

Herkenhoff, Philipp; Krautheim, Sebastian; Sauré, Philip

Working Paper

A Simple Model of Buyer-Seller Networks in International Trade

CESifo Working Paper, No. 9124

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Herkenhoff, Philipp; Krautheim, Sebastian; Sauré, Philip (2021) : A Simple Model of Buyer-Seller Networks in International Trade, CESifo Working Paper, No. 9124, Center for Economic Studies and Ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/236666>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

A Simple Model of Buyer-Seller Networks in International Trade

Philipp Herkenhoff, Sebastian Krautheim, Philip Sauré

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

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

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

Poschingerstr. 5, 81679 Munich, Germany

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

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

A Simple Model of Buyer-Seller Networks in International Trade

Abstract

The recent literature on firm-to-firm trade has documented salient empirical regularities of the buyer-seller network. We propose a simplistic re-interpretation of the classical Krugman (1980) model that accounts for surprisingly many of the empirical regularities. This re-interpretation relies on randomized bundling of Krugman-varieties into heterogeneous firms, economically neutral ‘sales units’ that import foreign varieties but belong to local firms, and a statistical reporting threshold that applies to firm-to-firm transactions. We argue that our model provides an important benchmark for the assessment of theoretical models that aim to identify the determinants of firm-to-firm networks.

JEL-Codes: F100, F120, F140.

Keywords: firm-to-firm, buyer-seller, trade, network, random matching.

Philipp Herkenhoff
University of Mainz / Germany
philipp.herkenhoff@uni-mainz.de

Sebastian Krautheim
University of Passau / Germany
sebastian.krautheim@eui.eu

Philip Sauré
University of Mainz / Germany
philip.saure@uni-mainz.de

July 9, 2021

We would like to thank Roc Armenter, Facundo Alborno, Ariel Burstein, Alejandro Cunat, Andreas Fischer, Carsten Eckel, Peter Egger, Peter Eppinger, Philipp Harms, Miriam Kohl, Wilhelm Kohler, Andrea Lassmann, Andrei Levchenko, Valeria Merlo, Gernot Müller, Gianmarco Ottaviano, Frank Stähler, Thierry Verdier and seminar participants at the University of Mainz (Sept. 2020) and the University of Tübingen (Feb. 2021) for helpful comments and discussions.

1 Introduction

International trade takes place through a complex network of transactions between buyers and sellers. Firm-to-firm networks shape key firm characteristics and macroeconomic aggregates alike by determining, among others, the geography of international trade (Chaney, 2014, and Chaney, 2018), the labor market reactions to international trade (Eaton et al., 2019), the distribution of firm size (Bernard et al., 2021, Panigrahi, 2021), the business cycle (Lim, 2018, Huneus, 2018, Taschereau-Dumouchel, 2020), and economic growth (Acemoglu and Azar, 2020).

In recent years, a dynamic literature on firm-to-firm trade has uncovered a set of salient empirical regularities concerning the number of connections between firms and the distribution of trade values among these connections.¹ In parallel, the literature has proposed intriguing theoretical models to study the factors shaping firm-to-firm connections.²

In this paper we argue that caution is warranted when evaluating theoretical models based on the salient empirical patterns of the firm-to-firm network, because some of the most prominent patterns emerge mechanically from a simplistic model. We develop such a simplistic model of randomized firm-to-firm trade to ask: which stylized facts of firm-to-firm trade networks are generated by purely stochastic firm-to-firm trade? It turns out that the answer to this question is: surprisingly many.

To develop our benchmark model, we go ‘back to square one’ in terms of modelling firms in international trade. For us (and probably most trade economists), this means that we turn to the model of Krugman (1980), where each country produces a mass of identical varieties, each of which, being subject to variable transport costs only, is shipped to and consumed in all destinations. We suggest a re-interpretation of this canonical model along with a re-interpretation of the notion of a ‘firm’, thus providing a new perspective on the informational content of firm-level trade data. Importantly, our re-interpretation leaves unchanged all economic activities of Krugman (1980), which are thus fully characterized by the classical (multi-country) setup.

Our interpretation deviates from the canonical view on Krugman (1980) in three central elements. First, we make a semantic distinction between a ‘plant’, a ‘sales unit’ and a ‘firm’, which are not distinguished in Krugman (1980). In our terminology, ‘plants’ are homogeneous producers (and exporters) of a single differentiated variety –

¹See Bernard and Moxnes (2018) for a survey. This literature has expanded on the previous focus of the analysis on either the buyer *or* the seller (e.g., Antràs et al., 2017, Fally and Hillberry, 2018, Bernard et al., 2018a) to placing firm-to-firm connections at center stage.

²These factors include search frictions (Chaney, 2014, Eaton et al., 2016), information frictions (Eaton et al., 2014, Benguria, 2015, Monarch and Schmidt-Eisenlohr, 2017, Chaney, 2018), relationship-specific fixed costs (e.g., Bernard et al., 2018c, and Lim, 2018), switching cost (Monarch, 2016), and heterogeneity in consumers’ valuation of varieties (Carballo et al., 2018).

i.e., they are the entities called ‘firms’ in Krugman (1980). To sell a variety to consumers in a given country, a unique local ‘sales unit’ is required. These sales units, however, are economically neutral: they operate under perfect competition with zero marginal cost, so that the Krugman (1980) setup remains effectively unchanged. The two entities – plants and sales units – are the economic agents in our model. These plants and sales units are bundled into legal entities called ‘firms’, so that observations at the ‘firm’-level in the data reflect the aggregate of these bundles’ activities.³

As the second element of our re-interpretation of the Krugman (1980) model, we introduce randomization in the bundling process. Specifically, plants and sales units are randomly assigned to ‘firms’ and the size of these firms, defined by the mass of plants and sales units they contain, is drawn from an exogenous (Pareto) distribution.⁴ The random nature of bundling implies that the link of a given plant with its foreign sales unit generates randomized trade relations between firms in the exporting and the importing country.⁵

As the third and final element of our re-interpretation of the Krugman (1980) model, we assume that firm-to-firm transactions of cross-border trade enter the model’s trade statistics only if their value exceeds a *reporting threshold*. This assumption captures a common feature of data collection by customs authorities and implies that firm-to-firm transactions are censored in the trade statistics.⁶

These three elements of our re-interpretation do not affect economic decisions. They have, however, implications for our interpretation of trade statistics. In particular, the distinction between the legal entity ‘firm’ and the economic entities active within its boundaries implies that we have to re-think the link between actual economic activity (the plant level) and the way it is represented in the data (at the firm level). Consider, e.g., firm A which has twice the mass of plants of firm B. While each plant produces and exports its own variety of the consumption good, in the data we observe two firms of different sizes. Moreover, exports of the smaller firm B are less likely to be recorded in the trade statistics than exports of firm A, simply because exports of the former are less likely to pass the reporting threshold.⁷

³The labels of ‘firms’ and their subsidiaries reflect our aim to connect our approach to the empirical work with firm-level data – see Section 4.3.1 for a discussion on this point. Throughout our description of the setup, we interpret the plant as a producer of the final product, and the sales unit as its distributing unit. We do so to adopt the standard of the literature, see, e.g., Chaney (2014), and thereby remain close to Krugman (1980). Section 4.3.2 offers generalizations of this interpretation, including trade in intermediate goods.

⁴This assumption implies that firms in our model sell multiple varieties. We distinguish, however, between a variety and a product and discuss in Section 4.3.3 how the product margin can be explicitly analyzed using a slight extension of our framework.

⁵This approach is reminiscent of the balls-and-bins model of trade by Armenter and Koren (2014). Section 4.1 discusses in detail the similarities and differences between their approach and ours.

⁶We discuss the reporting threshold and its role for our results in detail in Section 4.1.

⁷We remain agnostic about the origins of the heterogeneity of firm size. For example, our framework

Our paper’s main contribution is to show that a set of salient empirical patterns of international buyer-seller data emerge from our re-interpreted multi-country Krugman (1980) model. For the selection of these empirical patterns, we draw on recent studies by Blum et al. (2010) (who document firm-level trade patterns between Argentina and Chile), Bernard et al. (2018c) (for Norway), Bernard et al. (2018b) (for Colombia), Carballo et al. (2018) (for Costa Rica, Ecuador, and Uruguay) and Bernard and Moxnes (2018). To replicate – and in some cases expand on – their key stylized facts on buyer-seller networks in international trade, we use Colombian import data. We find that our simplistic model replicates a number of the most robust and salient empirical patterns. First, the buyer margin and the seller margin are driven by gravity forces: the number of exporters an importer connects to (seller margin) as well as the number of importers an exporter connects to (buyer margin) are increasing in the size of the partner country and decreasing in bilateral distance. Second, a firm’s aggregate trade volume is proportional to the number of its partners: the firm’s import value increases in the number of foreign suppliers with unit elasticity; by symmetry, the same holds for a firm’s export value and the number of its foreign buyers. Third, the share of local firms with more foreign sellers (buyers) decreases in the number of a firm’s sellers (buyers) with unit elasticity. For example, the share of local firms connected to more sellers than firm k is decreasing in firm k ’s number of sellers with unit elasticity. Fourth, large, well-connected firms dominate the firm-to-firm trade network: few firms have many partners and many firms have few partners. Fifth, negative assortative matching rules firm-to-firm connections: the better connected an importing (exporting) firm is, the less well-connected is its average buyer (seller). Sixth, the conditional sales distribution is stable: the size and the connectedness of a firm do not affect its sales to its median (or any other percentile) buyer. Seventh, firms tend to follow a hierarchical pecking order: an importer’s set of exporters is a subset of any larger importer’s set of exporters.⁸

At the current juncture, as an emerging literature analyzes increasingly disaggregated trade data and seeks to formulate intricate theories of firm-to-firm connections, it is appropriate to partition the empirical patterns into two sets: a first set with which conventional models do come to terms with and another set these models do not. Our simplistic model with random firm-to-firm matching delivers such a partition and helps

is compatible with an interpretation of Melitz-type productivity differences. Thus, we implicitly allow the size of a firm to be determined by factors other than the usual differences in productivity (as in Bustos, 2011) or quality (as in Fieler et al., 2018). Differences in firm sizes may also reflect (historic) differences in wealth across individuals which, in times of highly imperfect capital and financial markets mapped into heterogeneous firm sizes (as in Bonfiglioli et al., 2019) or emerge from heterogeneous management quality (Bloom et al., 2020) and signalling motives (e.g., Amaldoss and Jain, 2015). Those factors, to the extent that they can be incorporated in the Krugman setup, are immediately applicable to our standard framework.

⁸In addition to these core results of our baseline model, Section 4.2 discusses variations of our model that allow to capture other regularities of the data.

to answer the original question of Armenter and Koren (2014): which facts are useful for testing new theories? Since many of the salient empirical patterns cited in the literature emerge from our simplistic model, they should *not* be interpreted as empirical support for sophisticated microeconomic modelling of firm-to-firm interactions.⁹

We argue that our paper’s message is less destructive than it may seem, as it indicates which empirical patterns *can* help researchers to select modelling approaches for firm-level decisions. Generally, any significant deviation from the predictions generated by random matching of buyers and sellers suggests that firms engage in directed economic activity.¹⁰ This observation applies, in particular, to the dynamics of firm trade that do not only mirror the growth rates of firms but, instead, reflect strategic firm decisions. Also, complementarities of firms’ input suppliers (as in Halpern et al., 2015) or geographic clustering of export destinations due to the structure of search costs (as in Chaney, 2014) point at patterns that go beyond those predicted by randomized firm-to-firm trade.¹¹ With its underlying Krugman (1980) model, our model involves genuine economic decision making and provides a starting point for re-introduction of firms with economic content at both sides of the buyer-seller relationship.

Our paper connects to various literatures. Most importantly, we contribute to the growing literature that documents mutually consistent stylized facts of firm-to-firm networks in international trade, using different datasets from various countries. Prominent studies in this realm are Blum et al. (2009) (for Chilean exporters and Colombian importers), Blum et al. (2010) (Chilean importers and Argentinean exporters), Monarch and Schmidt-Eisenlohr (2017) (U.S. customs data identifying Chinese counterparts), and Carballo et al. (2018) (exporters from Costa Rica, Ecuador and Uruguay and their buyers). Bernard et al. (2018b) use Colombian customs data at the transaction-level identifying Colombian importers and foreign exporters.¹² Bernard et al. (2018c) analyze Norwegian exports in a transaction-level dataset that identifies, Norwegian exporters and their international buyers. This body of work produces a strikingly consistent set of

⁹For example, Bernard et al. (2018c) use a set of facts to motivate their model of trade in intermediate goods with heterogeneous buyers and sellers and match-specific fixed costs. We argue that successfully matching their set of facts does not improve on our simplistic model and thus does not constitute evidence in favor of (nor against) match-specific fixed costs. At the same time, we emphasize that our model is a merely benchmark and should not be read as a realistic description of economic reality.

¹⁰Examples of such deviations relate to the predicted unit elasticities in the patterns presented in Sections 3.2 and 3.3 or the invariant sales distribution from Section 3.6. Section 4.2 gives a more comprehensive list of the testable implications.

¹¹Further promising dimensions include firms’ engagement in directed search for upstream suppliers and downstream buyers along complex value chains (see, e.g., Bernard et al., 2021). Firms’ decisions regarding pricing and product composition may turn out to be relevant factors for the microeconomic analysis of the firm and firm-to-firm interactions. They are muted in Krugman- or Melitz-type settings as long as firms are infinitesimally small, but materialize prominently in Atkeson and Burstein, 2008, Eckel and Neary (2010), Arkolakis et al. (2010), Blum et al. (2019), or Auer et al. (2018).

¹²Their dataset is similar to ours, and we discuss the differences in Section A.1.

empirical patterns, despite the apparent differences of the respective economies in terms of economic development, production structure and per-capita income (for example Colombia in Bernard et al., 2018b, and Norway in Bernard et al., 2018c). Bridging the literature of domestic and international firm networks, Bernard et al. (2021) study the universe of buyer-seller relationships in Belgium and find that the number of customers determines heterogeneity in firm size.¹³

With the central element of randomized firm-to-firm connections, our paper is tightly connected to the ‘balls-and-bins’ paper by Armenter and Koren (2014), who show that the gravity patterns of international trade, including an operating extensive margin, emerge mechanically from a purely stochastic model of ‘balls’ falling into ‘bins’.¹⁴ Conceptually, we apply their approach of randomized trade relations to the firm-to-firm trade network. Instead of excluding optimizing agents from the model, however, we strip only the legal entity ‘firm’ of its economic content, while keeping the underlying economic transactions of a fully microfounded model.¹⁵ Parallel and independent research by Bernard and Zi (2021) complements our analysis by applying the purely stochastic approach with a discrete number of balls and bins to firm-to-firm trade. Under general firm size distributions, the authors investigate the information content of data under a large class of statistical transformations and at different aggregation levels. On the one hand, this work confirms the generality of our main message and shows that our basic insights do not hinge on our distributional assumptions and the particularities of the Krugman (1980) model. On the other hand, comparing the two complementary approaches highlights that the specific structure of our framework yields a long list of sharp, testable predictions that relate directly to the model’s structural parameters.¹⁶ The approach of randomized trade connections is also reminiscent of Chaney (2018), who spells out theory-independent sufficient conditions for the well-known effect of geographical distance on trade to emerge.

