Journal of Economic Dynamics & Control 187 (2026) 105313

Contents lists available at ScienceDirect

Journal of Economic Dynamics and Control

journal homepage: www.elsevier.com/locate/jedc

Firm-level production networks: What do we (really) know? ⋆

Andrea Bacilieri [a,b], András Borsos [a,c,d], Pablo Astudillo-Estévez [a,c,e], Mads Hoefer [a,f], François Lafond [a,f]

[a] *Institute for New Economic Thinking, University of Oxford, Oxford, UK*
[b] *Oxford Martin School, University of Oxford, Oxford, UK*
[c] *Complexity Science Hub Vienna, Vienna, Austria*
[d] *Department of Financial Systems Analysis, Central Bank of Hungary, Budapest, Hungary*
[e] *Colegio de Economía, Universidad San Francisco de Quito USFQ, Quito, Ecuador*
[f] *Smith School of Enterprise and Environment, University of Oxford, Oxford, UK*

ARTICLE INFO

JEL classification:
C81
D57
D85

Keywords:
Production networks
Input-output analysis
Firm-level data

ABSTRACT

Are standard production network properties similar across all available datasets, and if not, why? We provide benchmark results from two administrative datasets (Ecuador and Hungary), which are exceptional in that they have no reporting threshold. We compare these networks with a leading commercial dataset (FactSet) and published results on national firm-level production networks. Administrative datasets with no reporting thresholds have remarkably similar quantitative properties, while a number of important properties are biased in datasets with missing data.

1. Introduction

Almost a century after Leontief's *The Economy as a Circular Flow* (1928), national input-output (I-O) tables are available for the large majority of advanced economies, have been harmonised and extended to international tables and serve as the basis for environmentally-extended national accounts. These datasets continue to power the development of major macro-econometric, general equilibrium and agent-based models used by policymakers across the world.

⋆ We would like to thank the Servicio de Rentas Internas (SRI) and its Centro de Estudios Fiscales for providing the raw data of Ecuador for research purposes. András Borsos would like to thank INET for its hospitality. We would like to thank audiences at CCS 2022, ENBES, CSH Vienna, Harvard Growth Lab, ETH Zürich, Conference on Network Science and Economics, ESCoE 2023, Bank of Italy-CEPR-EIEF-OECD Conference 2024, and Andrew Bernard, Alexandra Brintrup, Vasco Carvalho, Banu Demir, Christian Diem, Doyne Farmer, Kerstin Hötte, Leonardo Ialongo, Glenn Magerman, Luca Mungo, Anton Pichler and Kjartan van Driell for comments, Charlotte Wargniez for excellent research assistance, and Louise Guillouët for sharing statistics. This work was supported by Baillie Gifford, the Oxford Martin Programme on the Post-Carbon Transition and the Institute for New Economic Thinking at the Oxford Martin School. This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via contract no. 2019-1902010003. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorised to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. FactSet had the opportunity to review the FactSet data description. Our code is available at https://github.com/andreabacilieri/what_do_we_really_know_replication. We have used modern AI tools for literature search, language editing and coding tasks.

E-mail addresses: andrea.bacilieri@oxfordmartin.ox.ac.uk (A. Bacilieri), borsosa@mnb.hu (A. Borsos), pastudillo@usfq.edu.ec (P. Astudillo-Estévez), mads.hoefer@ouce.ox.ac.uk (M. Hoefer), francois.lafond@inet.ox.ac.uk (F. Lafond).

https://doi.org/10.1016/j.jedc.2026.105313
Received 3 July 2025; Received in revised form 20 January 2026; Accepted 4 March 2026
Available online 25 March 2026
0165-1889/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

While these achievements are remarkable, these datasets remain highly aggregated, covering as few as 56 sectors in the World Input-Output Database (WIOD) and a maximum of 405 industries for the most disaggregated tables published by the US Bureau of Economic Analysis (BEA). In comparison, there are about 200 million firms in the world and 6 million in the US.[1] A host of recent papers have started to explore firm-level data on production networks, demonstrating its importance for understanding, e.g., stock co-movement, the propagation of shocks, international trade or aggregate fluctuations (see Table 1 for a comprehensive list of references).

In principle, firm-level data are likely to be much more useful than aggregate data, since aggregation can create substantial biases (Morimoto, 1970). Within the same industry, firms differ in the extent to which they buy and sell from other industries, so shock propagation models at the industry level will generally lead to biased results, especially because shocks do not affect all firms within an industry in the same way (Diem et al., 2023).

Data on firm-level production networks are thus very useful. But what data are available and how good are they? Are there generic properties of firm-level production networks that hold across all datasets? We address these questions through a detailed analysis of three important datasets: administrative VAT data from Ecuador and Hungary, and a leading commercial dataset covering large firms in the global supply network (FactSet). We complement these results with an extensive synthesis of the literature.

One of our goals is to provide a benchmark for researchers who only have a single dataset and worry that their results may be driven by peculiarities of that dataset. For *internal validity*, researchers must establish that the properties they study are not simply artefacts of that dataset's underlying data collection method. We find that several network properties are sensitive to the data collection method and to the reporting threshold. Raising the threshold changes those statistics systematically, while other properties appear robust within the range we study. For *external validity*, results from one country are more likely to generalise when key structural features of production networks are similar across countries. We document that many, though not all, of these features are comparable across our countries. Importantly, our evidence suggests that differences in reporting thresholds produce larger variation than cross-country heterogeneity, so comparative studies should harmonise reporting thresholds and possibly attempt to correct for truncation effects when using harmonised data is not possible.

For all years in Ecuador and for the last year in Hungary, there are no reporting thresholds, so, in theory, we observe the population of domestic firm-to-firm transactions; we call these our *complete* datasets. As Fig. 1 makes clear, the change in the reporting threshold in Hungary had a dramatic effect on the number of transactions in the network. Throughout the paper, we exploit the comparison between the early and recent years of the Hungarian dataset to understand the impacts of the reporting threshold. We develop this analysis further in Appendix D, where we quantitatively examine its influence by synthetically truncating the Ecuadorian network.

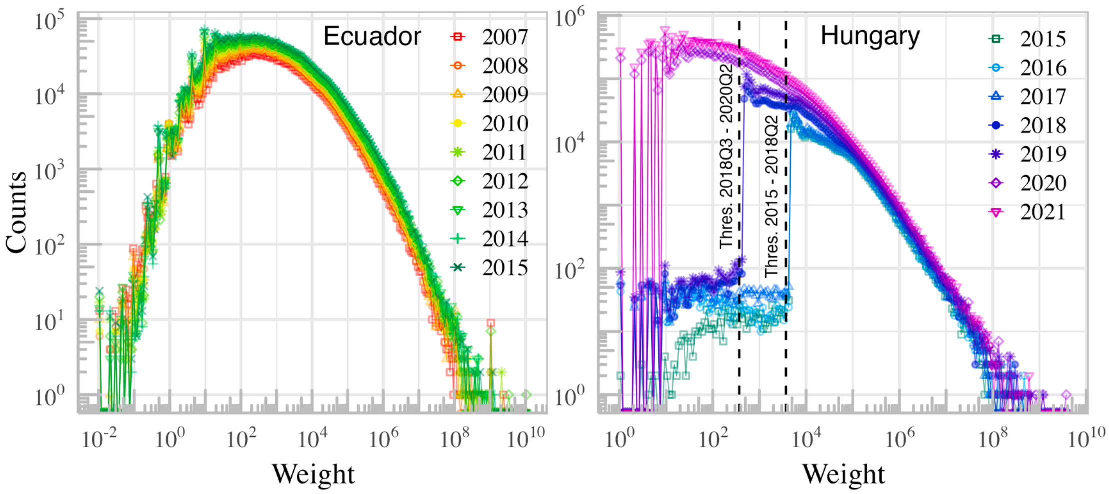

Fig. 1. Distribution of the weights for Ecuador (left) and for Hungary (right) over time. For Hungary, the two vertical lines mark the changes in the reporting threshold (i.e. "Thres."). The first threshold was in force from 2015 until the second quarter of 2018. The second threshold was effective up to the second quarter of 2020, after which it was removed. The values indicated by the vertical lines apply to the majority of the firms in the economy, but deviations arise depending on the tax rate to which firms are subject; see Appendix A.2.3 for more information. We used 200 log-spaced bins for both datasets. The values are in USD for Ecuador and in 1000 HUF (≈ \$2.8) for Hungary.

We find a remarkable similarity in key network statistics across complete datasets, providing us with a credible benchmark of what we "really know" – properties of production networks that are very likely to be similar despite country heterogeneity. We

[1] OpenCorporates reports 201,708,765 as of 17 November 2021 (https://opencorporates.com/), while Statistica reports around 210 million for 2018, 2019 and 2020 (https://www.statista.com/statistics/1260686/global-companies/). For the number of US firms, we used the Statistics of US Businesses dataset provided by the US Census Bureau (https://www.census.gov/programs-surveys/susb.html).

compare these results with what we observe in incomplete datasets and interpret the difference as the bias due to less exhaustive reporting.

To give an example, the mean number of partners (mean degree) in complete datasets is around 40 (weakly increasing with the number of nodes), but in incomplete datasets, it can be much lower, so the reporting threshold strongly biases the observed mean degree downward. In fact, we show that this bias is predictable: given the reporting threshold, we can predict the observed mean degree fairly well. By contrast, the tail exponent of the distribution of transaction values (weights) does not appear to be affected by changes in reporting requirements in Hungary, as seen in Fig. 1. These two examples are somewhat trivial, but others are far from obvious.

For instance, the tail exponent of the in-degree distribution is biased, whereas the tail exponent of the out-degree distribution is not. Several important papers have recognised biases in production network data, especially in Compustat, and have attempted to mitigate them (Atalay et al., 2011; Herskovic et al., 2020; Taschereau-Dumouchel, 2025). We are the first to use credibly complete domestic B2B networks to establish a ground truth for key properties, benchmark these against a comprehensive review of the literature, study synthetically truncated networks and industry heterogeneity, and conclude with clear predictions about the impact, or lack thereof, of the reporting threshold.

There is a lively literature in econometrics showing that observing networks only partially leads to biased results in a number of settings (see e.g. Boucher and Houndetoungan, 2025, for recent results and a brief review). We anticipate that our results will be directly useful to that literature, as the sign and magnitude of the bias are typically hard to anticipate (and correct) unless one knows precisely how the data have been truncated and knows the values of specific statistics in the full, population-level sample (Marray, 2025).

The key message from our empirical investigation is that while firms can have a very large number of customers, they do not have a similarly large number of suppliers. We discuss this asymmetry throughout the paper. We provide additional details on the joint distributions of in- or out-strength (B2B expenses and sales) and in- or out-degree (number of suppliers and customers), documenting new facts. For example, while the number of customers increases with sales, this relationship is strongly heteroscedastic: larger firms tend to have more customers, but they can also have an exceptionally high or surprisingly low number of customers.

Overall, we choose what properties to report based on several factors. Since our main goal is to establish whether complete datasets exhibit similar properties, and whether and how incomplete datasets are biased, we adopt an inclusive approach and report properties examined by many previous studies, thereby allowing for comparisons across a large and diverse sample. We draw inspiration from the network science literature, which has identified a number of structural properties that are typically of broad interest, for instance, as signatures of the network formation process or as predictors of the dynamics of quantities flowing through the network. Rather than listing all possibly interesting properties from a network point of view, however, we pay special attention to quantities that have been emphasised in recent economic literature, and discuss this literature throughout the paper.

Taken together, our results provide the first comprehensive picture of the most fundamental statistics on production networks at the firm level and provide a crucial benchmark for all economists and statisticians putting together these data.[2] In contrast to other studies which interpret facts *qualitatively* (e.g. “firms with higher sales have more customers”), we go beyond “stylised” facts and systematically provide clear *quantitative* estimates. Our results are thus helpful to all researchers who do not have access to administrative data but need key moments to calibrate their macroeconomic models or create synthetic datasets.

The paper is organised as follows. In Section 2, we provide a taxonomy of datasets that can be used to study firm-level production networks. In Section 3, we present results on the binary structure of firm-level production networks and in Section 4, we report findings on the weighted network. Section 5 discusses our results and Section 6 concludes.

2. Datasets

In an ideal case, data on firm-level production networks would be time-stamped, transaction-level data with a distinction between price and quantities. While prices and quantities are available in rare cases (such as Belgium, see e.g. Duprez and Magerman, 2018), most datasets contain either a money flow (how much firm j spends on inputs provided by firm i) or simply a binary indicator that i is a supplier of j. Table 1 shows examples of these different types of datasets, along with reference papers.

National datasets. The main source of national datasets is *data collected for VAT purposes*, which mainly cover supplier-customer relations between firms registered in the country. These datasets usually record money flows. Often, there is a threshold below which transactions need not be reported (see Table 2). As we will see, this threshold does affect some properties of the network. Taking Belgium as an example, the reporting threshold is €250 (even though smaller transactions might be declared), and firms and operations exempt from VAT declarations include microenterprises, medical and socio-cultural activities, and any financial transactions (Dhyne et al., 2015). Dhyne et al. (2015) report that for 2012, the revenues of firms in the network represent 95% of national gross output.

[2] Aside from independent national initiatives, there exists at least two projects aiming at computing key moments of national datasets in an harmonised way: workstream 2 of the “Challenges for Monetary Policy Transmission in a Changing World” at the European Central Bank (https://www.ecb.europa.eu/pub/research-networks/html/champ.en.html) and the “Leveraging InterFirm Transaction Data” at the OECD (https://www.oecd.org/en/about/projects/leveraging-inter-firm-transactions.html).

Table 1
Taxonomy of production network datasets, with examples.

Type and examples	Weighted	Source(s)
National		
Data collected for VAT purposes		
Ecuador	Yes	This paper; Mungo et al. (2023)
Hungary	Yes	This paper; Diem et al. (2022, 2023)
Belgium	Yes	Dhyne et al. (2015), Magerman et al. (2016), Dhyne et al. (2016, 2021, 2022), Duprez and Magerman (2018), Bernard et al. (2022)
Bulgaria	Yes	Alexopoulos et al. (2025)
Chile	Yes	Grigoli et al. (2023), Huneeus (2020)
China	Yes	Egger et al. (2025)
Costa Rica	Yes	Alfaro-Urena et al. (2018, 2022)
Dominican Republic	Yes	Cardoza et al. (2025)
Estonia	Yes	Criscuolo et al. (2024), Masso and Vahter (2021)
Italy	Yes	Pessina (2020)
Kenya	Yes	Chacha et al. (2024)
Rwanda	Yes	Spray and Wolf (2018)
Spain	Yes	Peydró et al. (2020)
Turkey	Yes	Demir et al. (2022, 2023)
Uganda	Yes	Spray and Wolf (2018), Spray (2017)
Uruguay	Yes	Gadenne et al. (2019a)
West Bengal	Yes	Kumar et al. (2021), Gadenne et al. (2019b)
5 Indian states	Yes	Panigrahi (2023)
Data from payment systems		
Brazil	Yes	Silva et al. (2020)
Estonia	Yes	de la Torre et al. (2016)
Japan	Yes	Fujiwara et al. (2021)
Netherlands	Yes	Ialongo et al. (2022)
Data collected for providing business services		
Japan (Tokyo Shoko Research)	No	E.g. Saito et al. (2007), Konno (2009), Ohnishi et al. (2009, 2010), Fujiwara and Aoyama (2010), Carvalho et al. (2021), Inoue (2016), Furusawa et al. (2017), Lu et al. (2017), Zhigang et al. (2018), Kichikawa et al. (2019), Bernard et al. (2019)
Japan (Teikoku Databank)	No	Mizuno et al. (2015)
Northern Italy (Infocert)	Yes	Guichon et al. (2024)
South Korea	No	Lee et al. (2016), Sang (2016)
US (Billtrust)	Yes	Costello (2020)
Global		
Data collected from financial reporting requirements		
FactSet	Yes/No	This paper and König et al. (2022); Taschereau-Dumouchel (2025) for the US
Bloomberg	Yes/No	E.g. Wu and Birge (2014), Wu (2016a)
Capital IQ (S&P)	Yes/No	E.g. Chakraborty and Ikeda (2020)
Compustat (S&P)	Yes/No	E.g. Cohen and Frazzini (2008), Atalay et al. (2011), Herskovic et al. (2020), Atalay et al. (2014), Carvalho and Voigtländer (2014), Barrot and Sauvagnat (2016), Wu and Birge (2014)
Shipment data		
FactSet	Yes	This paper
S&P	Yes	Wu (2016a)
Import-export data		
All countries	Yes/No	Examples where matched to national networks: Dhyne et al. (2021), Duprez and Magerman (2018), Spray (2021), Demir et al. (2022), Huneeus (2020)

Notes: We exclude studies that have no information on intra-national production networks, product-level datasets, and firm-level datasets focused on a single industry.

A second major source of data is *payment systems*, which track monetary flows, even when these do not correspond to a payment for goods or services (e.g. loans and subsidies). Brazil gathers firm-to-firm transactions through two real-time gross settlement systems (Sistema de Transferência de Reservas and Sistema de Transferência de Fundos) provided and operated by the central bank of Brazil. These two datasets collect customer-supplier relations via wire transfers made by firms through their banks, with no threshold. In 2014, the value of all recorded transactions was 20 times the value of national GDP (Silva et al., 2020). Similarly, Fujiwara et al. (2021) constructs a network from payments between Japanese firms that hold an account at the Shiga regional bank, while Ialongo et al. (2022) use payments made between clients of the same bank (they study two banks separately: ABM AMRO Bank NV and ING Bank NV).

A third source of data comes from *credit rating companies*. A prominent example is the production network data for Japan, collected by two private companies for company credit rating and reports (Tokyo Shoko Research, Ltd, and Teikoku Databank, Ltd). When rating and advising firms, these companies collect information on their suppliers and customers but do not keep track of the money flows. Depending on the credit rating company, firms are asked to list up to 24 or 60 of their suppliers and customers, so the in-

Table 2
Reporting thresholds by country.

Dataset	Year	Transaction size threshold		Firm size threshold		Source
		Raw	/GDPpp	Raw	/GDPpp	
Belgium	2002–2014	250 EUR	0.7			Bernard et al. (2022)
Costa Rica	2008–2015	2,500,000 CRC	40.56			Alfaro-Urena et al. (2022)
Domin. Rep.	2012–2017	0 DOP	0	0 DOP	0	Cardoza et al. (2025)
Chile	2003–2011	0 CLP	0	250m CLP	2816	Huneeus (2020)
Ecuador	2008–2015	0 USD	0	0 USD	0	This paper
Estonia	2015–2021	1,000–12,000 EUR	4.23–50.79	40,000 EUR	169	Criscuolo et al. (2024)
Hungary	2015–2018	3,703,703 HUF	83.45	0	0	This paper
Hungary	2018–2020	370,370 HUF	7.46	0	0	This paper
Hungary	2021	1000 HUF	0.02	0	0	This paper
Kenya	2015–2021	0 KES	0	5m KES	2488	Chacha et al. (2024)
Spain	2008–2009	3005 EUR	13.05			Peydró et al. (2020)
Turkey	2010–2014	5000 TRY	19.01			Demir et al. (2022)

Notes: The table shows the official reporting thresholds as gathered from the literature, omitting details of each country's idiosyncratic rules. The thresholds on the value of transactions and on firm sizes are shown in absolute and relative terms, expressed as percentage of GDP per person ("/GDPpp"), using World Bank data for the most recent year in the "Year" column. The table does not consider thresholds imposed by researchers (e.g. removing small firms). In Estonia, the threshold is €1000 per month, although there are a lot of voluntary registrations below the threshold. We determined an effective annual firm-level transaction threshold as follows. The maximum a pair can trade while staying below the threshold is €999 in each of the 12 months, which we round to €12,000. The minimum a pair needs to trade be subject to mandatory reporting is €1,000, occurring in only one of the 12 months. Thus, the effective threshold is somewhere between €1000 and €12,000. A similar but more complex issue arises for Hungary (see Appendix A.2.3), where we report the threshold as if it were always on an annual basis. Fig. 1 shows that this works well, in the sense that the data exhibits a strong shift at that value.

and out-degrees have an artificial cut-off. Another example is Credit2B (acquired by Billtrust), which also provides credit reporting services and collects supplier-customer transactions.

Our national datasets: Ecuador and Hungary. In this paper, we analyse two national datasets: Ecuador (2007–2015) and Hungary (2015–2021). Ecuador requires VAT filings from both firms and natural persons. We use data on firms only (see Appendix A.2.2 for more information). For these firms, we know some characteristics, such as their industry and location, but we do not have access to their revenues, total expenditure or any other financial variable. The monetary values of the transactions are in US dollars, the national currency of Ecuador.

Hungary's network is collected by the National Tax and Customs Administration of Hungary. Until the first half of 2020, Hungary required firms to report transactions that exceeded a threshold (which changed over time); for more information on the threshold and Hungary's dataset, see Appendix A.2.3. The data cover the period 2014–2021, but we dismiss the first year because the data quality is poor. For most Hungarian firms, we have access to financial statements, but do not make use of them. The transactions are expressed in HUF 1,000 ($\approx$ \$2.8).

The 2015 Ecuadorian and 2021 Hungarian networks serve as our reference examples of complete datasets throughout the paper. We highlight these in the tables using a dotted line.

Global datasets. The datasets with global coverage feature mainly listed firms, which account for a large portion of gross output. A key source for these data comes from US Financial Accounting Standards, which require publicly traded firms to report customers that account for 10% or more of their annual revenues (formally called *major customers*). Due to the data collection process, coverage is biased towards companies listed on US stock exchanges. Although companies might report customers that account for less than 10%, this threshold still skews the type of relations observed.

Standard and Poor (S&P) provides two datasets of this kind: Compustat and Capital IQ. Capital IQ provides information on over 60,000 publicly traded companies worldwide, while Compustat tracks the order of a thousand firms and links per year (Cohen and Frazzini, 2008). Compustat is solely based on relations derived from the disclosure of major customers. Other data providers such as Bloomberg and FactSet collect additional information on supply chains by looking at annual fillings/reports, investor presentations, company websites and press releases. As a result, these datasets are still biased in terms of the kind of transactions and companies they keep track of but are much more comprehensive. In a cross-sectional comparison of these datasets (excluding Capital IQ), Wu (2016b) finds Bloomberg to be the most comprehensive, although it is not possible to observe the network over time and access to bulk data remains very difficult. Importantly, the vast majority of these network data are not weighted.

Our global dataset: FactSet. In this paper, we use FactSet Supply Chain Relationships, which we merge with supply-chain relations derived from shipment data (Supply Chain Shipping Transactions). Shipment data are collected daily from the US Bill of Lading, which is required for all seaborne trade, and covers private and publicly traded firms. Our FactSet dataset covers the years 2014–2023, but we only present the results for the 2021 network in the main text because the number of links and firms drops in the last two years. See Appendix A.2.1 for more details.

Data cleaning. For FactSet, Hungary and Ecuador, we keep only firms in the largest weakly connected component (LWCC). Two firms are in the same weakly connected component if they are connected by at least one path, ignoring the direction of the edges. We retain the LWCC mostly because some properties are not well defined for networks with disconnected components (e.g. path length and clustering). This procedure leaves the number of firms or links virtually unaffected in our VAT datasets, but removes 8% of firms and 2% of edges in FactSet.

Comparing to national accounts aggregates. Aggregating firm-level data does not necessarily lead to quantities that are conceptually comparable with national accounts (see Appendix A.1). In Appendix A.2.3, we compare the sectoral composition of the Hungarian I-O table (intermediate sales + gross fixed capital formation) with that of the VAT dataset (network sales). In wholesale and retail trade, firm-to-firm sales are substantially higher than in the national accounts, as expected, because gross output in this industry excludes the value of the goods and services purchased for resale. In industries with important non-market activities (public administration and defence, education and health), the size of the industry in the VAT dataset is much lower than in the national accounts, again as expected. For other industries, such as manufacturing, which is under-represented in our VAT dataset, the reasons for the discrepancies are less clear.

3. Results on binary networks

In this section, we discuss binary network metrics, such as average degree, degree distributions, correlations between the in- and out-degree, and degree assortativity. We then describe local patterns: reciprocity, clustering coefficient and average path length.

Throughout the paper, our figures usually display all the years for which we have data. These show that most properties are very stable over time, except for Hungary due to changes in the reporting threshold. Consequently, when we report results in tables, we only report the most recent year for Ecuador, but 3 different years for Hungary, corresponding to years when different reporting thresholds were in place. For FactSet, we report a single year, as we do for Ecuador; however, we do not select the last year (2023). Instead, we report 2021 because the number of firms and relationships observed in 2022 and 2023 declines due to issues in data recording (see Appendix A.2.1).

3.1. Density and growth

How many suppliers and customers do firms have? As we will see in the next section, the number of suppliers or customers varies greatly across firms. However, before discussing the dispersion of these distributions, we provide detailed statistics on the average because it highlights the heterogeneity across datasets.

We define the *in-degree* k_i^{in} as the number of suppliers of firm i and the *out-degree* k_i^{out} as the number customers of firm i. The *average degree* is given by

$$\bar{k} = \frac{1}{N}\sum_{i=1}^{N} k_i^{\text{in}} = \frac{1}{N}\sum_{i=1}^{N} k_i^{\text{out}},$$

where N denotes the number of firms.

Mean degree is highly heterogeneous across datasets. We would not expect the average number of suppliers or customers of firms to differ dramatically across various economies. As a result, heterogeneity in the mean degree helps us to characterise heterogeneity across datasets arising from differences in data collection and data cleaning methods. Fig. 2 shows the average degree for all datasets for which we could find data in the literature, often with several data points per dataset corresponding to different years or studies. The mean degree varies from less than 3 to around 50, over an order of magnitude difference.

To a large extent, the observed variation in the mean degree appears to stem from differences in the data collection method. VAT-based datasets have a fairly high average degree. The years in which Hungary has a low average degree correspond to periods when a high reporting threshold was in place. When we study the effects of the reporting threshold (Appendix D), we find that there is indeed a clear relationship between the mean degree and the reporting threshold (Fig. D.15), which largely explains the differences in mean degree between countries with high reporting thresholds (Spain, Estonia, Costa Rica and early Hungary) and those with low or no thresholds (recent Hungary, Ecuador, Belgium and Dominican Republic; see Table 2).

Datasets collected by private companies (Japan, FactSet and the four smallest networks) tend to have a much lower average degree. The datasets based on transactions in two Dutch banks fall somewhere in between. The Dutch datasets only include transactions between accounts within the same bank, so although these banks are large, the data are still significantly truncated.

Mean degree tends to increase with network size in the time-series dimension. One might expect the mean degree to increase with the total number of nodes in the network, both in real data, for economic reasons (as firms might choose more partners when more partnering opportunities are available), and in the observed sample, for statistical sampling reasons (as observing a new node also increases the likelihood of uncovering a previously unobserved edge pointing to an existing node).

Fig. 2 (see also Table C.12) presents a mixed picture. The overall cross-sectional relationship is very noisy as small datasets tend to have a smaller average degree, but the average degree varies widely in larger datasets because non-VAT datasets can sample many

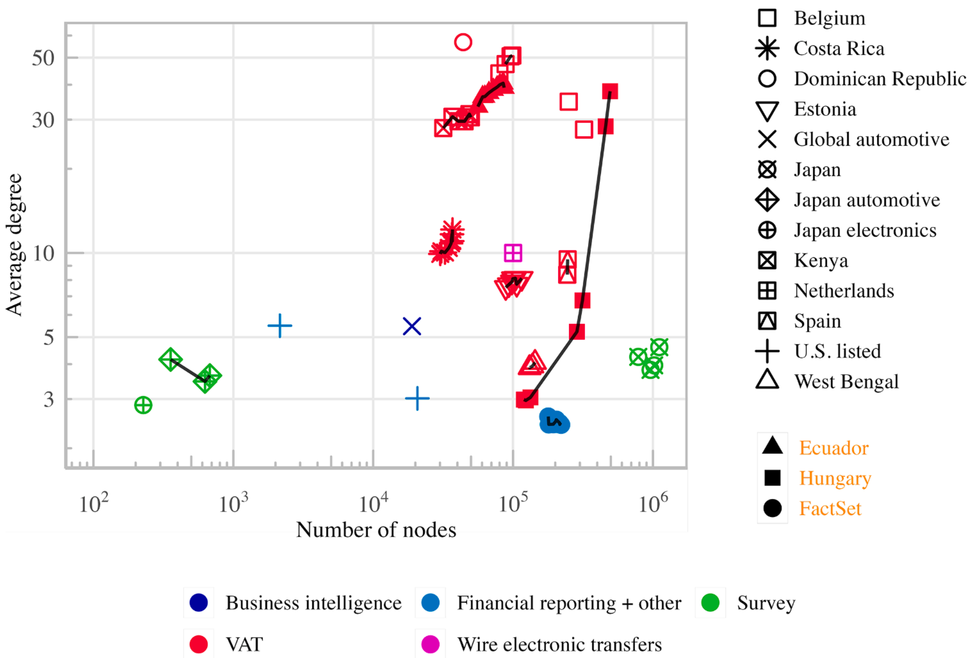

Fig. 2. Number of nodes and average degree over time on a log-log scale for Ecuador, Hungary, FactSet and the networks in the literature we reviewed. Colours refer to the data collection method. Names in orange correspond to the networks analysed in this paper and in black are the data taken from the literature. See Table C.12 for a list of the networks in the literature we reviewed. We did not include networks that were pooled over years. Connected dots belong to a dataset with a consistent cleaning procedure over time.

firms but relatively few edges. The time series dimension of each dataset, while very short, provides good evidence that the mean degree increases with size as

$$\bar{k}_t \sim N_t^{\eta}, \tag{1}$$

as commonly observed in growing networks (Dorogovtsev and Mendes, 2003).

