



Full length articles

Volatility in the small and in the large: The lack of diversification in international trade[☆]

Francis Kramarz^a, Julien Martin^b, Isabelle Mejean^{c,*}^a CREST-ENSAE and CEPR, Palaiseau, France^b Université du Québec à Montréal, CREST, and CEPR, Quebec, Canada^c CREST-Ecole Polytechnique and CEPR, Palaiseau, France

ARTICLE INFO

Article history:

Received 26 June 2018

Received in revised form 5 September 2019

Accepted 30 October 2019

Available online 22 November 2019

Research data related to this submission: http://isabellemejean.com/ReplicationPackage_KMM_JIE.html

ABSTRACT

How does international trade affect the risk exposure of firms and countries? Trade induces specialization, thus increasing economies' exposure to idiosyncratic supply shocks. But greater geographic diversification in trade destinations offers natural hedging properties against demand shocks. In this paper, we offer an integrated economic and econometric view of the impact of trade on firms and countries volatility. Exporters' volatility is shown to directly depend on the (lack of) diversification in their portfolio of clients. Indeed, most exporters, including the largest, have one or two main clients that dwarf the others. This structure of trade networks implies that individual exporters are strongly exposed to microeconomic demand shocks. The concentration of trade flows further implies that such risk does not wash out across firms, thus contributing to aggregate fluctuations.

© 2019 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

This paper presents an integrated analysis of the sources of volatility of firms and countries in international trade markets. Both individual and aggregate exports are shown to be strongly exposed to shocks hitting firms' foreign clients. Hence, micro-economic foreign shocks appear to be a key driver of the volatility of exports at the firm-level as well as in the aggregate.

The results contribute to the broad literature on the impact of international trade on the risk exposure of firms and countries. First, greater participation into international markets, by increasing the role of large firms as well as sectoral concentration, has been shown to magnify a country's vulnerability to idiosyncratic supply shocks (di Giovanni

and Levchenko, 2009, 2012). However, cross-country diversification dampens macro-economic volatility by reducing exposure to domestic demand shocks (Caselli et al., 2015). Focusing on micro-demand shocks, Kelly et al. (2013) document that firms with a more diversified customer base display a lower volatility. Hence, to fully understand the trade-volatility nexus, the structure of exporters' sales, within as well as across destinations, must be analyzed together with the various shocks that hit firms and countries in international markets.

This paper provides a step in this direction by modeling the universe of trade relationships between French exporters and their European partners, observed over a period of 15 years. The richness of the data together with a new empirical strategy allow us to provide a comprehensive decomposition of the different sources of trade volatility at various levels of aggregation.

We start with an economic model of firm-to-firm sales in which the growth of firm-to-firm exports is a combination of macro shocks and micro perturbations.¹ French sellers' competitiveness abroad varies

[☆] This paper is a thoroughly revised version of the paper entitled "Volatility in the Small and in the Large: Evidence from Trade Networks". The paper has benefited from comments at various seminars and conferences. We are particularly grateful to Ezra Oberfield, Felix Tintelnot, Glenn Magerman, the editor and two anonymous referees for their insightful suggestions. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714597). Our work is also supported by a public grant overseen by the French National Research Agency (ANR) as part of the "Investissements d'Avenir" program (Idex Grant Agreement No. ANR-11-IDEX-0003-02 / Labex ECODEC No. ANR-11-LABEX-0047). Martin thanks FRQSC for financial support (grant 2015-NP-182781).

* Corresponding author.

E-mail address: isabelle.mejean@polytechnique.edu (I. Mejean).

¹ In this paper, the distinction between macro and micro shocks is somewhat unconventional since the "macro" shocks are indeed allowed to vary across sectors. This is meant to emphasize the difference with "micro" shocks, which are assumed specific to an individual firm. Likewise, the distinction between supply and demand shocks is dictated by the particular market under study. In this market, French exporters represent the supply-side while European importers are at the root of a demand curve. "Demand" shocks are thus defined as shocks affecting foreign importers, a preference shock in our model, while "supply" shocks are those affecting individual exporters, e.g. cost shifters.