¹³There is a closely connected literature on purely domestic firm-to-firm networks that focuses on domestic production networks. Tintelnot et al. (2021) also use Belgian firm-to-firm data to document that, while only few firms export and import directly, the majority of firms does so indirectly through domestic linkages. The study also documents negative assortative matching, which is prevalent in datasets on international trade connections (in particular, Bernard et al., 2018c, and Bernard et al., 2018b). Using a detailed dataset for Japan, Bernard et al. (2019) document a strong link between a firm’s size and the number of its suppliers and highlight the role of within-country geographic distance for the number of firm connections. The same data are used by Carvalho et al. (2021) to study the effects of the 2011 Japanese earthquake on supply chains.

¹⁴This approach has inspired a number of studies. Eaton et al. (2013) pursue a similar approach to introduce firms with positive mass that impact economic aggregates. Head et al. (2017) develop a stochastic benchmark for trade involving the consumer side. Within the literature on firm-to-firm connections, Bernard et al. (2018c) argue that the balls and bins approach does not generate the salient empirical patterns of firm-to-firm trade. We discuss this point in detail in Section 4.1.

¹⁵Section 4.1 further discusses how our approach relates to the purely stochastic balls-and-bins approach.

¹⁶See also our discussion in Section 4.1.

Our paper also speaks to several closely related literatures, without explicitly including the corresponding modelling features. One of these literatures deals with the scope and boundary of the firm and analyzes the determinants of endogenous firm heterogeneity, such as firms’ product scope (e.g., in Eckel and Neary, 2010, and Bernard et al., 2011) and endogenous technology choice (as, e.g., in Lileeva and Trefler, 2010, and Bustos, 2011). Studies that are especially close to the literature on firm networks are those analyzing trade in intermediate inputs among firms (Amiti and Konings, 2007, Goldberg et al., 2010, Halpern et al., 2015, and Fieler et al., 2018).¹⁷ A further literature studies different aspects of firm dynamics: one-sided firm trade (Albornoz et al., 2012, and Ruhl and Willis, 2017) or firm-to-firm connections (see Blum et al., 2010, and Gimenez-Perales, 2021, in the context of international trade network and Carvalho and Voigtländer, 2014, in the domestic one). Chaney (2014) proposes a model of dynamic network formation between buyers and sellers in international trade. The model endogenizes firms’ trade costs, complementing and improving on previous work that investigates the determinants of (fixed) costs of international trade. Related studies focus on market entry costs (Arkolakis, 2010) and the role of financial frictions as impediments to trade (studied in Manova, 2013, Chaney, 2016, and Bonfiglioli et al., 2019). A strand of the literature with a marcoeconomic approach analyzes the sources of aggregate fluctuations through the input-output network (Acemoglu et al., 2012, Di Giovanni et al., 2014, and Lim, 2018). Other studies investigate the role of substitution elasticities among single products and their heterogeneity across firms (see Oberfield, 2018, but also Halpern et al., 2015, and Gimenez-Perales, 2021, which link back to Chaney, 2008).

Some of the listed dimensions may be readily included in extensions of our framework.¹⁸ Other features, such as those determining network formation or financial and other frictions, may ultimately produce predictions that go genuinely beyond the predictions of our randomized model.¹⁹

The remainder of the paper is structured as follows: Section 2 outlines our baseline model, Section 3 uses Colombian transaction-level import data to highlight seven major stylized facts and shows one by one that our model matches them all, Section 4 discusses

¹⁷The related issue of multinational firms and their organizational structure, analyzed in the literature based on Antràs (2003) and Antràs and Helpman (2004) and offshoring in general (Grossman and Rossi-Hansberg, 2008), is equally kept outside of our model.

¹⁸Our our model can be readily expanded to produce empirical regularities in the product dimension – see Section 4.3.3.

¹⁹One example could be third country effects in international trade, as identified in Chaney (2014). In a model of the dynamic formation of a cross-border buyer-seller network with search frictions, he uses French firm-level data to confirm that a firm exporting to country A is more likely to start exporting to country B if A is close to B , independently of the distance of B to the exporting country. Such effects are outside our model and can provide a starting point for identifying relevant implications of trade networks that go beyond the predictions of our randomized model.

the relation of our paper to the balls and bins framework of Armenter and Koren (2014) and points at directions for future research. Section 5 concludes.

2 A re-interpretation of Krugman (1980) with random matching

The aim of our analysis is to determine how far a simple re-interpretation of Krugman (1980) can go in rationalizing key patterns on the international trade network. We start by summarizing and interpreting the features of a multi-country version of Krugman (1980) that will be central for our analysis.

Heavily drawing on previous work with standard modelling framework, we only specify the parts of the model that are essential for our re-interpretation. Doing so, we leave the demand side of the model entirely unchanged. The supply-side, on the other hand, is subject to a re-interpretation that will shape our approach to the data but is inessential for economic outcomes.

2.1 Demand

Consumers obtain utility from the consumption of differentiated varieties of a final consumption good. Demand in country i for a variety ω originating from country j is

$$q_{ji}(\omega) = (p_{ji}(\omega)/P_i)^{-\sigma} Y_i / P_i, \quad (1)$$

where P_i is the ideal price index, Y_i total income of country i and σ stands for the constant substitution elasticity stemming from consumers' utility. Defining Ω_j as the set of varieties produced in j and sold in i , with homogeneous production plants, $p_{ji}(\omega) = p_{ji} \forall \omega \in \Omega_j$ is the local retail price in country i of varieties produced in country j including trade costs. Standard profit-maximization implies

$$p_{ji} = \frac{\sigma}{\sigma - 1} \tau_{ji} w_j, \quad (2)$$

where τ_{ji} are bilateral variable trade costs, w_j is the wage in country j and marginal unit labor requirements are normalized to one.

2.2 Supply – the ‘firm’ under the microscope

We now turn to our re-interpretation of the Krugman (1980) model and its link to the firm-to-firm trade data. First, we make an important semantic distinction between a ‘plant’, a ‘sales unit’ and a ‘firm’, which are not distinguished in Krugman (1980).

2.2.1 Plants

A *plant* produces a single variety of the differentiated final consumption good. It chooses quantities, prices and export destinations. All plants share the same productivity level and make positive operating profits, which just cover entry costs. So far, our plant is akin to a ‘firm’ in Krugman (1980). We do however, explicitly model an activity that is implicitly carried out in Krugman (1980): distribution to the final consumer.

2.2.2 Sales units

Distribution of varieties to consumers is carried out by *sales units*. When a plant seeks to sell its variety in a given market, it must do so through a local sales unit. Sales units can distribute goods to consumers at zero cost and operate under perfect competition. These assumptions keep the activity of sales units economically neutral and keep our model observationally equivalent to Krugman (1980).

In the absence of fixed cost of exporting, all plants in country j export to all destinations. For each destination i , the plant is randomly matched to a unique sales unit that is the exclusive distributor of its variety in that market.

2.2.3 Firms

We define a firm as a collection or a bundle of local plants and sales units.²⁰ While the economic activity takes place at the level of plants and sales units, the ‘firm’ is the legal entity that provides an umbrella for its economic activities.

Two features of our firms deserve special attention – heterogeneous firm size and the assignment of plants and sales units to firms.

Firm heterogeneity. A firm’s size is determined by the mass of varieties produced within its boundaries.²¹ The firm’s size follows a distribution, which we take as exogenous and specify by the cdf $F(\mu)$ and the according pdf $f(\mu)$. We assume that firm size follows identical distributions in all countries and consequently do not index F or f by j or i . We also remain agnostic about the sources of firm heterogeneity. Throughout the paper, we will refer to a firm in country j with mass μ_j as “firm μ_j ”.

For our core analytical exercises, we work with the most commonly used firm size distribution, the Pareto distribution

$$F(\mu) = 1 - (\mu/\underline{\mu})^{-\theta} \tag{3}$$

²⁰This means that we rule out multinational firms.

²¹This setup is a notional change of the Krugman (1980) model only. As long as firms are negligibly small relative to the economy, no incentives to intervene in production or pricing arise. The setup obviously coincides with the original interpretation when each firm has one plant only.

which has the pdf $f(\mu) = \theta \mu^{-\theta-1} (\underline{\mu})^\theta$, the expected value $\underline{\mu}\theta/(\theta - 1)$ and the χ^{th} percentile, defined through $F(\mu) = \chi$,

$$\mu_\chi = (1 - \chi)^{-1/\theta} \underline{\mu}. \quad (4)$$

Random assignment of plants and sales units to firms. We assume that the total mass of plants and sales units within a country is assigned to the domestic firms according to *identical and independent randomization*.²² The random bundling, in connection between domestic plants and foreign sales units generate random cross-border connections between firms. Since each firm comprises a mass of plants, it connects to all foreign firms by the law of large numbers. However, larger firms mechanically have “broader” or more intensive connections, as they comprise more plants and sales units within their boundaries.²³

2.3 ‘Firm’-to-‘firm’ trade observed in the data

‘Firm’-to-‘firm’ connections. Each pair of domestic plant and foreign sales unit has the same ex-ante probability to be matched. This assumption implies that the probability of a plant to be matched to a specific foreign firm is proportional to the number of sales units within this foreign firm and thus to this firm’s size. Conversely, the probability of a foreign sales unit to be matched to a domestic firm is proportional to that firm’s size. Applying finally the law of large numbers, the mass of plants exporting from one firm to the other is proportional to the product of both firms’ sizes.

Formally, the assumption that importing and exporting firms are randomly matched is reflected by independence of the joint distribution, which describes the probability that a producing firm of size μ_j and a selling firm of size μ_i are matched. This joint

²²The random nature of this bundling process reflects the fact that the Krugman-varieties in a given economy are all identical. However, the assignment of sales units is equally randomized and independent, so that the probability of a plant and a sales unit being linked is independent of how many other plants and sales units of the same firms are linked.

²³We superimpose a structure on the Krugman (1980) model, in which plants and sales units are randomly matched and randomly bundled into firms. This constitutes a similarity of our approach to the purely stochastic balls and bins approach along the lines of Armenter and Koren (2014), which we discuss in Section 4.1. The intuition for the connection is as follows. The ‘balls’ (the plant-sales unit connections) originate from ‘exporter-bins’ and ‘fall’ into ‘importer-bins’. Larger exporter-bins launch more balls, which in turn are more likely to fall into larger importer-bins. On the one hand, our model is a continuous version of that idea, since we deal with a continuum of firms (bins) and a continuum of plant-sales unit connections (balls) (instead of a discrete number as in Armenter and Koren, 2014); on the other hand, we build a bridge between their purely stochastic approach and standard economic modelling by “microfounding” the ‘balls’ as the exports of Krugman (1980) plants. In Section 4.1 we discuss the relation between the continuous and the discrete case as well as how sparsity (the fact that in the data only a small fraction of possible links is observed) emerges in our continuous model.

distribution (cdf) is independent²⁴

$$F(\mu_j, \mu_i) = F(\mu_j)F(\mu_i) \quad (5)$$

and the according pdf is

$$f(\mu_j, \mu_i) = f(\mu_j)f(\mu_i). \quad (6)$$

In the absence of fixed costs of exporting, all varieties in the Krugman (1980) model are exported, and demand is given in (1) and all varieties are exported to all countries. Since an exporter exports all of its μ_j varieties to firms in destination i , the subset of its varieties sold to a specific firm μ_i is $\mu_j\mu_i/(E(\mu)N_i)$. Combined with demand in equation (1) and our assumption $E(\mu)N_i = L_i$ above, the value of firm μ_j 's exports to firm μ_i is then:

$$X(\mu_j, \mu_i) = \frac{\mu_i\mu_j}{L_i}(p_{ji}/P_i)^{1-\sigma}Y_i. \quad (7)$$

Reporting threshold. The prediction that all varieties are exported to all markets is clearly at odds with the data on firm-level trade but stems from the fact that we tie our hands to the Krugman (1980) model. Instead of moving away from Krugman (1980), we turn to a possible explanation of the discrepancy between model and data that concerns the data collection process.

Specifically, we postulate a *reporting threshold*, \bar{t} , for firm-to-firm trade and assume that firm-to-firm transactions are recorded only if the monetary value of all varieties sold by an exporting firm μ_j to an importing firm μ_i is above \bar{t} , i.e. $X(\mu_j, \mu_i) > \bar{t}$.²⁵ Firm μ_i is thus registered as a buyer of firm μ_j and firm μ_j is registered as a seller to firm μ_i if and only if

$$\mu_i\mu_j > \bar{\mu}_{ji} := \frac{\bar{t}}{(p_{ji}/P_i)^{1-\sigma}Y_i/L_i}. \quad (8)$$

²⁴Note that the distributions of the importer and the exporter size are identical and equal to $F(\mu)$. This is a consequence of two assumptions: first, the distribution of firm size is identical for all countries and second, plants and sales units are randomly bundled to firms. The second assumption, in combination with the law of large numbers, implies that the mass of (foreign) sales units over the mass of plants is identical for all firms within a country. Of course, the two assumptions may be relaxed.

²⁵It is common practice of customs administrations to approximate aggregate trade value of small transactions and collect little underlying information on the trading parties. The United Nations (2004), Chapter 3.5, Paragraph 69 (“Reporting Threshold and Retention of Records”) specifies that goods “...can be declared in less detail or be made exempt from reporting requirements [...] when the value (or quantity) is below a certain customs-defined threshold...” and that, further, “[c]ompilers may also establish a threshold for statistical purposes, i.e., set a value below which transactions may not be processed and included in the detailed trade statistics, or may be included in the trade statistics based on a sampling approach.” The manual also mentions a specific example, stating that “in the United States, most import transactions valued at less than USD 1,500 may be reported ‘informally’, with only minimal information report”.

The threshold in (8) is decreasing in the trade-promoting variables, i.e., the wage (expenditure) of the destination $w_i = Y_i/L_i$ and its price index P_i but increasing in the trade-impeding variables, i.e., the wage in the source country w_j and the iceberg trade costs (compare (2)).

Connections-extensive margin. With all important assumptions in place, Figure 1 visualizes our model setup. Firm μ_j on the horizontal axis represents, say, an exporting firm and μ_i on the vertical axis represents an importing firm. $\underline{\mu}$ on each axis represents the minimum value of the Pareto distribution of firm sizes. Along the hyperbola, we have that $\mu_i\mu_j = \bar{\mu}_{ji}$ so that only trade between firm sizes located above and to the right of the line are recorded. For an importing firm of size μ_i , only connections with exporters larger than $\bar{\mu}_{ji}/\mu_i$ are recorded. As the importer's size increases, it thus reaches deeper into the pool of exporters, expanding along the connections-extensive margin. For the indicated low levels of the threshold $\bar{\mu}_{ji}$, an importing firm with sizes above μ_0 connects with all potential exporters including smallest firms in the market. For all importers with sizes $\mu_i < \mu_0$, however, the connections-extensive margin is active. That means, any change in $\bar{\mu}_{ji}$ (induced by changes in trade costs or a market conditions of the partner country) affect the mass of exporters this firm connects to.