Table 3 shows estimates of η, taking only years where the datasets are comparable (similar reporting thresholds and cleaning) for a given country. Although we do not want to draw conclusions from time series with sometimes only three observations, Eq. (1) with $0 < \eta < 1$ appears as a good hypothesis for administrative datasets. In contrast, this relationship does not apply to FactSet, where N_t varies very little. In the last column, we pool all the datasets except FactSet, Belgium and West Bengal, and run a standard panel regression with individual fixed effects, leading to a 95% confidence interval for η equal to [0.19, 0.42]. For $0 < \eta < 1$, the mean degree increases with the size of the network, even though the network is sparse, in the sense that the density goes to zero as the network size goes to infinity.

Table 3
Mean degree and network size.

	Dependent variable: $\log \bar{k}$								
	Belgium	Costa Rica	Ecuador	Estonia	FactSet	Hungary	Kenya	West Bengal	Fixed eff.
$\log N$	0.66	0.75***	0.37***	0.25*	-0.17	0.20	0.17*	0.46**	0.33***
	(0.11)	(0.19)	(0.05)	(0.11)	(0.13)	(0.12)	(0.08)	(0.08)	(0.06)
Constant	-3.69	-5.50**	-0.48	-0.79	2.98	-1.22	1.60	-4.08**	
	(1.21)	(1.99)	(0.57)	(1.24)	(1.65)	(1.42)	(0.83)	(0.89)	
Obs.	3	8	9	7	8	3	6	4	35
R^2	0.98	0.72	0.88	0.52	0.21	0.73	0.54	0.95	0.54

*p<0.1; **p<0.05; ***p<0.01
Notes: See Fig. 2 and Appendix C.1 for the data sources. For Hungary, we run the regression only for 2015–2017, when the reporting threshold did not change. For West Bengal, the regression is run quarterly for the year 2016. The fixed effects model excludes FactSet, West Bengal and Belgium but includes Spain, for which we have 2 data points. We remove Belgium because observations are not over consecutive years, FactSet because it has little variation and West Bengal because it is not a nation. If we remove Spain and/or introduce Belgium, the estimated coefficient does not change. Standard error in parenthesis.

An interesting model that requires η for calibration is that by Herskovic et al. (2020). In their model, a diversification argument suggests that firms with more partners should be less volatile. As a result, a higher η should reduce firm-level and aggregate fluctuations, conditional on N. The key insight of Herskovic et al.'s model, formalised in their Equation (11), is that firm volatilities follow a common factor, which is driven by the dispersion in firm sizes. When firm sizes are highly dispersed, firms are more likely to have customer bases dominated by a small number of very large firms, which raises concentration and amplifies firm-level and aggregate volatility. The extent to which each firm responds to this common factor depends directly on η. A higher η implies more trading links and therefore greater diversification, which weakens a firm's response to changes in the dispersion of firm sizes.

Herskovic et al. (2020) estimate $\eta = 0.13$ using Compustat data. Table 3 shows that this parameter is hard to measure on single-country datasets due to the limited number of observations. Instead, the panel estimator provides a credible estimate of 0.33. In a calibration exercise like that of Herskovic et al. (2020), this would lead to loadings several times smaller, with the exact magnitude depending on N.

A number of papers have discussed the impact of network density, although mostly in the context of aggregate I-O tables, where the number of industries is usually fixed over time. Specifically, density increases income levels (Gloria et al., 2024), decreases volatility (Miranda-Pinto, 2021) and increases business cycle asymmetry (growth skewness, Miranda-Pinto et al., 2023). Acemoglu and Azar (2020) also note an increase in the average number of suppliers in the 61-industry I-O table for the US between 1960 and 2000. They develop a theory where higher density is a symptom of cascades of input adoptions, driven by optimal choice following technology improvements, and is therefore associated with growth. Here we show that in firm-level networks, density is highly biased by reporting thresholds, and that production networks are fundamentally sparse (density $\to 0$ as $N \to \infty$, although mean degrees do increase with N). This makes it clear that, to confirm industry-level cross-country results using firm-level networks cross-country comparisons (an agenda we expect to arise in the near future), researchers will have to be particularly careful in accounting for the extent to which density and mean degree are affected by the reporting threshold rather than genuine cross-country differences.

Since most of the available datasets have a very short time-series dimension and coverage may change over time, we will focus on snapshots and avoid interpreting time-series patterns. Having looked at the average degree, we now turn to degree heterogeneity and the correlation between in- and out-degrees.

3.2. Degree distributions, correlation and assortativity

3.2.1. Degree distributions

Like many other networks, production networks exhibit very broad degree distributions. Fig. 3 shows the empirical complementary cumulative distribution function (CCDF) for the in- and out-degree of Ecuador, Hungary and FactSet. Visually, it appears that all distributions have heavy tails.

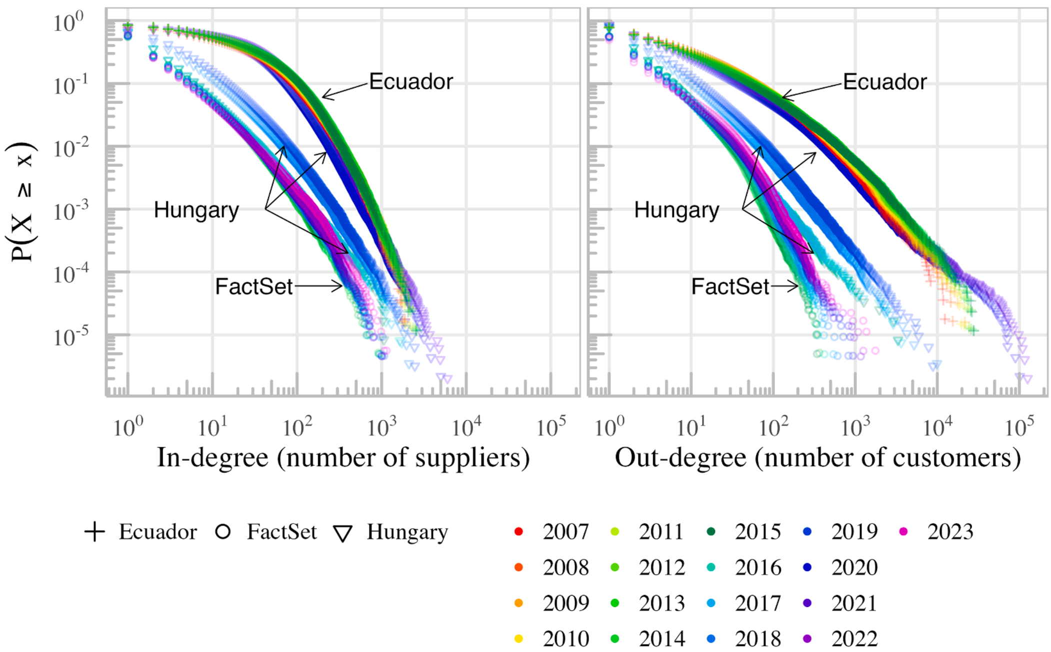

Fig. 3. Empirical CCDF of the number of suppliers (left) and customers (right) over time for the three networks we study. We compute the CCDF as $\bar{F}_n(x) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}(X_i \geq x)$, where $\mathbf{1}$ is the indicator function.

There is a striking difference between the distribution of the number of suppliers (left), where the maximum for our complete datasets is around 10^3–10^4, and the distribution of the number of customers (right), where the maximum is an order of magnitude higher (10^4–10^5), which is just an order of magnitude below the number of nodes in the network (10^5–10^6, see Fig. 2). In other words, firms with the highest number of customers sell to a very large portion of firms in the economy, whereas firms with the largest number of suppliers buy from a much smaller share of firms.

These findings confirm an industry-level result discussed in Carvalho and Tahbaz-Salehi (2019): it is uncommon to find industries with many suppliers, but some industries provide almost universal inputs. At the firm level, several studies illustrate that the range of the distribution of the number of customers tends to be much wider than that of the number of suppliers. This is evident in Figure 1a in Grigoli et al. (2023) for Chile, Figure A1 in Cardoza et al. (2025) for the Dominican Republic and Figure 2 in Alfaro-Urena et al. (2018) for Costa Rica. In these cases, as in our data (with the exception of FactSet), the maximum number of customers is roughly an order of magnitude higher than the maximum number of suppliers. In contrast, Luo and Whitney (2015), using an automotive dataset, find the opposite pattern, with a substantially heavier tail in the in-degree distribution. Based on Japanese and FactSet data, Bernard et al. (2019, Figure 2) and Taschereau-Dumouchel (2025, Figure 8), respectively, report tail behaviour and maxima that are similar for in- and out-degrees. These results illustrate how analyses based on partial or industry-specific datasets can lead to different conclusions compared to complete datasets.

To characterise the distributions more quantitatively, we estimate their tail exponent.[3] In practice (see Appendix B.1 for details), we treat degree distributions as regularly varying distributions; that is, distributions in which the share of nodes with a degree greater than k has the form

$$\text{Prob}(X > k) = \ell(k)k^{-\gamma}, \tag{2}$$

where $\ell(k)$ is a slowly varying function. Regularly varying distributions ("power laws"), have infinite variance if $\gamma \leq 2$. From a statistical point of view, the tail exponent is the key quantity used to characterise the behaviour of regularly varying distributions. It is also the key statistic of interest in many applications (e.g. Acemoglu et al., 2012).

Throughout the main body of the paper, whenever we report an estimated power-law exponent, we use the estimator from Clauset et al. (2009), which we call `plfit`, because it is standard in the literature. However, we verify all our results using state-of-the-art implementations of the tail index estimators for Generalized Extreme Value Distributions (GEVD) provided by Voitalov et al. (2019). This approach allows us to (mostly) avoid a debate on the relative quality of the fit between the power law and other distributions that may have heavy tails. See Appendix B.1 for details.

Table 4 shows estimates of the power-law exponents of the degree distributions and confirms the difference we observed between the in- and out-degree. Focusing on our complete datasets (Ecuador and Hungary 2021), the results are highly consistent, with exponents of around 2.5 for in-degree and 1.5 for out-degree. Although Appendix C.3 shows exceptions for some years and estimators, most notably for Ecuador in the early years.

Table 4
Power-law fit of the degree distributions.

Dataset	Year	In-degree	Out-degree	Estimation method	Data type	Source
Ecuador	2015	2.38	1.59	`plfit`	VAT	This paper
Hungary	2021	2.69	1.42	`plfit`	VAT	This paper
Hungary	2019	1.83	1.62	`plfit`	VAT	This paper
Hungary	2015	1.62	1.46	`plfit`	VAT	This paper
FactSet	2021	1.73	2.39	`plfit`	Financial reporting	This paper
Japan	2005	1.37	1.46	OLS_{CCDF}	Business services	Bernard et al. (2019)
Japan	2005	1.37	1.25	`plfit`	Business services	Ohnishi et al. (2010)
Japan	2006	1.35	1.26	`plfit`	Business services	Fujiwara and Aoyama (2010)
Dutch bank 1	2019	1.44	1.28	`plfit`	Payment system	Ialongo et al. (2022)
Dutch bank 2	2019	1.77	1.31	`plfit`	Payment system	Ialongo et al. (2022)
Chile	2019	0.28	0.40		VAT	Grigoli et al. (2023)
Dominican Rep.	2019	0.30	0.45	`paretofit`	VAT	Cardoza et al. (2025)
Costa Rica	2008	0.58	0.73		VAT	Alfaro-Urena et al. (2018)
US listed	04/2012- 06/2013	2.76	1.88	`plfit`	Financial reporting	Wu and Birge (2014)
US listed	1979–2007	1.00			Financial reporting	Atalay et al. (2011)
US listed	1978–2013	1.25	1.44		Financial reporting	Barrot and Sauvagnat (2016)
FactSet US	2016	0.97	0.83	Rank 1/2 estimator	Financial reporting	Taschereau-Dumouchel (2025)

Notes: The in-degree is the number of suppliers and the out-degree the number of customers. Most studies use `plfit` (Clauset et al., 2009). Bernard et al. (2019) regress the log CCDF on the log degree (OLS_{CCDF}). Cardoza et al. (2025) fit a pure Pareto using Stata's `paretofit`. A few other studies also appear to fit a pure Pareto (Barrot and Sauvagnat, 2016; Alfaro-Urena et al., 2018; Grigoli et al., 2023). Taschereau-Dumouchel (2025) uses the rank 1/2 estimator of Gabaix and Ibragimov (2011). For the data type taxonomy, see Table 1.

Why do other studies report different results? This is due to a combination of data types (i.e. VAT or not, reporting threshold or not) and estimation method, with inherent cross-country variation playing a comparatively minor role.

In Appendix D, we examine how the power-law exponents change as the reporting threshold increases. We find a clear pattern whereby the exponent of the distribution of the number of customers is *not* affected, whereas the exponent for the number of suppliers is biased towards smaller values (i.e. fatter tails); this can be seen in the Table 4 for the early years in Hungary. This finding explains why Hungary's in-degree distribution exhibits an exponent smaller than 2 in the early years, a consequence of reporting threshold bias.

[3] Appendix C.2 also shows that the number of nodes with in- or out-degree equal to zero can be relatively high (up to about 40%).

Next, and still focusing on VAT studies, we find that Chile, Dominican Republic and Costa Rica report very low exponents, below 1. The method used to estimate the exponents is not always clearly described and a value lower than 1 is highly implausible,[4] so we do not give further consideration to these estimates. Note also that these papers, like ours, report a higher maximum for the out-degree distribution than for the in-degree. This observation is difficult to reconcile with their finding of a smaller tail exponent for in-degree than for out-degree, since a smaller exponent implies fatter tails, which would typically be associated with more extreme realisations. In our estimates, the higher maximum for out-degrees is consistent with a fatter-tailed out-degree distribution, as reflected in a smaller exponent. Finally, studies with non-VAT datasets typically report exponents between 1 and 2, with some exceptions. We have already seen that non-VAT datasets differ substantially from complete VAT datasets, with, for instance, a mean degree of an order of magnitude smaller than in complete datasets.

We conclude that the power-law exponents for the degree distributions of complete, nationwide production networks are around 1.5 for the out-degree and around 2.5 for the in-degree. The out-degree distribution is much broader than the in-degree distribution: the maximum out-degree is about an order of magnitude larger than the maximum in-degree. We are not aware of any models that have successfully predicted these patterns quantitatively, although several models do feature an asymmetry between in- and out-degrees. Acemoglu and Azar (2020) predict that the out-degree distribution is substantially more unequal than that of the in-degree and can exhibit power-law tails. The intuition underlying this asymmetry is that certain inputs are affected by cost-reducing technological improvements, which lower their prices and induce downstream adoption, leading successful suppliers to accumulate a large number of customers. In the model of Carvalho and Voigtländer (2014), central inputs (in the network-proximity sense) are more likely to be widely adopted over time; this self-reinforcement mechanism leads to a power-law distribution of out-degree. These contributions point to a technological explanation for the asymmetry between in- and out-degrees, whereby some products or industries arise as near-universal suppliers or general-purpose technologies. Finally, Bernard and Zi (2022) offer a benchmark balls-in-bins model, matching buyers of a given size with sellers of a predetermined attractiveness (defined as the probability of receiving any given purchase order). The model can generate degree asymmetry by assuming different variances for the distributions of buyers' size and sellers' attractiveness.

As noted by Carvalho and Tahbaz-Salehi (2019), while both industry- and firm-level networks exhibit highly heterogeneous out-degree distributions, industry-level networks display much weaker heterogeneity in in-degrees than firm-level networks do. Here, we provide precise estimates of the tail exponents of the in- and out-degree distributions, confirming that firm-level in-degrees are substantially more heterogeneous than industry-level in-degrees, but remain markedly less heterogeneous than firm-level out-degrees.

Correlation between in- and out-degrees. Do firms with more customers also tend to have more suppliers? Fig. 4 shows that yes, in- and out-degrees are positively correlated. Firms with many customers also tend to have many suppliers, although, as noted previously,

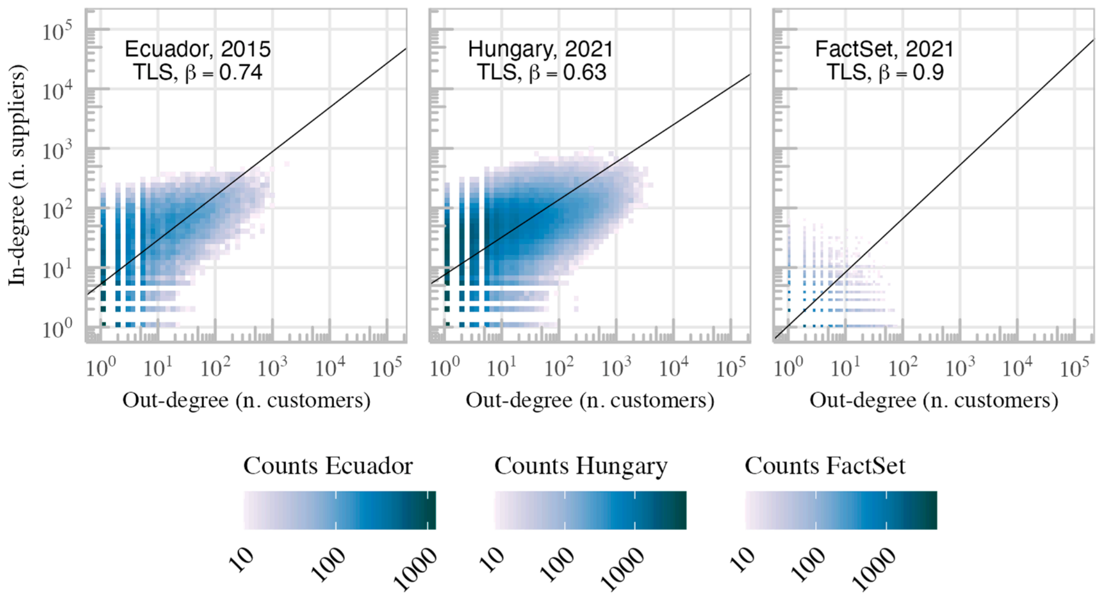

Fig. 4. 2D histogram for the number of customers (x-axis) and suppliers (y-axis). We divide each axis into 60 log-spaced bins and then count the number of data points falling in each square. We do not show squares with less than 10 observations. TLS stands for total least squares and β is the estimated coefficient. The black line shows the TLS fit.

[4] Power laws with an exponent less than 1 are extremely rare, especially in economics, although the exponent is around 1 for city and firm sizes. Applying three GEVD estimators to the in- and out-degree distributions of 49 networks, Voitalov et al. (2019, Table II) report only a handful of cases with $\gamma < 1$ and no case where all the 3 estimators suggest $\gamma < 1$.

they hardly have as many suppliers as customers. This suggests that for every doubling in the number of customers, we should see less than a doubling in the number of suppliers. In other words, the slope of the in-degree vs out-degree relation should be less than 1 or, equivalently, the slope of the out-degree vs in-degree relation should be greater than 1. To quantify this slope, we use Total Least Squares (TLS), which finds the line that minimises the squared residuals, measured as the perpendicular distance from the data point to the regression line. Unlike Ordinary Least Squares, TLS is symmetric (see Appendix B.2.1, where we also provide covariance matrices from which regression coefficients and R^2 values can be derived).

We find that the slope is indeed substantially less than 1. Taking the value of 0.63 for Hungary, firms with 10% more customers have only 6.3% more suppliers. Since firms with more sales have approximately proportionately higher expenses (Appendix C.5), if we are prepared to make a time series interpretation of our cross-sectional results, this may suggest that firms grow at the extensive margin on the customer side and at the intensive margin on the supplier side. However, Appendix B.2.4 warns against thinking in terms of a representative firm, since firms become increasingly heterogeneous in out-degree as they grow.

If the asymmetry between in- and out-degrees is due to different extensive and intensive margin behaviour, this may well have an impact on how supply and demand shocks propagate. Taschereau-Dumouchel (2025) offers a model where the extensive margin responds strongly to the business cycle (productivity shocks), showing that these topological adaptations likely weaken the impact of idiosyncratic shocks on aggregate fluctuations. Because Taschereau-Dumouchel relies on the FactSet dataset, which displays broadly similar in- and out-degree distributions, he calibrates the joint degree distribution using a bivariate power law with a common exponent. It would be interesting to explore how these results change with a calibration based on asymmetric degree distributions.

3.2.2. *Assortativity*

An interesting hypothesis in the literature is that supply chains are characterised by negative degree assortativity, meaning that highly connected firms tend to be connected to less connected firms (Bernard et al., 2019; Fujiwara and Aoyama, 2010; Bernard et al., 2022; Alfaro-Urena et al., 2018). Given the highly heterogeneous degree distributions, negative degree assortativity can be the symptom of nestedness,[5] whereby large firms connect to all types of firms, but small firms connect only to large firms.

We compute degree assortativity, as defined in Newman (2003), as the Pearson correlation coefficient between the degrees of firms at opposite sides of the same edge. Newman's metric, which we label r, is easy to interpret: it varies from -1 for perfectly disassortative networks to 1 for perfectly assortative networks, and equals zero when there is no correlation between the degree of connected nodes. For directed networks, there are four degree assortativity measures, each of which combines the in- and out-degrees of the suppliers and customers.

To give some context, social networks are frequently characterised by assortative mixing ($r > 0$), while technological and biological networks often have disassortative mixing ($r < 0$, Newman, 2003). Canonical random graphs such as Erdős-Rényi (ER) or Barabási-Albert models have zero assortativity in the limit of a large number of nodes, but mild dissassortativity can emerge in models with no specific force leading to it. For instance, Boguná et al. (2004) show that when degree distributions are very heterogenous, a configuration model corrected to have only simple edges would have mild dissassortativity. Bernard and Zi (2022) show that a simple balls-in-bins model can generate dissassortative networks without requiring specific sorting or matching mechanisms.

For Ecuador and Hungary, we find a weakly negative assortativity between -1.5% and -13% (Table 5), depending on the type of assortativity measure. Similar values are reported in the literature. In contrast, for FactSet, we find an assortativity mildly positive or close to zero. Of all the possible ways to compute assortativity, the largest in magnitude is the correlation between the out-degree of suppliers and the in-degree of customers; that is, suppliers with many customers tend to sell to customers with few suppliers. Put simply, firms that have few suppliers tend to connect with large suppliers.

Table 5
Assortativity coefficients.

Dataset	Year	$r_{k,k}$	$r_{k^{in},k^{out}}$	$r_{k^{out},k^{in}}$	$r_{k^{in},k^{in}}$	$r_{k^{out},k^{out}}$	Source
Ecuador	2015	−12.0	−3.8	−13.0	−10.5	−5.3	This paper
Hungary	2021	−7.6	−1.5	−8.9	−5.6	−2.4	This paper
Hungary	2019	−4.4	−1.5	−5.6	−3.1	−2.7	This paper
Hungary	2015	−5.5	−2.9	−7.0	−5.3	−3.5	This paper
FactSet	2021	0.9	1.7	−0.3	−0.6	3.1	This paper
Japan	2006	−7.5	negative	negative			Fujiwara and Aoyama (2010)
Japan listed	2016	−21					Krichene et al. (2019)
West Bengal	2016Q4	−6.2					Kumar et al. (2021)

Notes: Assortativity coefficients as defined in Newman (2003). $r_{k^{in},k^{out}}$ denotes the correlation between the suppliers' in-degrees and the customers' out-degrees, where each edge is a data point; other columns are interpreted in a similar way. All values are multiplied by 100.

In Appendix D, we show that assortativity becomes less negative as the reporting threshold increases, which explains the difference between the various years in Hungary.

[5] Nestedness means that the rows and columns of the adjacency matrix can be rearranged so that the upper left part is mostly full of positive values, while the rest is mostly full of zeros. Mariani et al. (2019) define a network as perfectly nested if "the degree of i is smaller than the degree of j, then the neighbourhood of i is contained in the neighbourhood of j", and note that nested networks are usually disassortative.

Bernard et al. (2022), Bernard et al. (2019), Alfaro-Urena et al. (2018) and Cardoza et al. (2025) use a different measure of assortativity. Their downstream assortativity measures how the average number of suppliers of i's customers changes as i's number of customers changes. Likewise, upstream assortativity measures the change in the average number of customers of i's suppliers as i's number of suppliers changes. Regardless of the assortativity measure used, they find a negative degree assortativity for Belgium (Bernard et al., 2022), Japan (Bernard et al., 2019), the Dominican Republic (Cardoza et al., 2025) and Costa Rica (Alfaro-Urena et al., 2018).

3.3. Reciprocity, clustering and path lengths

In this section, we report standard binary network quantities. We start by documenting the substantial extent to which links are reciprocal. Then, using the undirected version of the network, we show that the prevalence of closed triangles among all possible triples (global clustering) is low and can mostly be explained by degree heterogeneity. In contrast, the average proportion of a node's neighbours that are themselves connected (local clustering) is much higher, and higher than a random benchmark that preserves the degree distribution. Finally, we demonstrate that the shortest paths between pairs of nodes are very small, typically around 3 steps. For most of these properties, however, we show that non-administrative datasets or datasets with a high reporting threshold provide biased results.

3.3.1. Reciprocity

Reciprocity is the probability that an existing edge is reciprocated. In social networks, it can be very high. For example, in friendship networks in US schools, the reciprocity is between 0.3 and 0.5 (Ball and Newman, 2013). For firm-level production networks, we find that reciprocity is much lower but still much higher than expected in an equivalent ER random graph, where it is very close to zero. Table 6 shows the empirical values of the reciprocity in Ecuador (around 5%), Hungary (4–9% depending on the threshold) and FactSet (around 2%).

Table 6
Reciprocity, clustering and average path length.

Dataset	Year	Reciprocity	C_g		$\bar{C}$		Average path length			Source
		Emp.	Emp.	CM	Emp.	CM	Emp.	ER	CM	
Ecuador	2015	4.6	2.5	2.8	28.0	13.3	2.8	2.9	2.8	This paper
Hungary	2021	3.9	0.5	0.5	19.6	8.1	2.9	3.4	3.0	This paper
Hungary	2019	6.7	1.2	0.6	11.4	1.1	4.1	5.1	3.9	This paper
Hungary	2015	8.7	1.1	0.7	12.9	1.3	4.8	6.7	4.3	This paper
FactSet	2021	2.5	1.7	0.2	3.2	0.3	6.1	8.0	4.9	This paper
FactSet US	2016		2.4				4.8			Taschereau-Dumouchel (2025)
Japan	2005						4.6	7.0		Ohnishi et al. (2010)
Japan	2006		0.2	1.8	4.6		5.6	6.9		Fujiwara and Aoyama (2010)

Notes: All values are in percent except for the average shortest path length ("Average path length"). C_g and $\bar{C}$ are the global and average local clustering coefficients. "Emp." stands for empirical, ER and CM for Erdos-Renyi and configuration model. For both the ER and CM, we run 100 simulations for the clustering coefficients, and 10 for path lengths; each simulation samples 10^4 pairs at random.