because of shocks to unit costs, assumed to be common across firms in an industry, and because of idiosyncratic shocks to their productivity. These supply-side perturbations are at the root of variations in the demand expressed by all foreign partners to which the firm is connected. Firm-to-firm trade also varies due to demand shocks that are either aggregate in nature (i.e. common to all buyers within a destination) or idiosyncratic, hitting the importing firm or eventually the match it forms with a particular French exporter. We show how to identify these shocks using various moment conditions dictated by the model. Importantly, the flexible estimation does not place constraints on heterogeneity of idiosyncratic shocks, be they seller-, buyer- or match-specific. Their variance can be heterogenous in the cross-section because of the identity of the firm (eg. wholesaler versus manufacturing firm) or the nature of transaction (eg. arm's length vs intra-firm trade).

These theoretically-justified orthogonality conditions constitute the first element used to econometrically identify the model's shocks. Another identification condition, of econometric nature, comes from exploiting the structure of our data set, a bipartite graph with (French) exporting firms on one side and (European) importers on the other side. In the cross-section, the largest connected component of this graph encompasses the vast majority of trade flows because exporters tend to simultaneously serve several foreign buyers with the same product while foreign buyers often purchase various products to different French suppliers. Such connectedness is useful inasmuch as it helps us separately identify the model's shocks using high-dimensional fixed effects estimators.

Having estimated the different shocks hitting firms and countries in trade, these are combined with the micro-structure of trade flows to quantify the importance of various shocks for the volatility at different levels of aggregation. The aggregation exercise is straightforward. Volatility is a weighted average of the variances and covariances in firm-to-firm growth rates. The weights depend on the structure of trade networks and the level of aggregation (firm-level or aggregate level here). Their distribution has natural hedging properties against some of the shocks hitting firms in international markets.²

At the level of individual firms, we show that aggregate shocks, whether they hit the supply or the demand side of the market, do not generate much volatility. Their variance is indeed small, in comparison with microeconomic risks. By contrast, both microeconomic supply and demand shocks generate a substantial amount of volatility. In the baseline regression, they respectively represent 45 and 50% of the average firm's volatility. These contributions are stable across various empirical models based on slightly different identification assumptions that correspond to varying market structures. That microeconomic supply shocks are an important source of firm-level volatility is consistent with existing evidence and is not surprising as such shocks cannot be diversified (see eg. Comin and Mulani, 2006). What is more surprising is the substantial contribution of microeconomic demand shocks to firm-level volatility, which could have naturally faded away when the firm diversifies its portfolio of clients. These shocks are quantitatively important because i) their estimated volatility is high; and ii) exporters'

² In doing so, we take the structure of trade networks as given, although potentially time-varying. Hence, we do not intend to tackle the question of how shocks to firms in international markets affect firm-to-firm relationships at the extensive margin. This question is treated in models of endogenous trade networks such as Bernard et al. (2019) who study the impact of a long-term shock, namely the opening of a high-speed line, on the creation of buyer-seller relationships. Instead, we condition on the realized structure of networks and study how (temporary) shocks affect the dynamics of trade relationships, at the firm-to-firm level, for individual exporters and in the aggregate. From that point of view, we follow the methodology used in the empirical granularity literature, e.g. Gabaix (2011) or di Giovanni et al. (2014).

sales are highly skewed towards their main partner.³ Interestingly, such a lack of diversification is observed in supposedly well-integrated markets, namely good markets within the European Union. Since a firm's number of foreign clients is a decreasing function of the destination's market potential (Bernard et al., 2018), we conjecture that the lack of diversification, and its impact on the volatility of exports, is even stronger outside of the European Union.

Selection into exports implies that large firms on average serve more markets and more buyers within a country (Mayer and Ottaviano, 2008; Bernard et al., 2018). As a consequence, they are better diversified and should display less volatility. We confirm this intuition in our sample. The impact of demand shocks is half as large in the top decile of the firm size distribution as it is in the first decile. This contributes to explaining why large firms' sales display less volatility, on average.⁴ However, this negative correlation between size and volatility is less pronounced than could be expected. Microeconomic demand shocks continue to matter substantially at the top of the distribution because large firms also have skewed export sales, concentrated on a small number of partners.