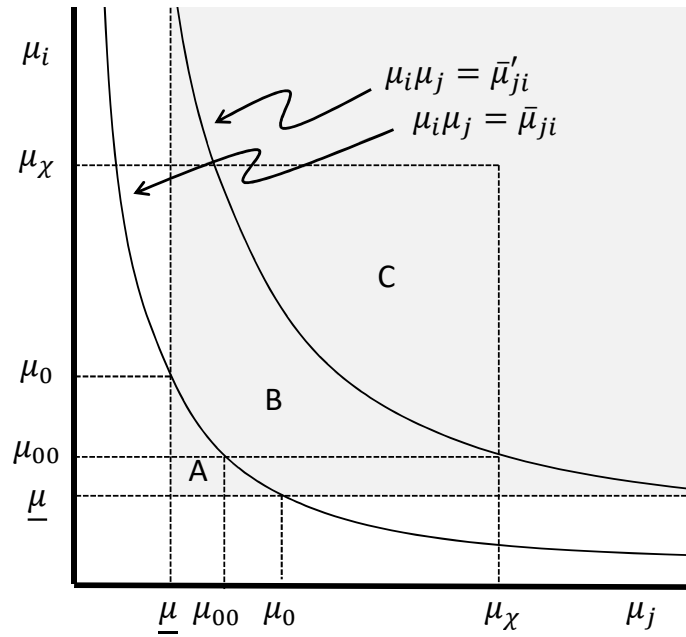
Empirically, the active connections-extensive margin seems to be the relevant case. Based on matched importer-exporter data from Chile and Colombia, Blum et al. (2009) show in Table 3.1 of their paper that in 2006, there was a total of 823 Colombian firms importing from Chile. However, the Chilean exporter at the 99th percentile sold only to 19 importers and the largest Chilean exporter sold to 30 Colombian firms. Based on this evidence, we focus on the case featuring the active connections-extensive margin when we present our results in Section 3.²⁶

3 The Data through the Lens of the Model

In this section we assess which of the salient regularities of international buyer-seller networks highlighted in the literature can be captured by our simplistic re-interpretation of the Krugman (1980) model. Our finding is: surprisingly many. Specifically, we provide a list of seven empirical regularities, which may seem intricate but mechanically emerge from our setup.

²⁶A firm-size distribution with no positive lower bound or a finite upper bound would readily generate the feature that no firm connects to all potential partners. We explore an alternative version of our model with a truncated Pareto distribution (as used in Helpman et al., 2008) and show that all the results presented in this section continue to hold in approximation. The results are available upon request.

Figure 1: Range of Recorded Connections and Large Firms – Pareto Distribution



Note: The grey area represents the set of all possible and thus of all active firm connections. The hyperbola defined by $\mu_i \mu_j = \bar{\mu}_{ji}$ partitions this set into those connections that have low trade volume and remain unrecorded in trade statistics (lower left: set A) and those that have low trade volume and are recorded in trade statistics (upper right: union of B and C). As the cutoff value ($\bar{\mu}_{ji}$) increases from $\bar{\mu}_{ji}$ to $\bar{\mu}'_{ji}$, the hyperbola shifts to the upper right and more trade connections remain unrecorded (A and B).

We have argued in Section 2.3 that the case with an active connections-margin, i.e., where $\bar{\mu}_{ji}/\mu_j > \underline{\mu}$, is the empirically relevant one. Accordingly, our results are based on this case. Throughout the section, whenever we refer to ‘connections’ or ‘trade flows’, we mean the recorded ones, i.e., those that are above our postulated reporting threshold.

While our main point is a theoretical one, we use Colombian transaction-level import data to replicate – and in some cases expand on – several regularities of the firm-to-firm trade network, which previous literature has found for different countries. We describe the data in Appendix A.1.²⁷

3.1 The Firm Margin and Gravity

Trade data exhibit a strong link between a firm’s number of firm-connections in a foreign country with size and distance of that partner country. In particular, Bernard et al. (2018b) show in their Table 5 that the number of an importer’s connections as well as the value of its imports from each foreign firm correlate with the established gravity variables (i.e., the economic size of the partner country and geographic distance) in the usual way. Bernard et al. (2018c) find the same pattern for Norwegian exporters (Table 5 in their online appendix).

Based on our data, we replicate these patterns for Colombian importing firms in Table 1, assessing their number of connections and total imports per partner country. Columns (1) - (3) document that the classical gravity variables, i.e., exporter GDP and bilateral distance, show a significant conditional correlation with the number of connections with usual sign and about the magnitude expected from the literature.²⁸

Turning to our model, we write the mass of firms in country j recorded as an exporter (or seller) to a firm of size μ_i as

$$S(j, \mu_i) = N_j \int_{\bar{\mu}_{ji}/\mu_i}^{\infty} f(\mu) d\mu = N_j \left(\frac{\mu_i \underline{\mu}}{\bar{\mu}_{ji}} \right)^\theta = N_j \left(\frac{\mu_i \underline{\mu} (p_{ji}/P_i)^{1-\sigma} Y_i}{\bar{t}L_i} \right)^\theta, \quad (9)$$

where the lower integration limit follows from equation (8) and indicates that connections of a firm of size μ_i are only recorded if the exporting firm in j supplies a sufficiently high value of its exports. Notice that the fraction $f(\mu_j)$ is multiplied by the total mass of exporters N_j .

²⁷We use the dataset presented and analyzed in Gimenez-Perales (2021), which is a variant of the one used in Bernard et al. (2018b).

²⁸In Appendix Table A.2, we also report evidence on a decomposition of aggregate Colombian imports into various margins, as presented in Bernard et al. (2018c) (but also for one-sided firm trade in Blum et al., 2019, and Gomtsyan and Tarasov, 2020). The table highlights the importance of the exporter margin in explaining Colombian imports.

Table 1: Firm-Level Gravity

	log(# Connections)	log(Imports)	log($\frac{\text{Imports}}{\text{Connection}}$)	log(# Connections)	log(Imports)	log($\frac{\text{Imports}}{\text{Connection}}$)
log(GDP)	0.351*** (0.053)	0.443*** (0.066)	0.092*** (0.019)	0.377*** (0.052)	0.477*** (0.060)	0.100*** (0.016)
log(GDP p.c.)				-0.147** (0.065)	-0.191** (0.079)	-0.044* (0.022)
log(Distance)	-0.169** (0.071)	-0.309*** (0.098)	-0.140*** (0.045)	-0.210*** (0.057)	-0.362*** (0.080)	-0.152*** (0.045)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.203	0.290	0.449	0.212	0.295	0.450
Observations	84011	84011	84011	84011	84011	84011

Note: Regression using OLS. Robust standard errors clustered at the country level. * p<0.05, ** p<0.01, *** p<0.001.

Similarly to (9), we can compute the recorded import value of firm μ_i from country j , $M(j, \mu_i)$, by integrating the values of firm-to-firm trade (7) over all exporters μ_j for which the reporting threshold is exceeded:

$$M(j, \mu_i) = N_j \int_{\bar{\mu}_{ji}/\mu_i}^{\infty} X(\mu_j, \mu_i) f(\mu_j) d\mu_j = N_j \bar{t} \frac{\theta}{\theta - 1} \left(\frac{\mu_i \underline{\mu} (p_{ji}/P_i)^{1-\sigma} Y_i}{\bar{t} L_i} \right)^{\theta}. \quad (10)$$

Equations (9) and (10) show that the number of connections per importer, $S(j, \mu_i)$, as well as the import values, $M(j, \mu_i)$, are decreasing in bilateral trade costs τ_{ji} and in the exporting country's wage (see (2)). At the same time, both variables are increasing in the mass of firms in exporting country (N_j), which is proportional to the exporting country's total effective labor units (L_j).

We can formulate the according expressions for the exporting firm. Realizing that the direction of trade flows changes when switching the two arguments in $X(\mu_i, \mu_j)$ and replacing $\bar{\mu}_{ji}$ by $\bar{\mu}_{ij}$, we compute the mass of importers (or buyers) in country j per exporter μ_i as

$$B(j, \mu_i) = N_j \int_{\bar{\mu}_{ij}/\mu_i}^{\infty} f(\mu) d\mu = N_j \left(\frac{\mu_i \underline{\mu} (p_{ij}/P_j)^{1-\sigma} Y_j}{\bar{t} L_j} \right)^{\theta}, \quad (11)$$

For the recorded exports of firm μ_i to country j , we get

$$X(j, \mu_i) = N_j \int_{\bar{\mu}_{ji}/\mu_i}^{\infty} X(\mu_i, \mu_j) f(\mu_j) d\mu_j = N_j \bar{t} \frac{\theta}{\theta - 1} \left(\frac{\mu_j \underline{\mu} (p_{ij}/P_j)^{1-\sigma} Y_j}{\bar{t} L_j} \right)^\theta. \quad (12)$$

With these expressions, we obtain the following

Proposition 1 *The Firm Margin and Gravity.*

- (i) *A firm's aggregate import value from foreign market j and the number of its suppliers in market j are increasing in the foreign market size, N_j , the domestic price level, P_i , and the domestic wage, w_i , but decreasing in the foreign wage, w_j , and in variable trade costs, τ_{ji} .*
- (ii) *Similarly, a firm's aggregate export value to a foreign market j and the number of its buyers in market j are increasing in the foreign market size, N_j , the foreign price level, P_j , and the foreign wage, w_j , but decreasing in the domestic wage, w_i , and in variable trade costs, τ_{ij} .*

Proof The first part follows from (9) and (10) in connection with (2), the second part from (11) and (12) with (2).

The proposition relates directly to the empirical patterns summarized in Table 1 above. Specifically, if we read the exporting country's economic size (GDP_j) as a proxy for the country's total number of firms and notice bilateral trade costs by distance (τ_{ji}) enters the price p_{ji} as in (2), then Columns (1) and (2) confirm the association of our gravity variables with both, the number of connections per importer and the firm-level import values as predicted in equations (9) and (10).²⁹ Guided by both equations, we then add to the regressions the explanatory variable exporter GDP per capita. The results, reported in Columns (4) - (6) of Table 1, show that the number of connections per firm is indeed decreasing in per-capita GDP of the source, while the exporting country's GDP and bilateral distance remain statistically significant with the expected sign.³⁰

Our model also predicts that the average import value per connection (reported in the third column) is unrelated to any characteristic at the exporter-level. While still significant in two of three cases, all coefficients drop substantially in magnitude compared to the first two columns and are much closer to zero.

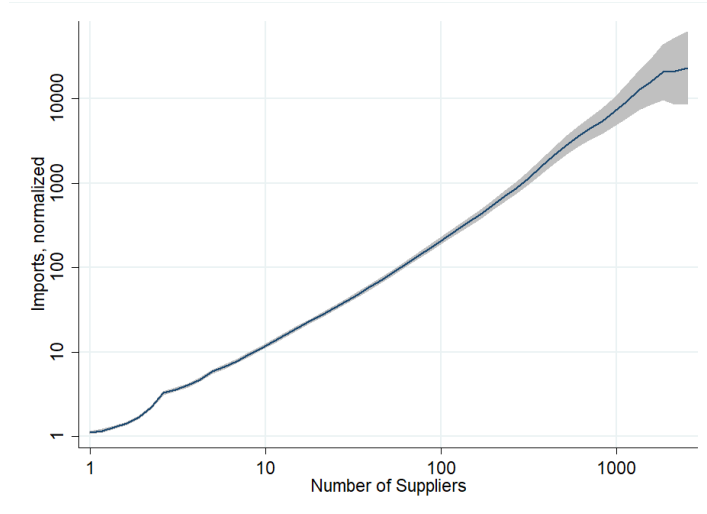
²⁹The firm fixed effects absorb the variation from μ_i and there is country-level variation of the importer i , since we use Colombian import data.

³⁰We stress that, in order to move closer to our model, we would ideally replace the exporter's wage w_j with unit labor cost and exporter GDP with the number of (exporting) firms N_j . However, we are not aware of comprehensive datasets that allows for a similarly broad country coverage.

3.2 Firm-Level Trade and the Number of Connections

Closely in line to the first fact above, trade data exhibit an approximately one-to-one relation between the value of firm-level imports and the number of exporters a firm connects to. This pattern is qualitatively illustrated by Blum et al. (2010) for Chilean data, where import from countries with the lowest trade values are imported by the largest Chilean firms. This regularity is presented for importers in Figure 5 in Bernard et al. (2018b) and for exporters in Figure 6 in the supplemental material of Bernard et al. (2018c).

Our own Figure 2 confirms these patterns, plotting the relationship between the number of suppliers to a Colombian importer on the horizontal axis and the normalized imports of that firm on the vertical axis, both on log scales.³¹



Note: The figure shows the fitted line from a kernel-weighted local polynomial regression of firm-origin log imports on firm-origin log number of suppliers. Axes are scaled in logs. Imports are normalized by mean imports of one-supplier firms.

Figure 2: Log Imports and the Number of Suppliers

Turning to our model, we simply combine equations (9) with (10) to express firm-level imports as a function of the number of connected exporters as

$$M(j, \mu_i) = \bar{t} \frac{\theta}{\theta - 1} S(j, \mu_i). \quad (13)$$

Equation (13) shows that, for a given threshold and Pareto shape parameter, an importer's number of exporters is a sufficient statistic for its total imports and we

³¹For the construction of the graph, we use a kernel-weighted local polynomial regression of log normalized imports on the log number of suppliers. The shaded gray area is the 95% confidence interval. We normalize firm-level imports by the mean imports of firms with only one supplier. In a linear regression of log imports on the log number of suppliers, we find an elasticity of 1.171 with a standard error of 0.006.

formulate the following

Proposition 2 *Firm-Level Trade and the Number of Connections.*

- (i) *A firm's import value from foreign market j increases in the number of the firm's suppliers from market j with unit elasticity.*
- (ii) *Similarly, a firm's export value to foreign market j increases in the number of the firm's buyers in market j with unit elasticity.*

Proof The first part follows directly from (13). The second part follows, equivalently, by combining (11) and (12).

The positive relation between imports and the number of connections at the firm level follows from the fact that for a given reporting threshold, \bar{t} , larger firms import more through more sales units and, simultaneously reach deeper into the pool of potential exporters. Under the Pareto-distributed firm size, the increase in the number of connected exporters is proportional to the increase in trade through the intensive margin and the relation is log-linear with unit slope.

The analogous intuition applies to firm-level exports.

3.3 The Distribution of the Number of Partners

Trade data exhibit a strong link between a firm's number of foreign firm-connections and the share of local firms with more foreign firm-connections. This relationship is shown in Bernard et al. (2018b) in Figure 3(a) for importers and in Figures 3(b) and 4 for exporters.³²

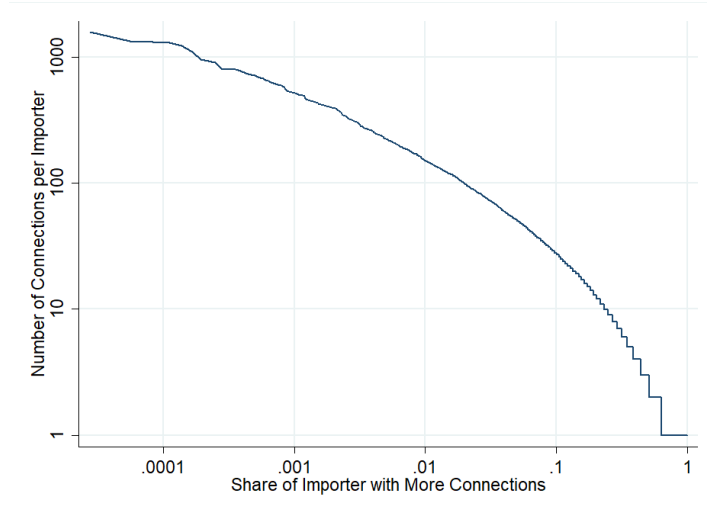
Using our data, Figure 3 replicates this pattern by plotting the total number of connections of an importing firm on the vertical axis against the share of importers in Colombia with more connections on the horizontal axis.³³ Figure 4 plots the same relationship separately for each of the top five Colombian sourcing locations.³⁴ The figures show that a few firms, either exporters or importers, have large numbers of connections while large numbers of firms have just one or two foreign partners.