3.3.2. Clustering

In social networks, it is very common that two of a person's friends are also themselves friends. Do we observe a similar pattern among firms? That is, if a firm transacts with two other firms, are these two firms likely to have a supply-chain relationship? We convert the networks to undirected networks by assuming that each directed edge is an undirected edge and we remove duplicated edges that arise from reciprocal edges. We then look at two standard metrics: the local and the global clustering coefficient.

The *global clustering coefficient* gives information about the density of triangles in the entire network. It gives the share of paths of length two that are closed. The most common way to write its definition is (Newman, 2018)

$$C_g = \frac{\text{number of triangles} \times 3}{\text{number of connected triples}},$$

where the factor of three in the numerator corrects for the fact that a triangle gets counted three times when we count the number of connected triples in the network. By contrast, the *local clustering coefficient* is the property of a single node:

$$C_i = \frac{\text{number of pairs of neighbours of } i \text{ that are connected}}{\text{number of pairs of neighbours of } i}.$$

Note that C_i is undefined for firms with degree 1 since they do not have a single pair of neighbours ($C_i = 0/0$); we exclude these firms from the average. A low local clustering coefficient is an indicator of centrality, in the sense that firms with low clustering coefficients are, by definition, bridging pairs of firms that are themselves not connected.

Table 6 shows the average local clustering coefficients $\bar{C}$ and the global clustering coefficient C_g in our three networks and in the literature. Both the global and the average local clustering coefficients are substantially larger in Ecuador than in other networks, although the complete Hungarian (2021) network also features high local clustering. We compare these results to a configuration model (CM), a random benchmark that preserves the nodes' degrees but is otherwise random.[6] The complete administrative datasets provide a relatively clear picture: global clustering is comparable to the random benchmark, while local clustering is much higher, although the absolute numbers are quite different between Ecuador and Hungary 2021.

The patterns in the incomplete datasets are different, with global and local clustering higher than in the random benchmark. However, the magnitude of the local clustering, both the empirical and simulated ones, is considerably lower than in the complete datasets.

As is well known, the difference between local and global clustering coefficients is partly due to the fact that the degree distribution is highly heterogeneous and that there is a negative correlation between node-level local clustering and degree. To see this, consider a firm with a very high degree. It usually has a very low local clustering, otherwise the network would be dense. Thus, a large degree firm has a huge number of *pairs* of partners that are not connected. Each of these pairs contributes to the overall number of triangles that are not closed, leading to a low global clustering. However, the average local clustering coefficient is an average where each firm weighs equally. So large firms contribute little and the average is driven by the many small degree firms, which can easily have fairly high local clustering coefficients.

Overall, the excess local clustering compared to a CM shows that matching is not only determined by degree. An intuitive explanation could be geography, as Bernard et al. (2019) and Mungo et al. (2023) find that firms tend to connect with firms that are closer in space, which would make reciprocal links and triads more likely. More generally, to explain the presence of excess clustering, a successful class of models is the one based on hidden geometries, where nodes are more likely to be connected if they are close in some underlying metric space (Serrano et al., 2008).

3.3.3. Paths

As in Section 3.3.2, we convert the directed networks into undirected ones. We then consider the *shortest path length* between two nodes; that is, the fewest number of hops one has to make to go from one to the other.

In production networks, we would expect that short path lengths imply that shocks at the firm level can reach most firms in the network more quickly and, potentially, more strongly. For example, Carvalho et al. (2021) study the impact of the 2011 great Japan earthquake and find that indirect suppliers and customers of directly affected firms were also impacted, but the effect decays substantially with network distance. While studies of firm-level production networks typically do not report statistics of path lengths, we provide these because we think that models of endogenous formation of production networks should try to match them.

Fig. 5 shows the distribution of the length of the shortest paths for our three networks. First, the distribution is stable over time except for Hungary, where we see the strong effect of the reporting threshold. In the early years, Hungary's distribution is closer to

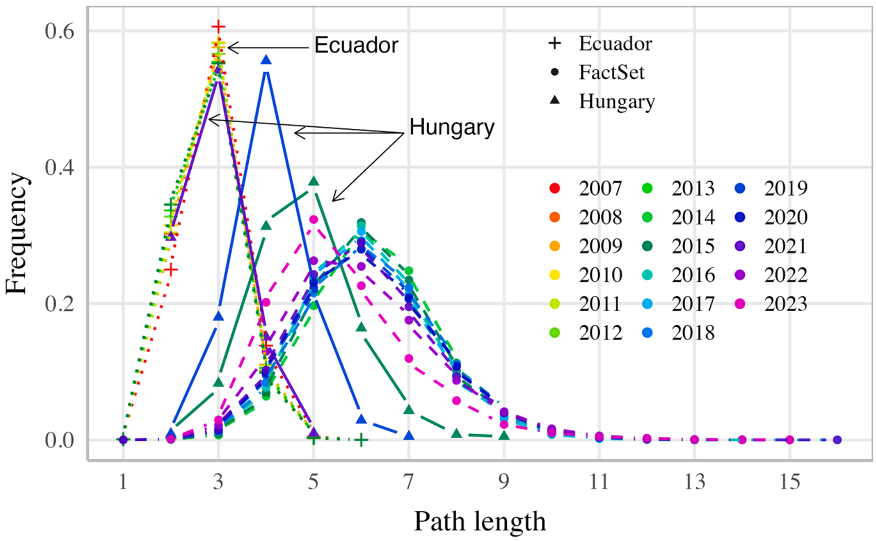

Fig. 5. Distribution of the length of the shortest paths in Ecuador, Hungary and FactSet over time. We convert the networks to undirected networks by assuming that every directed edge is undirected and we remove duplicated edges. Due to the long computation time, we compute the frequency of shortest path lengths in a random sample of 10^4 pairs for all our networks. Additionally, for Hungary, we compute the frequency of shortest path lengths only for the three years in which the reporting threshold does not change during the year (2015, 2019 and 2021).

[6] Consider a graph G that we wish to compare to the random benchmark. In theory, the CM should provide the set of all possible random graphs that have the exact same degree sequence as G. To compare G with the random benchmark, one should draw graphs uniformly from this set of random graphs and compute the metric of interest. In practice, it is very hard to sample uniformly from the set of *simple* (no self-loops and no duplicated edges) *connected* graphs. We use the fast algorithm `sample_degseq(…, method = 'simple')` from the `igraph` package in `R`, but with the disadvantage that it allows duplicated edges and self-loops, and can return disconnected graphs. The prevalence of self-loops, duplicated edges and disconnected nodes is small. We omit the comparison with ER graphs, where both clustering coefficients are very close to zero.

FactSet's, which is missing many firms and relationships. After the threshold is removed, the distribution in Hungary is astonishingly similar to that of Ecuador. Second, the mode for our two complete networks is 3 and the average is between 3 and 4 (Table 6).

As a point comparison, Carvalho (2014) reports an average path length of 4 in the US BEA I-O tables (417 industries). This is a surprising result as we might have expected a higher distance in a very large firm-level network than in a small, aggregated industry network.[7] This can be due to differences between the US economy and Hungary/Ecuador, but we think that it is more likely due to national accounting conventions; for example, because national accounts do not show links from wholesalers and retailers to their suppliers of goods destined to be resold.

Are these results surprising given the density of the network and the degree distributions reported in the previous sections? Given the large number of nodes, one could have expected that if we pick two firms at random, it would typically take many steps to connect them. It turns out that in most networks, the average shortest path length is very small, a phenomenon known as the *small-world* effect (Milgram, 1967; Watts and Strogatz, 1998). This effect is relatively well understood since even simple models of random network formation produce fairly short path lengths. In Table 6, we compare the average shortest path length in our networks with those expected from an ER model and CM.[8] We also include the results reported by Fujiwara and Aoyama (2010) and Ohnishi et al. (2010) for Japan and compute the ER benchmark for them using their published data on the number of nodes and edges. In complete datasets, the empirical average shortest path length is very close to that of the CM. For incomplete datasets, there is a difference, but we refrain from interpreting this further.

4. Results on weighted networks

Due to a lack of data, far less is known about weighted networks compared to binary networks. The Belgian network is probably the most studied firm-level network with information on the monetary values of firm-to-firm transactions. Recently, however, a number of other datasets have appeared, mostly from VAT data but also from payment systems (Table 1). In this section, we provide a detailed analysis of the distribution of key quantities of weighted networks.

We do not have data for the weights in FactSet, but we do have data for our other two networks, Ecuador and Hungary. We first discuss the distribution of weights, finding a power-law exponent slightly above 1. We then show that the strength distributions have an exponent very close to 1, confirming studies of firm sizes. Next, we look at the relationships between strengths and degrees: the relationship between network sales and the number of customers, and between network expenses and the number of suppliers.

While classic I-O analysis and more modern models focus on technical coefficients and the Leontief inverse, we do not have the necessary data to compute these quantities for Ecuador. To sidestep this issue, as a final result, we consider the distribution of the influence vector, a centrality measure motivated by the benchmark Cobb-Douglas model with uniform final demand shares. In this case, each element of the influence vector gives the elasticity of aggregate output to a shock to the total factor productivity (TFP) of each firm (Acemoglu et al., 2012; Carvalho and Tahbaz-Salehi, 2019).

4.1. *Distributions of the value of transactions*

A small number of studies report summary statistics for the value of the transactions (or "network weights") and find them to be heavy-tailed (Dhyne et al., 2015; Magerman et al., 2016; Bernard et al., 2022; Huneeus, 2020); Fig. 1 in Section 1 confirms these findings. More quantitatively, we find that the estimated power-law exponents are remarkably similar and around 1.1–1.2 for both Ecuador and Hungary, over time and regardless of the estimation method used (see Table 7 and C.16 for a more detailed account).

Table 7
Tail exponents for weighted network quantities.

Dataset	Year	Weight	In-strength	Out-strength	Influence	Source
Ecuador	2015	1.14	0.88	0.92	1.28	This paper
Hungary	2021	1.18	1.01	1.02	1.23	This paper
Hungary	2019	1.14	0.99	1.00	1.42	This paper
Hungary	2015	1.15	1.05	0.92	1.43	This paper
Belgium	2012				1.12	Magerman et al. (2016)
Dutch bank 1	2019		1.03	1.05		Ialongo et al. (2022)
Dutch bank 2	2019		0.69	0.72		Ialongo et al. (2022)

Notes: Parameters estimated using `plfit`.

The total value of the transactions (or the total weight of the network) is 78.8 billion for Ecuador in 2015 and 81.7 billion for Hungary in 2021.

[7] In the ER model, the average path length scales with size as $\ln N$, while in growing networks that lead to a power-law degree distribution with $1 < \gamma < 2$, the average path length increases much more slowly (Cohen and Havlin, 2003), scaling as $\ln \ln N$, because the presence of hubs shortens the distance between most pairs of nodes.

[8] In the ER model, each possible edge exists with probability p. We calibrate p to the empirical density of the (undirected) network. In the ER model, when p is such that the mean degree is less than 1, the network likely consists of disconnected clusters, and when $p > 1$, a "giant" component emerges. In our case, the mean degree is always substantially above 1, but a draw from the ER ensemble still almost always contains a small number of nodes outside of the LWCC. We remove these before computing shortest path lengths.

4.2. Distributions of network sales and expenses (strengths)

Our data is only on *network* sales and expenses (out- and in-strengths),[9] but we would expect that they are highly correlated with other indicators of a firm's size. It is well established that the distribution of firms' revenues has an exponent close to 1, with similar exponents also found when size is measured in terms of employees or receipts (Axtell, 2001).

Fig. C.12 shows the distribution of network sales and expenses over time for Ecuador and Hungary. Both are markedly stable across years, except, of course, for Hungary, where the distributions exhibit a clear break at the reporting threshold. While the threshold applies to the weights, many firms have an in- or out-degree equal to 1, so their strength is equal to the weight of their single edge. Table 7 shows the estimated power-law exponents, which are close to 1 (Tables C.17 and C.18 show that the GEVD estimators tend to suggest exponents slightly above 1).[10]

As expected, firms with higher network sales also tend to have higher network expenses, although there is considerable dispersion at lower values; see Fig. C.13 and Fig. B.9.

4.3. Strength-degree relationships

Do firms with more customers have higher network sales? Do firms with more suppliers have higher network expenses? Intuition suggests that yes, but the exact value of the elasticities is important. For example, we may think that the marginal customer is less important than the average customer, so that growing the customer base by 1% would result in less than a 1% increase in sales.

Fig. 6 shows the 2D histogram for our two datasets. Again, the 2015 Ecuadorian network and the 2021 Hungarian network appear remarkably similar. There is a clear positive relationship between strength and degree in all four cases.

An interesting observation is that there appears to be a considerable number of large firms (in terms of network sales or expenses) with only a few partners. In Fig. B.9, we document heteroskedasticity more precisely and find that regressions of strength on degree are fairly homoskedastic, whereas regressions of degree on strength are highly heteroskedastic, due to the presence of small-degree firms that have high strengths.[11]

While Figure S1a in Ialongo et al. (2022) provides similar visual evidence for two different datasets of firm-to-firm transactions that occur between clients of a bank, we are not aware of any papers discussing or explaining this pattern. We think it is consequential in the context of systemic risk. For instance, if a firm's importance depends positively on its size and its vulnerability to shocks depends negatively on its ability to diversify its supplier or customer base (e.g. Herskovic et al., 2020), then large firms with few partners are both important and vulnerable. The upper left quadrant of each panel in Fig. 6 shows that there are many such firms.

Table 8 quantitatively investigates the data from Fig. 6. While some papers report regressions of strength on degree, others report regressions of degree on strength, so we show both. We find good consistency among our datasets and with the literature, despite some methodological differences in the choice of variables (including industry fixed effects does not affect the results, see Appendix E.2). At first sight, the elasticities do not seem to change dramatically with the change in the reporting threshold, although the out-degree elasticity of network sales appears to be more affected than the in-degree elasticity of network expenses.

A more thorough investigation (Appendix D) shows that the reporting threshold tends to increase the value of the coefficient, except in the case of the regression of in-strength on in-degree, where there is no clear bias. While the reporting threshold leads to only "small" increases in the measured slope coefficients, this can represent a serious bias for studies that need quantitatively precise estimates.

To give one key example, Bernard et al. (2022) highlight the fact that in Belgium, the regression of the log of sales on the log of the number of customers has a coefficient lower than one, suggesting that a firm's marginal customer represents fewer sales than its average customer. Table 8 shows that this fact only holds in Belgium and Ecuador, and fails in several other instances, potentially questioning the premise of Bernard et al.'s theory (Bernard et al., 2022, use variables demeaned at the NACE 2 level, but as we show in Appendix E.2, this does not affect substantially any of the elasticities we report). However, the analysis of all coefficients against the reporting threshold, in Appendix D.4, makes it clear that the other coefficients (greater than 1) are typically untrustworthy because they are biased. Indeed, Fig. D.18 shows that the coefficients obtained in truncated networks are biased in the same direction and with a similar magnitude as one would expect based on our synthetic truncation of the Ecuadorian network. Thus, while our estimate for the complete Hungarian dataset is still slightly above 1, this careful analysis of the impact of the reporting threshold suggests that Costa Rica and Estonia would likely feature a coefficient around or less than one if there were no reporting threshold.

Beyond the example of Bernard et al. (2022), these elasticities are often central to the papers reporting them, as they are used to assess the relative importance of the intensive and extensive margins. In the models of Arkolakis et al. (2023) and Demir et al. (2023), for instance, firms expand their network as they grow in size and these elasticities are used to calibrate the model.

[9] Network expenses exclude labour costs, imports, taxes and subsidies. Network sales exclude sales to final demand, exports and taxes. Sales and expenses on capital goods are still likely to be present. In Appendix A.2, as a benchmark, we use intermediate sales plus GFCF from national accounts, without including exports or other final demand components.

[10] By examining qq-plots, we also find that, in contrast to other distributions reported in the paper, the lognormal provides a good fit, a point also noted by Ialongo et al. (2022).

[11] Fig. 6 gives the visual impression that the strength-degree relationship may be heteroskedastic, but it is important to bear in mind that typically, the sample maximum increases and the sample minimum decreases with sample size. In other words, the reason why we observe a narrower range of values for strength conditional on high degrees, compared to conditional on low degrees, is simply a consequence of there being fewer high-degree observations from which to sample, rather than a lower variance. Fig. B.9 in the Appendix suggests heteroskedasticity mostly in the out-degree–out-strength relationship.

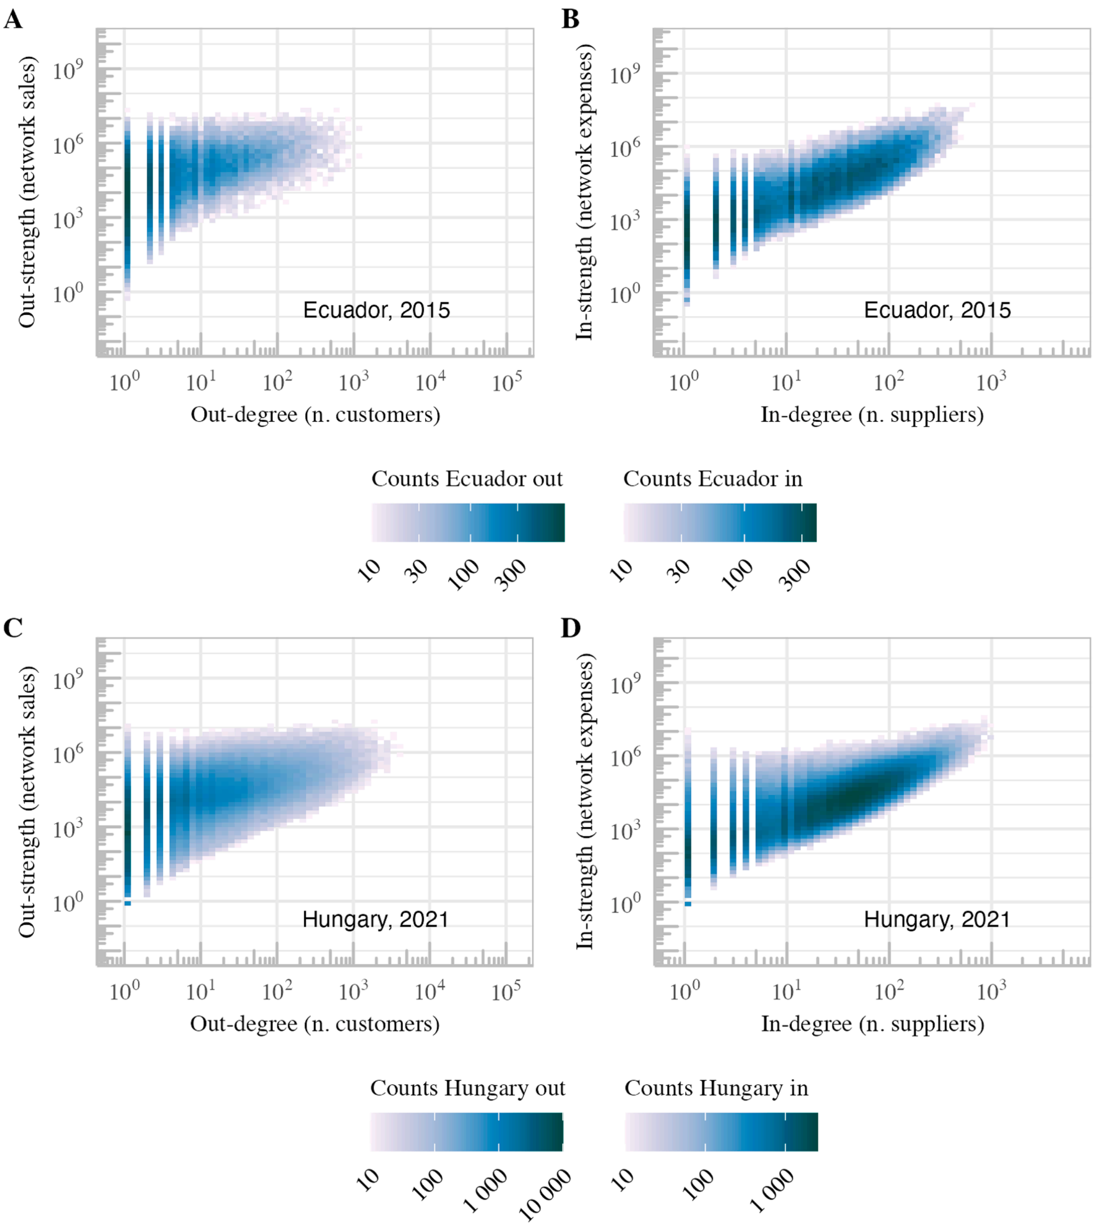

Fig. 6. **A**, **C**, 2D histograms for the degrees (number of customers or suppliers) against strength (network sales or expenses), for Ecuador in 2015 and Hungary in 2021. We bin both axes into 60 equally-spaced bins and we then count the number of observations in each square. We do not color squares that have less than 10 observations.

4.4. The influence vector

We conclude by characterising the distribution of the *influence vector*, which is defined as

$$\boldsymbol{v} \equiv \frac{\alpha}{N}\left[\boldsymbol{I} - (1-\alpha)\boldsymbol{P}^{\text{in}}\right]^{-1}\mathbb{1}\,,$$

where $\alpha \in (0,1]$ is a parameter, $\boldsymbol{P}^{\text{in}}$ is the (column-stochastic) matrix of input shares computed as $P^{\text{in}}_{ij} = Z_{ij}/\sum_i Z_{ij}$ where Z_{ij} is the expenses by j on input i, N is the number of firms, $\boldsymbol{I}$ is an identity matrix and $\mathbb{1}$ is a vector of ones.

The influence vector corresponds to PageRank, applied to the appropriately transposed weighted matrix with a damping factor of $1-\alpha$. This formulation enables the use of a fast implementation from the `igraph` package in `R`, with a well-understood error tolerance.

PageRank, a classic measure of centrality, was popularised in the context of production networks by Acemoglu et al. (2012), and was subsequently computed by Magerman et al. (2016) using firm-level data (using $\alpha = 0.2$) and by Carvalho (2014) on US I-O tables (using $\alpha = 0.5$). While the Domar weights are often the key statistics of interest, for instance, in assessing the impact of micro-level

Table 8
Strength-degree elasticities.

Dataset	Year	$\ln s \| \ln k$	$\ln k \| \ln s$	R^2	Source
Network sales & number of customers					
Ecuador	2015	0.88	0.31	0.27	This paper
Hungary	2021	1.05	0.36	0.37	This paper
Hungary	2019	1.22	0.40	0.49	This paper
Hungary	2015	1.38	0.38	0.53	This paper
Belgium	2014	0.77		0.35	Bernard et al. (2022)
Chile	2018–2019		0.42	0.46	Arkolakis et al. (2023)
Costa Rica	2008–2015	1.20			Alfaro-Urena et al. (2018)
Estonia	2021	1.10	0.46	0.51	Criscuolo et al. (2024)
Estonia	2015	1.14	0.47	0.53	Criscuolo et al. (2024)
Turkey	2015		0.44	0.33	Demir et al. (2023)
Network expenses & number of suppliers					
Ecuador	2015	1.55	0.41	0.63	This paper
Hungary	2021	1.35	0.44	0.60	This paper
Hungary	2019	1.43	0.47	0.67	This paper
Hungary	2015	1.39	0.45	0.63	This paper
Chile	2018–2019		0.45	0.20	Arkolakis et al. (2023)
Costa Rica	2008–2015	0.89			Alfaro-Urena et al. (2018)
Estonia	2021	1.47	0.50	0.74	Criscuolo et al. (2024)
Estonia	2015	1.46	0.51	0.74	Criscuolo et al. (2024)
Turkey	2015		0.58	0.61	Demir et al. (2023)
Japan	2005–2010		0.33		Bernard et al. (2019)
Sales & number of partners					
Chile	2014–2020		0.33	0.25	Miranda-Pinto et al. (2023)
Japan	2005	1.30			Watanabe et al. (2013)

Notes: OLS regressions of either the log of strength on the log of degree (column $\ln s \,|\, \ln k$) or the log of degree on the log of strength (column $\ln k \,|\, \ln s$). All the observations equal to zero are dropped. Japan is the only non-VAT dataset. Bernard et al. (2022) use network sales and add 4-digit industry fixed effects. Alfaro-Urena et al. (2018) use total sales demeaned by industry and keep only firms in the 3 highest volume industries. Arkolakis et al. (2023) use total sales and include year and state fixed effects. Demir et al. (2023) consider manufacturing firms and use sales for both the in-degree and out-degree regressions. Miranda-Pinto et al. (2023) use total sales and firms with more than 5 employees. In Miranda-Pinto et al. (2023) and Watanabe et al. (2013), the degree is the number of suppliers and customers.

shocks on aggregate fluctuations, this measure of centrality can be thought of as isolating the role of the network (Carvalho and Tahbaz-Salehi, 2019).

Fig. 7 shows the distribution of the influence vector for Ecuador and Hungary. These clearly display heavy tails with a constant slope over three orders of magnitude and an overall shape very similar to that reported for US industries by Carvalho (2014).

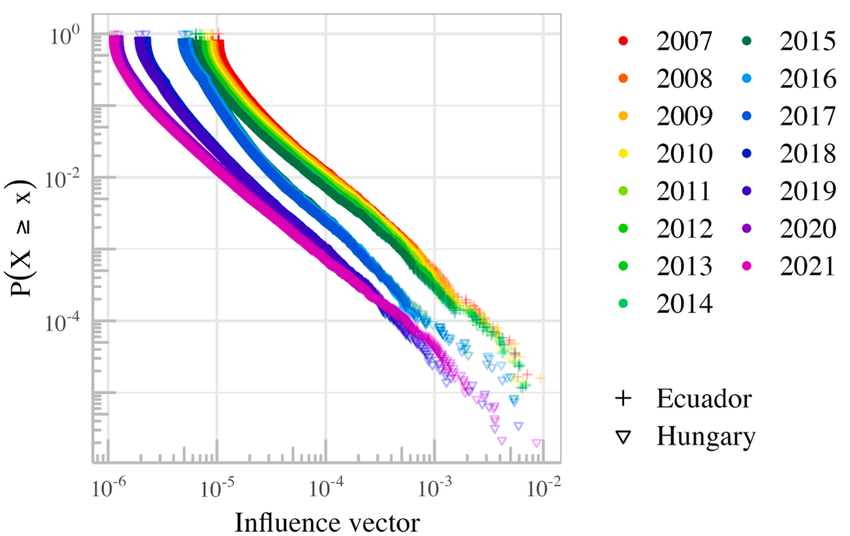

Fig. 7. Empirical CCDF of the influence vector for Ecuador (cross) and Hungary (triangle) over time.

Table 7 reports the estimated power-law exponents for our networks (see Table C.19 for more years and estimators) and for Belgium (Magerman et al., 2016). The firm-level datasets display fatter tails compared to the industry-level estimate of 1.48 by Carvalho (2014). The estimates for Hungary are higher in the early years, suggesting that they are sensitive to the reporting threshold; we confirm this in Appendix D. Although these differences are relatively small, indicating good overall consistency, a slight change in a tail exponent can result in a variation of economically meaningful magnitude.

5. Discussion

Table 9 summarises our results. It provides a quantitative summary of the main properties we examined and proceeds with a qualitative assessment of agreement across complete datasets and the nature of the bias in incomplete datasets.

Table 9
Summary of results.

