aligned}\pi_i &= \bar{p}_i y_i - c_i y_i \\ &= (\mu_i - 1) c_i y_i \\ &= \frac{\mu_i - 1}{\alpha_i} w L_i\end{aligned}$$

where we used the fact that the labor cost share satisfies $\alpha_i = wL_i/c_i y_i$ in the last line. Aggregate income is then

$$X = wL + \sum_i \pi_i = wL + \sum_i \frac{\mu_i - 1}{\alpha_i} w L_i$$

Invoking the assumption from the main text that $\alpha_i = \alpha$ and $\mu_i = \mu$, we get

$$wL + \sum_i \pi_i = \left(1 + \frac{\mu - 1}{\alpha}\right) wL.$$

Aggregate income X is proportional to labor income wL .

C.4 Forward fixed point

Total firm sales are $S_i = \mathcal{F}_i + \sum_{j \in \mathcal{C}_i} m_{ij}$. We first derive expressions for final demand and then demand from other firms.

Final demand. Using equation (3) and defining $\tilde{P}_i \equiv P_i^{1-\sigma}$ and $\tilde{\mathcal{P}} \equiv \mathcal{P}^{1-\sigma}$, the final demand price index is

$$\mathcal{P}^{1-\sigma} = \sum_i \left(\frac{\tau_i c_i}{\phi_i} \right)^{1-\sigma} = \sum_i P_i^{(1-\sigma)\gamma_i(1-\alpha_i)} \tilde{g}_i Z_i.$$

Using equation (3), final demand is

$$\begin{aligned}\mathcal{F}_i &= \left(\frac{\phi_i}{\tau_i c_i} \right)^{\sigma-1} \mathcal{P}^{\sigma-1} X \\ &= \tilde{g}_i Z_i P_i^{(1-\sigma)\gamma_i(1-\alpha_i)} \frac{X}{\mathcal{P}^{1-\sigma}}.\end{aligned}$$

Firm Demand. Using (3), (8) and (4), firm demand is

$$\begin{aligned} \sum_{j \in \mathcal{C}_i} m_{ij} &= \sum_{j \in \mathcal{C}_i} \left(\frac{\phi_{ij}}{p_{ij}} \right)^{\sigma-1} P_j^{\sigma-1} \gamma_j M_j \\ &= \tilde{g}_i Z_i P_i^{(1-\sigma)\gamma_i(1-\alpha_i)} \sum_{j \in \mathcal{C}_i} \frac{\gamma_j (1-\alpha_j)}{\mu_j} \frac{S_j}{P_j^{1-\sigma}} \omega_{ij}, \end{aligned}$$

where we used the fact that $M_i = S_i (1 - \alpha_i) / \mu_i$. Combining the two sources of demand, we get total sales:

$$S_i = \tilde{g}_i Z_i P_i^{(1-\sigma)\gamma_i(1-\alpha_i)} \left(\frac{X}{\mathcal{P}^{1-\sigma}} + \sum_{j \in \mathcal{C}_i} \frac{\gamma_j (1-\alpha_j)}{\mu_j} \frac{S_j}{P_j^{1-\sigma}} \omega_{ij} \right). \quad (20)$$

C.5 Change of variables

A problem with the expression for S_i in equation (20) is that it depends on $\tilde{g}_i = \bar{\psi} \bar{\theta}^{\gamma_i(1-\alpha_i)} \bar{\omega}$, which is a function of the unidentified parameters $\bar{\psi}$ and $\bar{\theta}$. This section solves this indeterminacy.