Turning to our model, this pattern follows, from the Pareto distribution of firm sizes once more in a very tractable manner. Recall that $S(j, \mu_i)$ in Equation (9) is the number

³²Bernard et al. (2018c) confirm those results in Figures 4 and 5 of their online appendix.

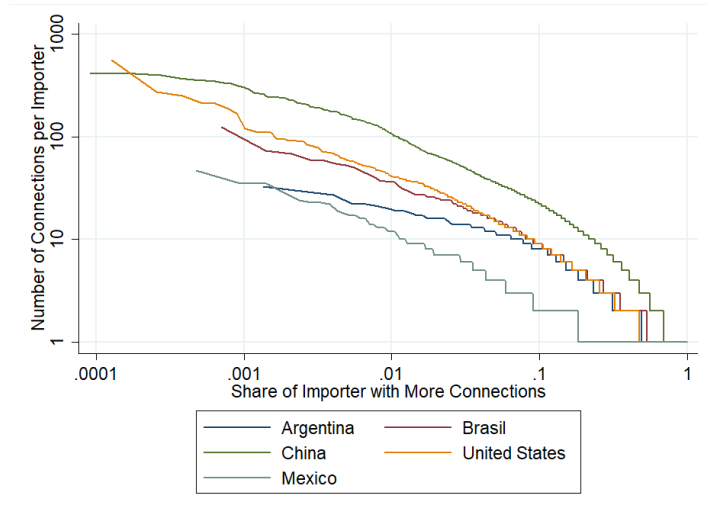
³³A linear regression of the log number of connections on the log share of importers with more connections yields an elasticity of -1.31 with a standard error of 0.006.

³⁴The corresponding elasticities are given by -0.95 (standard error 0.006) for the United States, -1.18 (0.010) for China, -0.48 (0.008) for Mexico, -0.93 (0.015) for Brazil, and -0.83 (0.023) for Argentina.



Note: The elasticity estimated in a linear regression of the log number of connections on the log share of importers with more connections is -1.31 with a standard error of 0.006.

Figure 3: The Distribution of Connections across Importers - All Countries



Note: The elasticities estimated in a linear regression of the log number of connections on the log share of importers with more connections are -0.95 (std. error 0.006) for the United States, -1.18 (0.010) for China, -0.48 (0.008) for Mexico, -0.93 (0.015) for Brazil, and -0.83 (0.023) for Argentina.

Figure 4: The Distribution of Connections across Importers - Top 5 Source Countries

of firms in j that export to a firm of size μ_i . This expression is obviously increasing in the importing firm's size μ_i – the statement applies to any firm size distribution. Therefore, the share of firms with at least as many connections as μ_i is equal to the fraction of firms that is larger than μ_i . In the Pareto case, we express this fraction as

$$\Pr[\mu \geq \mu_i] = 1 - F(\mu_i) = \mu_i^{-\theta} \underline{\mu}^\theta. \quad (14)$$

Solving for μ_i^θ and plugging the result into equation (9) yields the following relation

between the number of exporters for an importer of size μ_i and the fraction of importers that has more connections than μ_i as

$$S(j, \mu_i) = N_j \underline{\mu}^{2\theta} \bar{\mu}_{j_i}^{-\theta} \Pr[\mu \geq \mu_i]^{-1}. \quad (15)$$

This equation yields the following

Proposition 3 *The Distribution of the Number of Partners.*

- (i) *An importing firm's number of sellers in a foreign market j is decreasing in the fraction of local firms connected to more sellers in market j with unit elasticity.*
- (ii) *Similarly, an exporting firm's number of buyers in a foreign market j is decreasing in the fraction of local firms connected to more buyers in market j with unit elasticity.*

Proof The first part follows from (15), the second by applying the same manipulations that lead to (15) to $B(j, \mu_i)$ from (11).

3.4 The Role of Large Firms in the Trade Network

Trade data show a prevalence of large firms with more than one connection in aggregate trade. For example, Blum et al. (2009) report that more than half of the Chilean exporters sell to a single Colombian importer, which is, however, large in import volume. Carballo et al. (2018) show that exports from Costa Rica, Ecuador and Uruguay are mainly driven by few exporting firms with multiple buyers. A similar pattern is shown in Bernard and Moxnes (2018) for Norwegian export data, where most connections involve large firms on at least one side of the trade relation and these account for the largest part of bilateral trade.

Table 2: The Role of Large Firms in Firm-to-Firm Trade

Imports	(1) One-to-one	(2) Many-to-one	(3) One-to-many	(4) Many-to-many
Share of Value (%)	2.8%	5.0%	30.7%	61.6%
Share of Counts (%)	3.8%	5.0%	35.0%	56.2%

Note: The unit of observation is a firm-destination. Firms with one connection in each of two markets are counted as single-connection firms. Column (1) indicates matches, in which both importer and exporter have only one connection in a market, column (2) refers to matches in which the exporting firm has one connection and the importer has multiple connection, column (3) indicates the reverse case and in column (4), both importer and exporter have multiple connections.

Based on our data, Table 2 shows that over half of all firm-to-firm connections involve firms with many connections on both sides and these connections account for more than 60% of total trade. Our summary statistics in Table A.1 also show that 50% of importers connect with more than 3 exporters and that 50% of exporters connect with more than one importer. Recalling further from Figure 2 that firms with more connections are larger in terms of sales and purchases. Taken together, these observations point at the dominant role for large firms in firm-to-firm trade.

Turning to our model, we gauge the role of large firms in our model by computing the mass of connections that involves at least one large firm as a share of total recorded connections. To that aim, we define large firms as those above a certain size percentile χ (e.g., $\chi = 0.9$ for the 90th percentile). We then define the mass of all recorded connections as³⁵

$$M_0 = \int_{\underline{\mu}}^{\infty} \int_{\underline{\mu}}^{\infty} f(\mu_i, \mu_j | \mu_i \mu_j > \bar{\mu}_{ji}) d\mu_i d\mu_j. \quad (16)$$

(This is the share of connections that is located to the top right of the hyperbola represented in Figure 1.)

A firm, either the exporter or the importer, that is located at the χ th size percentile has size μ_χ . The mass of connections involving at least one large firm, i.e., one that is above the χ th percentile of the size distribution, is thus defined as

$$M_1^\chi = \int_{\underline{\mu}}^{\infty} \int_{\underline{\mu}}^{\infty} f(\mu_i, \mu_j | \mu_i \mu_j > \bar{\mu}_{ji} \ \& \ (\mu_i \geq \mu_\chi \ | \ \mu_j \geq \mu_\chi)) d\mu_i d\mu_j. \quad (17)$$

In the appendix, we show that under the Pareto distribution, M_0 is given by

$$M_0 = (\bar{\mu}_{ji}/\underline{\mu}^2)^{-\theta} \left[1 + \theta \ln(\bar{\mu}_{ji}/\underline{\mu}^2) \right]. \quad (18)$$

We also show that if χ is large such that $\mu_\chi \geq \mu_0$, M_1^χ is given by

$$M_1^\chi = 1 - \chi^2, \quad (19)$$

while it is

$$M_1^\chi = (\bar{\mu}_{ji}/\underline{\mu}^2)^{-\theta} \left[2 + 2\theta \ln(\bar{\mu}_{ji}/(\underline{\mu}\mu_\chi)) \right] - (\mu_\chi/\underline{\mu})^{-2\theta} \quad (20)$$

³⁵This mass is expressed as the share of all possible connections $N_i N_j$. We will normalize accordingly in equation (17). Since we then take ratios, suppressing the factor $N_i N_j$ is irrelevant.

if χ is small such that $\mu_\chi < \mu_0$. Here, μ_0 is defined as

$$\mu_0 = \bar{\mu}_{ji}/\underline{\mu}. \quad (21)$$

(The firm size μ_0 is marked in Figure 1 in Section 2.) For both cases we can show that the share of connections involving at least one large firms,

$$m^\chi(\bar{\mu}_{ji}) = M_1^\chi/M_0, \quad (22)$$

is increasing in $\bar{\mu}_{ji}$, or

$$\frac{d}{d\bar{\mu}_{ji}} m^\chi(\bar{\mu}_{ji}) > 0. \quad (23)$$

By the definition of $\bar{\mu}_{ji}$, equation (23) implies the following

Proposition 4 *Large Firms and Connections.* *For any percentile χ , the share $m^\chi(\bar{\mu}_{ji})$ of recorded connections with at least one firm larger than the χ^{th} percentile*

(i) is continuously increasing in $\bar{\mu}_{ji}$,

(ii) equals one if $\bar{\mu}_{ji} \geq \mu_\chi^2$.

Proof See the Appendix.

The proposition obviously implies that, for χ large enough, the share of connections involving at least one large firm is arbitrarily close to one and the role of large firms in international trade becomes dominant.

As a corollary to the proposition, $m^\chi(\bar{\mu}_{ji})$ is increasing in the reporting threshold \bar{t} , trade costs, τ_{ji} , and the exporting country's wage, w_j , but decreasing in the importing country's price index, P_i , and its per-capita GDP, w_i .

Using the Colombian data, we calculate the share of all connections involving firms of which either the importer or the exporter is above a given percentile in the distribution of the number of connections per firm and plot it against the respective percentile in Figure 5. Defining 'large' firms as those in the upper tenth percentile (one percent) of the size distribution, Figure 5 shows that more than 85% (about half) of all connections involve at least one large firm. This observation illustrates the dominant role of large firms in the network. The figure also plots the ratio $m^\chi(\bar{\mu}_{ji})$ implied by the model as a function the percentile χ calibrated to the data.³⁶ We do not seek to apply a

³⁶We set the distribution's lower bound, $\underline{\mu}$, to 1 and take the value of θ from the literature. Specifically, we take the estimate based on firm sales from Table 1 of Di Giovanni et al. (2011) and assume $\sigma = 4$, giving us a value of $\theta = 3.051$. Then we choose the value of $\bar{\mu}_{ji}$ such that the sum of squared distances between the model and the data is minimized.

formal metric to assess our model’s fit with the data but observe that the shape of our calibrated function is close to the data and the qualitative match seems reasonably successful.

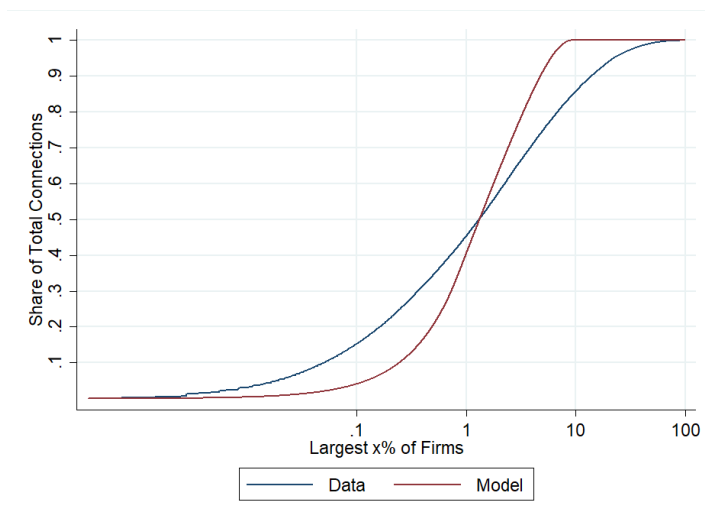


Figure 5: The Share of Connections Involving at Least One Large Firm - Data and Model

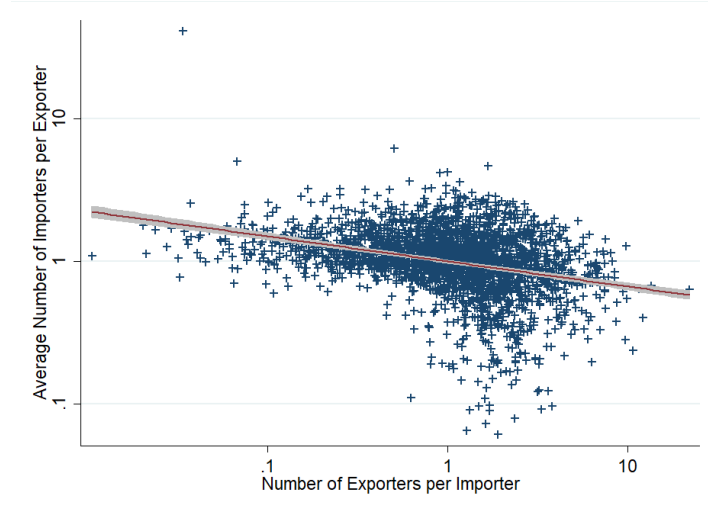
3.5 Negative Assortative Matching

Trade data exhibit negative assortative matching (NAM) in firm-to-firm networks, i.e., large firms tend to match with small ones. For example, Blum et al. (2010) report that small Argentinean exporters tend to connect to one large Chilean importer. Bernard et al. (2018b) show in Figure 7 that Colombian importers that have more exporters connect to exporters that sell to fewer importers on average. An analogous relationship is shown for Norwegian exporters in Figure 1 of Bernard et al. (2018c).

Using our data, Figure 6 plots an importer’s number of trade partners in a country (horizontal axis) to these trade partner’s average number of connections (vertical axis).³⁷ The variables exhibit a negative and statistically significant relationship with an elasticity of -0.18 and a standard error of 0.010. The point (10, 0.1) in the graph can be interpreted as follows: the connected exporters of an importer with ten times the average number of connections in a source country, have on average one tenth of the average number of connections with importers in Colombia. Taking into account

³⁷Axes are scaled in logs. In the construction of the graph we follow Bernard et al. (2018c) and first calculate the number of importers connected to by each exporter as well as the number of exporters in each country connected to by a Colombian importer. We then calculate the mean number of connected importers for each *observed number* of connected exporters by country, thereby pooling Colombian importers with the same number of connections in a source country. We then take log of both variables and demean them by country.

the strong correlation between firm imports and the number of connections presented in Section 3.2, NAM means that larger firms have on average smaller trading partners.



Note: The slope coefficient is estimated to be -0.18 (std. error 0.010).

Figure 6: Negative Assortative Matching

Turning to our model, we consider an importer of size μ_i and relate it to the size of its exporter at the χ^{th} percentile, μ_χ . Negative assortative matching implies that the χ^{th} percentile consists of smaller exporters when the importer is larger, i.e. $\partial\mu_\chi(\mu_i)/\partial\mu_i < 0$. This case also implies that, e.g., the median exporter of a larger importer is smaller than the median exporter of a small importer. The cumulative distribution functions of exporters, conditional being recorded as an exporter to firm μ_i , is given by

$$F(\mu|\mu \geq \bar{\mu}_{ji}/\mu_i) = 1 - (\mu\mu_i/\bar{\mu}_{ji})^{-\theta}. \quad (24)$$

The according χ^{th} percentile, defined through $F(\mu|\mu_\chi \geq \bar{\mu}_{ji}/\mu_i) = \chi$, is

$$\mu_\chi(\mu_i) = (1 - \chi)^{-1/\theta} \bar{\mu}_{ji}/\mu_i. \quad (25)$$

This identity directly implies that $\partial\mu_\chi(\mu_i)/\partial\mu_i < 0$ for any χ and thus proves negative assortative matching between exporting and importing firms in our model.