Property	Results	Consistency in complete datasets	Bias in incomplete datasets	Section
Binary network				
Mean degree	$\approx [30, 50]$	High	Downward	3.1
Mean degree ~ #firms	Elasticity $\approx 1/3$	Medium		3.1
Out-degree distribution	Tail exponent $\approx [1.4, 1.6] < 2$	Very high	None	3.2.1
In-degree distribution	Tail exponent $\approx [2, 3] > 2$	Very high	Downward	3.2.1
In-degree ~ out-degree	TLS slope $\approx [0.6, 0.8] < 1$	High		3.2.2
Degree assortativity (all)	$\approx -[0.2, 0.01] < 0$	High	Towards zero	3.2.2
Reciprocity	Much higher than random, $\approx [3, 5]\%$	High		3.3.1
Global clustering	Low, comparable to CM	Moderate		3.3.2
Average local clustering	$\approx [19, 28]\%$, higher than CM	High	Downward	3.3.2
Average path length	slightly below 3	Very high	Upward	3.3.3
Weighted network				
Weights	Tail exponent $\approx [1.1, 1.2]$ <2	Very high	None	4.1
Strengths	Tail exponent close to 1	Very high	None	4.2
Degree-Strength correlations	Higher for in- than out-relations	High		4.2
Out-strength ~ out-degree	OLS slope $\approx$ [0.77, 1.05]	High	Upward	4.2
Out-degree ~ out-strength	OLS slope $\approx$ [0.31, 0.36]	High	Upward	4.2
In-strength ~ in-degree	OLS slope $\approx$ [1.35, 1.55]	High		4.2
In-degree ~ in-strength	OLS slope $\approx$ [0.41, 0.44]	High	Upward	4.2
In-strength ~ out-strength	TLS slope ≈ 1	High		4.2
Influence vector ($\alpha = 0.5$)	Tail exponent $\approx [1.2, 1.3] < 2$	High	Upward	4.4

Notes: See the relevant section for the definition of the properties and evidence for the reported results. The edge direction is from a supplier to a customer, so the in-degree is the number of suppliers and the out-degree is the number of customers. The "Results" column presents intervals summarising the quantitative values obtained from our complete datasets. These intervals are intentionally rounded to indicate that they reflect our judgment based on the precise evidence presented in the main text and appendix, rather than being derived from a formal procedure for aggregating the data, or constructing confidence or credible intervals. The column "Consistency" provides our qualitative evaluation of the extent to which the reported results are similar across complete administrative datasets (Ecuador 2015, Hungary 2021 and, where available, Belgium and/or the Dominican Republic). The column "Bias" presents our qualitative assessment of whether the incomplete datasets exhibit a clear and systematic deviation from the complete datasets, based on the tables in the main text and the appendices, particularly Appendix D. If we think the evidence is lacking, we leave the "Bias" column blank. Properties marked $y \sim x$ refer to OLS or TLS estimates for the log-transformed values. All the results are very persistent over time. CM stands for configuration model and TLS for Total Least Squares.

We base our assessment on the tables reported in the main text, as well as the tables and results in the Appendices, particularly those covering other years and alternative estimators. We consider the complete datasets to be Hungary 2021, Ecuador, the Dominican Republic and Belgium. Chile is excluded, as it has a low reporting threshold for transactions but a high reporting threshold for firm sizes. We also refer to Appendix D, which contains a quantitative analysis of the effect of the threshold on several key properties.

Overall, we find that for most properties, there is a strong, quantitative agreement between complete datasets: there are properties of production networks that we think are solid enough to be considered "really known".

Incomplete datasets sometimes feature a clear deviation from complete ones, allowing us to identify the direction of the bias introduced by incomplete reporting. In some cases, the results are obvious (e.g. a downward bias in the mean degree), while in others, they are highly non-trivial (e.g. a bias in the in-degree tail exponent but not in the out-degree tail exponent).

A typical paper in the literature proceeds by analysing a single network, identifying a new key fact and then developing a sophisticated analytical theory to explain it. Our results here provide good news for this agenda, but also a warning. The good news is that results developed in one context indeed have strong potential for external validity, given the similarity between complete datasets. However, our results also provide a warning that incomplete datasets can lead to biases in key properties. The behaviour of quantitative models may or may not depend very sensitively on estimated statistics. Table 9 provides a go-to benchmark for researchers in the field to, first, compare the properties of the network they analyse and, second, if their network is biased, to develop a sense for which moments are likely to be robust enough to base their model on.

A potentially important limitation of this study is that while all the datasets we discuss have a great variety in terms of geographical location (Latin America, Western and Eastern Europe, Africa and the Middle East) and development levels, they do tend to have similar sizes, with a number of nodes spanning roughly one order of magnitude only (30,000–500,000). However, we generally expect network size to influence certain aspects of network structure. For instance, as shown in Table 3, the mean degree weakly increases with the number of nodes, $\bar{k} \sim N^{1/3}$. Thus, while we believe that most of the properties, like tail exponents, would hold in larger countries, other properties, like mean degree, would likely differ. Although there are currently no complete VAT datasets for large countries, a careful study of scaling effects in countries for which we do have complete data would be a very relevant avenue for further research.

6. Conclusions

There is a broad consensus that modern macroeconomics should be bottom-up, data-rich and account for interactions. However, this agenda is hindered by the fact that we know very little about firm-level production networks, raising concerns that observed differences across datasets may reflect variations in data collection methods rather than genuine cross-country differences.

In this paper, we have made the first systematic attempt to summarise what is known about the structure of production networks. What data are available? Are there generic properties of firm-level production networks that hold across different countries and over time? Do discrepancies between datasets arise from the data collection and cleaning methods?

As expected, some properties of production networks hold across all datasets, at least qualitatively; for example, sparsity, heavy-tailed degree and strength distributions, high local clustering and small average path length. However, our analysis shows that we can be considerably more precise, thanks to the fact that many quantities are remarkably similar across complete datasets. Calibrating models using incomplete datasets can lead to targeting clearly biased moments.

In addition to our systematic attempt to compare datasets, we have also established or confirmed a few facts of economic significance. For example, the distribution of the number of customers exhibits a much heavier tail than that of the number of suppliers. We have also found that many large firms (in terms of network sales or expenses) have very few customers or suppliers. This suggests the existence of very large firms with limited downstream and upstream diversification, with potentially important implications for systemic risk.

There are several limitations to this work, which we regard as just one key step in an important research agenda. A first line of research will need to delve more deeply into data collection methods and assess the comparability of firm-level datasets with classical national accounts objects. A second line of research should look at more sophisticated properties, perhaps driven by theoretical research. For example, we have refrained from calculating quantities that rely on geographical locations, even though this is a clear avenue for applications.

To conclude, this paper serves as a reference point for those interested in datasets of firm-level production networks, with two main objectives. First, it makes key moments and statistics publicly available, which should be useful for guiding theoretical research and for researchers who lack access to administrative data but require key moments to calibrate their macroeconomic models or to create synthetic datasets. Second, it contributes to an emerging agenda of developing standards for data collection, cleaning and matching for micro-level production network data around the world.

Declaration of interests

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Bacilieri reports financial support was provided by Intelligence Advanced Research Projects Activity. Lafond reports was provided by Intelligence Advanced Research Projects Activity. Lafond reports financial support was provided by Baillie Gifford. Bacilieri reports was provided by Baillie Gifford. Francois Lafond reports a relationship with Macrocosm Inc that includes: consulting or advisory and equity or stocks. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Data

In this appendix, we describe our three datasets. For Hungary, we also provide an illustrative comparison with national accounts quantities. Before this, we briefly review a number of reasons why we would not expect aggregate firm-level networks to match national accounts.

A.1. Differences between firm-level networks and national I-O tables

We provide a non-technical discussion of the differences between firm-level data and the Supply-Use (SU) or the I-O tables framework from national accounts. For a detailed handbook on the compilation of SU and I-O tables, see United Nations (2018, especially p. 43–44, 105–107 and 163) and Eurostat (2008). We omit a discussion of missing firms and missing transactions.

Investment vs intermediate consumption. I-O tables are central to national accounts as they make it possible to compute GDP (i.e. total value added) by subtracting intermediate consumption from gross output (i.e. total sales). In national accounts, investment spending (i.e. gross fixed capital formation) is classified as part of final demand, thus separating it from intermediate consumption, whose transactions are recorded in the inter-industry transaction matrix. In contrast, firm-to-firm transaction networks include both intermediate and investment transactions without distinction.

In practice, the differential treatment of intermediate and investment transactions between the two data types can introduce a substantial bias, as the aggregate value of the transactions observed in the firm network should, in principle, exceed that of the inter-industry transaction matrix. If investment is 25% of GDP[12] and intermediate transactions are about as large as GDP, then the network transactions should be 25% higher than in the I-O table. This bias should be highly heterogeneous across industries: Vom Lehn and Winberry (2022) reconstruct the investment network for 37 industries in the US, showing that a few industries (i.e. construction, machinery, professional and technical services, and motor vehicles) represent a very high share of investment goods.

Wholesale trade, retail trade and transport. In national accounts, the convention is that wholesale, retail and transport should be better thought of as "pass-through sectors" rather than producing and consuming in a similar way as the other sectors do. More precisely, national accounts consider that the output of these industries is not their total sales but the *margins* they apply over the goods they buy and sell or transport.

When industry j buys goods produced by industry i through a wholesaler k, I-O tables record the flow of goods directly between industries i and j. The total cost paid by industry j is then split into the sales proceeds for industry i and a "trade margin" received by k. Another way to think about this is to consider that the wholesaler is a service provider – its true inputs are, say, labour, electricity and real estate, but not the goods that it buys only for reselling.

In sharp contrast, in firm-to-firm transaction data, we would observe the wholesaler buying the goods and reselling them, and we would not observe a transaction between industry i and j. Therefore, total sales of wholesale, retail and transport should be much higher in firm-level data than in the I-O tables (roughly 5 times higher if margins are 20%). Furthermore, we also expect the structure of the inter-industry matrix to be substantially different.

This distinction between I-O tables and firm-level networks can have important implications for the analysis of shock propagation. For example, if a supermarket were hit by a natural disaster and forced to shut down, a firm-level model could simulate the complete cessation of that firm's buying and selling activities, thereby also capturing the resulting cascading effects throughout the network. In contrast, in I-O tables, its buyers and suppliers would still have a direct link (even if at an aggregate level), so a naive model would allow trade to happen and therefore fail to estimate the full impact. In other cases, one may prefer a dataset that treats margin industries as in the national accounts; for instance, disruptions to specific suppliers of a wholesaler affect its customers in a way that is mapped correctly by national accounts and would be missed in firm-level data.

Financial institutions and financial services. The measurement of financial services in national accounts is complex. Additionally, financial institutions usually obey specific accounting rules and regulations. As a result, it is customary to remove financial firms when analysing firm-level datasets. Despite the decision to retain all the firms in our raw data, we expect that datasets of firm-level networks may or may not include financial firms. If financial firms are included, we expect quite different flows in and out of these firms, depending on the type of data source used (e.g. public balance sheets, VAT records, surveys or payment systems). Reconciling this with national accounts should proceed on a case-by-case basis.

Unit of analysis and industrial classification. The underlying principle is that of having a unit of observation that performs the same kind of production (known as the local kind-of-activity unit, Eurostat, 2008), which means that a firm or one of its establishments is divided into smaller units based on the principle of homogeneous production. Establishments tend to be the preferred statistical unit as they have a more homogeneous production and data collection is still practical.

However, the unit in micro-level network data is typically a firm, the sector of which is determined by the firm's primary activity, which can cause significant issues as multi-establishment and multi-product firms are common. Having multi-product firms implies that to aggregate firm-level data into proper I-O tables, one needs to split a firm's output into various industries and then decide how to break down the firm's inputs into the different outputs, a well-known problem for constructing symmetric I-O tables from SU tables (Miller and Blair, 2009). Another problem is that firm-level datasets may use a different classification system.

For datasets based on transaction-level records, such as those that can be obtained from banks or payment providers, another issue arises when firms hold multiple accounts. Consider, for example, a large multi-product firm operating in several regions and that buys legal services from a large legal services firm with offices across the country. It is possible that the customer firm centralises its payments, so we would see a transaction from one headquarter to another, rather than multiple firm-to-firm links.

Prices and volume measures. National I-O tables use different concepts of prices, depending on whether the price includes VAT, and trade and transport margins. In firm-level data, we think it is more likely that we observe amounts actually paid, which would probably include trade and transport margins, and which may be quoted with or without VAT in VAT datasets.

[12] This is roughly the ratio at the world level according to World Bank data (available at https://data.worldbank.org/indicator/NE.GDI.TOTL.ZS).

Timing of transactions. In principle, national accounts, firms' financial accounts and VAT payment datasets are compiled on an accrual basis, meaning that they record "flows at the time economic value is created, transformed, exchanged, transferred or extinguished. This means that flows that imply a change of ownership are entered when the change occurs, services are recorded when provided, output at the time products are created and intermediate consumption when materials and supplies are being used" (Nations, 2010, p. 55, par 3.166). This can create substantial inconsistencies with firm-to-firm transaction datasets compiled from direct money flows, such as payment systems, due to the prevalence of trade credit.

International trade. Multinational firms typically file their accounts (and taxes) in various countries so that national accounts can ultimately try to separate the contributions of foreign firms domestically and domestic firms abroad. When using firm-level data, the ability to reconstruct tables close to national accounts would depend on the ability to access detailed financial accounts. Again, the specifics of the data collection method would matter.

Taxes and government sector. In most countries, the public sector accounts for a large share of GDP. National accounts can represent this fairly accurately by classifying government activities according to their purpose. Thus, the SU tables for industries such as health or education would typically show aggregates of public and private activities. In general, we would expect it to be difficult to reconcile firm-level datasets and national accounts for non-market activities.

Informal sector. In most countries, national accounts estimate the value of the informal economy, which is unlikely to have a counterpart in tax-based administrative data. Wiedemann et al. (2024) offer a method to augment the VAT dataset from Kenya with reconstructed links for the informal economy, using non-VAT data sources that do include the informal economy.

All considered, reconstructing or reconciling national accounts with firm-level datasets is a serious challenge, which we do not attempt here. Buda et al. (2022) provide a proof of concept that this can be done for consumption using payment data, but we are not aware of any study that has done this for network data, which is more difficult. Having acknowledged these issues, we proceed to describe the datasets we use and provide a comparison of our Hungarian dataset with the relevant national accounts.

A.2. Datasets

Throughout the paper, we keep firms in the largest connected component. A network is *connected* if there is at least one path between all pairs of firms. The network is directed, so we keep firms in the largest *weakly* connected component (LWCC). This data truncation is very small on our VAT networks, but not insignificant for FactSet.

We now describe each of our 3 datasets in detail. For Hungary, we will provide a comparison with national accounts. We choose Hungary instead of Ecuador because we can use years with different thresholds.

A.2.1. FactSet

We use two data sources provided by FactSet: Supply Chain Relationships and Supply Chain Shipping Transactions, both downloaded on 23 April 2024. The Supply Chain Relationships data come partly from companies' filings required by US Federal Accounting Standards,[13] and partly from information on supply chain relationships released in investor presentations, company websites and press releases. The second source (Supply Chain Shipping Transactions) records shipment declarations at ports from the US Bill of Lading. FactSet collects this information from the US Customs and Border Protection.

The supply chain dataset starts in 2003, while the shipment dataset in 2013. We do not use years prior to 2014 because FactSet changed its data collection methodology in 2013, improving the quality of the dataset. We use data up to 2023, but we only report results for the 2021 network in the main text because the number of firms and edges in the network decreases in 2022 and 2023 (see Table C.12) – this is due to an increase in entities with generic identifiers, which we remove, as discussed below.

Due to the nature of the data collection process, coverage is biased towards companies listed on US stock exchanges, large firms and large transactions.

The monetary values of customer-supplier transactions are rarely available. When they are, they are reported as a percentage of the seller's revenues. However, it is unclear to which reported revenue figure this percentage refers to, e.g. quarterly or annual income statement. Similarly, in the shipment dataset, the cumulative value of the goods shipped is not always reported and with valuation methods that are not necessarily consistent with balance sheet information. Therefore, we use only the binary topology.

In the raw dataset, links report the year, month, day and hour. The start and end dates correspond to when the record was first published and when the end was announced. We consider a relationship to exist in a given year if it exists at any time during that year. We remove all links involving unknown shippers and consignees, which are represented by generic entity identifiers (0JWCDG-E and 0JWCDH-E). We remove duplicate links and self-loops, and we keep only firms in the LWCC.

[13] The Statement of Financial Accounting Standards No. 131 requires publicly traded firms on US stock exchanges to report customers that account for 10% or more of their annual revenues, formally called *major customers*.

A.2.2. Ecuador

The Internal Revenues Service (IRS) of Ecuador collects data through VAT filings (Astudillo-Estevez, 2021). In principle, it includes information on every transaction between all the entities, aggregated per year. Firms need to report both their suppliers and customers.

The VAT records are divided into three categories: VAT-exempt transactions, 0% VAT rate transactions and taxable transactions. Exempt transactions, where firms neither collect nor report VAT, include in-kind contributions, business sales with asset-liability transfers, corporate mergers and donations to charitable entities. The 0% VAT rate applies to essential goods within the VAT system, such as unprocessed agricultural products, fish, milk, bread, medications, fertilisers and agricultural equipment. All other transactions are subject to standard VAT rates, including non-essential consumer goods like clothing, electronics and cosmetics. This categorisation does not affect the total value of the transaction reported in our dataset.

Sometimes, the value of the transaction reported by the customer and by the supplier may differ, and there may be a mismatch between reporting forms. The IRS takes care of cleaning possible mismatches. The IRS is particularly concerned to detect possible fraud in transactions with large firms, which are identified using the weighted degree centrality. In the first couple of years of the data collection, numbers were reported manually, so the latest years are more reliable.

We also have information on industry classification: ISIC Rev. 4 codes.

We use a version that excludes natural persons, foreign companies that are not registered in Ecuador, and where an implausibly large value for a transaction was replaced with its value in the previous year. We then remove edges associated with a transaction value of zero (there are no negative transactions to remove) and self-loops, which represent transfers among establishments of the same firm (these transactions are not taxed and are purely for accounting purposes within the firm). Finally, we keep only firms in the LWCC.

A.2.3. Hungary

Source and reporting threshold. Hungary's network is collected by the National Tax and Customs Administration of Hungary. Before the first half of 2020, firms had to report a supply chain relationship if the tax content of their cumulative trades or invoices exceeded a certain value during the reporting period. There are exceptions, however, depending on the sector the firm operates in or the type of goods traded. The length of the reporting period can be annual, quarterly or monthly, depending primarily on the amount of VAT declared by the firm in previous periods.

From 2015 until the second quarter of 2018, the threshold was HUF 1 million. It is calculated on the tax content of the sum of the transactions between two firms in a given reporting period. Given that the tax rate is 27% (although there are exceptions as noted above), the value of the transaction (excluding tax) above which reporting is required is HUF 3,703,703. During this period, firms had to report both directions, i.e. both their purchase and their sales transactions that were above the threshold. We use the network constructed from the information reported by the buyers, as they have a clear incentive to report (claiming back VAT) and the network appears much more complete.

From the third quarter of 2018 to the second quarter of 2020 there were three important changes. First, the threshold was lowered to 100,000 HUF. Second, it became interpreted at the invoice level regardless of the reporting frequency (only invoices with a tax content of more than 100,000 HUF had to be reported). Third, only the purchases had to be reported. Finally, there is no threshold since the second half of 2020, so all invoices must be reported.

We also have information on industry classification: NACE Rev. 2 codes.

The dataset covers the period 2014–2021. However, we dismiss the first year because the quality of the data is poor; this might be due to the inexperience of both the authorities and firms with the new reporting requirements. For an in-depth description of the Hungarian dataset, we refer to Borsos and Stancsics (2020).

Consistent with the other datasets, we exclude edges corresponding to transactions with zero or negative values, as well as self-loops. We retain only firms that belong to the LWCC.

Sectoral composition. Fig. A.8 shows the sectoral composition of the Hungarian economy according to the national I-O tables and our firm-level network aggregated at the sector level. We choose the year 2015 and 2020 because we have I-O tables for these years and the reporting thresholds are very different. The top and bottom panels have different x-axes ranges, allowing us to show separately large and small industries.

We compute industry size as intermediate sales plus GFCF in the I-O table and we use firms' network sales in the VAT data. Using the Corporate Registry of Hungary, we aggregate firms into industries using NACE (Rev. 2) codes at the 1-digit level, which match the economic sectors of the national I-O table. We then aggregate both to the 1-digit level. Some firms do not have to report their sector (e.g. freelancers and non-profit organisations); these account for 11% of total firms' network sales in 2015 and 16% in 2020.

The sectoral composition differs between the firm-level network and the national accounts. Manufacturing stands out, accounting for roughly half of total sales in the I-O tables, but only about a third in the VAT data. The industries with a difference greater than 80% (in absolute terms) are wholesale and retail; electricity and gas; public administration and defence; human health and social work activities; arts, entertainment and recreation. For "margin" industries (wholesale and retail trade) and industries with high public sector involvement (i.e. non-market activities, such as public administration, defence, education and health), this is not surprising in view of the discussion in Section A.1, but for other industries the reasons for these discrepancies are less clear.

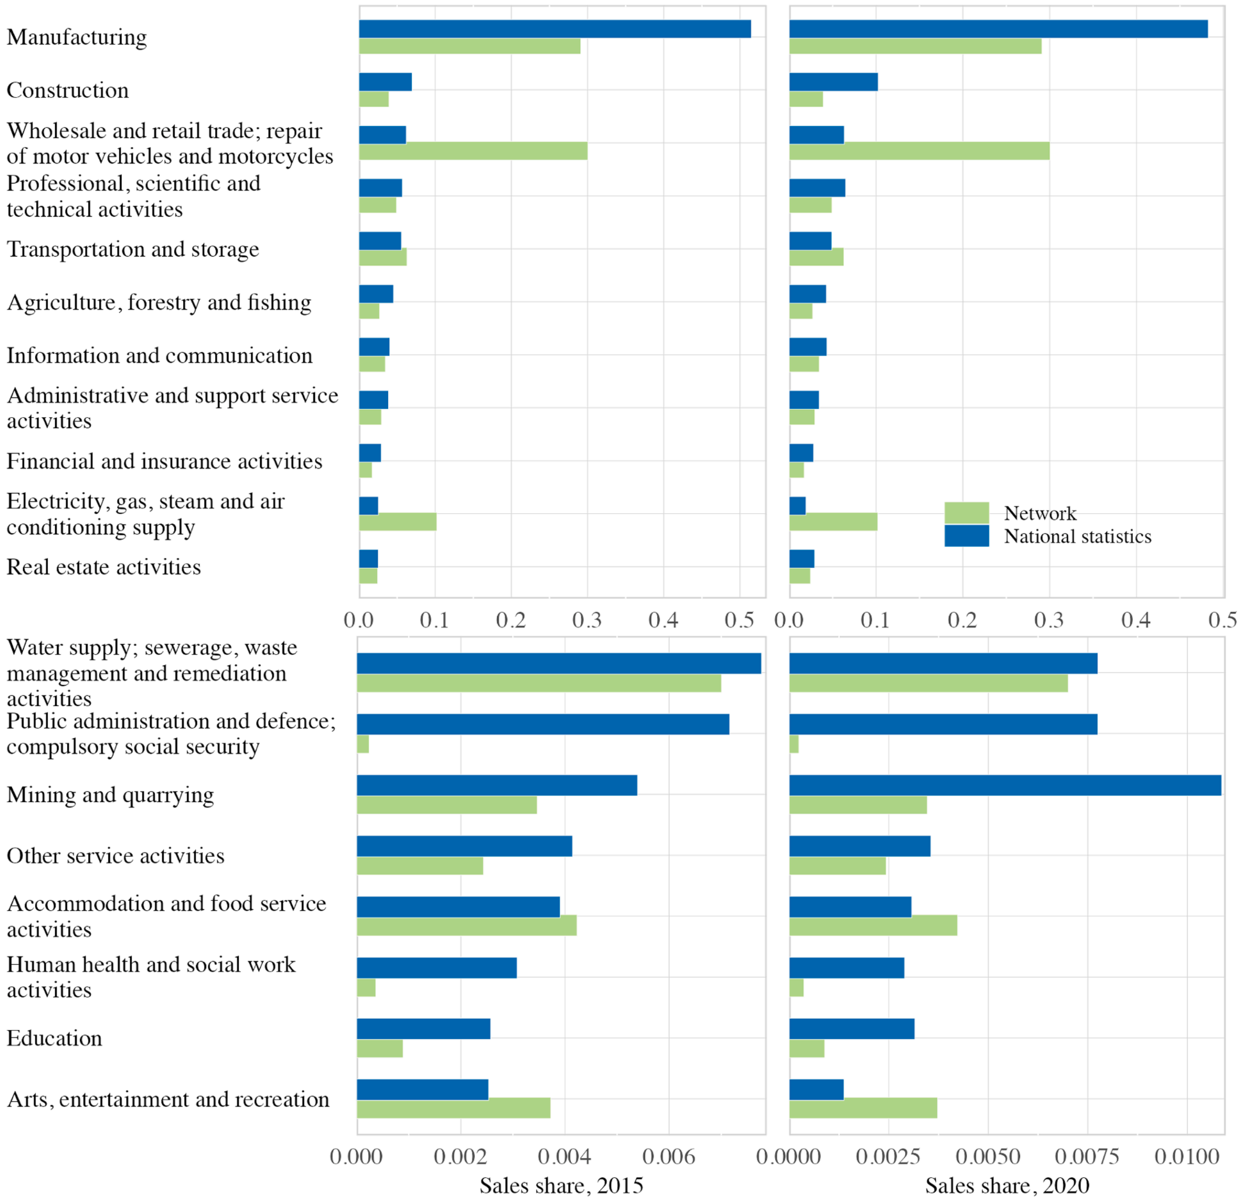

Fig. A.8. Sectoral composition in the national I-O table at the sector level (blue) and in our firm-level dataset (green) in 2015 (left) and 2020 (right). We compare the sales of firms in the network (aggregated at the sector level) with the sum of intermediate sales and GFCF in the national I-O table. We use NACE Rev. 2 codes at the 1-digit level. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Appendix B. Characterizing distributions

B.1. Estimating power-law exponents

Estimating power-law exponents is a difficult task (Nair et al., 2022). Real-world data almost never follow pure Pareto distributions, but instead exhibit power-law *tails*. When the power-law behaviour occurs only in the tail, using methods that assume a pure Pareto distribution are typically misleading; thus, identifying where the power-law behaviour "starts" is a critical step. In the remainder of this section, we discuss these issues in detail and justify our choice of estimators. For a comprehensive treatment, see Nair et al. (2022) and Voitalov et al. (2019).

Power laws as regularly varying distributions. Power laws are sometimes defined as distributions with a "Pareto tail", which, taken literally, would mean that beyond a threshold $k_{\min}$, the tail of the (complementary cumulative) distribution is exactly proportional to

$k^{-\gamma}$. In applications where the tail of the distribution is of interest, it is better to consider the larger class of *regularly varying distributions*, which have a complementary cumulative distribution function (CCDF) of the form $\bar{F}(k) = \ell(k)k^{-\gamma}$, where $\ell(k)$ is a slowly varying function, that is, it satisfies

$$\lim_{x\to\infty} \frac{\ell(tx)}{\ell(x)} = 1 \, .$$

The presence of the slowly varying function implies that the shape of the distribution can deviate noticeably from a pure power law for low and moderately high degrees (the body of the distribution). This is expected for any real-world phenomenon, where noise, heterogeneity and other factors are present. However, the distribution maintains the key feature we are interested in: the extreme tail behaviour. For example, in models of the granular origins of aggregate fluctuations (Gabaix, 2011; Acemoglu et al., 2012), what matters is that the variance of a given distribution diverges, which is the case for all regularly varying distributions with $\gamma < 2$.

How can we test whether a distribution is regularly varying? And if it is, how can we estimate γ? Extreme value theory (the Fisher–Tippett–Gnedenko theorem) tells us that the asymptotic distribution of the (suitably normalized) maximum of a sample of i.i.d. random variables, if it exists, is the Generalized Extreme Value Distribution (GEVD) with extreme value index ξ,

$$\Pr(V < v) = \exp\left(-(1+\xi v)^{-1/\xi} \right) .$$

There are three subfamilies, characterised by the value of ξ. For $\xi < 0$, the GEVD is the Weibull distribution; for $\xi = 0$, it is the Gumbel distribution; and for $\xi > 0$, it is the Fréchet distribution. It turns out that the maximum domain of attraction (MDA) of the Fréchet distribution is exactly the set of all regularly varying distributions. Thus, for any regularly varying distribution with tail index γ, the distribution of the suitably normalised sample maximum is the GEVD with tail index $\xi = 1/\gamma > 0$.

Estimating power-law exponents. Estimating tail exponents requires two choices: (1) a choice of the number of order statistics (a threshold above which the power-law behaviour begins), and (2) a choice of the estimator for the tail exponent.

The standard estimator in the literature is the method from Clauset et al. (2009), which uses maximum likelihood estimation (MLE, the Hill estimator) with the threshold chosen to minimise the Kolmogorov-Smirnov distance between the estimated and empirical CCDFs. We use R igraph's implementation, which finds the threshold such that the p-value of the Kolmogorov-Smirnov test between the fitted distribution and the original sample is the largest.