To make progress, we invoke the assumption in the main text that $\alpha_i = \alpha$, $\gamma_i = \gamma$ and $\mu_i = \mu$. Define the fixed point

$$\tilde{P}_i = \sum_{k \in \mathcal{S}_i} \tilde{P}_k^{\gamma(1-\alpha)} Z_k \omega_{ki}. \quad (21)$$

It is then straightforward to show that $P_i^{1-\sigma} = \tilde{g}^{1/[1-\gamma(1-\alpha)]} \tilde{P}_i$, where $\tilde{g} = \bar{\psi} \bar{\theta}^{\gamma(1-\alpha)} \bar{\omega}$. After some algebra, one can then rewrite equation (20) to

$$S_i = Z_i \tilde{P}_i^{\gamma(1-\alpha)} \left(\frac{X}{\sum_j \tilde{P}_j^{\gamma(1-\alpha)} \tilde{z}_j} + \frac{\gamma(1-\alpha)}{\mu} \sum_{j \in \mathcal{C}_i} \frac{S_j}{\tilde{P}_j} \omega_{ij} \right), \quad (22)$$

which is independent of $\bar{\psi}$ and $\bar{\theta}$.

C.6 Existence and Uniqueness

We prove existence and uniqueness by showing that the general equilibrium belongs to the class of models analyzed by Allen et al. (2016).

Allen et al. (2016) consider the following system of equations:

$$\prod_{h=1}^K (x_i^h)^{\gamma_{kh}} = c_i^k + \sum_{j=1}^N K_{ij}^k \prod_{h=1}^K (x_j^h)^{\beta_{kh}},$$

where $i, j \in \{1, \dots, N\}$ are firms/sectors, x_i^h is the type h equilibrium variable, c_i^k is a constant and K_{ij}^k are exogenous linkages between i and j . With $K = 1$ this reduces to

$$x_i^\gamma = c_i + \sum_{j=1}^N K_{ij} x_j^\beta. \quad (23)$$

The backward fixed point in equation (21) can be written in the form of equation (23) with $\gamma = 1$, $c_i = 0$, $\beta = \gamma(1 - \alpha)$ and $K_{ij} = Z_k \omega_{ki}$. Using their notation, \mathbf{A} is simply $\gamma(1 - \alpha)$ and therefore the maximum eigenvalue of \mathbf{A}^p is also $\gamma(1 - \alpha) < 1$. According to their Theorem 2(i), there exists a unique and strictly positive solution to the backward fixed point.

The forward fixed point in equation (22) can be written in the form of equation (23) with $\gamma = 1$, $c_i = Z_i \tilde{P}_i^{\gamma(1-\alpha)} X / \left(\sum_j \tilde{P}_j^{\gamma(1-\alpha)} \tilde{z}_j \right)$, $\beta = 1$ and $K_{ij} = \gamma(1 - \alpha) \omega_{ij} / \left(\mu \tilde{P}_j \right)$. Using their notation, \mathbf{A} is 1 and therefore the maximum eigenvalue of \mathbf{A}^p is also 1. According to their Theorem 2(ii.a) there exists at most one strictly positive solution to the forward fixed point.

D Additional Results and Robustness

D.1 Exogenous mobility test

We report 81 separate hypothesis tests across all buyer-seller pair bins in Table 18. Each column refers to the change from buyer decile t to $t + 1$, and each row refers to the change from seller decile t to $t + 1$. For example, the cell (3-2,2-1) reports the difference $\overline{\ln m_{3,2}} - \overline{\ln m_{3,1}} - \left(\overline{\ln m_{2,2}} - \overline{\ln m_{2,1}} \right)$.

Table 18: Exogenous mobility test.

		Buyer decile								
		2-1	3-2	4-3	5-4	6-5	7-6	8-7	9-8	10-9
Seller decile	2-1	0.01*	0.02*	0.02*	0.03*	0.02*	0.01	0.01*	0.01	-0.04*
	3-2	0.04*	0.01*	0.01	0.00	0.00	0.01*	0.00	0.01	0.02*
	4-3	0.01	0.01*	0.02*	0.00	0.01	0.00	0.01	0.00	0.03*
	5-4	0.00	-0.01	0.00	0.01*	0.00	0.00	0.00	0.00	0.08*
	6-5	0.02*	0.02*	0.01	0.00	0.01	0.01	0.02*	0.00	0.02*
	7-6	0.04*	0.02*	0.02*	0.01*	0.01*	0.01	0.01	0.03*	0.00
	8-7	0.00	0.01*	0.00	0.00	0.03*	-0.01	0.00	-0.01	0.01
	9-8	-0.01	0.02*	0.02*	0.01	-0.01	0.02*	0.01	0.05*	0.08*
	10-9	0.03*	0.04*	0.01	0.02*	0.02*	0.02*	0.02*	0.04*	0.19*