Proposition 5 *Negative Assortative Matching.*

- (i) *The size of an importer's χ^{th} percentile seller is decreasing in the importer's size with unit elasticity.*
- (ii) *Similarly, the size of an exporter's χ^{th} percentile buyer is decreasing in the exporter's size with unit elasticity.*

Proof The first part follows from (25); the second from the observation that the importer and exporter distributions are identical and by replacing $\bar{\mu}_{ji}$ with $\bar{\mu}_{ij}$ in (24) and (25).

The proposition is formulated in terms of firm size, but the statement immediately translates to the number of connections and thus to Figure 6. In particular, in combination with equations (9) and (11) that establish the link between firm size and the number of recorded connections, the first part of Proposition 5 shows that the more exporters an importing firm connects to, the less connections has its average exporter. Similarly, the second part of Proposition 5 shows that the more importers an exporting firm connects to, the less connections has its average importer.

The intuition for this result is that a large importer can compensate for a smaller export partner to generate trade volumes that are large enough to be recorded. Larger firms therefore reach deeper into the pool of potential trade partners than smaller firms, thus reducing the size of the average partner.

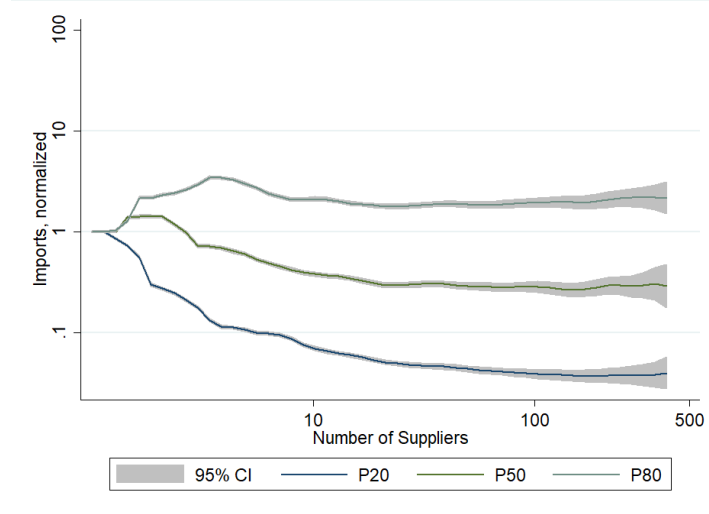
3.6 Conditional Sales Distribution

Trade data show that the distribution of an exporter’s sales or an importer’s purchases across its connections is very stable as the number of connections increases. This stylized fact is documented for Colombian importers in Figure 6 of Bernard et al. (2018b). Bernard et al. (2018c) present according evidence on the export side, using data for Norwegian exporters (Figure 7 of their online appendix).

Using our data, Figure 7 replicates the figures from the reference literature, plotting the source-country-level number of suppliers to Colombian importers on the horizontal axis and the normalized source-country-level imports by Colombian importers on the vertical axis. Both axes are scaled in logs and imports are normalized by mean imports of firms with only one supplier in a market. Including the 20th, 50th, and 80th percentile of the conditional imports distribution, the figure shows that the distribution remains stable as the number of connections of the importers increases. In particular, importers with 100 suppliers do not buy more from their median exporter than importers with 10 suppliers (although in the aggregate, importers with 100 suppliers do of course import more). In Figure 8 we repeat the exercise for the top five source countries of Colombian imports and find conditional sales distributions that are stable or even decreasing in the number of connections.

This result indicates that large firms export and import more not because they sell more to each partner, but because of the number of partners they have. This underscores the importance of the partner-extensive margin at the firm level for explaining large trading firms – and by the concentration of trade among those largest firms – also

for explaining aggregate trade flows.



Note: The figure shows the fitted lines from kernel-weighted local polynomial regressions of the 20th, 50th, and 80th percentile of firm-origin log imports on firm-origin log number of suppliers. Axes are scaled in logs. Imports are normalized by mean imports of one-supplier firms.

Figure 7: Conditional Sales Distribution

Turning to our model, the value of trade between importer μ_i and exporter μ_χ located in j is given by (7), which is proportional to $\mu_i\mu_\chi$. Combining equations (7), (8) and (4), the firm-to-firm trade value is given by

$$X(\mu_i, \mu_\chi(\mu_i)) = (1 - \chi)^{-1/\theta} \bar{t}. \quad (26)$$

The identity shows that the absolute import volume of an importer from its χ^{th} percentile exporter is independent of the importer's size.

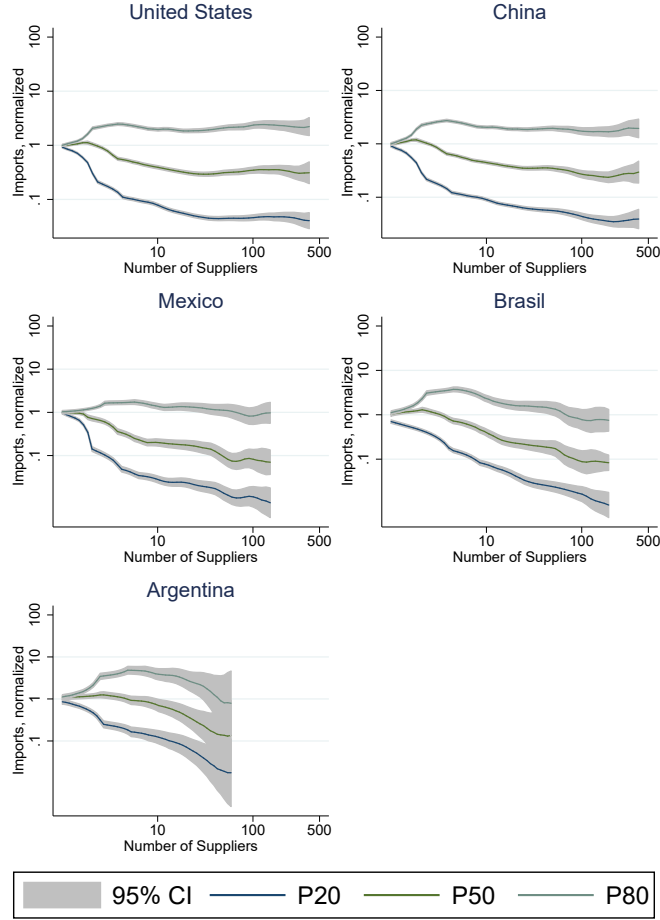
Normalizing further by the median sales of *any* importing firm μ_{i_0} (e.g., the 'smallest' importer) thus yields

$$\frac{X(\mu_i, \mu_\chi(\mu_i))}{X(\mu_{i_0}, \mu_{1/2}(\mu_i))} = (2(1 - \chi))^{-1/\theta}. \quad (27)$$

These observations give rise to the following

Proposition 6 *Conditional Distribution of Trade.*

- (i) *A firm's distribution of import values from market j across the percentiles of its suppliers in market j is independent of the importing firm's size.*
- (ii) *Similarly, a firm's distribution of export values to market j across the percentiles of its buyers in market j is independent of the exporting firm's size.*



Note: The figure shows the fitted lines from kernel-weighted local polynomial regressions of the 20th, 50th, and 80th percentile of firm-origin log imports on firm-origin log number of suppliers. Axes are scaled in logs. Imports are normalized by mean imports of one-supplier firms.

Figure 8: Conditional Sales Distribution

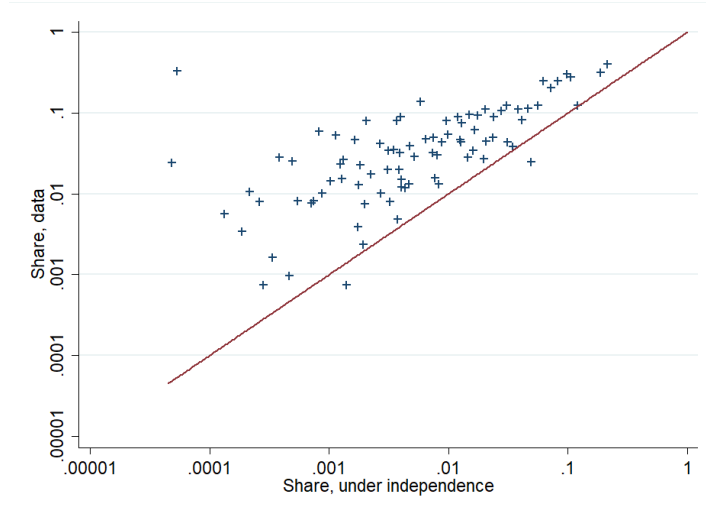
Proof The first part follows from (27). The second part follows from the parallel expressions, derived by switching arguments in $X(\mu_i, \mu_j)$ and computing $X(\mu_\chi(\mu_i), \mu_i)$.

3.7 Hierarchy of Connections

In trade data, the connections of importers and exporters are governed by a hierarchy. Bernard et al. (2018c) show in Figure 9 of their online appendix that in the majority of destinations, the share of Norwegian exporters connecting with foreign importers in the order of their connectedness to the Norwegian market is larger than expected under the statistical benchmark of independence of connection probabilities.

Using our data, we follow the approach of Bernard et al. (2018c) and compare, for each source country, the share of Colombian importers, who buy from the exporter with the most Colombian connections, the exporter with the second-most Colombian connec-

tions, and so on, with the share of Colombian importers, who would do so if the probabilities of a Colombian importer buying from a particular exporter were independent. Denote by r_S the rank of an exporter in a country based on its number of connected importers in Colombia. The share of all Colombian importers from that country the exporter with rank r_S connects to is denoted by π_{r_S} . If connections are independent, for a Colombian importer, the probability of connecting only to the most-connected exporter in a country is given by $p_1 = \pi_1 \prod_{r=2}^S (1 - \pi_r)$, with S being the aggregate number of exporters in a market. Accordingly, the probability of connecting only to the most and second-most connected exporters is given by $p_2 = \pi_1 \pi_2 \prod_{r=3}^S (1 - \pi_r)$ and so on. In Figure 9, each datapoint is a country, for which we plot the share of firms in the data that follow the hierarchy on the vertical axis against the share that would be expected under the assumption of independence of connections, i.e. $\sum_{k=1}^S \pi_k$. We find that for a large number of source countries, the share of firms following the hierarchy is larger than would be expected under the independence assumption.



Note: Axes are scaled in logs. Exporting countries with more than 20 exporters and importers are considered. The vertical axis shows the share of importers in all importers from a market, who connect with the most connected exporter, the second-most connected exporter, and so on, in each market. The horizontal axis shows those shares that would be expected under independent connection probabilities.

Figure 9: Hierarchy Across Exporters

Turning to our model, the set $\mathcal{S}(\mu_i)$ of exporters an importer of size μ_i connects to is given by

$$\mathcal{S}(\mu_i) = \{\mu_j | \mu_j \geq \bar{\mu}_{ji}/\mu_i\}. \quad (28)$$

Clearly, as the importer size, μ_i , increases, the set $\mathcal{S}(\mu_i)$ expands and contains ever smaller exporting firms μ_j , i.e., $\mathcal{S}_{\mu_i} \subset \mathcal{S}_{\mu'_i}$ for all $\mu'_i > \mu_i$.

The similarly defined set of importers, exporter μ_i connects to is

$$\mathcal{B}(\mu_i) = \{\mu_j | \mu_j \geq \bar{\mu}_{ij}/\mu_i\}. \quad (29)$$

(where only the indices of $\bar{\mu}_{ij}$ are switched relative to (28)) defines the set of importers an exporter of size μ_i connects to.

These equations lead to the following

Proposition 7 *Hierarchy of Connections.* *Firms-to-firm connections follow a strict hierarchy that is governed by firm size:*

- (i) *An importer's set of suppliers from foreign market j is a strict subset of any larger importer's set of suppliers from foreign market j .*
- (ii) *Similarly, an exporter's set of buyers in foreign market j is a strict subset of any larger exporter's set of buyers in foreign market j .*

Proof The first part follows from (28), the second one from (29).

Importers always connect to the largest exporter and extend their reach downwards towards smaller exporters. Intuitively, the restriction that trade flows above a certain threshold are recorded is less binding for larger importers, so that their connections with the smaller exporters are still recorded.

4 Discussion

The previous section has documented how salient patterns of firm-to-firm trade data emerge from our simple Krugman-based model. The model crucially rests on the cross-border links of plants to sales units, and the random bundling of those actors into local firms of Pareto-distributed sizes. These features, in combination with a reporting threshold for firm-to-firm trade, generate our surprisingly rich set of predictions.

This section offers a critical interpretation on some of our model's features and discusses further perspectives as well as testable implications and possible extensions. First, we expand on the stochastic element of our model, discussing the main similarities and differences to the balls and bins model of trade by Armenter and Koren (2014), our central reference in this regard. Second, we take a look at the empirical patterns that our model does not fully match but also at the predictions that may be explored in future empirical work. Third, we lay out our understanding of the terms *variety* and *plant* and how these definitions of our model relate to data. In this context, we also discuss how our model may be applied to firm-to-firm trade in inputs and point at possible extensions for further dimensions, such as *products* and product categories.

4.1 Balls and Bins and Sparsity in Trade Connections

The central role of the stochastic bundling in our model constitutes a conceptual similarity to the balls and bins model of trade by Armenter and Koren (2014). These authors show that a purely stochastic model with a discrete number of balls falling into a discrete number of bins delivers the well-known gravity patterns of international trade and generates a rich set of predictions for the extensive margin of trade that has robust empirical backing.

How does our re-interpretation of Krugman (1980) relate to the stochastic balls and bins model? We think of our model as a hybrid: we apply the stochastic approach at the crucial part, where firm-to-firm connections are formed, but we maintain an underlying structure of economic decision making that models and explains the ‘balls’: in our model, the ‘balls’ are the economic transactions emerging from the Krugman (1980) model. Only the ‘bins’ (firms) are free of economic meaning and constructed as purely legal and accounting entities. This approach allows us to apply the logic of the balls and bins model and show that the empirical patterns emerge mechanically when firms of different size trade with one another. At the same time, we maintain a fully-fledged Krugman (1980) model at the plant-level. Our model thus builds a bridge between the stochastic approach and economic modelling: on the one hand, it allows to directly study the effect of *removing* all microeconomic complexity at the firm level and the firm-to-firm level; on the other hand, it invites to *re-introduce* economic modelling choices at the firm level.