Voitalov et al. (2019) provide a review of the extremal estimators for the tail index and an implementation of the double bootstrap procedure for selecting the threshold. The double bootstrap method takes two estimators proven to be consistent, applies them to different sample sizes and selects the value of the threshold that brings the two estimators into closest agreement. These methods have been shown to be consistent; see Voitalov et al. (2019) and Nair et al. (2022) for further details and references.

For discrete distributions, such as the degree distributions, Voitalov et al. (2019) approximate the degrees to continuous reals by adding small symmetric noise sampled uniformly at random in the interval $[-0.5, 0.5]$. They show that adding such noise (a) does not substantially affect the estimated value of the tail index, provided that the distribution is regularly varying with Fréchet as the MDA; and (b) that the noise improves the convergence and stability of the estimators.

In the main text, we report the estimates from the method of Clauset et al. (2009) to make our results comparable with published results and because it has been shown to be fairly robust to finite-size effects, typically observed in network data.[14] We do not report the standard errors from Clauset et al.'s (2009) MLE because we do not think that pure Pareto tails are the correct benchmark.[15] However, we verify all our results using the state-of-the-art implementations of Voitalov et al. (2019).

Classifying distributions and hypothesis testing. With regularly varying distributions, it is not possible to do hypothesis testing (Voitalov et al., 2019). Regularly varying distributions do not form a parametric class of distributions. Instead, they are non-parametric with infinite degrees of freedom due to the unspecified slowly varying function $\ell(k)$. Importantly, there is an infinite number of regularly varying distributions for which a sampled sequence of finite length does not appear to be regularly varying. Likewise, there is an infinite number of distributions that are not regularly varying for which a sampled sequence of finite length may appear to be regularly varying. Therefore, the best strategy one can adopt is to consider all the available consistent estimators and check for agreement on the ranges of the estimated γ's.

As a result, we do not use a formal criteria for classifying distributions as regularly varying or not, but informally we are guided by the classification scheme adopted by Voitalov et al. (2019), where a distribution is *not a power law* if at least one of the extreme value estimators returns $\xi \le 0$; *hardly a power law* if $\xi > 0$ for all the extreme value estimators and at least one $\xi \le 1/4$ ($\gamma \ge 4$); *a power law* if for all the extreme value estimators $\xi > 1/4$ ($\gamma < 4$); and *a power law with divergent second moment* if for all the extreme value estimators $\xi \ge 1/2$ ($\gamma \le 2$). In practice, for all the distributions we fit, we, almost always, find $\gamma < 4$.

[14] In some data-generating processes, including certain canonical models of growing networks, the asymptotic distribution is a power law, but for finite sizes, the distribution is a power law with an exponential cutoff (which has finite moments). Serafino et al. (2021) show that the estimator of Clauset et al. (2009) performs well in retrieving the true power-law exponent even in finite-size networks. See also Figure 8, panels c, h and m in Voitalov et al. (2019), showing the same result but for finite-size i.i.d sequences drawn from a power law with an exponential cutoff, $k^{-\gamma}e^{-k/n^{\omega}}$, where n is the sample size (Voitalov et al., 2019, Equation D4). This distribution has finite moments for fixed n, but has pure power-law behaviour asymptotically. We find that, in practice, the estimator of Clauset et al. (2009) performs relatively well in estimating γ for reasonable sample sizes ($10^3 - 10^5$), particularly for small exponents.

[15] Furthermore, the MLE standard errors are based on the assumption that the data are i.i.d., which is unlikely to be a good assumption for network data. In practice, the MLE standard errors for the distributions we study are very small and we think deceptively so.

The 1/4 threshold is chosen because if ξ is positive but very small, it is not possible to test whether $\xi = 0$. If $\xi = 0$, then the distribution is in the Gumbel MDA, which includes both light-tailed distributions and heavy-tailed distributions that are not regularly varying. The value 1/4 is somewhat subjective and we may have wanted to make it depend on the number of observations. For instance, Dorogovtsev and Mendes (2003, Figure 3.32) provide a heuristic argument: to estimate a power law with a reasonable degree of precision, one needs data that span at least 2 or 3 orders of magnitude and, given a sample size, the range of the data needed is heavily affected by γ.[16] For example, power laws with $\gamma > 4$ would require a tremendous amount of data to span enough orders of magnitude to be adequately measured. This phenomenon is apparent in our data, where the variation across years and estimators is much higher for in-degrees (which have $\gamma \approx 2.5$) than for out-degrees (which have $\gamma \approx 1.5$; compare Tables C.14 and C.15).

In reporting our detailed estimates (Tables C.14, C.15, C.17, C.18, C.16, and C.19), we show the estimated values of γ and the number of data points used to estimate the tail exponent, i.e. the number of order statistics, which is determined by the double-bootstrap method for the estimators based on GEVD and by minimising the K-S distance for `plfit`. The number of order statistics is interesting because it shows that some estimators use much more data than others, providing an additional robustness check.

Lognormals may seem power laws. The lognormal distribution has a finite second moment and is among the heavy-tailed distributions that are in the Gumbel MDA, so it is not regularly varying. However, when the lognormal distribution has a high variance, it can be easily mistaken for a power law. This can be seen from the probability density function:

$$\log p(x) = -\log x - \log(\sigma\sqrt{2\pi}) - \frac{(\log x - \mu)^2}{2\sigma^2} ,$$

where as $\sigma \to \infty$, the quadratic term tends to zero. Therefore, when the variance of the lognormal is very high, the distribution can look very similar to a power law.

Sornette (2006) shows that the lognormal can be rewritten as

$$p(x) = (x_0\sqrt{2\pi\sigma^2})^{-1}(x/x_0)^{-1-m(x)} ,$$

where $x_0 = \exp(\mu)$ and $m(x) = \log(x/x_0)/(2\sigma^2)$ is a slowly varying function of x. When σ^2 is large enough, there is a large range of values x such that $m(x)$ is very small and the lognormal looks like a power law in this region.

In our case, we have found by examining qq-plots that lognormal fits are good for the distribution of strengths but not for other quantities.

B.2. Characterizing joint distributions

In this appendix, we collect a number of technical details and empirical results related to the characterisation of the joint distributions.

B.2.1. Total least squares

In many instances, we are interested in characterising the direction of the relationship between two variables. For instance, we expect that firms with large sales also have large expenses, so we can hypothesise the deterministic relationship $s_{\text{in}} \propto s_{\text{out}}^{\theta}$, perhaps with $\theta \approx 1$. Regressing (log) sales on (log) expenses only characterises the slope of the *conditional* relationship. Therefore, the estimate of θ will differ if we regress sales on expenses or the other way around, unless they are perfectly correlated.[17]

To characterise the “slope”, we use the Total Least Squares (TLS) estimator, which is well-known as a specific errors-in-variables estimator (see also demeaning regressions). Essentially, it minimises the squared *perpendicular* distances (rather than horizontal or vertical) between each point and the regression line, exactly as in principal component analysis. In fact, in our bivariate case, the TLS slope is equal to the ratio of the first two entries of the leading eigenvector of the covariance matrix. In practice, we demean the data, estimate the TLS slope $\hat{b}$ and find the intercept as $\hat{a} = \bar{y} - \hat{b}\bar{x}$, where $\bar{y}$ and $\bar{x}$ are sample averages.

B.2.2. Covariances

Here, we report the covariance matrices for the strengths and degrees of Ecuador (2015) and Hungary (2021). Since we are interested in log-transformed values, we need to drop the zeros. When a node has an in-strength of zero, it always has an in-degree of zero; similarly, for out-strength and out-degree. However, it is possible for a node to have suppliers but not to have any customers, or the other way around (Table C.13). As a result, we need to report two covariance matrices.

Table B.10 shows the variances and covariances computed by removing only the nodes that have a value of zero for a specific metric (mean and variance) or pair of metrics (covariances). Instead, Table B.11 reports the covariance matrix computed after removing all the nodes with at least one zero value. These tables allow the reader to compute results that we do not report explicitly in the main text. We provide 3 examples.

[16] In the case of pure Pareto tails, the formula for the standard errors is $\sigma \approx \gamma/\sqrt{N}$ (Clauset et al., 2009, Equation 3.2), which makes it clear that for a given sample size, fatter tails are more precisely estimated.

[17] If β is the coefficient of the OLS regression of y on x, $\beta = \text{Cov}(y,x)/\sigma_x^2$, and $\tilde{\beta}$ is the coefficient of the regression of x on y, we have $\beta = \rho^2(1/\tilde{\beta})$, where ρ^2 is the squared correlation coefficient; that is, the R^2 from both regressions. Thus $\beta = (1/\tilde{\beta})$ iif $\rho = 1$ or -1.

Table B.10
Covariance matrix keeping only nodes with pairwise positive values.

	Ecuador (2015)				Hungary (2021)			
	k^{out}	k^{in}	s^{out}	s^{in}	k^{out}	k^{in}	s^{out}	s^{in}
k^{out}	2.83	1.37	2.49	2.39	2.73	1.14	2.85	1.92
k^{in}	1.37	2.36	1.98	3.66	1.14	2.00	1.20	2.71
s^{out}	2.49	1.98	8.17	4.34	2.85	1.20	7.99	2.98
s^{in}	2.39	3.66	4.34	8.99	1.92	2.71	2.98	6.10
Mean	1.85	2.83	10.62	9.91	1.67	3.05	8.61	9.25

Notes: All variables are log-transformed. The row 'Mean' shows the average of the log-transform of the positive values.

Table B.11
Covariance matrix keeping only nodes with positive values for all four metrics.

	Ecuador (2015)				Hungary (2021)			
	k^{out}	k^{in}	s^{out}	s^{in}	k^{out}	k^{in}	s^{out}	s^{in}
k^{out}	2.89	1.37	2.12	2.39	2.75	1.14	2.19	1.92
k^{in}	1.37	2.06	1.98	3.07	1.14	1.65	1.20	2.19
s^{out}	2.12	1.98	6.97	4.34	2.19	1.20	5.90	2.98
s^{in}	2.39	3.07	4.34	7.63	1.92	2.19	2.98	5.22
Mean	2.07	3.19	11.08	10.53	2.04	3.30	9.38	9.64

Notes: All variables are log-transformed.

Example 1: Total least squares. In Fig. C.13, we report the TLS estimate of the relationship between in- and out-strengths as 1.08. This is the ratio of the first two entries of the leading eigenvector of the covariance matrix. For Ecuador, the matrix (using Table B.11, because Table B.10 contains variances excluding only observations equal to zero in each variable separately) $\begin{pmatrix} 6.97 & 4.34 \\ 4.34 & 7.63 \end{pmatrix}$ has leading eigenvector $(0.679, 0.733)$, yielding a TLS estimate of $0.733/0.679 = 1.08$, as reported in Fig. C.13.

Example 2: Least squares. In Table 8, we report regressions of (log) strength on (log) degree. For instance, for the regression of out-strength on out-degree for Hungary in 2021, using Table B.10, the coefficient is $\hat{\beta} = \frac{\text{Cov}(\ln s^{\text{out}}, \ln k^{\text{out}})}{\text{Var}(\ln k^{\text{out}})} = 2.85/2.73 = 1.044$, as reported (up to rounding errors).

Example 3: Large variances. In Section 4.2, we report that the strength distributions can also be well fitted by a lognormal distribution. It is well known that, when the lognormal scale parameter is large, lognormal distributions are very difficult to distinguish from distributions with regularly varying tails (Section B.1). Sornette (2006, p. 95), for instance, uses a value up to $\sigma = 3$ (holding the lognormal location parameter, i.e. the median, at 1) to make this point. We can calculate that fitting a lognormal to the in-strength distribution of Ecuador (Table B.10) leads to $\hat{\sigma} = \sqrt{8.99} \approx 3$.

B.2.3. Heteroskedasticity and non-linearities

To investigate whether the conditional relations feature non-linearities and/or heteroskedasticity, we use binned scatter plots. To gauge non-linearities, we employ GLM, and to gauge heteroskedasticity, quantile regressions at the 10th and 90th levels, both as implemented by Cattaneo et al. (2019) and using 100 bins. Fig. B.9 shows the results for Ecuador (2015) and Hungary (2021), considering the two conditional relationships within each possible combination pair of in- and/or out-, and strength and/or degree joint relationships.

Overall, linearity appears to hold fairly well over large ranges. However, it is interesting that the deviations from linearity (computed by simple OLS) are almost systematically the same both in Ecuador and in Hungary.

Regarding homoskedasticity, the variance in the in-strengths is noticeably smaller when conditioning on intermediate values of the number of partners (say, around 100; top row); this effect is less pronounced for the out-strengths. For the degree-strength relationships (second row), there is a clear trend of increasing variance in the number of customers as we condition on higher and higher sales, particularly for the out-degrees. In other words, firms with very high sales may have many customers, and on average they do, but it is not uncommon to find very large firms having just a few customers. This fact does not appear to have been noticed in the literature. It could also be due to the fact that we only observe intermediate domestic customers.

B.2.4. Mean degree and size

In this section, we document two facts in more detail concerning degrees and the relationship between degrees and firm size. Fig. B.10 shows the degree distributions for Ecuador and Hungary, reproducing some of the information in Fig. 3 but displaying both the in- and out-degrees on the same chart. First, Fig. B.10 shows that Ecuador and Hungary (when complete) have extremely similar degree distributions, with the Ecuadorian distributions tailing off slightly earlier because of the substantially smaller number of firms. Second, there exists a degree value, k^*, at which the fraction of nodes with degree greater than k^* is the same for both the in- and

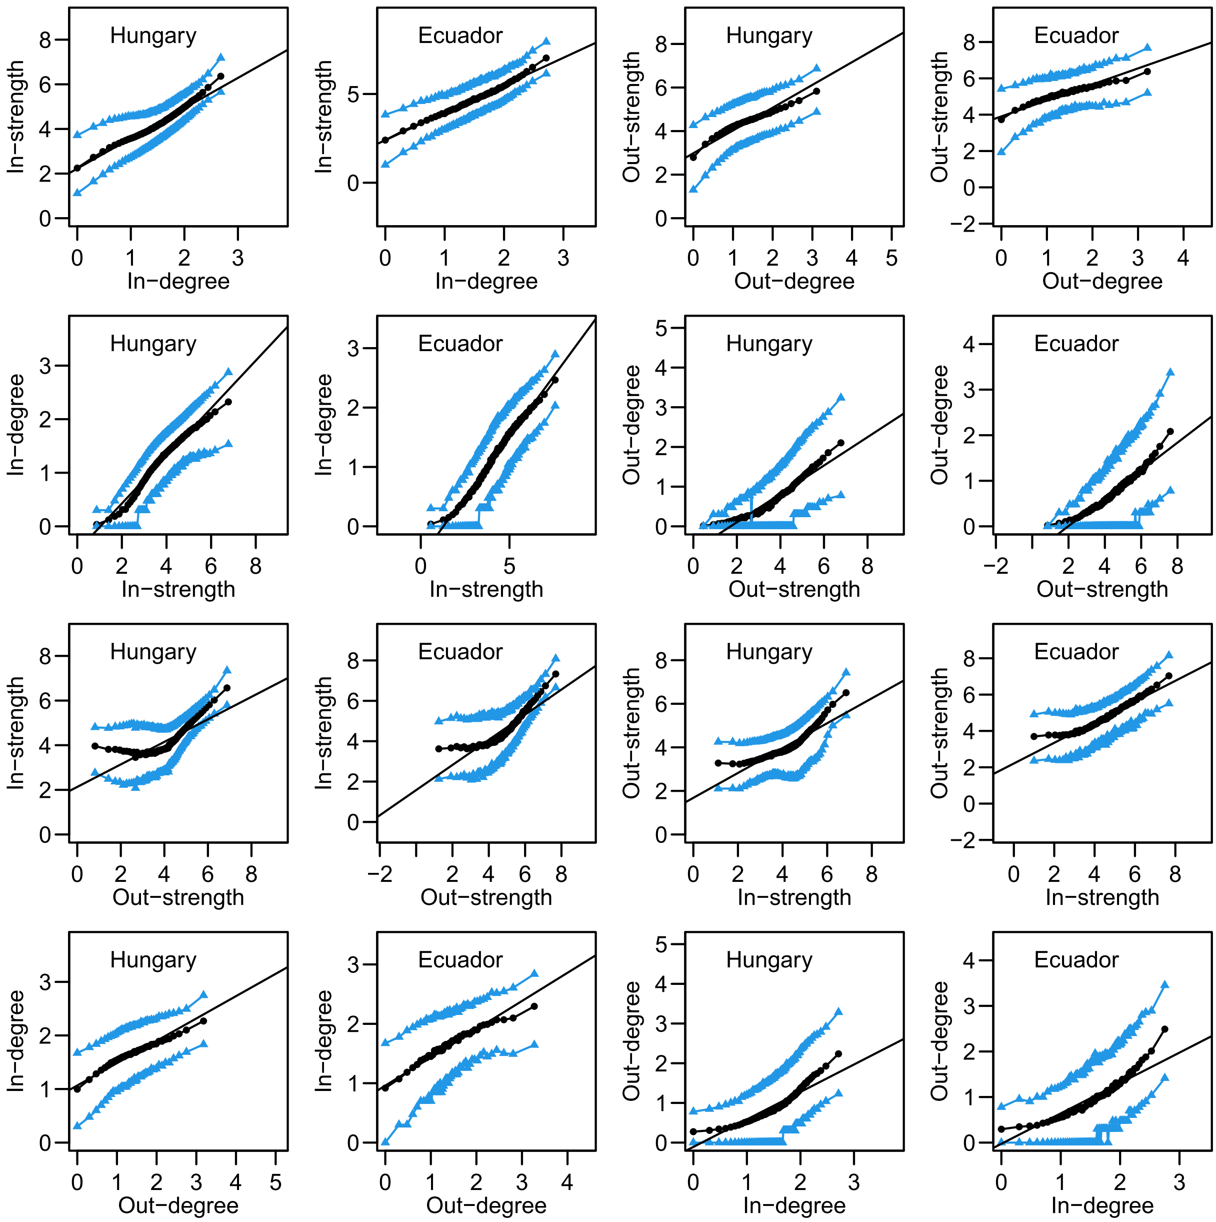

Fig. B.9. Binned scatter plots of the conditional relationships among strengths and degrees for Ecuador (2015) and Hungary (2021). The black dots show the non-parametric GLM estimates and the blue dots show the non-parametric 10th and 90th quantile regressions, using the implementation of Cattaneo et al. (2019) and 100 bins. The black line is the simple OLS regression. The ranges of the x- and y-axis are determined by the range of the data to highlight that the last bins in the tails reflect a large range of data. All axes are in log base 10. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

out-degree. Above k^*, the fraction of firms with out-degree greater than k exceeds the fraction of firms with in-degree greater than k, whereas below k^*, the reverse holds.

Fig. B.10 suggests that large firms, characterised by high in- and out-degrees and high in- and out-strengths, may have more customers than suppliers, which in turn suggests that small firms may instead have more suppliers than customers. However, this pattern is not directly evident in the figure, as it depends in complex ways on the joint distribution of these four variables. To examine this, we define firm size as the average of the in- and out-strength and, for firms of similar sizes, we compute three statistics of the distributions of in- and out-degrees: mean, median and 90^{th} percentile. Fig. B.11 shows the results.

For both countries, the median out-degree is smaller than the median in-degree for a large range spanning five orders of magnitude before the tail data become noisy. By contrast, the mean in- and out-degree exhibit a crossover. For Ecuador, it is unique and occurs at

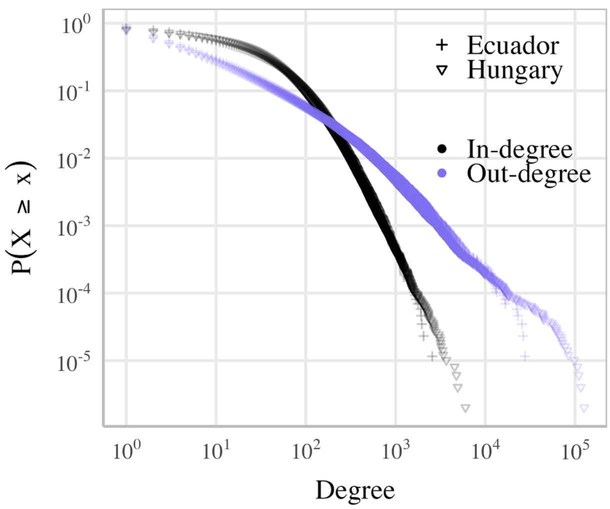

Fig. B.10. Empirical CCDF of the number of suppliers (in-degree) and the number of customers (out-degrees) for Ecuador in 2015 (cross) and Hungary in 2021 (triangle). In-degree is shown in black and out-degree in purple. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

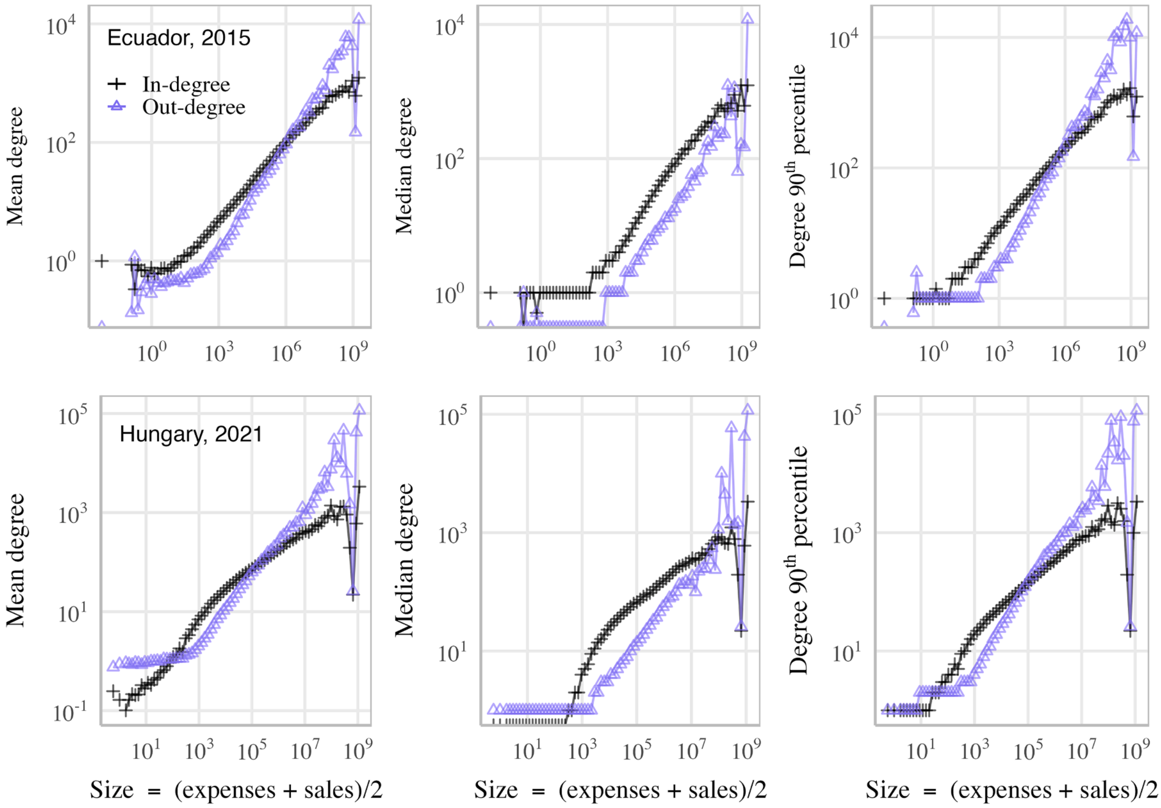

Fig. B.11. Binned scatter plots of the conditional relationships among size and in- and out-degrees for Ecuador in 2015 (top row) and Hungary in 2021 (bottom row). In-degrees (number of suppliers) are marked with a black cross and out-degrees (number of customers) with a purple triangle. We divide the x-axis into 80 log-spaced bins and compute the metric of interest in each bin. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

a relatively large size, around \$1-2 million and a degree greater than 100. For Hungary, the broad picture remains, with a crossover point around HUF 350 million, which is in the ballpark of \$1 million (HUF 1000 ≈ \$3) and, again, at a degree slightly above 100. We also find another crossover point for very small firms, below approximately HUF 100,000. Firms in that region have a mean in- or out-degree of 1 or less, indicating that numerous firms have no customers or suppliers.

To explain the different behaviour between the mean and median in- and out-degrees, in the right-most panels we show the 90th percentile, which behaves like the mean. In other words, the mean is driven by the upper tail of the distribution, with increasingly extreme values of the out-degree for larger and larger firms. Very large firms with substantially more customers drive these statistics. By contrast, for in-degrees, higher percentiles do not diverge from the median as size increases, as shown already in the second row of Fig. B.9, where heteroskedasticity appears clearly for out-degrees but not, or less so, for in-degrees.

Overall, and with the caveat regarding the very small firms in Hungary, the results are again remarkably consistent and noteworthy. To our knowledge, the facts shown in Fig. B.11 have remained undocumented and thus unexplained in the literature.

Appendix C. Additional results

C.1. Network size

Table C.12 shows the number of firms (N), supply-chain relationships (E) and the average degree ($\bar{k}$) for all the firm-level networks we analyse in this paper and in the literature we reviewed. Below, we report the details of how we construct the table. The asterisks indicate that the value has been inferred by us rather than explicitly reported in the cited paper.

For the Belgian network, the data from Dhyne et al. (2015) were extracted from Section 2 and Section 4, reporting, respectively, an approximate number of nodes and an exact number of edges. Always for Belgium, Kikkawa et al. (2019) report the number of nodes and edges in Table 1 and Table 2 of Section 2.1, Magerman et al. (2016) in Section 3.3, Dhyne and Duprez (2015) in Section 1.2, Dhyne et al. (2016) in Section 1.1, and Bernard et al. (2022) in Table 5.

For Brazil, Silva et al. (2020) report an approximate number of nodes and edges in Section 3.

For Chile, Huneeus (2020) report the number of edges in Table F.15, Appendix F, which is in millions and with two decimals. Miranda-Pinto et al. (2023) report the mean total degree (in- plus out-degree), calculated using the directed network, in Table 1; hence, the undirected average degree, shown in Table C.12, is obtained by dividing the mean total degree by 2.

For Costa Rica, we infer the number of edges. Alfaro-Urena et al. (2018) report the number of firms in Table 8. In Table 6, they record the number of firms with at least one supplier (the set of customer firms) and the customers' average degree. We obtain the number of edges by multiplying the number of customers by the customers' average degree ($N_c \cdot \bar{k}_c$), where the average degree is adjusted for the total number of firms in the network. This adjustment is obtained by dividing the number of customers by the total number of firms and then multiplying by the customers' average degree ($\frac{N_c}{N} \cdot \bar{k}_c$).

In Table 4, Alfaro-Urena et al. (2018) also report the number of suppliers and the suppliers' average degree. Inferring the total number of edges and the average degree using the supplier-side information yields slightly different results than using the customer-side information reported above. In Table C.12, we report the statistics calculated using the customer-side information (Table 6 in Alfaro-Urena et al., 2018).

For the Dominican Republic, Cardoza et al. (2025) report the annual average number of nodes and edges for the period 2012–2019. We attribute those averages to 2019.

For Estonian, de la Torre et al. (2016) report the number of nodes and edges in Table 3. The data for Criscuolo et al. (2024) are taken from Figure 1; we thank the authors for providing the underlying statistics.

For globally listed firms, Wu (2016a) reports the number of nodes and edges in the cleaned and raw network in Section 2.2.

For the global automotive network, the number of nodes and edges are taken from Brintrup et al. (2015), Section 3A.

For Japan, Bernard et al. (2019) report the number of nodes and edges in Section 3, Ohnishi et al. (2010) in Section 2, Fujiwara and Aoyama (2010) in Section 3B, Mizuno et al. (2015) in Footnote 4, Inoue (2016) in Section 2 and Lu et al. (2017) in Section 4.

For the Japanese networks of the automotive and electronics industries, the data of Luo et al. (2012) can be found in Table 2.

For Kenya, Chacha et al. (2024) report the number of nodes and edges in Table A2.

For the Netherlands (the two Dutch banks), data from Ialongo et al. (2022) can be found in the Data Section.

Data for Rwanda and Uganda are in Spray and Wolf (2018), Chapter 17, Section 2. In Section 3, Spray (2017) also reports data for Uganda, although for a different network.

For Spain, data are taken from Section 3 in Peydró et al. (2020).

For Turkey, Demir et al. (2022) report the number of nodes and edges only approximately in Section 3.1.