Large firms' lack of diversification combined with the connectedness of the largest foreign buyers implies that microeconomic demand shocks contribute substantially to the volatility of aggregate exports. In the European Union, we show that they explain 33% of the variance of export growth, to be compared with a contribution of individual supply shocks of 34% and of aggregate shocks of 62%.⁵ Hence, individual shocks to foreign clients do not vanish in the aggregate. Instead, they contribute substantially to the level of aggregate volatility of bilateral and multilateral French exports. Their quantitative impact also depends on the microstructure of trade networks. Differences in exporters' size and diversification as well as heterogeneity in the connectedness of foreign buyers explain about a third of the dispersion in the volatility of exports across destinations. This points to introducing this source of fluctuations in more sophisticated models of trade openness and volatility.

Related literature. Our paper is related to several strands of literature. As mentioned earlier, the impact of trade openness on macroeconomic volatility is an important topic in the macro-development literature, see e.g. Koren and Tenreyro (2007), Caselli et al. (2015), di Giovanni and Levchenko (2009), di Giovanni and Levchenko (2012). By estimating a rich structure of shocks, we believe we contribute to this literature through a precise quantification of the respective contributions of various sources of risks encountered by firms in their export

³ For demand shocks to induce a substantial amount of volatility, both i) and ii) need to prevail in the data. Having a diversified portfolio of clients is neither a necessary nor a sufficient condition for having little exposure to microeconomic demand shocks. Individual exporters might end up with little exposure to such shocks while interacting with a small number of clients if those clients display little volatility in purchases and/or if they face negatively correlated shocks. Our empirical analysis is ex-ante agnostic about this possibility since we do not impose any structure on the cross-sectional correlations nor on the amount of volatility and its heterogeneity across individuals. Ex-post, we estimate a negative correlation between the volatility induced by microeconomic demand shocks and the skewness in firms' sales. This correlation implies that, if firms compensate a skewed portfolio of sales by the interaction with less volatile clients, the strategy is not pronounced enough so that it annihilates the mechanical impact of a poorly diversified portfolio on the firms' exposure to microeconomic demand risk.

⁴ This is consistent with evidence in Kurz and Senses (2016), obtained from different data using a different empirical strategy. They show that firm-level employment is significantly less volatile in i) exporting firms as opposed to non-exporting firms; and ii) firms that export to a larger number of destinations. Both results are consistent with a reduction in volatility for firms with a more diversified portfolio of clients.

⁵ Note that the contributions do not sum to 100% because of empirical covariance terms. The relative contributions of various types of shocks vary a bit when we estimate the structural shocks using slightly different identification assumptions as commented in the text. In all specifications, the contribution of microeconomic demand shocks is substantial though.

activities.⁶ In that respect, the exercise is similar to Caselli et al. (2015) who also compare the contribution of various shocks, although their model does not have heterogeneous firms thus neglecting the potential role of microeconomic risks emphasized in di Giovanni and Levchenko (2012).⁷ Microeconomic shocks have been shown to be an important source of aggregate fluctuations in the granularity literature initiated by Gabaix (2011) and Acemoglu et al. (2012). These seminal findings find strong support in the data (see eg. di Giovanni et al., 2014; Magerman et al., 2016). While this literature emphasizes the role of idiosyncratic supply shocks, our decomposition also takes into account microeconomic demand risk. Moreover, we tackle the impact of such shocks on aggregate fluctuations, as well as on firm-level volatility.

We also emphasize the role of demand shocks for individual firms' volatility. In doing so, we contribute to the literature on firm-level volatility. This literature documents a large amount of volatility in firm-level data, irrespective of the measure of performance adopted.⁸ Recent papers further document systematic differences in the volatility of domestic and exporting firms (Buch et al., 2009; Vannoorenberghe, 2012; Vannoorenberghe et al., 2016). While our analysis focuses on already exporting firms, the empirical strategy makes it possible to offer new results regarding the heterogeneity in exporters' exposure to a variety of shocks.

The use of firm-to-firm trade flows naturally creates a connection with the recent trade literature exploiting comparable data and trying to rationalize the structure of trade networks observed in the data (see Carballo et al., 2018; Bernard et al., 2018). Our analysis takes a different track. Instead of studying the structure of trade networks, we take the structure as given and study what it implies in terms of their exposure to shocks. A consequence of this is the purely positive nature of the exercise. Given the structure of trade networks and the amount of natural hedging properties it entails, we study the resilience of trade to various shocks. Another interesting question which the paper does not answer is why does the structure of trade networks displays such high degree of concentration.