Note: The table shows the double difference from equation (14) in the main text. Significance: * < 5%, ** < 1%, *** < 0.1%. t values are based on Welch's t-test.

D.2 Variance decompositions

This section derives statistical properties of the baseline variance decomposition. Consider the following identity:

$$s \equiv \sum_k a_k.$$

The variance of s is

$$\text{var}(s) = \sum_k \sigma_{kk} + \sum_k \sum_{i \neq k} \sigma_{ki}, \quad (24)$$

where $\sigma_{ki} = \text{cov}(a_k, a_i)$. In the baseline decomposition, we regress each element a_k on s . By the properties of OLS, the estimate is

$$\beta_k = \frac{\text{cov}(a_k, s)}{\text{var}(s)} = \frac{1}{\text{var}(s)} \left(\sigma_{kk} + \sum_{i \neq k} \sigma_{ki} \right). \quad (25)$$

Note that the sum of all β_k 's equals one,

$$\sum_k \beta_k = \frac{1}{\text{var}(s)} \left(\sum_k \sigma_{kk} + \sum_k \sum_{i \neq k} \sigma_{ki} \right) = 1.$$

Also note that in the case with only two components, the covariance term in equation (25) is split equally among components:

$$\begin{aligned} \beta_1 &= (\sigma_{11} + \sigma_{12}) / \text{var}(s) \\ \beta_2 &= (\sigma_{22} + \sigma_{12}) / \text{var}(s). \end{aligned}$$

D.3 A Monte Carlo Simulation

This section performs a Monte Carlo simulation to assess the statistical properties of the estimating model. We simulate a random network, draw random shocks for Z_i and ω_{ij} , and solve the model. This entails solving the full general equilibrium, which is described in Section 7. Given simulated data on (i) network linkages and (ii) $\ln m_{ij}$, we estimate buyer and seller fixed effects from equation (1) as described above. We then test whether we can recover the model parameters \tilde{P}_i and Z_i (from $P_i^{1-\sigma} = \gamma_i M_i / (\theta_i \bar{\theta})$ and $Z_i = \psi_i (\theta_i / (\gamma_i M_i))^{\gamma_i (1-\alpha_i)}$, see Sections 4.3 and 4.4) as well as the error term ω_{ij} .

We summarize key parameters of the model in Table 19. Figure 13 shows the estimated and true values for $\ln Z_i$, $\ln P_i^{1-\sigma}$ and $\ln \omega_{ij}$ on the horizontal and vertical axes, respectively. The correlations between true and estimated values are between 0.95 and 0.99, showing that the empirical model recovers the true model parameters with high precision.

Table 19: Summary of Parameters.

Parameter	Definition	Value
<u>Network activity:</u>		
$\ln Z_i$	Production capability	Standard normal
$\ln \omega_{ij}$	Pair-wise sales shock	Standard normal
α	Labor cost share	0.4
μ	Markup	1.33
γ	Network input share	1
N	Number of firms	1,000
<u>Network formation:</u>		
$\ln \epsilon_{ij}$	Matching shock	Standard normal
Match if $\ln Z_i + \ln \tilde{z}_j + \ln \epsilon_{ij} > 2.5$		

Figure 13: Monte Carlo Simulation.





Figure 14: Model Fit. Exogenous Final Demand.



Note: The figure shows the equilibrium $\ln S_i$ calculated from equation (17) with exogenous final demand on the horizontal axis and $\ln S_i$ from balance-sheet data on the vertical axis.