For the networks of US-listed firms, Atalay et al. (2011) report the number of nodes and edges in the Data Section (Supporting Information), Cohen and Frazzini (2008) in Section 2 and Wu and Birge (2014) Table 1. Barrot and Sauvagnat (2016) record only the number of nodes in Section A3, Online Appendix.

For West Bengal, Kumar et al. (2021) report the network statistics in Table 1.

For the 5 Indian states (i.e. Gujarat, Maharashtra, Tamil Nadu, Odisha, and West Bengal), in Section 2.1, Panigrahi (2023) provides data on the number of nodes and edges; these numbers are approximate.

Table C.12
Binary network statistics.

Country	Year	N	E	$\bar{k}$	Source(s)
Ecuador	2007	56,058	1,873,023	33.4	This paper
	2008	60,492	2,189,163	36.2	This paper
	2009	62,781	2,277,980	36.3	This paper
	2010	67,159	2,514,144	37.4	This paper
	2011	72,200	2,774,900	38.4	This paper
	2012	76,828	3,013,282	39.2	This paper
	2013	80,519	3,222,500	40	This paper
	2014	84,937	3,439,953	40.5	This paper
	2015	86,345	3,372,929	39.1	This paper
Hungary	2015	119,469	356,788	3	This paper
	2016	123,371	366,048	3	This paper
	2017	133,166	405,075	3	This paper
	2018	285,906	1,493,764	5.2	This paper
	2019	313,117	2,116,912	6.8	This paper
	2020	458,814	13,022,589	28.4	This paper
	2021	493,616	18,710,235	37.9	This paper
FactSet	2014	193,697	472,217	2.4	This paper
	2015	201,277	501,903	2.5	This paper
	2016	203,417	513,570	2.5	This paper
	2017	203,658	515,678	2.5	This paper
	2018	216,491	529,095	2.4	This paper
	2019	215,164	529,973	2.5	This paper
	2020	215,581	524,751	2.4	This paper
	2021	220,508	535,187	2.4	This paper
	2022	180,003	438,134	2.4	This paper
	2023	178,491	463,667	2.6	This paper
Belgium	2002–2012	400,000	88,437,335	221.1*	Dhyne et al. (2015)
	2002	88,301	4,187,000	47.4*	Kikkawa et al. (2019)
	2007	95,941	4,848,000	50.5*	Kikkawa et al. (2019)
	2012	98,745	5,026,000	50.9*	Kikkawa et al. (2019)
	2012	79,788	3,505,207	43.9*	Magerman et al. (2016)
	2012	250,000	8,700,000	34.8*	Dhyne and Duprez (2015)
	2014	321,824	8,900,000	27.7*	Dhyne et al. (2016)
	2014	94,334			Bernard et al. (2022)
Brazil	2003–2014	6,200,000	410,000,000	66.1*	Silva et al. (2020)
Bulgaria	2017	312,762			Alexopoulos et al. (2025)
Chile	2005		5,670,000		Huneeus (2020)
	2008		6,830,000		Huneeus (2020)
	2011		6,580,000		Huneeus (2020)
	2014–2020			20.3*	Miranda-Pinto et al. (2023)
Costa Rica	2008	30,153	299,204*	9.9*	Alfaro-Urena et al. (2018)
	2009	30,708	311,355*	10.1*	Alfaro-Urena et al. (2018)
	2010	32,583	327,007*	10*	Alfaro-Urena et al. (2018)
	2011	34,690	361,583*	10.4*	Alfaro-Urena et al. (2018)
	2012	36,009	389,742*	10.8*	Alfaro-Urena et al. (2018)
	2013	36,702	404,769*	11*	Alfaro-Urena et al. (2018)
	2014	36,572	426,750*	11.7*	Alfaro-Urena et al. (2018)
	2015	36,869	447,366*	12.1*	Alfaro-Urena et al. (2018)
Dominican Republic	2019	44,000	2,500,000	56.8*	Cardoza et al. (2025)
Estonia	2014	16,613			de la Torre et al. (2016)
	2015	88,603	667,133	7.5*	Criscuolo et al. (2024)
	2016	93,164	717,574	7.7*	Criscuolo et al. (2024)
	2017	98,022	772,090	7.9*	Criscuolo et al. (2024)
	2018	99,655	806,150	8.1*	Criscuolo et al. (2024)
	2019	103,742	842,572	8.1*	Criscuolo et al. (2024)
	2020	106,186	821,328	7.7*	Criscuolo et al. (2024)
	2021	115,281	935,432	8.1*	Criscuolo et al. (2024)
Global listed	1994–2015	23,059	2,257,761	97.9*	Wu (2016a)
Global listed cleaned	1994–2015	10,930	1,007,998	92.2*	Wu (2016a)

Notes: Values denoted by an asterisk (*) are inferred; see the text above this table.

Table C.12
Continued.

Country	Year	N	E	$\bar{k}$	Source(s)
Global automotive	10/2013–01/2014	18,942	103,602	5.5*	Brintrup et al. (2015)
Japan	2005	785,939	3,338,319	4.2*	Bernard et al. (2019)
	2005	961,318	3,667,521	3.8*	Ohnishi et al. (2010)
	2006	1,019,854	4,041,442	4*	Fujiwara and Aoyama (2010)
	2008	552,145			Mizuno et al. (2015)
	2009	541,816			Mizuno et al. (2015)
	2010	518,565			Mizuno et al. (2015)
	2011	520,087			Mizuno et al. (2015)
	2012	525,836			Mizuno et al. (2015)
	2012	1,109,549	5,106,081	4.6*	Inoue (2016)
	2010	1,160,000			Lu et al. (2017)
	2013	1,610,000			Lu et al. (2017)
Japan automotive	1983	356	1480	4.2*	Luo et al. (2012)
	1993	679	2473	3.6*	Luo and Magee (2011), Luo et al. (2012)
	2001	627	2175	3.5*	Luo et al. (2012)
Japan electronics	1993	227	648	2.9*	Luo and Magee (2011), Luo et al. (2012)
Kenya	2015	31,684	886,940	28*	Chacha et al. (2024)
	2016	36,920	1,134,159	30.7*	Chacha et al. (2024)
	2017	40,677	1,204,754	29.6*	Chacha et al. (2024)
	2018	44,997	1,332,150	29.6*	Chacha et al. (2024)
	2019	48,697	1,528,410	31.4*	Chacha et al. (2024)
	2020	49,955	1,528,109	30.6*	Chacha et al. (2024)
Netherlands	2019	100,000	1,000,000	10*	Ialongo et al. (2022)
Rwanda	2009–2014	65,193			Spray and Wolf (2018)
	2010	18,714			Spray and Wolf (2018)
	2014	32,330			Spray and Wolf (2018)
Spain	2008	245,524	2,328,908	9.5*	Peydró et al. (2020)
	2009	243,936	2,040,869	8.4*	Peydró et al. (2020)
Turkey manufacturing	2010–2014	600,000	6,000,000	10*	Demir et al. (2022)
Uganda	2010	29,274			Spray and Wolf (2018)
	2014	41,578			Spray and Wolf (2018)
	2010–2015	100,428			Spray and Wolf (2018)
	2009–2016	83,000	420,000	5.1*	Spray (2017)
U.S. listed	2016	20,702	62,474	3*	Taschereau-Dumouchel (2025)
	04/2012–06/2013	2152	11,819	5.5*	Wu and Birge (2014)
	1979–2007	39,815	14,204	0.4*	Atalay et al. (2011)
	1980–2004	30,622	11,484	0.4*	Cohen and Frazzini (2008)
	1980–2009		48,839		Herskovic et al. (2020)
	1978–2013		21,528		Barrot and Sauvagnat (2016)
West Bengal	2016 Q1	131,148	510,026	3.9*	Kumar et al. (2021)
	2016 Q2	131,303	509,404	3.9*	Kumar et al. (2021)
	2016 Q3	137,482	541,254	3.9*	Kumar et al. (2021)
	2016 Q4	143,518	582,045	4.1*	Kumar et al. (2021)
5 Indian states	2011–2016	2,500,000	103,000,000	41.2*	Panigrahi (2023)

Notes: Values denoted by an asterisk (*) are inferred; see the text above this table.

C.2. Number of customer-only or supplier-only firms.

Table C.13 shows the share of firms which are supplier only ($k^{\mathrm{in}} = 0$) or customer only ($k^{\mathrm{out}} = 0$). These proportions differ across datasets, with no obvious patterns. It is possible that the share of customer-only firms is higher than the share of supplier-only firms, but Hungary is a notable exception.[18] Taken at face value, this would mean that the share of firms with no domestic non-labour inputs is smaller than the share of firms with no domestic business customers. We find a similar result for Ecuador but not for Hungary. For non-VAT datasets, the shares of customer-only and supplier-only appear comparable.

[18] An exception is also Uganda, where Spray (2021) reports 87,000 suppliers but only 13,000 customers. We do not report Uganda in Table C.13 because we do not know the total number of firms.

Table C.13
Share of customer-only ($k^{out}=0$) or supplier-only ($k^{in}=0$) firms.

Dataset	Year	Supplier-only	Customer-only	
Ecuador	2015	15.3	20	This paper
Hungary	2021	19.4	13.7	This paper
Hungary	2019	28.1	18.2	This paper
Hungary	2015	36.6	21.1	This paper
FactSet	2021	40.8	45.5	This paper
Belgium	2012	0.1	15.4	Magerman et al. (2016)
Costa Rica	2008–2015	9.7	30.4	Alfaro-Urena et al. (2018)
Estonia	2015–2021	11	35–41	Criscuolo et al. (2024)
Spain	2009	8	24	Peydró et al. (2020)
US listed	04/2012–06/2013	31.1	27.3	Wu and Birge (2014)

C.3. Tail exponents of degree distributions

Tables C.14 and C.15 show the power-law exponents of the in- and out-degree distributions (CCDFs) estimated using the method of Clauset et al. (2009), marked γ^{plfit}, and the three estimators of Voitalov et al. (2019) based on extreme value theory.

Table C.14
Tail exponents for in-degree distributions.

	plfit		Hill		Moment		Kernel	
	γ	κ	γ	κ	γ	κ	γ	κ
Ecuador								
2007	2.07	2018	3.10	229	6.52	377	2.79	11,230
2008	2.06	2391	4.24	13	7.17	125	9.72	60,492
2009	2.16	2239	3.55	77	5.47	141	7.58	62,781
2010	2.25	2403	3.17	123	5.49	283	2.89	38,593
2011	2.16	3353	3.20	142	4.32	444	3.25	17,606
2012	2.36	2462	3.21	178	3.05	4945	2.90	48,982
2013	2.33	2734	3.27	216	3.06	4858	3.44	17,473
2014	2.70	1127	2.72	849	3.05	4019	3.60	33,354
2015	2.38	2900	2.68	547	3.00	4873	3.58	37,695
Hungary								
2015	1.62	1162	1.66	718	1.76	873	1.39	41,147
2016	1.66	836	1.65	466	1.73	832	1.36	44,674
2017	1.35	6663	1.70	352	1.77	981	1.38	50,558
2018	1.66	3916	2.09	153	2.43	551	2.00	8950
2019	1.83	2696	2.22	131	2.02	4684	2.13	8794
2020	2.51	3545	2.51	3524	2.68	12,622	2.71	118,412
2021	2.69	2246	2.71	1550	2.86	13,398	2.83	118,268
FactSet								
2014	1.72	3168	2.21	177	1.99	3147	2.07	7854
2015	1.78	2036	1.82	1403	2.01	2787	2.10	5310
2016	1.77	1922	1.80	1269	1.95	3157	2.08	7543
2017	1.84	1112	1.89	705	1.95	3376	2.06	7014
2018	1.74	2114	2.55	16	1.92	3770	2.07	8014
2019	1.80	1217	2.41	14	1.93	2947	2.04	4486
2020	1.71	1936	2.61	18	1.88	3345	2.01	7416
2021	1.73	1674	2.81	18	1.87	4030	1.97	7966
2022	0.96	112,103	1.60	1270	1.66	4920	1.72	11,179
2023	1.02	121,456	2.35	46	3.74	86	1.58	8939

Notes: Parameters estimated using `plfit` (Clauset et al., 2009) and the three tail-index estimators for generalized extreme value distributions of Voitalov et al. (2019). κ is the number of data points.

Table C.15
Tail exponents for out-degree distributions.

	plfit		Hill		Moment		Kernel	
	γ	κ	γ	κ	γ	κ	γ	κ
Ecuador								
2007	1.26	1963	1.83	68	2.03	127	1.64	2500
2008	1.82	190	1.77	62	2.02	85	1.63	2021
2009	1.38	934	1.58	67	1.84	132	1.41	5304
2010	1.13	3126	1.63	73	1.86	113	1.34	7566
2011	1.40	944	1.60	90	1.93	127	1.80	867
2012	1.36	972	1.70	66	1.82	171	1.76	1089
2013	1.65	210	1.59	102	1.86	259	1.72	946
2014	1.64	228	1.59	119	1.84	175	1.58	2184
2015	1.59	220	1.60	96	1.75	193	1.55	1809
Hungary								
2015	1.46	2771	1.46	2020	1.45	10,054	1.44	48,007
2016	1.43	3739	1.45	2786	1.47	8109	1.42	44,310
2017	1.45	4752	1.45	3746	1.49	5181	1.49	28,501
2018	1.61	1687	1.60	1116	1.65	2766	1.72	5155
2019	1.62	1444	1.63	1266	1.65	2569	1.72	7901
2020	1.43	3577	1.41	4135	1.39	12,200	1.47	25,514
2021	1.42	4081	1.42	3848	1.40	10,188	1.45	26,971
FactSet								
2014	2.69	864	2.68	326	3.32	1304	3.37	6303
2015	2.83	602	2.78	333	3.62	1529	3.90	4177
2016	2.56	1054	2.70	93	3.28	1368	3.71	2098
2017	2.51	967	2.79	124	3.21	1319	3.20	8129
2018	2.39	1071	2.75	206	3.01	727	3.51	4076
2019	2.22	1033	2.45	330	2.68	1180	2.99	4715
2020	2.39	466	2.35	286	2.59	736	2.81	4728
2021	2.39	522	2.36	292	2.60	1529	2.84	6005
2022	0.85	95,519	2.45	456	2.60	1916	3.04	5701
2023	0.85	88,693	2.45	364	2.55	2573	2.77	10,710

Notes: Parameters estimated using `plfit` (Clauset et al., 2009) and the three tail-index estimators for generalized extreme value distributions of Voitalov et al. (2019). κ is the number of data points.

C.4. *Weight distributions*

Table C.16 shows the estimated power-law exponents of the weight distributions (CCDFs) estimated using the method of Clauset et al. (2009), marked γ^{plfit}, and the three estimators of Voitalov et al. (2019) based on extreme value theory. To speed up calculations, we supply only values greater than 10^4 to the algorithm.

Table C.16
Tail exponents for weights distributions.

	plfit		Hill		Moment		Kernel	
	γ	κ	γ	κ	γ	κ	γ	κ
Ecuador								
2007	1.17	1405	1.14	1901	1.13	4579	1.15	7761
2008	1.20	1547	1.21	997	1.22	1512	1.25	3527
2009	0.95	17,442	1.24	351	1.31	1115	1.34	2582
2010	1.09	5403	1.12	2373	1.16	4653	1.19	8083
2011	1.12	5073	1.19	908	1.22	1754	1.20	7232
2012	1.12	2233	1.08	4346	1.11	4663	1.12	8368
2013	1.12	4661	1.13	3097	1.16	4939	1.18	8123
2014	1.16	3046	1.17	1975	1.18	4545	1.21	8345
2015	1.14	5093	1.20	1272	1.21	3493	1.22	7577
Hungary								
2015	1.15	15,095	1.16	8847	1.16	22,412	1.17	65,026
2016	1.17	10,589	1.18	6361	1.18	13,879	1.17	66,030
2017	1.15	6480	1.13	12,149	1.14	22,977	1.15	97,111
2018	1.18	6146	1.18	3920	1.21	5238	1.12	164,953
2019	1.14	10,879	1.15	7459	1.17	10,695	1.18	17,904
2020	1.13	11,259	1.13	9219	1.13	16,855	1.12	64,158
2021	1.18	15,069	1.17	12,530	1.18	23,555	1.20	51,491

Notes: Parameters estimated using `plfit` (Clauset et al., 2009) and the three tail-index estimators for generalized extreme value distributions of Voitalov et al. (2019). κ is the number of data points.

C.5. Strength distributions

Fig. C.12 shows the distribution of in- and out-strengths for Ecuador and Hungary, while Fig. C.13 shows the 2D scatter of in- and out-strengths with the TLS estimates.

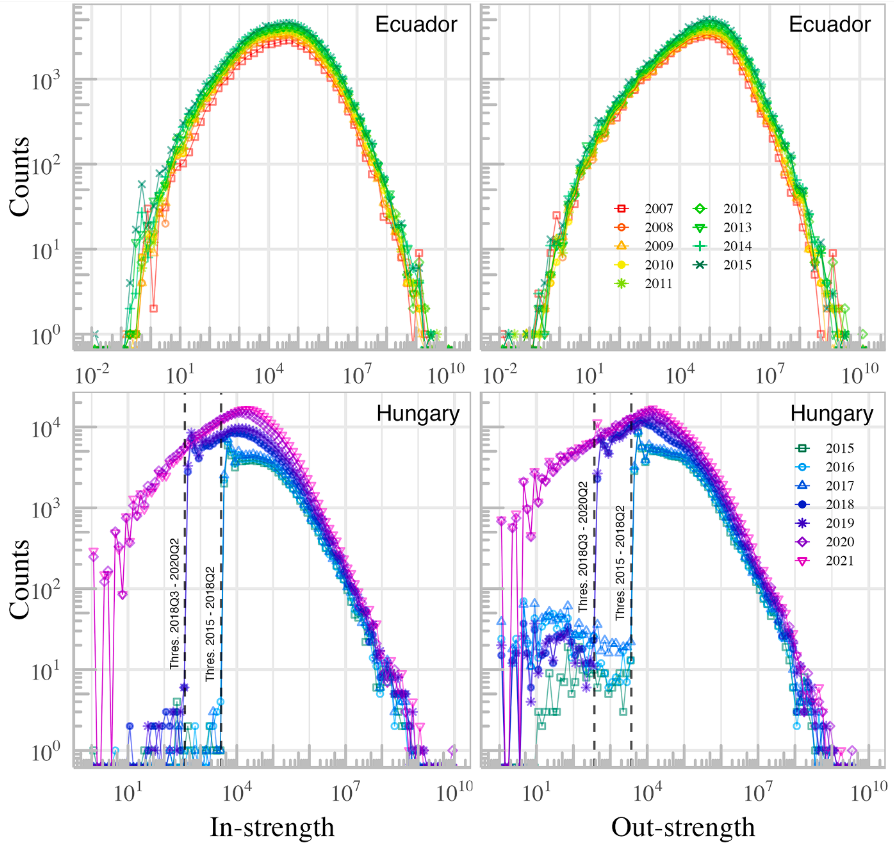

Fig. C.12. Distribution of the in- and out-strengths for Ecuador (left) and Hungary (right) over time. We used 80 log-spaced bins for Ecuador and 100 for Hungary. The two vertical lines for Hungary mark the reporting thresholds; see description of Fig. 1 and Appendix A.2.3. The values are in USD for Ecuador and in 1,000 HUF for Hungary.

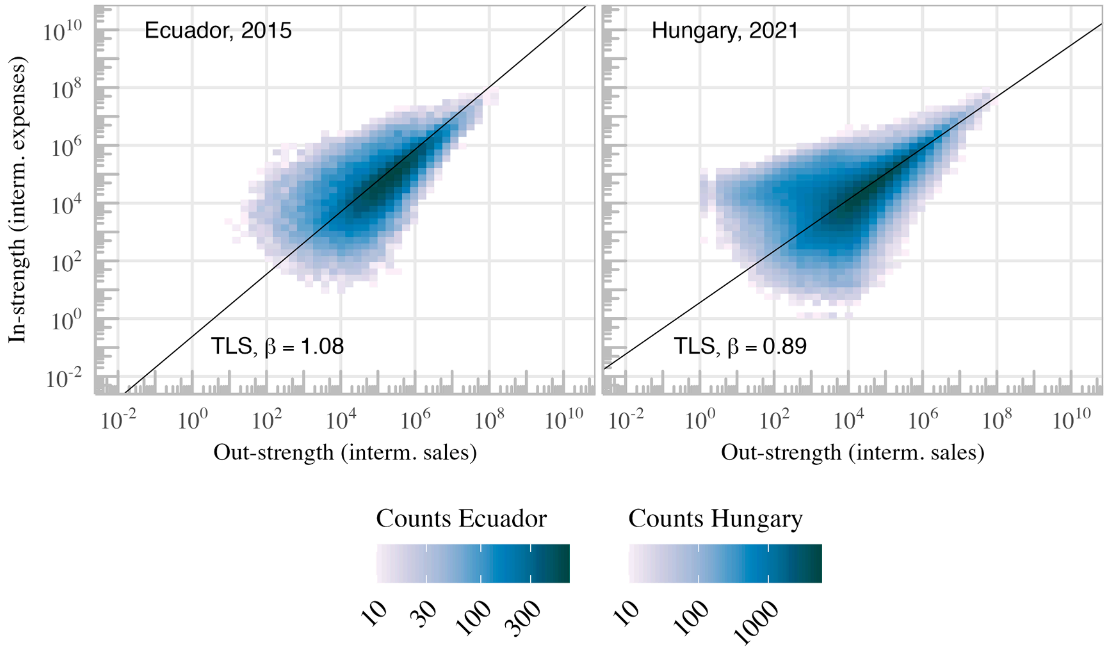

Fig. C.13. 2D histograms for network expenses and network sales for Ecuador in 2015 and Hungary in 2021. We divide both axes into 60 equally-spaced bins and count the number of data points in each square. We do not show squares that have less than 10 observations. The counts are log-transformed. The black line shows the TLS fit. β is the slope of the (TLS) regression line. Ecuador is in 2015 USD and Hungary in 2021 1,000 HUF.

Tables C.17 and C.18 show the estimated power-law exponents using the three estimators of Voitalov et al. (2019) based on extreme value theory and the estimator of Clauset et al. (2009).

Table C.17
Tail exponents for in-strength distributions.

	plfit		Hill		Moment		Kernel	
	γ	κ	γ	κ	γ	κ	γ	κ
Ecuador								
2007	0.94	1143	0.96	518	0.96	1535	0.97	2670
2008	0.91	2010	0.93	1393	0.98	1807	1.00	3659
2009	0.88	2426	1.05	447	1.02	1679	1.29	644
2010	1.01	885	1.04	426	1.08	1076	1.18	746
2011	1.06	584	1.08	355	1.14	474	1.11	1592
2012	0.95	1326	0.96	945	1.05	881	1.04	1628
2013	1.08	489	1.11	276	1.23	332	0.94	8671
2014	0.86	3973	1.12	303	1.21	350	0.93	8559
2015	0.88	3202	1.09	257	1.14	589	0.91	11,880
Hungary								
2015	1.05	1402	1.05	834	1.12	1052	1.40	401
2016	0.89	7218	1.11	380	1.17	708	0.95	18,376
2017	0.87	10,346	1.06	713	0.92	10,057	0.92	31,876
2018	0.91	10,143	0.93	8000	0.96	10,372	0.96	29,557
2019	0.99	4120	1.00	2893	1.06	2898	1.00	24,199
2020	0.97	13,733	0.98	7605	0.99	14,344	1.00	42,998
2021	1.01	13,968	1.01	10,252	1.02	19,370	1.03	60,318

Notes: Parameters estimated using `plfit` (Clauset et al., 2009) and the three tail-index estimators for generalized extreme value distributions of Voitalov et al. (2019). κ is the number of data points.

Table C.18
Tail exponents for out-strength distributions.

	plfit		Hill		Moment		Kernel	
	γ	κ	γ	κ	γ	κ	γ	κ
Ecuador								
2007	1.03	656	1.05	455	1.05	257	0.93	8924
2008	0.87	4145	1.13	323	1.23	306	0.90	16,924
2009	0.87	3575	1.08	378	1.21	527	0.91	16,456
2010	0.88	4291	1.22	28	1.23	451	0.92	15,639
2011	0.88	4751	1.14	206	1.20	307	0.92	16,628
2012	0.88	4296	1.09	320	1.08	327	0.91	15,601
2013	0.91	3843	1.24	151	0.96	4819	0.98	8276
2014	0.93	3738	1.45	35	0.99	3782	1.01	7125
2015	0.92	3742	0.95	2476	1.52	127	1.01	7678
Hungary								
2015	0.92	10,489	1.42	66	1.54	68	0.97	25,624
2016	0.96	5052	1.44	56	1.54	83	0.98	24,674
2017	0.95	5307	0.92	9095	0.96	12,447	0.98	24,761
2018	0.99	6660	1.00	4450	1.01	10,375	1.03	15,452
2019	1.00	5981	1.45	79	1.02	10,154	1.04	18,008
2020	0.99	6271	0.98	5835	1.00	10,385	1.01	22,283
2021	1.02	7823	1.03	5874	1.05	9693	1.07	11,139

Notes: Parameters estimated using `plfit` (Clauset et al., 2009) and the three tail-index estimators for generalized extreme value distributions of Voitalov et al. (2019). κ is the number of data points.

C.6. Input and output shares

The input shares are computed as $P_{ij}^{\text{in}} = Z_{ij} / \sum_i Z_{ij}$ and the output shares as $P_{ij}^{\text{out}} = Z_{ij} / \sum_j Z_{ij}$, where Z_{ij} is the payment from j to i.

Magerman et al. (2016) calculate the input shares of Belgian firms. They find the distribution to be "heavily skewed", with a mean of 0.02, a median of 0.003 and a standard deviation of 0.08. A supplier thus accounts for 2%, on average, of a firm's intermediate input mix. Similar findings are reported by Kikkawa et al. (2019), characterising the distribution of input shares as (roughly) lognormal.

Fig. C.14 shows the distributions of Ecuador and Hungary's input and output shares, displaying roughly a bell-shaped pattern for the log shares. In all the distributions of our complete networks, there is a clear mode around 0.1%. Although small input shares are the most common, it is not rare that a supplier or customer represents a large fraction of costs or sales.

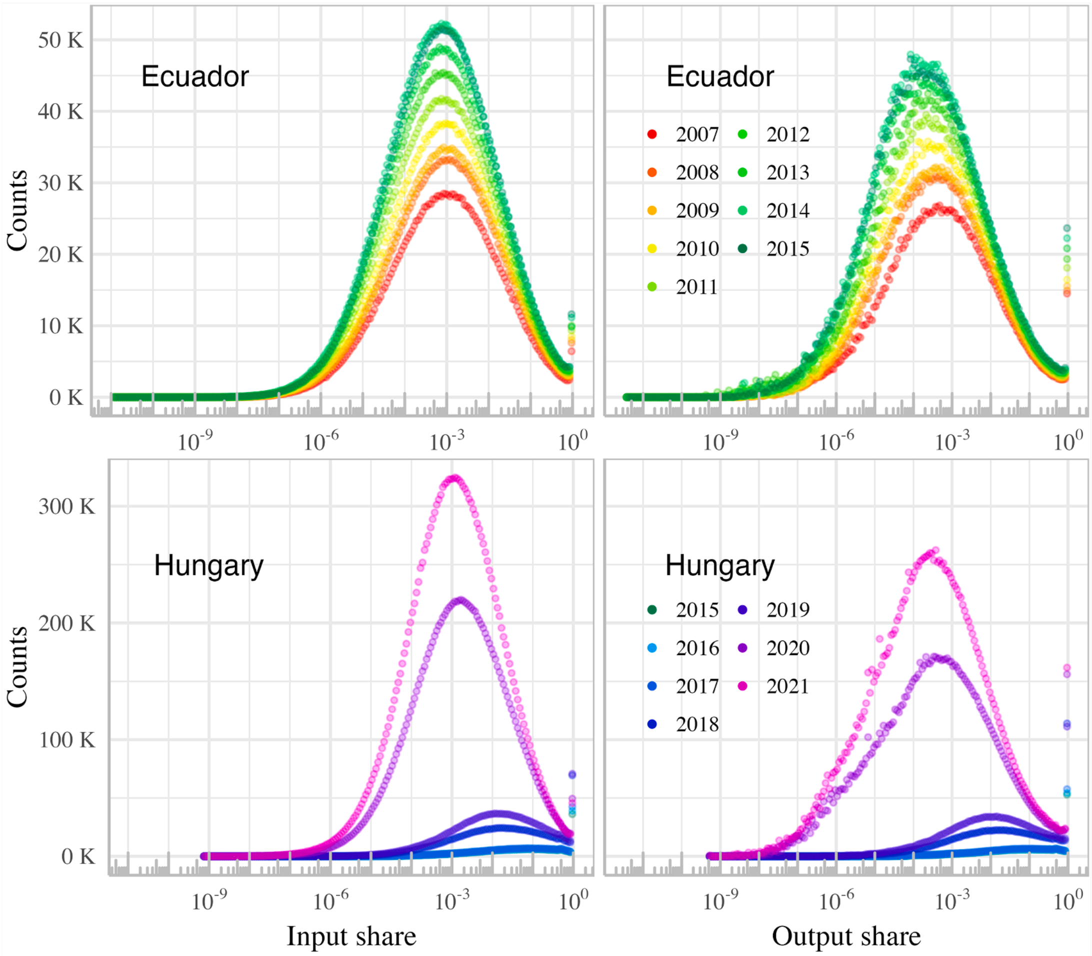

Fig. C.14. Distribution of the input shares (left) and the output shares (right) for Ecuador (top) and Hungary (bottom) over time on a semi-log scale. We binned the data into 200 log-spaced bins.