There is also an important technical difference between the approach in Armenter and Koren (2014) and ours. While Armenter and Koren (2014) postulate transactions of discrete (and identical) size, in our model we consider a continuum of firms, each of which comprises a continuum of plants and thus exports a continuous flow of goods. Our choice has the advantage that we can rely on the law of large numbers when solving the model. At the same time, the assumption of a continuum of varieties within a firm implies that all firms are connected with all other firms, which is clearly at odds with the *sparsity* of connections found in the data. Indeed, bilateral country-to-country trade flows are typically *sparse* (the number of observed positive trade flows is low compared to the number of possible connections, which is the motivating observation for Armenter and Koren, 2014) and, not surprisingly, firm-to-firm connections are sparse as well. Our model generates the feature of sparse firm-to-firm trade through the reporting threshold: since trade flows grow arbitrarily small as firms grow smaller the firms or as trade costs increase, the reporting threshold censors the smallest connections and thereby generates sparsity.³⁸ Of the universe of possible firm-to-firm connections, only a small

³⁸We have motivated the introduction of a reporting threshold in Section 2. United Nations (2004),

fraction is reported to be active.³⁹ While the balls and bins approach by Armenter and Koren (2014) thus discretizes the trade flows and delivers continuous probabilities of observing bilateral trade flows, our approach, by contrast, keeps trade flows continuous but delivers discrete probabilities of zero or one from observing firm-to-firm trade flows.

Alternatively, we could move closer to Armenter and Koren (2014) in generating sparsity in our model without recurring to the reporting threshold by adopting the idea of discrete trade flows in our model. In particular, we could sample a discrete number of observations from our model, i.e., set a number n_1 of plants and randomly assign each plant one sales unit in each country. Then, we could chose a number n_2 of firms, drawing a size for each of them from our Pareto distribution. Finally, within each country, all plants and sales units would be randomly bundled into firms, with a plant’s probability of being assigned to a firm being proportional to the firm’s size draw.

The thought experiment of discretizing our model along those lines would naturally generate some features of the data: this approach mechanically delivers sparsity, as a given local firm connects to very small foreign firms with very low probability only. At the same time, it would remove the implication of our continuous model that some very large firms should have trade flows above the reporting threshold with all foreign firms.⁴⁰ Clearly, for any finite number of draws, the largest firm will still be of finite size and only connect to a finite number of foreign firms.

Interestingly, Bernard et al. (2018c) discuss a balls and bins approach in their Online Appendix Section C. They argue that it cannot generate the key patterns in the firm-to-firm trade data. This result seems to be in stark contrast to our findings. The reason for this discrepancy lies in their definition of a ‘bin’. In their approach, the ‘bin’ is the connection between a given buyer and a given seller, i.e. the exporting firm-importing firm tuple defines a bin. Our take is that this definition is misfit for a direct

the manual mentioned in Footnote 25 above, also specifies in Paragraph 238 (“Treatment of low-value transactions”): “When the threshold for considering low value transactions is kept low, more complete and higher quality trade statistics are possible but only at the expense of a larger data-processing load. Whatever the threshold, estimates of trade below the threshold level should be made.” Also note that there is an obvious difference between a firm’s total bilateral trade volume and the size of its transactions. As shown in Kropf and Sauré (2014), however, the frequency and the size of transactions are increasing in the overall trade value, so that the statement that small firms do not pass the accounting threshold remains valid under endogenous decisions of transaction size. Our model is thus proof to the critique in Blum et al. (2016).

³⁹An obvious alternative to the reporting threshold would be a fixed cost of exporting at the firm level. The introduction of such a cost at the firm level would, however, make our approach shift away from Krugman (1980). Our aim is to suggest a benchmark model without any economic decision taken at the firm level. The reporting threshold is therefore our preferred modelling choice. We will see below that a discretization of the model is an alternative way to go.

⁴⁰In our theoretical analysis in Section 2.3 we simply disregard this, pointing out in Footnote 26 that the feature is avoided when truncating the Pareto distribution, in which case all of our results hold as an approximation. Bernard et al. (2018c) deal with the same issue by considering a limiting case where the minimum firm size of sellers goes to zero, combined with linking the minimal seller size inversely to the total measure of sellers.

application of the original balls and bins approach to the firm-to-firm network, as it replaces the connection between two bins with a bin itself. Importantly, the approach entirely disregards the firm size of the trading partners and thereby a central element of the firm-to-firm network. It is hence unsurprising that the balls and bins analysis in Bernard et al. (2018c) does not match the patterns in the data. We argue that the suitable definition should rely on the identification of the exporting firm with the export-bin, the importing firm with the import-bin, and the transaction of a plant to a sales unit with a ball. Our distinction between plants and firms, in combination with the mechanics of the reporting threshold, reintroduces this element and implies that the probabilistic approach yields a successful reflection of the data.

In an ongoing research project, Bernard and Zi (2021) recognize the problematic definition of a ‘bin’ in Bernard et al (2018c) and pursue a purely stochastic balls and bins approach. Similar to our paper, they argue that models with rich microeconomic structure at the firm-to-firm level should be benchmarked against a purely probabilistic framework. Their work complements our analysis in studying a generalized framework, where sparsity of firm-to-firm connections, negative assortativity and hierarchical matching (i.e., the patterns discussed in our Sections 3.5 and 3.7) emerge under arbitrary firm size distributions. The authors also investigate the information content of data under different statistical transformations and at varying aggregation levels. Their findings underscore the generality of the stochastic elements of our approach. Despite the synergies between both approaches, our approach generates a number of unique results. Specifically, by building on the empirically relevant Pareto distribution (see Di Giovanni et al., 2011, or Görg et al., 2017), our model does not only yield sharp structural predictions regarding the role of gravity variables (Propositions 1 and 4), the relation between trade values and connections (Proposition 2), and negative assortativity (Proposition 5), but also predicts additional patterns regarding the distributions of connections and trade values (Propositions 3 and 6).

4.2 Further Perspectives and Implications

Closely in line with Armenter and Koren (2014), we understand the general message of our stochastic model as twofold. We show that the list of salient empirical patterns generated by our model “will be consistent with a very large class of models” (Armenter and Koren, 2014, p. 2129). A theoretical model potential to match these patterns is thus no indication of empirical support for its economic mechanisms or modelling approach. For the pragmatist, this observation implies that a cheap way to explain the observed empirical regularities is by micro-founding a Pareto distribution of firm size (e.g., as in Luttmer, 2007).

At the same time, our paper offers constructive messages as well. First, the model delivers genuinely new predictions that are, to the best of our knowledge, not tested at present. For example, our model yields refinements to the standard, ad-hoc empirical approaches explored, e.g., in Section 3.1 regarding the impact of the classical gravity variables on firm imports and the number of firm connections. Also, equations (9) - (12) suggest that the economic size (GDP_j) picks up the effects of the mass of firms of the partner country, N_j . It thus suggests to include the number of foreign firms N_j directly and send it into a horse race with the classical gravity variables, in particular, with GDP_j . Similarly, the unit labor costs, currently proxied by w_j , can be measured more precisely and tested within the current gravity setup (see Proposition 1). Also, the models predicts asymmetries in the roles of gravity variables of the importer and the exporter, as highlighted by the pairs of equations (9) and (11) as well as (12) and (10). Datasets with firm-level information for several importing and exporting countries (such as the World Bank’s Exporter Dynamics Database) may be used to jointly test, for example, the predictions of per-capita expenditure, w_i , the importing country’s price level P_i , the exporting country’s production costs w_j and bilateral trade costs for firm-to-firm trade flows (see, e.g. equation (9)).⁴¹ Alternatively, when approaching equations (9) - (12) from a structural angle, log-linearized, econometric versions of these equations will yield estimates for the critical Pareto shape parameter θ (e.g., through the estimated coefficients of $\ln(Y_i/L_i)$ in equation (9)) and on its interaction with the demand elasticity (through the coefficient of the exporter’s unit labor cost). To the extent that our model’s mechanics operate at the sector-level, sector-specific θ may also be estimated from firm-to-firm trade. Investigating these dimensions would connect our work closely to existing work on the size distribution of firms, e.g., by Di Giovanni et al. (2011), Head et al. (2014) or Görg et al. (2017). Another predictions that is, to the best of our knowledge, not yet tested at present relate to the prevalence of large firms in connections (see Proposition 4). Thus, the share of connections with at least one large exporter on either side of the connection is predicted by equation (23) to vary systematically with the respective trade costs τ_{ji} , the exporting country’s wage w_j , the importing country’s price index P_i and its per-capita GDP w_i . Finally, the negative assortative matching illustrated in Figure 6 is predicted to involve a country-to-country dimension through the variable $\bar{\mu}_{ji}$, as shown in equation (25) (see also Proposition 5). Our model thus predicts that the corresponding country factors induce part of the noise affecting Figure 6. The influence of these country-level factors may be tested as well.

Our model does leave some empirical regularities and statistics unexplained and

⁴¹These predictions link our model to the literature on the role of per-capita income in international trade. Early studies include Hufbauer (1970) and Anderson (1979). More recent research is done, e.g., by Fieler (2011), Sauré (2012), Caron et al. (2014), and Fajgelbaum and Khandelwal (2016).

thus opens the possibility for the theory “to distinguish itself from the competition. . .” (Armenter and Koren, 2014, p. 2129). Thus, moving beyond the model’s current, somewhat narrow, specification, we can venture educated guesses about the model’s performance in further dimensions. For example, we may introduce differences in the distributions of selling and buying firms, or across countries and potentially across sectors (introducing the corresponding indices for the CDF in equation (5)). These differences, which then play out in the equations (9) - (12) and beyond, will deliver country or sector-specific estimates, e.g., of the Pareto shape parameter θ . We may also move beyond the standard setup of the Pareto distribution, adopting, e.g., a truncated Pareto à la Helpman et al. (2008), which generates an endogenous set of purely domestic firms if the reporting threshold is sufficiently large.⁴²

Time is also an important and relevant dimension: the work on the patterns firm dynamics mentioned in the introduction points at promising extensions of our model. Thus, work on firm exports (Albornoz et al., 2012, and Ruhl and Willis, 2017) and the dynamics of firm-to-firm connections (Blum et al., 2010 and Gimenez-Perales, 2021) documents empirical regularities that are likely to be informative about the factors that systematically shape firm network in international trade beyond our probabilistic approach. An obvious point in case is the role of intermediaries documented on Blum et al. (2010), which tend to specialize on specific destination countries, thus indicating economies of scale of organizing shipments by region.

These generalizations may be explored, in particular, to investigate the underlying factors for some subtle discrepancies between the theory and the empirics the previous sections also exposed. Thus, the deviations from the Pareto distribution documented in the previous literature may explain the small but statistically significant deviation of the slope in Figure 3 from the unit elasticity predicted in equation (13) (see Proposition 2). Also, the underlying firm size distributions may relate to the apparent concavity in Figures 4 and 3, which constitutes a deviation from our baseline theory (see Proposition 3).

In sum, our model not only delivers mechanical patterns that have their robust correspondence in the data. In addition, it shows that some parts of the empirical patterns may indeed inform us about the validity of modelling economic decisions at the firm level that go above and beyond the mere size of firms, i.e., the mechanical scale effect at the firm level. We have pointed at some of these patterns – deviations from the unit elasticity predicted in Proposition 2 or differences from the log-linear pattern predicted in Proposition 3 – but also at some novel predictions that are related with but not identical to traditional gravity variables (Propositions 1 and 4). Future empirical work on these dimensions may help to further narrow down the set of empirical patterns that

⁴²We read our Pareto distribution in our current setting as a proxy for this case.

are apt to test the economics of firm-to-firm relations with international trade data.

Finally, we stress that with its underlying Krugman (1980) model, our model provides a natural starting point for future research that re-introduces firms that face genuine economic incentives to form buyer-seller relationships.

4.3 Interpretations and Possible Extensions

In this last part of the section, we present some generalizations of the parsimonious approach presented in our model and also point at possible extensions of our model.

4.3.1 Firms vs. Plants

For our re-interpretation of the Krugman (1980) model in Section 2, we have distinguished between firms, plants and sales units and subsequently relied on these definitions when analyzing our model. While we introduced the terminology for expositional purposes, we do not want the reader to take it literally and we certainly want to avoid misunderstandings arising from our labels. When contemplating our re-interpretation, one may argue that trade transactions are organized, observed and recorded at the plant level so that plants' transactions (instead of firms' transactions) are subject to an accounting threshold.⁴³ From that point of view, our identification of a 'plant' with a variety would imply that the model's margins cease to operate in an interesting way: since all plants (varieties) are identical, either all or none would be recorded to export to a given destination. Our perception of the modelling setup clearly differs from such a narrow view. In particular, we understand firms merely as the accounting entities that ship their products overseas. We would thus re-label them as plants when working with plant-level data.

4.3.2 Trade in Inputs

We may also alter our previous interpretation of the 'variety' that is produced under a firm's roof and demanded and consumed by an individual. Specifically, we may read such a variety as an intermediate input, which is distributed to a foreign final good producer who, in turn, aggregates its varieties to a final consumption good. This reading of the model would bring us closer to most of the literature on production networks and firm-to-firm trade, which models links between input suppliers and final good producers, as discussed in the Introduction.⁴⁴

⁴³Such a misunderstanding may arise, in particular, in view of the widespread use of plant-level data in related work in international trade (see, e.g., Amiti and Konings, 2007 and Kasahara and Rodrigue, 2008 for an early and closely related example).

⁴⁴With an additional adjustment to the interpretation, the model can also be applied to domestic trade networks. In our model, sales units are exclusive sellers to a group of consumers (final goods

This interpretation would bring us also in line with the view in Caliendo and Parro (2015), who assume that competitive firms indexed by j purchase subsets \mathcal{N}_j of (foreign or domestic) varieties in order to produce a composite of intermediate varieties (labelled ‘composite intermediates’ in Caliendo and Parro (2015) according to

$$X_j = \left(\sum_{n \in \mathcal{N}_j} x_n^{1-1/\sigma} \right)^{\sigma/(1-\sigma)}.$$

Following this view, we could read our economy as one where consumers purchase these composite intermediates and combine to the final consumption good

$$C = \left(\sum_{j=1}^J X_j^{1-1/\sigma} \right)^{\sigma/(1-\sigma)} = \left(\sum_{n \in \mathcal{N}} x_n^{1-1/\sigma} \right)^{\sigma/(1-\sigma)} \quad (30)$$

Since the competitive firms who produce the composite of intermediate varieties sell at zero surcharge, all prices and thus allocation would remain unchanged so that our model falls back into the original Krugman (1980) paradigm.

4.3.3 Multi-Product Firms

When further zooming in on the accounting entity of a ‘firm’, we can also extend our reading of varieties produced within this firm. While some of these varieties may constitute goods for final consumption, others can represent intermediate goods, which are assembled via a CES-aggregator before they leave the factory gate. In this specific interpretation of Krugman (1980), we would actually partition the mass of the firm’s varieties, \mathcal{N} , into a (finite or infinite) number of subsets, $\{\mathcal{N}_{j \in J}\}_j$. For example, a small product with a rather simple production process consists of a smaller number of these varieties, while a large and complex product consists of more.⁴⁵ Again, this interpretation leaves the formal model unchanged (as long as substitution elasticities are identical at all aggregation levels) since for the partition $\{\mathcal{N}_j\}_j$ of a full set of varieties \mathcal{N} , we know that, formally identical to equation (30)

$$Y^{\mathcal{N}} = \left(\sum_{n \in \mathcal{N}} x_n^{1-1/\sigma} \right)^{\sigma/(1-\sigma)} = \left(\sum_{j=1}^J X_j^{1-1/\sigma} \right)^{\sigma/(1-\sigma)}$$

producers) because the latter are in the same country. Within a country, clients could be grouped for different reasons (location, infrastructure, sectors, business associations etc.)