D.4 Model Fit

This section explores the source of prediction error between firm-level total sales S_i in the model and in the data, as illustrated in Figure 11 in the main text.

A potential source of error is sales to final demand, which in the model is determined by the first term in parentheses in equation (17). Prediction errors in final demand will have a direct and an indirect effect. The direct effect is that total sales are the sum of network sales and final demand sales, so that errors in final demand for firm i directly translate into errors in its total sales S_i . The indirect effect is that errors in the total sales of firm i (the direct effect) affect the network sales of all suppliers of firm i . The prediction errors will thus propagate through the production network.

We investigate this source of error by replacing the endogenous final demand component in equation (17), $Z_i \tilde{P}_i^{\gamma(1-\alpha)} X / \left(\sum_j \tilde{P}_j^{\gamma(1-\alpha)} \tilde{z}_j \right)$, with final demand as observed in the data, \mathcal{F}_i , and then re-calculating the forward fixed point. The new scatterplot between model-generated and observed total sales is shown in Figure 14. Prediction errors are significantly reduced compared to Figure 11, and the correlation coefficient between S_i in the model and in the data is now 0.96. The remaining sources of error are heterogeneity in labor shares, network input shares and markups (α_i , γ_i and μ_i).

Table 20: Firm Size Decomposition by Year (2002-2014).

Year	N	Upstream $\ln\psi_i$	Downstream $\ln\xi_i$	Final Demand $\ln\beta_i$
2002	81,399	.17***	.78***	.05***
2003	83,804	.17***	.78***	.05***
2004	85,161	.18***	.78***	.04***
2005	86,602	.17***	.78***	.04***
2006	88,693	.17***	.79***	.04***
2007	91,155	.18***	.79***	.03***
2008	92,451	.18***	.79***	.03***
2009	92,511	.17***	.79***	.04***
2010	92,885	.17***	.79***	.04***
2011	94,257	.18***	.80***	.03***
2012	95,534	.18***	.79***	.03***
2013	94,297	.18***	.80***	.02***
2014	94,330	.18***	.81***	.01**

Note: The table reports coefficient estimates from OLS regressions of a firm size margin (as indicated in the column heading) on total firm sales. All variables in logs. Significance: * < 5%, ** < 1%, *** <0.1%.

D.5 Additional Robustness

D.5.1 Results by Year

This section performs the exact firm size decomposition for every year from 2002 to 2014. We find that the sources of firm size heterogeneity have remained remarkably stable over time.

We perform the analysis from Section 3 separately for each year in the data, and report the results in Table 20. The importance of the upstream component has stood firmly at 17-18%. The downstream component has gradually risen from 78% to 81%, closely following a decline in final demand from 5% to 1%. We observe similarly stable patterns when we consider the lower-tier decomposition of the downstream and upstream components (available upon request). These findings suggest that there may be inherent drivers of the firm size distribution whose relative importance persists despite the rise in production fragmentation across firm and country boundaries over the last 15 years.

D.5.2 Results by sector

We have also explored the stability of our results across different sectors (available upon request) by performing the decomposition exercise separately for six broad sector groups.

Across the board, the estimated coefficients are relatively close to the baseline findings in the main paper. One exception is construction (NACE 41 to 43), where the final demand term β_i enters with a coefficient of -0.10. However, this is expected, as large construction firms typically sell relatively less to final demand compared to small construction firms.

D.5.3 Firm Growth

The baseline decomposition relates the variance of sales across firms to the variance of different sales margins. A related question is what explains the variance of firm sales growth. We proceed as follows. First, we estimate equation (7) on two cross-sections, the baseline year 2014 ($t = 1$) and year 2002 ($t = 0$). We then calculate the change in every demeaned variable in the decomposition. For example, the overall decomposition from (??) becomes

$$\Delta_T \ln S_i = \Delta_T \ln \psi_i + \Delta_T \ln \xi_i + \Delta_T \ln \beta_i,$$