The rest of the paper is organized as follows. Section 2 describes the theoretical framework and empirical strategy. Section 3 presents the data and some stylized facts. Section 4 presents the results, starting with a summary of the estimated shocks in subsection 4.1 and, then, a discussion of the relative contribution of shocks to the volatility of individual firms (the volatility "in the small") and the volatility of aggregate exports ("in the large") (Subsections 4.2–4.4). Finally, Section 5 concludes.

⁶ The strategy for identifying shocks can be interpreted as an extension of di Giovanni et al. (2014). In comparison, the additional firm-to-firm dimension allows us to go deeper into the analysis of the microeconomic origins of aggregate fluctuations and statistically separate seller-related and buyer-specific sources of risk. Such a strategy is in contrast with the burgeoning and complementary literature studying the propagation of shocks in firm-to-firm networks, by exploiting natural disasters to trace the propagation of well-identified shocks, see e.g. Boehm et al. (2019), Barrot and Sauvagnat (2016) and Carvalho et al. (2016). While this approach has important advantages when it comes to identifying the propagation of shocks, such a strategy does not attempt to compare the relative importance of several sources of risk, which our approach does.

⁷ Note that this literature is interested in the impact of trade on countries' overall volatility while our strategy forces us to focus on the volatility of exports. In the data, both objects are highly correlated. At the level of countries, the correlation between the volatility of exports and that of GDP growth rates is as high as 90% (Source: Penn World Tables). In our firm-level data set, the cross-sectional correlation between measures of firms' overall sales volatility and export sales volatility is lower, equal to 23%, but still positive and significant. Because EU exports represent the vast majority of exports for most firms in our data, the correlation of volatilities in firm-level sales and EU exports, the object of study in the paper, is very comparable, at 17%.

⁸ Comin and Philippon (2006) document the increase in firm volatility using real measures (sales, employment and capital expenditures) as well as financial data (equity returns). See also Comin and Mulani (2006) and Asker et al. (2014) on sales data, Thesmar and Thoenig (2011) based on sales and employment, Campbell (2001) using stock returns.

2. Empirical strategy

The empirical analysis involves two steps. First, we use panel data on firm-to-firm export growth to recover the shocks hitting firms in international markets. In practice, this amounts to estimating a high-dimensional fixed effects equation in the cross-section of growth rates, for each available year. Second, the recovered time-series of shocks are aggregated within and across firms to compute measures of volatility "in the small" and "in the large", and their components.

The underlying model is described in Section 2.1. The strategy to recover the model's shocks from the data is detailed in Section 2.2. The aggregation procedure can be found in Section 2.3.

2.1. Theoretical framework

To understand the economic nature of the orthogonality conditions at the root of our estimation method, this section develops a model of firm-to-firm growth. Its structure is inspired from Atkeson and Burstein (2008), which we complement with various shocks in order to derive predictions on the dynamics of firm-to-firm trade. These shocks are treated as deep structural parameters and are thus assumed orthogonal to each other in the model and the empirical strategy. The structure of the model itself however implies that their impact on firm-to-firm growth is not necessarily uncorrelated, a possibility that the empirical strategy takes into account as explained in Section 2.2.

The analysis is kept in partial equilibrium and is used to derive the growth of sales by a firm s to an importer b located in a country c ($b = i$). In the empirical analysis, all sellers s are located in France, which explains that the model's notations do not specify the country of origin of firm s .

The demand side of the model features a nested CES structure. At the aggregate level, a competitive firm in country i produces a final consumption basket using the output y_b of a continuum of goods $b \in [0, 1]$ as inputs of a CES production function:

$$c_i = \left[\int_0^1 (a_b y_b)^{\frac{\eta-1}{\eta}} db \right]^{\frac{\eta}{\eta-1}}$$

To simplify notation and without loss of generality, goods are assumed to be produced locally, i.e. $c(b) = i \forall b$. a_b denotes a preference parameter associated with the variety produced by firm b and will later be allowed to vary randomly. η denotes the elasticity of substitution between varieties. The representative consumer chooses quantities y_b given a vector of prices P_b and overall demand A_i . Shocks to A_i are the aggregate demand shifters of the model.