C.7. Influence vector distributions

Table C.19 shows the estimated power-law exponents for the CCDF of the influence vector over time, using the `plfit` estimator of Clauset et al. (2009) and the three estimators of Voitalov et al. (2019) based on extreme value theory.

Table C.19
Tail exponents for influence distributions.

	plfit		Hill		Moment		Kernel	
	γ	κ	γ	κ	γ	κ	γ	κ
Ecuador								
2007	1.37	3347	1.31	1935	1.37	8109	1.36	27,922
2008	1.31	3012	1.29	2527	1.34	9378	1.32	31,471
2009	1.28	2274	1.27	2124	1.34	9538	1.32	32,570
2010	1.26	2029	1.27	2372	1.32	9363	1.30	38,385
2011	1.29	3515	1.29	3121	1.32	9799	1.32	38,962
2012	1.28	2838	1.28	2847	1.32	10,011	1.30	45,742
2013	1.27	2827	1.27	2984	1.32	9760	1.32	43,119
2014	1.25	3276	1.25	3266	1.30	9638	1.30	50,265
2015	1.28	3472	1.28	3013	1.30	8893	1.33	48,382
Hungary								
2015	1.43	2033	1.39	1475	1.78	91,638	1.47	13,901
2016	1.38	1467	1.36	1231	1.77	93,152	1.77	67,536
2017	1.34	1089	1.34	1098	1.75	101,798	1.75	72,535
2018	1.36	2091	1.38	2375	1.67	233,203	1.44	42,586
2019	1.42	6760	1.34	2108	1.67	313,116	1.39	29,875
2020	1.27	2311	1.25	2538	1.26	5809	1.32	231,306
2021	1.23	3330	1.22	2764	1.24	6427	1.29	247,665

Notes: Parameters estimated using `plfit` (Clauset et al., 2009) and the three tail-index estimators for generalized extreme value distributions of Voitalov et al. (2019). κ is the number of data points.

Appendix D. Truncation effects

In this appendix, we study how the reporting threshold affects key network quantities. We do this by systematically plotting the network quantities against the reporting threshold, using the data we have collected for various countries, which have different thresholds.

To make reporting thresholds comparable between countries with different currencies and levels of development, we choose to express the threshold as a percentage of GDP per capita, using World Bank data (see Table 2). For Hungary, we only report three years (2015, 2019 and 2021), one for each threshold level (see Appendix A.2.3).

We complement this analysis with statistics based on versions of the 2015 Ecuadorian network constructed by removing all links below a given percentile of the distribution of transaction values (from 0% to 95%). We then remove disconnected nodes and keep the LWCC (this last step removes only very few nodes).

In most figures, the x-axis (reporting the threshold as a percentage of GDP per person) is on a log scale. To be able to show countries with a strictly zero reporting threshold, we assume their threshold to be the same as that of Hungary in 2021: around 0.02% of GDP per capita (corresponding to HUF 1000 or about \$2.8).

D.1. Mean degree

Fig. D.15 presents the results for the mean degree. This is the most important chart since the mean degree is commonly reported, so we can compare many countries. Although the mean degree is expected to decline with increasing reporting thresholds, a more striking feature of Fig. D.15 is that the threshold level is highly predictive of the observed mean degree. In other words, Fig. D.15 shows that complete datasets have a mean degree around 30–50 and that countries with similar reporting thresholds have similar mean degrees.

The empirical findings of Fig. D.15 suggest that the reporting threshold largely explains the cross-country differences in observed mean degrees, reinforcing the conclusion that production networks have a similar structure across countries. These results also indicate that at least some observed properties of the networks can be de-biased. For instance, although Spain and Estonia have a mean degree around 7–9, based on Fig. D.15, we predict that if the reporting threshold were lowered to near zero, we would observe a mean degree at least 30–50.

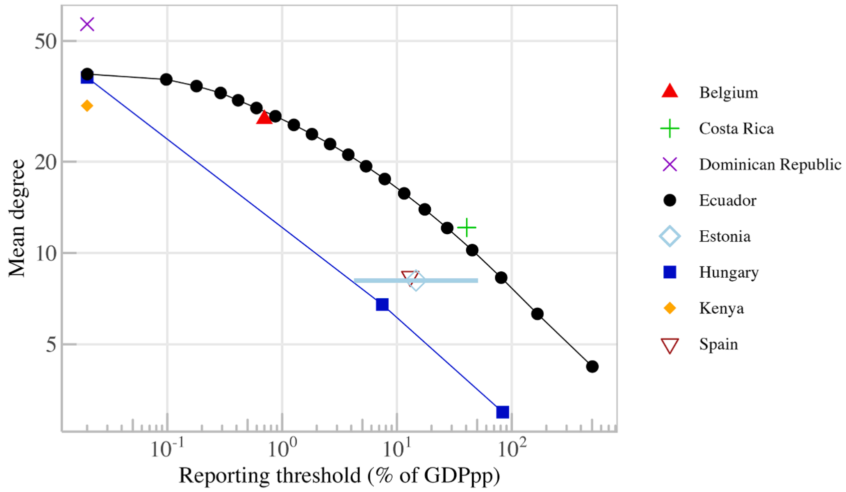

Fig. D.15. Mean degree as a function of the reporting threshold (expressed as a percentage of GDP per capita). In this and other figures in the Appendix, Estonia is shown as a horizontal line because its threshold is monthly and it is unclear how to translate this into a yearly threshold; see the caption of Table 2.

D.2. Assortativity and clustering

Fig. D.16 shows how assortativity and clustering vary with increasing reporting thresholds. The results for assortativity are clear: truncation makes assortativity less negative. Fig. D.16 suggests that truncation has a limited effect on the global clustering coefficient; specifically, the change induced by truncation is small, nonlinear and lower in magnitude than the difference observed between Ecuador and Hungary. By contrast, truncation has a substantial impact on the average local clustering coefficient.

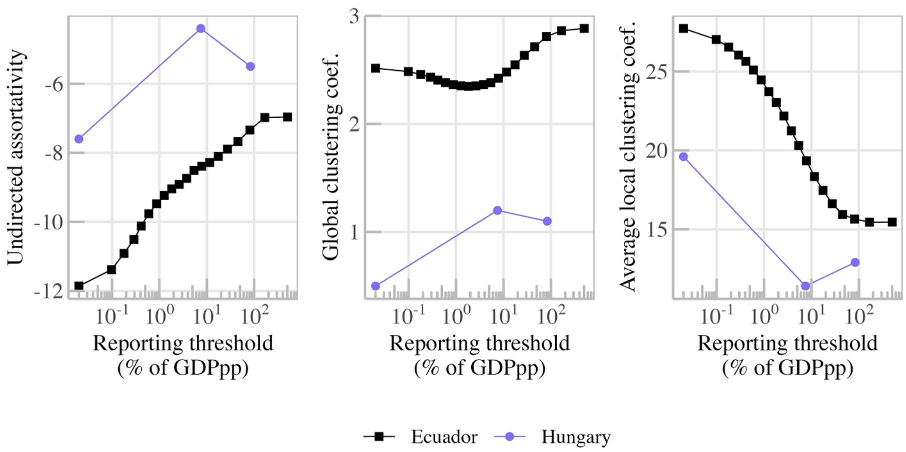

Fig. D.16. Assortativity, $r_{k,k}$ (left), global clustering coefficient, C_g (center), and average local clustering coefficient, $\bar{C}_i$ (right), as a function of the reporting threshold (expressed as a percentage of GDP per capita).

D.3. Degree distributions and the influence vector

Fig. D.17 shows that while the reporting threshold does not affect the power-law exponent of the out-degree distribution, leaving it around 1.5, it significantly impacts the measured power-law exponent of the in-degree distribution, shifting it from above to below 2.

These findings are non-trivial and potentially valuable. Similar to the mean degree, since the values of synthetically-truncated Ecuador and those of Hungary, across the three truncation regimes, overlap substantially, we can have good confidence in both the absolute level of these values (≈ 1.5 for out-degrees and 2.5 for in-degrees) and the effect of the reporting threshold (none for out-degrees and fatter tails for in-degrees). As a result, it may be possible to use this relationship to correct results obtained from a truncated, incomplete dataset.

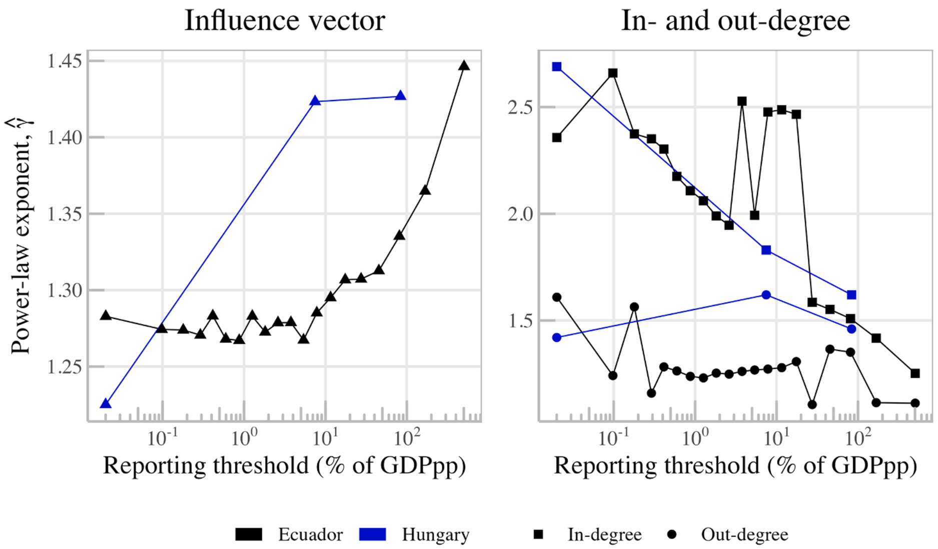

Fig. D.17. Power-law exponents for the influence vector (left) and the in- and out-degree distributions (right) as a function of the reporting threshold (expressed as a percentage of GDP per capita).

The results for the influence vector are also relatively clear, with Fig. D.17 indicating that higher reporting thresholds lead to a measured distribution that is less fat-tailed.

D.4. Strength-degree relationships

Fig. D.18 presents the results for the strength-degree relations. Similar to the average degree, this case is particularly informative because we have several data points. Despite some heterogeneity, the results overall provide fairly clear evidence of both an unambiguous effect of the reporting threshold and universality: complete datasets exhibit similar values.

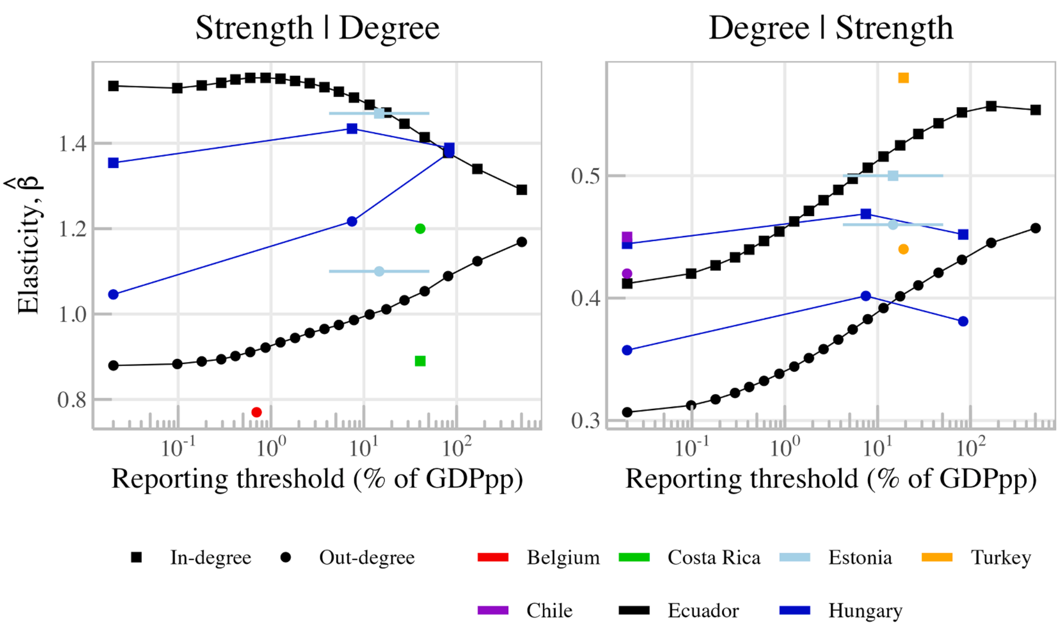

Fig. D.18. Strength-degrees elasticities as a function of the reporting threshold (expressed as a percentage of GDP per person). We always regress in-strength with in-degree (squares) and out-strength with out-degree (circles). To give an example, on the left panel, the square dots show the coefficients of the regressions of log in-strength on log in-degree.

First, in-relations have higher coefficients than out-relations, with Costa Rica as the only exception. Second, both within- and cross-country evidence suggests that truncation generally increases the value of the coefficient, except for the in-strength/in-degree relations (square dots in the left panel).

Appendix E. Industry-level analysis

In the main text, we do not make use of industry labels. The reason is that most of the heterogeneity is not driven by industry differences. We also worry that industrial codes are highly imperfect and not perfectly comparable across countries.

In this appendix, we investigate heterogeneity across industries. Our main goal is to show that most heterogeneity is not driven by industry differences and that our results are robust to controlling for industry heterogeneity. By establishing substantial within-industry differences, this appendix more generally motivates the use of firm- rather than sector-level data.

While we also take this as an opportunity to comment on inter-industry differences, we refrain from providing an overly detailed interpretation of the results, since we have only two datasets with complete data; we leave this to further research.

In principle, we could investigate edge-level quantities, such as reciprocity, assortativity and edge weights, but these are more convoluted because each edge connects nodes from potentially two different industries. We therefore limit our investigation to node-level quantities and we leave this to further research.

E.1. Is heterogeneity driven by industry differences?

To answer this question, we proceed with a variance decomposition of key node-level quantities. Tables E.20 and E.21 present the results for Ecuador and Hungary, respectively, at various levels of industry disaggregation. The tables report the number of observations, the total variance, the variance explained by the between component, its share of total variance and the overall mean of the metric. Note that we do not observe the industry affiliation of all firms. For Ecuador, this concerns only a few firms, so the total variance and mean match those in Table B.10, and the mean local clustering matches Table 6; for Hungary, the corresponding figures differ slightly.

Table E.20
Variance decomposition of key network metrics for Ecuador in 2015.

Metric	N	Var_{tot}	$\text{Var}_{between}$	$\text{Share}_{between}$	Mean
ISIC 1-digit					
In-degree (log)	73,138	2.36	0.20	8.54%	2.83
Out-degree (log)	69,048	2.83	0.35	12.38%	1.85
In-strength (log)	73,138	8.99	0.88	9.82%	9.91
Out-strength (log)	69,048	8.17	0.79	9.72%	10.62
Local clustering coefficient	76,491	0.05	0.00	4.11%	0.28
ISIC 2-digit					
In-degree (log)	73,138	2.36	0.27	11.25%	2.83
Out-degree (log)	69,048	2.83	0.42	14.91%	1.85
In-strength (log)	73,138	8.99	1.16	12.89%	9.91
Out-strength (log)	69,048	8.17	1.04	12.76%	10.62
Local clustering coefficient	76,491	0.05	0.00	5.30%	0.28
ISIC 3-digit					
In-degree (log)	73,138	2.36	0.30	12.75%	2.83
Out-degree (log)	69,048	2.83	0.51	17.98%	1.85
In-strength (log)	73,138	8.99	1.31	14.59%	9.91
Out-strength (log)	69,048	8.17	1.21	14.82%	10.62
Local clustering coefficient	76,491	0.05	0.00	6.24%	0.28
ISIC 4-digit					
In-degree (log)	73,138	2.36	0.34	14.34%	2.83
Out-degree (log)	69,048	2.83	0.54	19.15%	1.85
In-strength (log)	73,138	8.99	1.50	16.65%	9.91
Out-strength (log)	69,048	8.17	1.34	16.46%	10.62
Local clustering coefficient	76,491	0.05	0.00	7.29%	0.28

In all cases, the between-group component accounts for a relatively small share of the total variance, in the range of 3–20%, and increasing with industry disaggregation as expected. The results are also very similar across countries.

One result that is sufficiently large and consistent across countries to warrant interpretation is that the share of variance explained by the between component is higher for out-degree than for in-degree. We return to this point when examining mean degrees across industries in Section E.3.

These results strongly justify our approach in the main text of quantifying heterogeneity without appealing to industry differences. Nevertheless, they also suggest that further investigation could reveal interesting patterns in how specific industries contribute to heterogeneity.

Table E.21
Variance decomposition of key network metrics for Hungary in 2021.

Metric	N	Var_{tot}	$\text{Var}_{between}$	$\text{Share}_{between}$	Mean
NACE 1-digit					
In-degree (log)	294,490	2.13	0.20	9.29%	3.12
Out-degree (log)	270,781	2.95	0.34	11.62%	2.01
In-strength (log)	294,490	6.52	0.51	7.87%	9.35
Out-strength (log)	270,781	6.88	0.55	7.97%	9.36
Local clustering coefficient	292,001	0.03	0.00	3.08%	0.19
NACE 2-digit					
In-degree (log)	294,490	2.13	0.23	10.98%	3.12
Out-degree (log)	270,781	2.95	0.40	13.61%	2.01
In-strength (log)	294,490	6.52	0.62	9.50%	9.35
Out-strength (log)	270,781	6.88	0.94	13.67%	9.36
Local clustering coefficient	292,001	0.03	0.00	3.60%	0.19
NACE 3-digit					
In-degree (log)	294,490	2.13	0.27	12.78%	3.12
Out-degree (log)	270,781	2.95	0.53	18.04%	2.01
In-strength (log)	294,490	6.52	0.76	11.67%	9.35
Out-strength (log)	270,781	6.88	1.17	16.99%	9.36
Local clustering coefficient	292,001	0.03	0.00	4.40%	0.19
NACE 4-digit					
In-degree (log)	294,490	2.13	0.29	13.67%	3.12
Out-degree (log)	270,781	2.95	0.60	20.33%	2.01
In-strength (log)	294,490	6.52	0.87	13.31%	9.35
Out-strength (log)	270,781	6.88	1.29	18.83%	9.36
Local clustering coefficient	292,001	0.03	0.00	4.89%	0.19

E.2. Do strength-degree relations change when controlling for industry differences?

In the main text, we report regressions of degrees on strengths and vice versa. Although Arkolakis et al. (2023), Miranda-Pinto et al. (2023) and Demir et al. (2023), for instance, already showed that including industry fixed effects does not substantially affect the coefficients, we repeat this exercise here. Tables E.22 and E.23 show the results. Each column reports a different specification,

Table E.22
Strength–degree elasticities under alternative industry and year fixed-effects specifications for Ecuador.

Network sales & number of customers					
ln s | ln k					
ISIC 1-digit	0.88	0.95	0.90	0.94	0.95
ISIC 2-digit	0.88	0.96	0.90	0.96	0.96
ISIC 3-digit	0.88	0.98	0.90	0.98	0.98
ISIC 4-digit	0.88	0.99	0.90	0.98	0.98
ln k | ln s					
ISIC 1-digit	0.31	0.32	0.32	0.33	0.33
ISIC 2-digit	0.31	0.32	0.32	0.34	0.34
ISIC 3-digit	0.31	0.33	0.32	0.34	0.34
ISIC 4-digit	0.31	0.33	0.32	0.34	0.34
Network expenses & number of suppliers					
ln s | ln k					
ISIC 1-digit	1.55	1.52	1.53	1.49	1.49
ISIC 2-digit	1.55	1.51	1.53	1.48	1.48
ISIC 3-digit	1.55	1.51	1.53	1.48	1.48
ISIC 4-digit	1.55	1.51	1.53	1.48	1.48
ln k | ln s					
ISIC 1-digit	0.41	0.40	0.42	0.41	0.41
ISIC 2-digit	0.41	0.40	0.42	0.41	0.41
ISIC 3-digit	0.41	0.41	0.42	0.41	0.41
ISIC 4-digit	0.41	0.41	0.42	0.41	0.41
2015 only	Yes	Yes	No	No	No
Industry FE	No	Yes	No	Yes	Yes
Year FE	No	No	No	No	Yes

Table E.23
Strength–degree elasticities under alternative industry and year fixed-effects specifications for Hungary.

Network sales & number of customers					
ln s | ln k					
NACE 1-digit	0.89	0.94	0.90	0.95	0.95
NACE 2-digit	0.89	0.94	0.90	0.95	0.95
NACE 3-digit	0.89	0.95	0.90	0.95	0.95
NACE 4-digit	0.89	0.95	0.90	0.96	0.95
ln k | ln s					
NACE 1-digit	0.38	0.39	0.37	0.37	0.37
NACE 2-digit	0.38	0.40	0.37	0.39	0.39
NACE 3-digit	0.38	0.40	0.37	0.39	0.39
NACE 4-digit	0.38	0.40	0.37	0.38	0.38
Network expenses & number of suppliers					
ln s | ln k					
NACE 1-digit	1.40	1.40	1.40	1.39	1.39
NACE 2-digit	1.40	1.40	1.40	1.39	1.39
NACE 3-digit	1.40	1.40	1.40	1.39	1.39
NACE 4-digit	1.40	1.39	1.40	1.38	1.38
ln k | ln s					
NACE 1-digit	0.46	0.45	0.44	0.43	0.43
NACE 2-digit	0.46	0.45	0.44	0.43	0.43
NACE 3-digit	0.46	0.45	0.44	0.44	0.43
NACE 4-digit	0.46	0.45	0.44	0.44	0.44
2021 only	Yes	Yes	No	No	No
Industry FE	No	Yes	No	Yes	Yes
Year FE	No	No	No	No	Yes

for one year or multiple years, with or without industry and year fixed effects (for Hungary, we use at most 2020 and 2021). The results confirm that adding industry fixed effects or using multiple years, with or without year fixed effects, does not meaningfully affect the reported coefficients.

E.3. Where do we find interesting industry differences and are they consistent across countries?

Although inter-industry differences do not provide a dominant explanation for the large heterogeneity observed in firm-level properties, we do, in some cases, observe substantial differences across industries that are worth exploring.

We start with mean strengths, which are simpler to interpret and provide a foundation for the discussion of degrees. Fig. E.19 shows scatter plots of mean in- and out-strength for Ecuador in 2015 and Hungary in 2021. The relative closeness of the points to the unit line confirms that industries tend to sell as much as they buy within the network. However, given the logarithmic scale, visually small differences can be substantial.

Industries D (electricity), B (mining) and C (manufacturing) tend to have larger firms in both countries. There is no industry exhibiting a strong imbalance in both countries. As a result, we refrain from further discussion, as we focus only on quantitatively strong and highly convergent evidence across the two countries.

Fig. E.20 shows the same scatter plots but for degrees. In this case, heterogeneity is much higher and echoes the evidence reported in Table E.20, with out-degree exhibiting greater cross-industry heterogeneity than in-degree.

Some patterns appear plausible and sensible. Electricity (D) and wholesale (G), by their nature, buy in bulk and sell to many customers. In contrast, manufacturing (C) and agriculture (A) have many more suppliers than customers. It is plausible that, in these industries, firms aggregate inputs from many suppliers into products that they sell to large buyers, such as wholesalers. For other industries, while it is tempting to speculate, we refrain from further discussion until evidence from several countries confirms which patterns are genuinely common and associated with intuitive economic behaviour.

These differences warrant one additional robustness check. In the main text, we find that the TLS slope for the in-degree $\sim$ out-degree relation is substantially less than 1 and around 3/4. This suggests that, subject to the caveat that cross-sectional variation may not faithfully reflect time-series behaviour, firms scale their out-degree faster than their in-degree. We then ask whether this pattern also holds at the industry level.

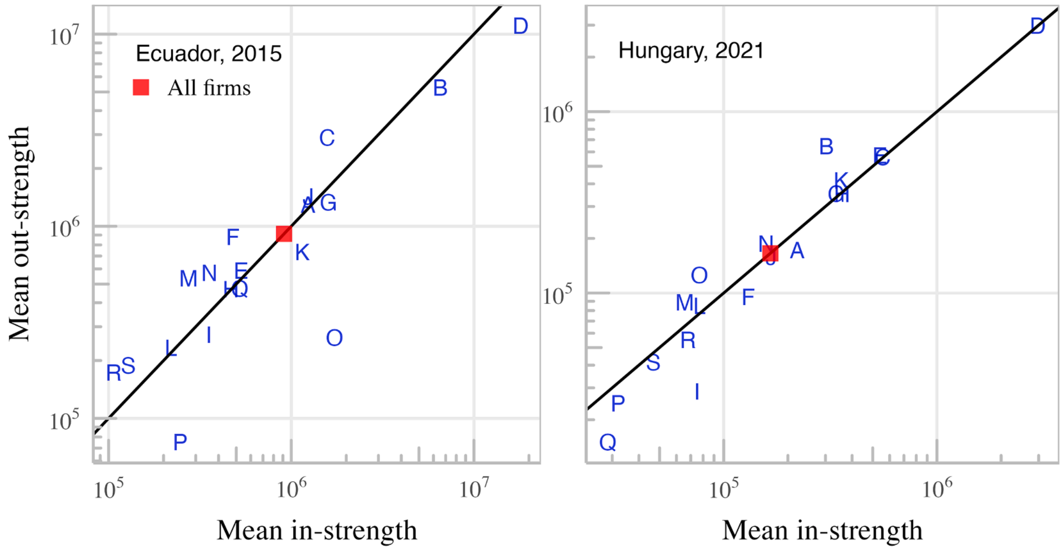

Fig. E.19. Mean in- and out-strength by industry for Ecuador 2015 (left) and Hungary 2021 (right). The red square marks the economy-wide value, which includes all firms. Letters indicate the 1-digit level NACE codes for Hungary and ISIC codes for Ecuador. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

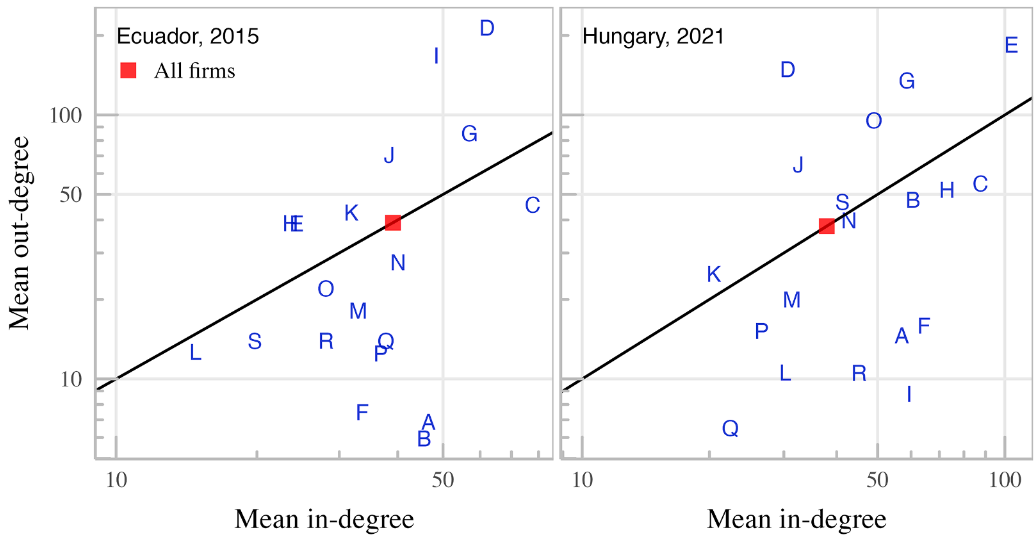

Fig. E.20. Mean in- and out-degree by industry for Ecuador in 2015 (left) and Hungary in 2021 (right). The red square marks the economy-wide value, which includes all firms. Letters indicate the 1-digit level NACE codes for Hungary and ISIC codes for Ecuador. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Fig. E.21 presents the estimated TLS coefficients from regressions of in-degree on out-degree. Most industries exhibit coefficients close to the economy-wide values of 0.74 for Ecuador in 2015 and 0.63 for Hungary in 2021. A few industries have exponents greater than 1, but this is inconsistent across the two countries. Some industries lie below the economy-wide value in both countries, such as wholesale, or above it in both countries, such as construction. While this may reflect genuine patterns observable in other countries, we refrain from over-interpreting these results given our limited sample.