⁴⁵An example of the former, say, a bar of soap would consist of a set \mathcal{N}_j with smaller mass than an example of a latter, $\mathcal{N}_{j'}$, that combine to, say, a windmill.

where

$$X_j = \left(\sum_{n \in \mathcal{N}_j} x_n^{1-1/\sigma} \right)^{\sigma/(1-\sigma)}$$

Importantly, this conception of the mass of varieties within a firm would be consistent with our re-interpretation of Krugman (1980) as long as the bundled varieties (\mathcal{N}_j) ‘fit’ into the according importing firm. In other words, no bundle \mathcal{N}_j can consist of a mass of varieties that exceeds the mass of sales units in the according importing firm.⁴⁶ In these re-interpretations of our model, the defining feature of a ‘firm’ remains the organization of cross-border trade transactions.⁴⁷

All of these possible interpretations of varieties within firms point at a relevant dimension that we have neglected throughout our analysis, but which the empirical literature on firm exports has explored: empirical regularities at the product level. Specifically, Table A.2 in the Appendix (and, e.g., Blum et al., 2009, and Bernard et al., 2018c) show that trade systematically expands along the product margins. That observation, too, is consistent with our model, as the following thoughts will show.

A direct way to introduce the product level to our model is by assuming that a variety produced by any firm is randomly assigned to the statistical product category k with probability γ_k (which is independent of the country and firm).⁴⁸ Our model then implies that each country’s expenditure share on product k is γ_k (where $\sum_k \gamma_k = 1$ must obviously hold).⁴⁹

Within this setup, a reporting threshold at the firm-firm-product level implies that equation (8) becomes product-specific and the condition for trade in a product between the two firms μ_i and μ_j to be recorded is

$$\mu_i \mu_j \geq \bar{\mu}_{jik} := \frac{\bar{t}}{(p_{ji}/P_i)^{1-\sigma} \gamma_k Y_i / L_i}. \quad (31)$$

Clearly, the lower the expenditure share on product k , the larger the product of the two trading firms’ size must be in order for trade to be recorded. Hence, there is a marginal expenditure share $\bar{\gamma}$, below which products remain unrecorded. The product margin is thus operating in this extension of the model as well and is predicted to

⁴⁶Formally, that requirement amounts to assuming a form of first order stochastic dominance of the distribution of foreign firms over the distribution of the size of bundles within exporting firms.

⁴⁷The work by Arkolakis et al. (2010) and Gomtsyan and Tarasov (2020) documents that firms do organize and ship multiple products at the time and that it is cost-saving to do so.

⁴⁸Here, we closely follow the empirical literature by defining firms as “bundles of establishments in the same or different industries” (Magyari et al., 2017).

⁴⁹An alternative way to introduce products to our model is by generalizing country i ’s utility to $U = \exp \left[\int_{k \in [0,1]} \gamma_k \ln(C_{jk}) dk \right]$, where $C_{jk} = \left[\int_{j \in \Omega_{jk}} c_{ijk}^{1-1/\sigma} dj \right]^{\sigma/(1-\sigma)}$ is the bundle of k -products consumed so that $\gamma_k \geq 0$ is the expenditure share on product k .

react systematically to the gravity variables that shape the buyer margin and the seller margin through equation (31). Specifically, the product margin will expand with trade-promoting factors (proximity, per-capita expenditure of the importing country) and contract with trade inhibiting factors (trade and production costs).⁵⁰

In sum, a direct extension of our model will feature an operating product margin, which can be conjectured to be consistent with the recently documented empirical patterns along this margin as well.

5 Conclusion

This paper proposes a model of firm-to-firm trade, which does not go beyond Krugman (1980) in economic content and features purposefully simplistic firms. In our terminology, the ‘firm’ is simply a legal entity with a random assignment of a mass of plants that produce Krugman-type varieties. It also comprises a proportional mass of sales units, which import foreign varieties and sell them to domestic consumers. Cross-border firm-to-firm connections arise as a result of unique linkages between plants and sales units. Finally, firm-to-firm trade is assumed to enter trade statistics only when it exceeds a fixed reporting threshold. This model generates a number of salient empirical regularities of the recent firm-to-firm trade literature that the recent literature has documented. We argue that any microeconomic modelling choice at the firm(-to-firm) level that prides itself on finding support in the data must outperform our stochastic trade model in terms of matching the data.

We acknowledge that our general message may not appear very constructive but hope that it will ultimately turn out to be very much so. In spirit very close to Armenter and Koren (2014), we intend to draw a dividing line between clear empirical support for a theory and minimal requirements of consistency with firm-to-firm trade data. Moreover, by opting for the tractable and flexible Krugman (1980) workhorse model as the basis of our approach, our framework readily lends itself to stepwise re-introduce economic activity at the firm level. It may therefore provide a useful tool for future research that aims at identifying microeconomic modelling choices at the firm or firm-to-firm level. Those microeconomic mechanisms, a selection of which we have listed in our paper, will need to demonstrate their value added relative to our benchmark model.

⁵⁰Under concrete assumptions on the distribution of expenditure shares γ , the model will deliver sharper results on the impact of those gravity variables on product margin.

References

- ACEMOGLU, D. AND P. D. AZAR (2020): “Endogenous Production Networks,” *Econometrica*, 88, 33–82.
- ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 80, 1977–2016.
- ALBORNOZ, F., H. F. C. PARDO, G. CORCOS, AND E. ORNELAS (2012): “Sequential Exporting,” *Journal of International Economics*, 88, 17–31.
- AMALDOSS, W. AND S. JAIN (2015): “Branding Conspicuous Goods: An Analysis of the Effects of Social Influence and Competition,” *Management Science*, 61, 2064–2079.
- AMITI, M. AND J. KONINGS (2007): “Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia,” *American Economic Review*, 97, 1611–1638.
- ANDERSON, J. (1979): “A Theoretical Foundation of the Gravity Equation,” *American Economic Review*, 69, 106–116.
- ANTRÀS, P. (2003): “Firms, Contracts, and Trade Structure,” *The Quarterly Journal of Economics*, 118, 1375–1418.
- ANTRÀS, P., T. FORT, AND F. TINTELNOT (2017): “The Margins of Global Sourcing: Theory and Evidence from U.S. Firms,” *American Economic Review*, 107, 2514–2564.
- ANTRÀS, P. AND E. HELPMAN (2004): “Global Sourcing,” *Journal of Political Economy*, 112, 552–580.
- ARKOLAKIS, C. (2010): “Market Penetration Costs and the New Consumers Margin in International Trade,” *Journal of Political Economy*, 118, 1151–1199.
- ARKOLAKIS, C., S. GANAPATI, AND M.-A. MUENDLER (2010): “The Extensive Margin of Exporting Products: A Firm-level Analysis,” Tech. rep., National Bureau of Economic Research.
- ARMENTER, R. AND M. KOREN (2014): “A Balls-and-Bins Model of Trade,” *American Economic Review*, 104, 2127–2151.
- ATKESON, A. AND A. BURSTEIN (2008): “Pricing-to-Market, Trade Costs, and International Relative Prices,” *American Economic Review*, 98, 1998–2031.
- AUER, R. A., T. CHANEY, AND P. SAURÉ (2018): “Quality Pricing-to-Market,” *Journal of International Economics*, 110, 87–102.
- BENGURIA, F. (2015): “The Matching and Sorting of Exporting and Importing Firms: Theory and Evidence,” *Technical Report*.
- BERNARD, A., E. DHYNE, G. MAGERMAN, K. MANOVA, AND A. MOXNES (2021): “The Origins of Firm Heterogeneity: A Production Network Approach,” *Journal of Political Economy*, forthcoming.

- BERNARD, A., J. JENSEN, S. REDDING, AND P. SCHOTT (2018a): “Global Firms,” *Journal of Economic Literature*, 56, 565–619.
- BERNARD, A., A. MOXNES, AND Y. SAITO (2019): “Production Networks, Geography, and Firm Performance,” *Journal of Political Economy*.
- BERNARD, A., S. REDDING, AND P. SCHOTT (2011): “Multiproduct Firms and Trade Liberalization,” *Quarterly Journal of Economics*, 126, 1271–1318.
- BERNARD, A. AND Y. ZI (2021): “Sparse Production Networks,” *mimeo, University of Oslo*.
- BERNARD, A. B., E. A. BØLER, AND S. DHINGRA (2018b): “Firm-to-Firm Connections in Colombian Imports,” *NBER WP 24557*.
- BERNARD, A. B. AND A. MOXNES (2018): “Networks and Trade,” *Annual Review of Economics*, 10, 65–85.
- BERNARD, A. B., A. MOXNES, AND K. H. ULLTVEIT-MOE (2018c): “Two-Sided Heterogeneity and Trade,” *Review of Economics and Statistics*, 100, 424–439.
- BLOOM, N., K. MANOVA, J. VAN REENEN, S. T. SUN, AND Z. YU (2020): “Trade and Management,” *Review of Economics and Statistics*, 1–44.
- BLUM, B. S., S. CLARO, K. DASGUPTA, AND I. J. HORSTMANN (2019): “Inventory Management, Product Quality, and Cross-Country Income Differences,” *American Economic Journal: Macroeconomics*, 11, 338–88.
- BLUM, B. S., S. CLARO, AND I. HORSTMANN (2010): “Facts and Figures on Intermediated Trade,” *American Economic Review*, 100, 419–23.
- BLUM, B. S., S. CLARO, AND I. J. HORSTMANN (2009): “Intermediation and the Nature of Trade Costs: Theory and Evidence,” *mimeo, University of Toronto*.
- (2016): “A balls-and-bins model of trade: Comment,” *American Economic Review*, 106, 843–51.
- BONFIGLIOLI, A., R. CRINÒ, AND G. GANCIA (2019): “Trade, Finance, and Endogenous Firm Heterogeneity,” *Journal of the European Economic Association*, 17, 79–130.
- BUSTOS, P. (2011): “Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms,” *American Economic Review*, 101, 304–340.
- CALIENDO, L. AND F. PARRO (2015): “Estimates of the Trade and Welfare Effects of NAFTA,” *Review of Economic Studies*, 82, 1–44.
- CARBALLO, J., G. I. P. OTTAVIANO, AND C. V. MARTINCUS (2018): “The Buyer Margins of Firms’ Exports,” *Journal of International Economics*, 112, 33–49.

- CARON, J., T. FALLY, AND J. R. MARKUSEN (2014): “International Trade Puzzles: A Solution Linking Production and Preferences,” *Quarterly Journal of Economics*, 129, 1501–1552.
- CARVALHO, V., M. NIREI, Y. SAITO, AND A. TAHBAZ-SALEHI (2021): “Supply Chain Disruptions: Evidence from the Great East Japan Earthquake,” *Quarterly Journal of Economics*, 136, 1255–1321.
- CARVALHO, V. M. AND N. VOIGTLÄNDER (2014): “Input Diffusion and the Evolution of Production Networks,” Tech. rep., National Bureau of Economic Research.
- CHANEY, T. (2008): “Distorted Gravity: the Intensive and Extensive Margins of International Trade,” *American Economic Review*, 98, 1707–1721.
- (2014): “The Network Structure of International Trade,” *American Economic Review*, 104, 3600–3634.
- (2016): “Liquidity Constrained Exporters,” *Journal of Economic Dynamics and Control*, 72, 141–154.
- (2018): “The Gravity Equation in International Trade: An Explanation,” *Journal of Political Economy*, 126, 150–177.
- DI GIOVANNI, J., A. A. LEVCHENKO, AND I. MEJEAN (2014): “Firms, Destinations, and Aggregate Fluctuations,” *Econometrica*, 82, 1303–1340.
- DI GIOVANNI, J., A. A. LEVCHENKO, AND R. RANCIERE (2011): “Power Laws in Firm Size and Openness to Trade: Measurement and Implications,” *Journal of International Economics*, 85, 42–52.
- EATON, J., M. ESLAVA, D. JINKINS, C. KRIZAN, AND J. TYBOUT (2014): “A Search and Learning Model of Export Dynamics,” *Society for Economic Dynamics Meeting Papers No. 1535*.
- EATON, J., D. JINKINS, J. TYBOUT, AND D. XU (2016): “Two-Sided Search in International Markets,” *Annual Meeting of the Society for Economic Dynamics*.
- EATON, J., S. KORTUM, AND F. KRAMARZ (2019): “Firm-to-Firm Trade: Imports, Exports, and the Labor Market,” *Society for Economic Dynamics Meeting Papers No. 702*.
- EATON, J., S. S. KORTUM, AND S. SOTELO (2013): “International Trade: Linking Micro and Macro,” *Advances in Economics and Econometrics: Tenth World Congress*, II.
- ECKEL, C. AND J. P. NEARY (2010): “Multi-Product Firms and Flexible Manufacturing in the Global Economy,” *The Review of Economic Studies*, 77, 188–217.
- FAJGELBAUM, P. D. AND A. KHANDELWAL (2016): “Measuring the Unequal Gains from Trade,” *The Quarterly Journal of Economics*, 131, 1113–1180.

- FALLY, T. AND R. HILLBERRY (2018): “A Coasian Model of International Production Chains,” *Journal of International Economics*, 114, 299–315.
- FIELER, A. C. (2011): “Nonhomotheticity and Bilateral Trade: Evidence and a Quantitative Explanation,” *Econometrica*, 79, 1060–1101.
- FIELER, A. C., M. ESLAVA, AND D. Y. XU (2018): “Trade, Quality Upgrading, and Input Linkages: Theory and Evidence from Colombia,” *American Economic Review*, 108, 109–46.
- GIMENEZ-PERALES, V. (2021): “The Dynamics of Importer-Exporter Connections,” *mimeo*, *University of Mainz*.
- GOLDBERG, P. K., A. K. KHANDELWAL, N. PAVCNIK, AND P. TOPALOVA (2010): “Imported Intermediate Inputs and Domestic Product Growth: Evidence from India,” *The Quarterly Journal of Economics*, 125, 1727–1767.
- GOMTSYAN, D. AND A. TARASOV (2020): “Exporting Costs and Multi-Product Shipments,” *ZEW-Centre for European Economic Research Discussion Paper*.
- GÖRG, H., P. HENZE, V. JIENWATCHARAMONGKHOL, D. KOPASKER, H. MOLANA, C. MONTAGNA, AND F. SJÖHOLM (2017): “Firm Size Distribution and Employment Fluctuations: Theory and Evidence,” *Research in Economics*, 71, 690–703.
- GROSSMAN, G. M. AND E. ROSSI-HANSBERG (2008): “Trading Tasks: A Simple Theory of Offshoring,” *American Economic Review*, 98, 1978–97.
- HALPERN, L., M. KOREN, AND A. SZEIDL (2015): “Imported Inputs and Productivity,” *American Economic Review*, 105, 3660–3703.
- HEAD, K., R. JING, AND J. RIES (2017): “Import Sourcing of Chinese Cities: Order versus Randomness,” *Journal of International Economics*, 105, 119–129.
- HEAD, K., T. MAYER, AND M. THOENIG (2014): “Welfare and Trade without Pareto,” *American Economic Review*, 104, 310–16.
- HELPMAN, E., M. MELITZ, AND Y. RUBINSTEIN (2008): “Estimating Trade flows: Trading Partners and Trading Volumes,” *The Quarterly Journal of Economics*, 123, 441–487.
- HUFBAUER, G. (1970): “The Impact of National Characteristics and Technology on the Commodity Composition of Trade,” in *The Technology Factor in International Trade*, ed. by R. Vernon, Columbia University Press.
- HUNEEUS, F. (2018): “Production Network Dynamics and the Propagation of Shocks,” *mimeo*, *Yale University*.
- KASAHARA, H. AND J. RODRIGUE (2008): “Does the Use of Imported Intermediates Increase Productivity? Plant-Level Evidence,” *Journal of Development Economics*, 87, 106–118.