where Δ_T denotes the change from $t = 0$ to $t = 1$, e.g. $\Delta_T \ln S_i = \Delta_S \ln S_{i1} - \Delta_S \ln S_{i0}$. We then demean all variables at the NACE 4-digit level. Finally, we regress each component, e.g. $\Delta_T \ln \psi_i$, on $\Delta_T \ln S_i$. This decomposition allows us to assess the importance of the network in explaining firm growth. Note that long differencing is only feasible for firms that are observed with non-missing sales as well as buyer and seller effects in both years, such that we cannot perform the decomposition on firms that enter or exit during the sample period. However, the decomposition accounts for the adding and dropping of customers and suppliers, i.e. the terms $\Delta_T \ln \psi_i$ and $\Delta_T \ln \xi_i$ may change because of extensive margin adjustments.

The results are summarized in Table 21. At a broad level, the contribution of each component is quite close to what we found in the baseline cross-sectional analysis, yet there are some notable differences in magnitudes. For example, the downstream component dominates in the overall decomposition, with the same contribution as in the cross-section of 81%. However, the upstream and final demand components are now equally important at 9-10%, while final demand played a trivial role of 1% before. On the upstream side, all growth over time comes from improvements in own production capability (100%). On the downstream side, the number of customers is the primary driver of variation in both firm growth and firm sales in the cross-section, but it generates 61% of the former compared to 71% of the latter. This is counterbalanced by a greater role for the customer covariance term in the growth decomposition (38%) relative to the levels decomposition (26%). Alternative long differences from 2002-2008 and 2008-2014 give very similar results (available upon request).

Table 21: Firm Growth Decomposition (2002-2014).

Firm Size Component		Sales	Downstream	Upstream
		S_i	ξ_i	ψ_i
S_i	Upstream	ψ_i	.09***	
	Downstream	ξ_i	.81***	
	Final Demand	β_i	.10***	
ξ_i	# Customers	n_i^c	.61***	
	Avg Customer Capability	$\bar{\theta}_i$.01***	
	Customer Covariance	Ω_i^c	.38***	
ψ_i	Own Production Capability	Z_i		.99***
	# Suppliers	n_i^s		-.02***
	Avg Supplier Capability	$\bar{\psi}_i$.01***
	Supplier Covariance	Ω_i^s		.02***
N		41,177	41,177	41,177

Note: The table reports coefficient estimates from OLS regressions of a firm size margin (as indicated in the row heading) on total firm sales (column 3), $\ln \xi_i$ (column 4) or $\ln \psi_i$ (column 5). All variables in logs. Significance: * < 5%, ** < 1%, *** < 0.1%.

We draw three conclusions about the sources of firm growth from these patterns. First, the vast heterogeneity in growth rates across firms stems from some firms successfully expanding their sales to downstream buyers in the production network. This entails adding more customers over time, but also effectively redirecting sales towards buyers that are both big and well matched. The greater importance of the latter margin for firm growth relative to the cross-section would be consistent with the presence of matching costs and ex-ante imperfect information about buyers. Bigger firms may be able to match with more buyers at a given point, as well as to more effectively reallocate sales among them as match qualities are revealed over time, with both forces contributing to faster sales growth.

Second, while the variation in final demand is not important for firms' relative performance in the cross-section, tapping final consumers helps surviving firms expand revenues to a greater degree. Note that in our data, this corresponds to a rise in sales to final domestic consumers as well as to foreign markets.

Finally, faster growing firms enhance their efficiency and/or product quality mainly by increasing their own production capability. While big firms do benefit from more effective input sourcing in the cross-section, firms' sales growth does not come from further optimizing their sourcing behavior. To the limited extent that fast-growing firms do adjust along this dimension, they reduce the number of input suppliers and shift purchases towards well-

matched suppliers in equal measure, without changing average supplier capability. The contrasting results for the upstream and downstream network components of firm growth indicate that firms may face different matching frictions and information asymmetries in their interactions with buyers and suppliers, which translate into different firm dynamics on the production and sales side.