Again, these findings justify our focus on aggregate quantities in the main text, while suggesting potentially interesting differences across countries at the industry level. Our analysis in this section lays the foundation for a research agenda exploring such differences.

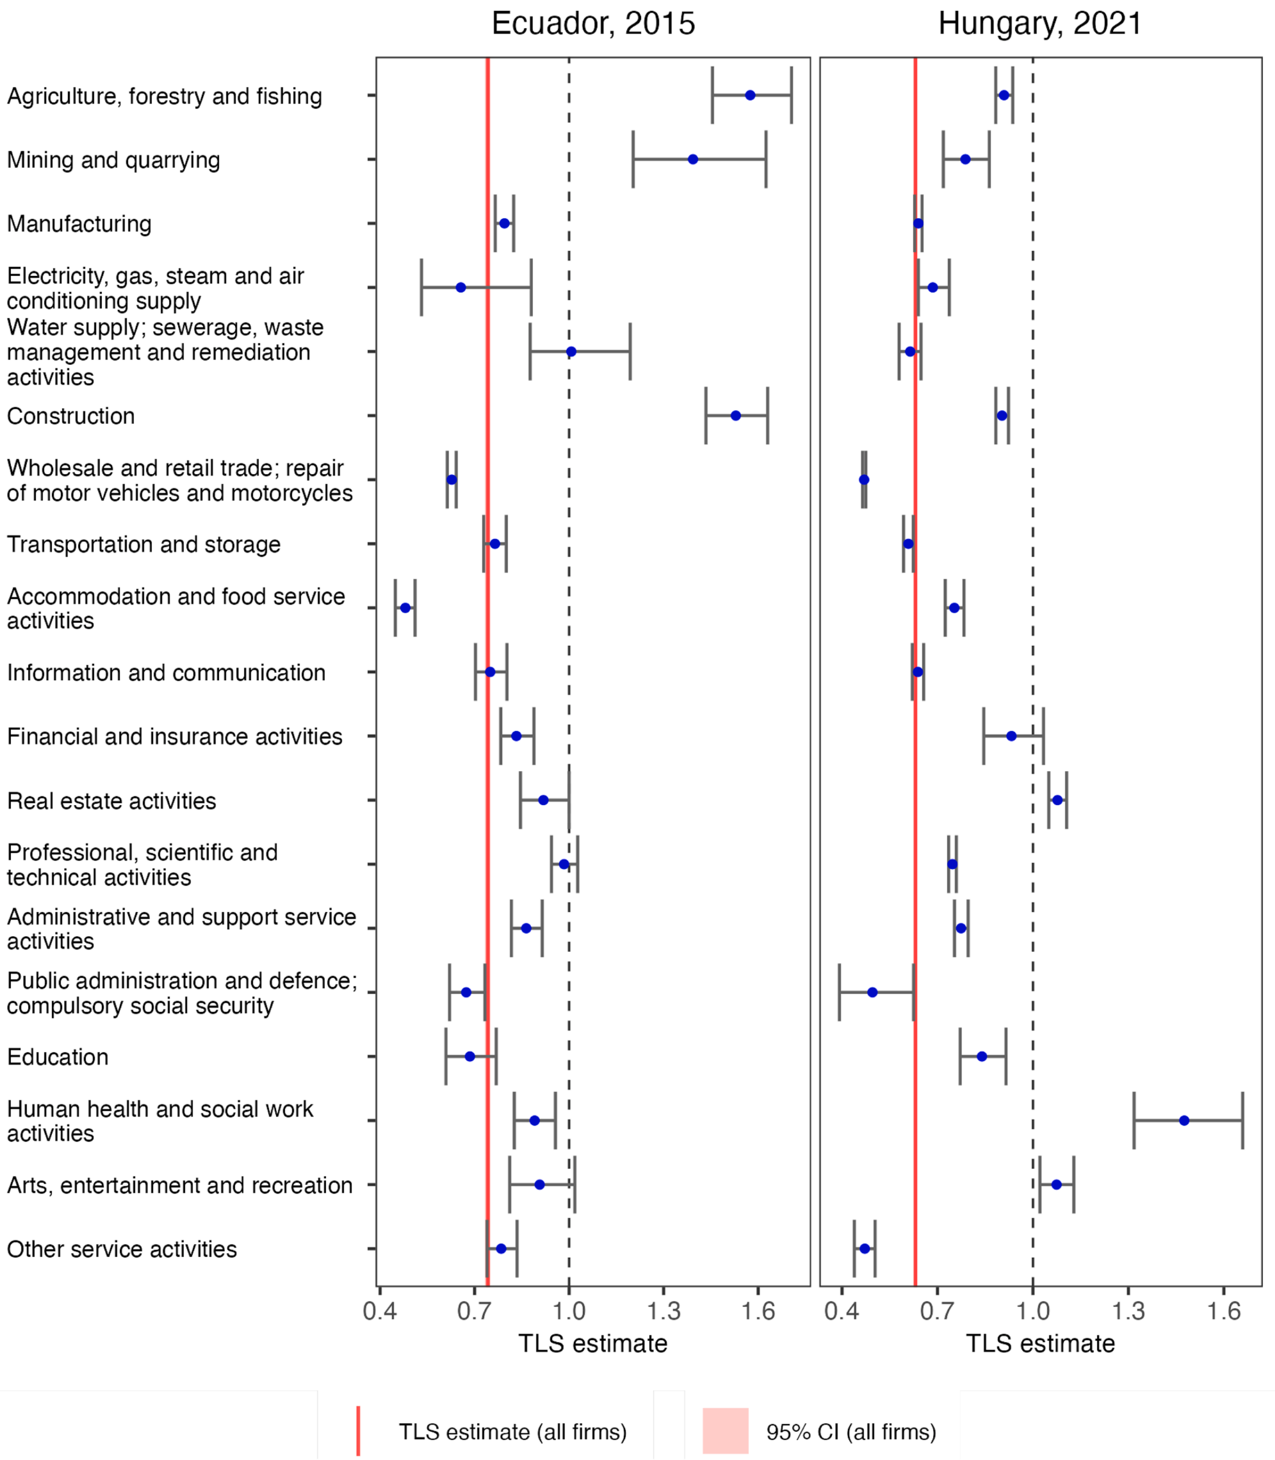

Fig. E.21. TLS estimates for the regression of in-degree on out-degree by industry (blue dot) for Ecuador in 2015 (left) and Hungary in 2021 (right). The grey bars show the 95% CI. The red line shows the TLS estimate for all firms in the dataset and the red-shaded area shows the 95% CI. We use the bootstrap method to estimate the confidence intervals. The black dashed line denotes the unit reference value. We use industrial codes at the 1-digit level; NACE codes for Hungary and ISIC codes for Ecuador. We do not show industry T, U and V due to the small number of observations. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

References

Acemoglu, D., Azar, P.D., 2020. Endogenous production networks. Econometrica 88 (1), 33–82.

Acemoglu, D., Carvalho, V.M., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. Econometrica 80 (5), 1977–2016.

Alexopoulos, A., Dellaportas, P., Gyoshev, S., Kotsogiannis, C., Olhede, S.C., Pavkov, T., 2025. A network and machine learning approach to detect Value Added Tax fraud. https://doi.org/10.48550/arXiv.2106.14005

Alfaro-Urena, A., Fuentes, M.F., Manelici, I., Vásquez, J.P., 2018. Costa Rican production network: stylized sacts. Central Bank of Costa Rica Research Documents No. 002.

Alfaro-Urena, A., Manelici, I., Vasquez, J.P., 2022. The effects of joining multinational supply chains: new evidence from firm-to-firm linkages. Q. J. Econ. 137 (3), 1495–1552.

Arkolakis, C., Huneeus, F., Miyauchi, Y., 2023. Spatial production networks. NBER Working paper 30954.

Astudillo-Estevez, P.A., 2021. Towards a Post-Oil Economy: A Complexity Approach to Understanding Natural Resource Dependency and Economic Diversification in Ecuador [Doctoral Dissertation]. Ph.D. thesis. University of Oxford.

Atalay, E., Hortacsu, A., Roberts, J., Syverson, C., 2011. Network structure of production. Proc. Natl. Acad. Sci. 108 (13), 5199–5202.

Atalay, E., Hortaçsu, A., Syverson, C., 2014. Vertical integration and input flows. A. Econ. Rev. 104 (4), 1120–1148.

Axtell, R.L., 2001. Zipf distribution of US firm sizes. Science 293 (5536), 1818–1820.

Ball, B., Newman, M. E.J., 2013. Friendship networks and social status. Network Sci. 1 (1), 16–30.

Barrot, J.-N., Sauvagnat, J., 2016. Input specificity and the propagation of idiosyncratic shocks in production networks. Q. J. Econ. 131 (3), 1543–1592.

Bernard, A.B., Dhyne, E., Magerman, G., Manova, K., Moxnes, A., 2022. The origins of firm heterogeneity: a production network approach. J. Polit. Econ. 130 (7), 1765–1804.

Bernard, A.B., Moxnes, A., Saito, Y.U., 2019. Production networks, geography, and firm performance. J. Polit. Econ. 127 (2), 639–688.

Bernard, A.B., Zi, Y., 2022. Sparse production networks. NBER Working Paper 30496.

Boguná, M., Pastor-Satorras, R., Vespignani, A., 2004. Cut-offs and finite size effects in scale-free networks. Eur. Phys. J. B 38 (2), 205–209.

Borsos, A., Stancsics, M., 2020. Unfolding the hidden structure of the Hungarian multi-layer firm network. MNB Working Papers 2020/139. Magyar Nemzeti Bank (Central Bank of Hungary).

Boucher, V., Houndetoungan, A., 2025. Estimating peer effects using partial network data. Rev. Econ. Stat. , 1–48.

Brintrup, A., Barros, J., Tiwari, A., 2015. The nested structure of emergent supply networks. IEEE Syst. J. 12 (2), 1803–1812.

Buda, G., Hansen, S., Rodrigo, T., Carvalho, V.M., Ortiz, Á., Mora, J. V.R., 2022. National accounts in a world of naturally occurring data: a proof of concept for consumption. Cambridge Faculty of Economics Working paper .

Cardoza, M., Grigoli, F., Pierri, N., Ruane, C., 2025. Worker mobility in production networks. Rev. Econ. Stud. 92 (6), 3682–3703. https://doi.org/10.1093/restud/rdae114

Carvalho, V.M., 2014. From micro to macro via production networks. J. Econ. Persp. 28 (4), 23–48.

Carvalho, V.M., Nirei, M., Saito, Y.U., Tahbaz-Salehi, A., 2021. Supply chain disruptions: evidence from the great East Japan earthquake. Q. J. Econ. 136 (2), 1255–1321.

Carvalho, V.M., Tahbaz-Salehi, A., 2019. Production networks: a primer. Annu. Rev. Econom. 11, 635–663.

Carvalho, V.M., Voigtländer, N., 2014. Input diffusion and the evolution of production networks. Working Paper 20025. National Bureau of Economic Research. https://doi.org/10.3386/w20025

Cattaneo, M.D., Crump, R.K., Farrell, M.H., Feng, Y., 2019. On binscatter. arXiv preprint arXiv: 1902.09608 .

Chacha, P.W., Kirui, B.K., Wiedemann, V., 2024. Supply chains in times of crisis: evidence from Kenya's production network. World Dev. 173. https://doi.org/10.1016/j.worlddev.2023.106363

Chakraborty, A., Ikeda, Y., 2020. Testing "efficient supply chain propositions" using topological characterization of the global supply chain network. PLoS ONE 15 (10), e0239669.

Clauset, A., Shalizi, C.R., Newman, M. E.J., 2009. Power-law distributions in empirical data. SIAM Rev. 51 (4), 661–703.

Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. J. Finance 63 (4), 1977–2011.

Cohen, R., Havlin, S., 2003. Scale-free networks are ultrasmall. Phys. Rev. Lett. 90 (5), 058701.

Costello, A.M., 2020. Credit market disruptions and liquidity spillover effects in the supply chain. J. Polit. Econ. 128 (9), 3434–3468.

Criscuolo, C., Dechezleprêtre, A., Guillouet, L., Lalanne, G., Manaresi, F., 2024. Estonia's firm-level production network: Lessons for industrial policy. OECD Science, Technology and Industry Working Papers No. 2024/13. OECD Publishing. Paris. https://doi.org/10.1787/e0f18b9c-en.

Demir, B., Fieler, A.C., Xu, D., Yang, K.K., 2023. O-ring production networks. J. Polit. Econ. forthcoming.

Demir, B., Javorcik, B., Michalski, T.K., Ors, E., 2022. Financial constraints and propagation of shocks in production networks. Rev. Econ. Stat. 106 (2), 437–454.

Dhyne, E., Duprez, C., 2015. Has the crisis altered the Belgian economy's DNA? Econ. Rev. ii, 31–43.

Dhyne, E., Duprez, C., et al., 2016. Three regions, three economies? Econ. Rev. iii, 59–73.

Dhyne, E., Kikkawa, A.K., Magerman, G., 2022. Imperfect competition in firm-to-firm trade. J. Eur. Econ. Assoc. 20 (5), 1933–1970.

Dhyne, E., Kikkawa, A.K., Mogstad, M., Tintelnot, F., 2021. Trade and domestic production networks. Rev. Econ. Stud. 88 (2), 643–668.

Dhyne, E., Magerman, G., Rubínová, S., 2015. The Belgian production network 2002–2012. NBB Working Paper No. 288. National Bank of Belgium, Brussels.

Diem, C., Borsos, A., Reisch, T., Kertész, J., Thurner, S., 2022. Quantifying firm-level economic systemic risk from nation-wide supply networks. Sci. Rep. 12 (1), 1–13.

Diem, C., Borsos, A., Reisch, T., Kertész, J., Thurner, S., 2023. Estimating the loss of economic predictability from aggregating firm-level production networks. arXiv preprint arXiv: 2302.11451 .

Dorogovtsev, S.N., Mendes, J. F.F., 2003. Evolution of Networks: From Biological Nets to the Internet and WWW. Oxford University Press.

Duprez, C., Magerman, G., 2018. Price Updating in Production Networks. Working Paper Research No 352. National Bank of Belgium. https://ideas.repec.org/p/nbb/reswpp/201810-352.html.

Egger, D., Faber, B., Li, M., Lin, W., 2025. Rural-Urban Migration and Market Integration. NBER Working Paper 34098. National Bureau of Economic Research. https://doi.org/10.3386/w34098

Eurostat, 2008. Eurostat Manual of Supply, Use and Input-Output Tables. Eurostat Methodologies and Working Papers 2008 edition. Office for Official Publications of the European Communities, Luxembourg.

Fujiwara, Y., Aoyama, H., 2010. Large-scale structure of a nation-wide production network. Eur. Phys. J. B 77 (4), 565–580.

Fujiwara, Y., Inoue, H., Yamaguchi, T., Aoyama, H., Tanaka, T., Kikuchi, K., 2021. Money flow network among firms' accounts in a regional bank of Japan. EPJ Data Sci. 10 (1), 1–26.

Furusawa, T., Inui, T., Ito, K., Tang, H., 2017. Global sourcing and domestic production networks. CESifo Working Paper Series No. 6658.

Gabaix, X., 2011. The granular origins of aggregate fluctuations. Econometrica 79 (3), 733–772.

Gabaix, X., Ibragimov, R., 2011. Rank- 1/2: a simple way to improve the OLS estimation of tail exponents. J. Bus. Econ. Stat. 29 (1), 24–39.

Gadenne, L., Rathelot, R., Nandi, T., 2019a. Linkages with Multinationals and Domestic Firms' Performance. Technical Note IDB-TN-01746. Integration and Trade Sector, Inter-American Development Bank. http://dx.doi.org/10.18235/0001848.

Gadenne, L., Rathelot, R., Nandi, T., 2019b. Taxation and Supplier Networks: Evidence from India. Working Paper 19/06. Saïd Business School, University of Oxford. https://oxfordtax.sbs.ox.ac.uk/wp-19/06-lucie-gadenne-tushar-k.-nandi-roland-rathelot-taxation-and-supplier-networks-evidence-from.

Gloria, J., Miranda-Pinto, J., Fleming-Muñoz, D., 2024. Production network diversification and economic development. J. Econ. Behav. Organ. 218, 281–295.

Grigoli, F., Luttini, E., Sandri, D., 2023. Idiosyncratic shocks and aggregate fluctuations in an emerging market. J. Dev. Econ. 160, 102949.

Guichon, J., Fatès, N., Contassot-Vivier, S., Amato, M., et al., 2024. Properties of B2B invoice graphs and detection of structures. In: Cherifi, H., Rocha, L.M., Cherifi, C., Donduran, M. (Eds.), Complex Networks & Their Applications XII. Springer Nature Switzerland, Cham. Vol. 1143, pp. 444–455. Series Title: Studies in Computational Intelligence. https://doi.org/10.1007/978-3-031-53472-0_37

Herskovic, B., Kelly, B., Lustig, H., Van Nieuwerburgh, S., 2020. Firm volatility in granular networks. J. Polit. Econ. 128 (11), 4097–4162.
Huneeus, F., 2020. Production network dynamics and the propagation of shocks. Draft https://www.fedehuneeus.com/s/JMP_FHL.pdf.
Ialongo, L.N., de Valk, C., Marchese, E., Jansen, F., Zmarrou, H., Squartini, T., Garlaschelli, D., 2022. Reconstructing firm-level interactions in the dutch input–output network from production constraints. Sci. Rep. 12 (1), 1–12.
Inoue, H., 2016. Controllability analyses on firm networks based on comprehensive data. arXiv preprint arXiv: 1604.01322.
Kichikawa, Y., Iino, T., Iyetomi, H., Inoue, H., 2019. Hierarchical and Circular Flow Structure of the Interfirm Transaction Network in Japan. Discussion papers 19063. Research Institute of Economy, Trade and Industry (RIETI). https://ideas.repec.org/p/eti/dpaper/19063.html.
Kikkawa, A.K., Magerman, G., Dhyne, E., 2019. Imperfect Competition in Firm-to-Firm Trade. http://dx.doi.org/10.2139/ssrn.3389836.
König, M.D., Levchenko, A., Rogers, T., Zilibotti, F., 2022. Aggregate fluctuations in adaptive production networks. Proc. Natl. Acad. Sci. 119 (38), e2203730119.
Konno, T., 2009. Network structure of japanese firms. scale-free, hierarchy, and degree correlation: analysis from 800,000 firms. Economics 3, 1–13.
Krichene, H., Fujiwara, Y., Chakraborty, A., Arata, Y., Inoue, H., Terai, M., 2019. The emergence of properties of the Japanese production network: how do listed firms choose their partners? Soc. Networks 59, 1–9.
Kumar, A., Chakrabarti, A.S., Chakraborti, A., Nandi, T., 2021. Distress propagation on production networks: coarse-graining and modularity of linkages. Physica A 568, 125714.
Lee, Y.J., Kim, S.D., Hong, J.P., Nbsp, H. G. C. a. S. M.Y., 2016. Industrial network analysis using inter-firm transaction data. Indian J. Sci. Technol. 9 (26), 1–9. https://doi.org/10.17485/ijst/2016/v9i26/97308
Leontief, W., 1928. The Economy as a Circular Flow. (in German), Archiv für Sozialwissenschaft und Sozialpolitik, 60, 577-623. Reprinted and translated into English in *Structural Change and Economic Dyamics*, 2 (1991), 181-212.
Lu, Y., Ogura, Y., Todo, Y., Zhu, L., 2017. Supply Chain Disruptions and Trade Credit. RIETI Discussion Paper Series 17-E-054. Research Institute of Economy, Trade and Industry (RIETI). https://ideas.repec.org/p/eti/dpaper/17054.html.
Luo, J., Baldwin, C.Y., Whitney, D.E., Magee, C.L., 2012. The architecture of transaction networks: a comparative analysis of hierarchy in two sectors. Industr. Corpor. Change 21 (6), 1307–1335.
Luo, J., Magee, C.L., 2011. Detecting evolving patterns of self-organizing networks by flow hierarchy measurement. Complexity 16 (6), 53–61.
Luo, J., Whitney, D.E., 2015. Asymmetry in -degree and out-degree distributions of large-scale industrial networks. arXiv preprint arXiv: 1507.04507 .
Magerman, G., De Bruyne, K., Dhyne, E., Van Hove, J., 2016. Heterogeneous firms and the micro origins of aggregate fluctuations. National Bank of Belgium Working Paper .
Mariani, M.S., Ren, Z.-M., Bascompte, J., Tessone, C.J., 2019. Nestedness in complex networks: observation, emergence, and implications. Phys. Rep. 813, 1–90.
Marray, K., 2025. Estimating spillover effects from sampled connections. http://dx.doi.org/10.2139/ssrn.5236701.
Masso, J., Vahter, P., 2021. Joining and Exiting the Value Chain of Multinationals and the Performance of Suppliers: Evidence from Inter-Firm Transaction Data. http://dx.doi.org/10.2139/ssrn.3989069.
Milgram, S., 1967. The small world problem. Psychol. Today 2 (1), 60–67.
Miller, R.E., Blair, P.D., 2009. Input-Output Analysis: Foundations and Extensions. Cambridge University Press.
Miranda-Pinto, J., 2021. Production network structure, service share, and aggregate volatility. Rev. Econ. Dyn. 39, 146–173.
Miranda-Pinto, J., Silva, A., Young, E.R., 2023. Business cycle asymmetry and input-output structure: the role of firm-to-firm networks. J. Monet. Econ. .
Mizuno, T., Souma, W., Watanabe, T., 2015. Buyer-supplier networks and aggregate volatility. In: The Economics of Interfirm Networks. Springer, pp. 15–37.
Morimoto, Y., 1970. On aggregation problems in input-output analysis. Rev. Econ. Stud. 37 (1), 119–126.
Mungo, L., Lafond, F., Astudillo-Estévez, P., Farmer, J.D., 2023. Reconstructing production networks using machine learning. J. Econ. Dyn. Control 148, 104607.
Nair, J., Wierman, A., Zwart, B., 2022. The fundamentals of heavy tails: Properties, emergence, and estimation. Vol. 53. Cambridge University Press.
Nations, U., 2010. System of National Accounts 2008. European Communities, International Monetary Fund, Organisation for Economic Co-operation and Development, United Nations, World Bank. https://doi.org/10.18356/4fa11624-en
Newman, M., 2018. Networks. Oxford University Press.
Newman, M. E.J., 2003. Mixing patterns in networks. Phys. Rev. E 67 (2), 026126.
Ohnishi, T., Takayasu, H., Takayasu, M., 2009. Hubs and authorities on japanese inter-firm network: characterization of nodes in very large directed networks. Progr. Theor. Phys. Supplement 179, 157–166.
Ohnishi, T., Takayasu, H., Takayasu, M., 2010. Network motifs in an inter-firm network. J. Econ. Inter. Coordinat. 5 (2), 171–180.
Panigrahi, P., 2023. Triad trade and small worlds of large spatial production networks. Working Paper Available at: https://piyushpanigrahi.github.io/TriadTrade_PiyushPanigrahi.pdf.
Pessina, L., 2020. Who writes the check to the government does matter: evidence from firm-to-firm links. In: Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association. Vol. 113. National Tax Association, pp. 1–40.
Peydró, J.L., Jiménez, G., Kenan, H., Moral-Benito, E., Vega-Redondo, F., 2020. Production and financial networks in interplay: crisis evidence from supplier-customer and credit registers. CEPR Discussion Paper .
Saito, Y.U., Watanabe, T., Iwamura, M., 2007. Do larger firms have more interfirm relationships? Physica A 383 (1), 158–163.
Sang, J. J. a. C.J., 2016. Structure and characteristics of transaction network in Korean non-financial industries. Indian J. Sci. Technol. 9 (26), 1–11.
Serafino, M., Cimini, G., Maritan, A., Rinaldo, A., Suweis, S., Banavar, J.R., Caldarelli, G., 2021. True scale-free networks hidden by finite size effects. Proc. Natl. Acad. Sci. 118 (2).
Serrano, M.Á., Krioukov, D., Boguná, M., 2008. Self-similarity of complex networks and hidden metric spaces. Phys. Rev. Lett. 100 (7), 078701.
Silva, T.C., Amancio, D.R., Tabak, B.M., 2020. Modeling economic networks with firm-to-firm wire transfers. arXiv preprint arXiv: 2001.06889.
Sornette, D., 2006. Critical Phenomena in Natural Sciences: Chaos, Fractals, Selforganization and Disorder: Concepts and Tools. Springer Science & Business Media.
Spray, J., 2017. Reorganise, Replace or Expand? The role of the supply-chain in first-time exporting. Cambridge-INET Working Paper Series No: 2017/18. Institute for New Economic Thinking, Department of Economics, Cambridge University, Cambridge.
Spray, J., 2021. Search externalities in firm-to-firm trade. IMF Working paper 2108.
Spray, J., Wolf, S., 2018. Industries without Smokestacks: Industrialization in Africa Reconsidered. Oxford University Press, Oxford, UK.
Taschereau-Dumouchel, M., 2025. Cascades and fluctuations in an economy with an endogenous production network. Rev. Econ. Stud. , rdaf036.
de la Torre, S.R., Kalda, J., Kitt, R., Engelbrecht, J., 2016. On the topologic structure of economic complex networks: empirical evidence from large scale payment network of Estonia. Chaos Solitons Fractals 90, 18–27.
United Nations, 2018. Handbook on Supply, Use and Input-Output Tables with Extensions and Applications. United Nations, New York, US.
Voitalov, I., van der Hoorn, P., van der Hofstad, R., Krioukov, D., 2019. Scale-free networks well done. Phys. Rev. Res. 1 (3), 033034.
Vom Lehn, C., Winberry, T., 2022. The investment network, sectoral comovement, and the changing US business cycle. Q. J. Econ. 137 (1), 387–433.
Watanabe, H., Takayasu, H., Takayasu, M., 2013. Relations between allometric scalings and fluctuations in complex systems: the case of Japanese firms. Physica A 392 (4), 741–756.
Watts, D.J., Strogatz, S.H., 1998. Collective dynamics of ’Small-world’ networks. Nature 393 (6684), 440–442.
Wiedemann, V., Kirui, B.K., Khandelwal, V., Chacha, P.W., 2024. Spatial inequality and informality in Kenya’s firm network. World Bank Policy Research Working Paper 10932.
Wu, D., 2016a. Shock spillover and financial response in supply chain networks: Evidence from firm-level data. Unpublished working paper. Available at: https://www.semanticscholar.org/paper/Shock-Spillover-and-Financial-Response-in-Supply-%3A-Wu/9bf39f16bf5aa2f663ff196dfeae6a5a8848af3c.

Wu, J., 2016b. Firm Performance and Risk in Supply Chain Networks. PhD thesis. The University of Chicago.
Wu, J., Birge, J.R., 2014. Supply chain network structure and firm returns. http://dx.doi.org/10.2139/ssrn.2385217.
Zhigang, L.I., Shang-Jin, W., Hongyong, Z., 2018. Production Chains, Exchange Rate Shocks, and Firm Performance. Discussion papers 18058. Research Institute of Economy, Trade and Industry (RIETI). https://ideas.repec.org/p/eti/dpaper/18058.html.