- KROPF, A. AND P. SAURÉ (2014): “Fixed Costs per Shipment,” *Journal of International Economics*, 92, 166–184.
- KRUGMAN, P. (1980): “Scale Economies, Product Differentiation, and the Pattern of Trade,” *American Economic Review*, 70, 950–959.
- LILEEVA, A. AND D. TREFLER (2010): “Improved Access to Foreign Markets Raises Plant-level Productivity... for some Plants,” *The Quarterly Journal of Economics*, 125, 1051–1099.
- LIM, K. (2018): “Endogenous Production Networks and the Business Cycle,” *mimeo*, *University of Toronto*.
- LUTTMER, E. G. J. (2007): “Selection, Growth, and the Size Distribution of Firms,” *Quarterly Journal of Economics*, 122, 1103–1144.
- MAGYARI, I. ET AL. (2017): “Firm Reorganization, Chinese Imports, and US Manufacturing Employment,” *mimeo*, *Columbia University*.
- MANOVA, K. (2013): “Credit Constraints, Heterogeneous Firms, and International Trade,” *Review of Economic Studies*, 80, 711–744.
- MONARCH, R. (2016): “‘It’s not you, it’s me’: Breakups in U.S.-China Trade Relationships,” *International Finance Discussion Papers 1165*.
- MONARCH, R. AND T. SCHMIDT-EISENLOHR (2017): “Learning and the Value of Trade Relationships,” *International Finance Discussion Papers 1218*.
- OBERFIELD, E. (2018): “A Theory of Input-Output Architecture,” *Econometrica*, 86, 559–589.
- PANIGRAHI, P. (2021): “Endogenous Spatial Production Networks: Quantitative Implications for Trade and Productivity,” *mimeo*, *University of California - Berkeley*.
- RUHL, K. J. AND J. L. WILLIS (2017): “New Exporter Dynamics,” *International Economic Review*, 58, 703–726.
- SAURÉ, P. (2012): “Bounded Love of Variety and Patterns of Trade,” *Open Economies Review*, 23, 645–674.
- TASCHEREAU-DUMOUCHEL, M. (2020): “Cascades and Fluctuations in an Economy with an Endogenous Production Network,” *mimeo*, *Cornell University*.
- TINTELNOT, F., K. KIKKAWA, M. MOGSTAD, AND E. DHYNE (2021): “Trade and Domestic Production Networks,” *Review of Economic Studies*, 88, 643–668.
- UNITED NATIONS (2004): *International Merchandise Trade Statistics: Compilers Manual*, United Nations, Department of Economic and Social Affairs, Statistics Division Ser. F., New York.

Appendix

A.1 Colombian Import Data

We use Colombian import data to replicate - and in some cases expand the analysis of - several regularities the previous literature has found for different countries. In this section we describe the data.⁵¹

Our dataset comes from the Colombian statistical authority DANE and contains 2.890.117 import transactions by Colombian firms in the year 2012. The Colombian importers are identified by a national tax code and for each of their transactions, we have transaction values in USD and an HS 10-digit product code. We use cif values throughout. The data also contain the country where the product was produced as well as the country and city from which the goods were shipped to Colombia, which may be different from the country of production.

Some of the recorded import transactions represent imports for non-commercial purposes, such as experimental or exhibition samples. Another set of import transactions originates from Colombian free trade zones. We ignore both types of transactions in our analysis, thereby dropping 0.65% of all observations, or 0.66% of the total import value.

One drawback of our data is the fact that we do not observe the identities of the exporting firms. We therefore proxy an exporting firm by a combination of the location from which the goods were supplied (the country-city combination) with the 4-digit product code of the transaction. In a prior cleaning step, we make sure that we do not double-count cities within an importing firm by unifying the spelling.⁵²

In Table A.1 we present some key statistics about our dataset. We report these statistics separately for aggregate Colombian imports and for the top 5 source countries in 2012.⁵³

Total Colombian imports in 2012 amount to USD 58.2 million, where the United States alone account for 24% of that value. There is a total of 36394 importing firms in Colombia connecting with 188550 exporters. This number exceeds the figure reported by Bernard et al. (2018b) by factor of about 2.3. Consistent with this difference, we find higher values for the mean and median number of exporters per importer (12.6 and 3 in our data vs. 5.3 and 2 in Bernard et al., 2018b), but very close values for the mean and median number of importers per exporter (2.4 and 1 in our data, 1.8 and 1 in Bernard et al., 2018b). Not surprisingly, we find that the mean (median) value of

⁵¹We describe the Colombian data as presented and analyzed in Gimenez-Perales (2021). The dataset is a variant of the one used in Bernard et al. (2018b).

⁵²Bernard et al. (2018b) are able to identify exporting firms in their data because they have direct access to the data from the Colombian tax and customs authority DIAN, which provides the names of the exporting firms. They employ a machine learning algorithm to group common misspellings or spelling variants in order to avoid double-counting exporters. Using finer levels of aggregation, we find that we can still replicate all the stylized facts reported below. We present our main results at the 4-digit level because at this level of aggregation, the total number of “exporters” in our data is closest to the number reported in Bernard et al. (2018b).

⁵³For the identification of the top 5 source countries, we use the variable representing the country of production. If we were to use the country from which the goods are supplied, Panama would show up as one of the top 5 source countries because Panama is an important hub for intermediaries. In the data, 98% of products supplied from Panama are actually produced in a different country.

Table A.1: Colombian Import Data - Summary Statistics

Country	All	United States	China	Mexico	Brasil	Argentina
Total value in USD (millions)	58217.70	13985.63	9816.75	6453.87	2859.67	2401.56
# Colombian importers	36293	15290	16799	5619	3657	1693
# foreign exporters	188550	59291	55951	10343	8187	1826
Mean value per importer-exporter (\$'000s)	127.08	112.47	57.59	339.39	180.41	473.31
Median value per importer-exporter (\$'000s)	4.02	3.47	2.86	5.65	6.14	13.10
Mean no. of exporters per importer	12.62	8.13	10.15	3.38	4.33	3.00
Median no. of exporters per importer	3.00	2.00	3.00	1.00	2.00	1.00
Mean no. of importers per exporter	2.43	2.10	3.05	1.84	1.94	2.78
Median no. of importers per exporter	1.00	1.00	1.00	1.00	1.00	1.00

Note: Data are for the year 2012. Exporters are defined as a combination of a 4-digit HS product code and the city from which the product is shipped.

trade per importer-exporter pair at USD 127080 (USD 4020) is only about one third (one sixth) of the value found in Bernard et al. (2018b).

This appendix describes the data, provides additional empirical results as well as proofs for some of our theoretical results.

A.2 Margins Decomposition

Table A.2 shows a margins decomposition of aggregate Colombian imports into the number of importers, the number of exporters, the number of products, average imports per importer-exporter-product and a density term defined according to Bernard et al. (2018c) as the share of active connections as a share of all possible connections.

Table A.2: Trade Margins and Aggregate Colombian Imports

	Importers	Exporters	Products	Density	Intensive
log(Imports)	0.502*** (0.018)	0.542*** (0.020)	0.485*** (0.016)	-0.932*** (0.031)	0.404*** (0.023)
Constant	-3.764*** (0.258)	-3.944*** (0.281)	-3.362*** (0.226)	6.689*** (0.445)	4.381*** (0.313)
R-squared	0.784	0.772	0.783	0.786	0.612
Observations	202	202	202	202	202

Note: * p<0.05, ** p<0.01, *** p<0.001

A.3 Proof of Proposition 4: Connections with at Least One Large Firm

In this theory appendix we show how we calculate the share of connections involving at least one large firm. We proceed in two broad steps:

1. Compute the mass of all recorded connections

$$M_0 = \int_{\underline{\mu}}^{\infty} \int_{\underline{\mu}}^{\infty} f(\mu_i, \mu_j | \mu_i \mu_j > \bar{\mu}_{ji}) d\mu_i d\mu_j \quad (\text{A.1})$$

2. Compute the mass all recorded connections including at least one large firm

$$M_1^X = \int_{\underline{\mu}}^{\infty} \int_{\underline{\mu}}^{\infty} f(\mu_i, \mu_j | \mu_i \mu_j > \bar{\mu}_{ji} \ \& \ (\mu_i \geq \mu_X \ | \ \mu_j \geq \mu_X)) d\mu_i d\mu_j \quad (\text{A.2})$$

Step 1: M_0 . We compute M_0 as the total mass of connections (the grey rectangle in Figure 1), which equals one minus the mass in the area A. To that aim, we define the cutoff value for μ_j above which even the smallest μ_i are registered as partners:

$$\mu_0 = \bar{\mu}_{ji} / \underline{\mu} \quad (\text{A.3})$$

To compute the mass of the area A in the figure, we need to integrate over the ranges defined by the borders of A, which are $[\underline{\mu}, \mu_0]$ for μ_j and $[\underline{\mu}, \bar{\mu}/\mu_j]$ for μ_i :

$$\begin{aligned}
M_0 &= 1 - \int_{\underline{\mu}}^{\mu_0} f(\mu_j) \left(\int_{\underline{\mu}}^{\bar{\mu}_{ji}/\mu_j} f(\mu_i) d\mu_i \right) d\mu_j \\
&= 1 - \int_{\underline{\mu}}^{\mu_0} f(\mu_j) \left(1 - (\bar{\mu}_{ji}/(\mu_j \underline{\mu}))^{-\theta} \right) d\mu_j \\
&= 1 - \left[1 - (\mu_0/\underline{\mu})^{-\theta} - \theta (\bar{\mu}_{ji}/\underline{\mu}^2)^{\theta} \int_{\underline{\mu}}^{\mu_0} 1/\mu_j d\mu_j \right] \\
&= (\mu_0/\underline{\mu})^{-\theta} + \theta (\bar{\mu}_{ji}/\underline{\mu}^2)^{\theta} \ln(\mu_0/\underline{\mu})
\end{aligned}$$

or

$$M_0 = (\bar{\mu}_{ji}/\underline{\mu}^2)^{-\theta} \left[1 + \theta \ln(\bar{\mu}_{ji}/\underline{\mu}^2) \right] \quad (\text{A.4})$$

where we used Equation (A.3) and $\int_{\underline{\mu}}^{\infty} f(\mu) d\mu = 1$, keeping in mind that $\underline{\mu} = (\theta - 1)/\theta$.

We note that this expression is decreasing in $\bar{\mu}_{ji}$:

$$\frac{d}{d\bar{\mu}_{ji}} M_0 < 0 \quad (\text{A.5})$$

To see this, observe that $x^{-1}(1 + \ln(x))$ is decreasing for $x > 1$ and substitute $x = (\bar{\mu}_{ji}/\underline{\mu}^2)^{\theta}$.

In Figure 1, an increase in $\bar{\mu}_{ji}$ amounts to a shift of the curved line to the upper right, decreasing the area that encompasses registered connections and thus their mass.

Step 2: M_1^X . There are two cases to distinguish: one where $\bar{\mu}_{ji}$ is relatively small so that $\mu_0 \leq \mu_\chi$ (the according line is marked by $\bar{\mu}_{ji}$ in Figure 1) and the other where $\bar{\mu}_{ji}$ is relatively large so that $\mu_0 > \mu_\chi$ (marked by $\bar{\mu}'_{ji}$ in Figure 1). In the first case, the mass of connections involving at least one large firm is simply one minus the mass involving no large firm (one minus the mass in areas A+B+C in Figure 1), i.e.,

$$M_1^X = 1 - \chi^2 \quad (\text{A.6})$$

The share of connections that involves at least one large firm is then M_1^X/M_0 . By our observation above that M_0 is increasing in $\bar{\mu}_{ji}$, this share must increase in $\bar{\mu}_{ji}$. In this case, the ratio

$$m^X = M_1^X/M_0$$

is obviously increasing in $\bar{\mu}_{ji}$ by (A.5). Thus, Proposition 4 holds for this first case.

In the second case, when $\mu_0 \geq \mu_\chi$ holds, we define the cutoff value below which μ_i registers trade only with large μ_j as

$$\mu_{00} = \bar{\mu}_{ji}/\mu_\chi \quad (\text{A.7})$$

The mass of registered connections that do not involve large firms is then the total mass of registered firms from Equation (A.4) minus the mass of firms within the range

indicated by area C Figure 1. Collecting the relevant borders of the integrals, we compute

$$\begin{aligned}
M_0 - M_1^X &= \int_{\mu_{00}}^{\mu_X} f(\mu_j) \left(\int_{\bar{\mu}_{ji}/\mu_j}^{\mu_X} f(\mu_i) d\mu_i \right) d\mu_j \\
&= \int_{\mu_{00}}^{\mu_X} f(\mu_j) \left(-(\mu_X/\underline{\mu})^{-\theta} + (\bar{\mu}_{ji}/(\mu_j\underline{\mu}))^{-\theta} \right) d\mu_j \\
&= -(\mu_X/\underline{\mu})^{-\theta} \left[-(\mu_X/\underline{\mu})^{-\theta} + (\mu_{00}/\underline{\mu})^{-\theta} \right] + \theta(\bar{\mu}_{ji}/\underline{\mu}^2)^{-\theta} \left[\int_{\mu_{00}}^{\mu_X} 1/\mu_j d\mu_j \right] \\
&= (\mu_X/\underline{\mu})^{-2\theta} - (\bar{\mu}_{ji}/\underline{\mu}^2)^{-\theta} + \theta(\bar{\mu}_{ji}/\underline{\mu}^2)^{-\theta} \ln(\mu_X^2/\bar{\mu}_{ji})
\end{aligned}$$

so that with Equation (A.4), we have

$$M_1^X = (\bar{\mu}_{ji}/\underline{\mu}^2)^{-\theta} \left[2 + 2\theta \ln(\bar{\mu}_{ji}/(\underline{\mu}\mu_X)) \right] - (\mu_X/\underline{\mu})^{-2\theta} \quad (\text{A.8})$$

The share of connections that involves at least one large firm is M_1/M_0 or

$$m^X = M_1^X/M_0 = \frac{2 \left[1 + \theta \ln(\bar{\mu}_{ji}/(\underline{\mu}\mu_X)) \right] - (\mu_X^2/\bar{\mu}_{ji})^{-\theta}}{1 + \theta \ln(\bar{\mu}_{ji}/\underline{\mu}^2)} \quad (\text{A.9})$$

This ratio is increasing in $\bar{\mu}_{ji}$:

$$\frac{d}{d\bar{\mu}_{ji}} m^X > 0 \quad (\text{A.10})$$

To see this, substitute $x = (\bar{\mu}_{ji}/\mu_X^2)^\theta$ (implying $x \in (0, 1]$) and rewrite the equation as m^X

$$m^X = 2 - \frac{2\theta \ln(\mu_X/\underline{\mu}) + x}{2\theta \ln(\mu_X/\underline{\mu}) + 1 + \ln(x)} \quad (\text{A.11})$$

and use that $(a + 1 + \ln(x))/(a + x)$ is increasing in x for the relevant case $x \in (0, 1]$ and $a > 0$.

The statement of Proposition 4 holds for this second case as well, which completes the proof.