At the lower-level of aggregation, firms produce using imperfectly substitutable inputs, that are either produced in-house or purchased from upstream firms s .⁹ Taking the firm's sourcing strategy as given, one can write firm b 's output as¹⁰:

$$y_b = \left[\sum_s (a_{sb} q_{sb})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$$

where q_{sb} is the quantity purchased from seller s , a_{sb} is a random preference parameter associated with variety s and σ is the elasticity of substitution between varieties. Shocks to a_{sb} and a_b are the microeconomic

⁹ Inputs produced in-house are thus such that $s = b$. Domestic inputs are such that s is located in country i . In our data, we only observe upstream firms located in France, i.e. we observe the component of input purchases sourced from French sellers.

¹⁰ See Antràs et al. (2017) for a theoretical model of firms' sourcing strategies. In the empirical analysis, the number of input providers will be allowed to vary over time but we will maintain the assumption that such extensive adjustments are independent of the shocks driving the dynamics of trade. See di Giovanni and Levchenko (2012) for a similar treatment of extensive margin adjustments.

demand shocks of the empirical framework. The price index associated with y_b defines as:

$$P_b = \left[\sum_s \left(\frac{p_{sb}}{a_{sb}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$$

The technology to produce inputs is assumed to be linear in (equipped) labor, the price ω of which is considered exogenous. Labor productivity is modeled as the product of an aggregate component Z and a firm-specific term z_s . Shocks to these productivity components are the supply shocks of the model. In practice, the aggregate TFP shock varies across sectors. The aggregate supply shock thus includes sectoral perturbations. Given an (endogenous) mark-up μ_{sb} , the price of inputs purchased by buyer b from country $c(b) = i$ to seller s can thus be written:

$$p_{sb} = \mu_{sb} \frac{\tau_i \omega}{z_s}$$

with τ_i , the trade cost, assumed constant in what follows.

Following Atkeson and Burstein (2008), we assume that; i) inputs are imperfect substitutes, $\sigma < \infty$; ii) inputs are more substitutable than goods, $1 < \eta < \sigma$; and iii) sellers play a static game of quantity competition. Namely, each seller chooses a quantity q_{sb} to be sold to buyer b taking as given the quantities chosen by the other firms, as well as the wage ω and aggregate demand A_i . As demonstrated in Atkeson and Burstein (2008), the vector of equilibrium prices implies mark-ups that are increasing in the firm's market share:

$$\mu_{sb} = \frac{\varepsilon(w_s^b)}{\varepsilon(w_s^b) - 1} \quad \text{where} \quad \varepsilon(w_s^b) = \left[\frac{1}{\sigma} (1 - w_s^b) + \frac{1}{\eta} w_s^b \right]^{-1}$$

and $w_s^b \equiv \left(\frac{p_{sb}}{a_{sb} P_b} \right)^{1-\sigma}$ is the share of seller s in buyer b 's input purchases.

Together, these assumptions imply that the growth of the demand for imports at the firm-to-firm level ($g_{sb,t} \equiv d \ln(p_{sb,t} q_{sb,t})$), the object of interest in the rest of the analysis, decomposes as follows:

$$g_{sb,t} = (1-\sigma)d \ln \frac{\omega_t}{Z_t} + d \ln A_{i,t} + (\sigma-1)d \ln z_{s,t} + (\eta-1)d \ln a_{b,t} + (\sigma-1)d \ln a_{sb,t} + (\sigma-\eta)d \ln p_{b,t} + (1-\sigma)d \ln \mu_{sb,t} \quad (1)$$

The first two lines in eq. (1) form a linear combination of the five (orthogonal) structural shocks fed into the model. The last two terms are endogenous variables that depend on the whole vector of (quality-adjusted) prices entering buyer b 's input purchases. Our objective is to recover the different components in eq. (1). We now explain how we set up the orthogonality conditions to achieve this objective.

2.2. Estimation strategy

The presence of the markup and the price index components in eq. (1) makes the decomposition a complicated function of the different shocks. To make this equation amenable to estimation, we impose additional assumptions regarding the market structure that imply a set of orthogonality conditions used to recover the shocks.

2.2.1. Benchmark model

In our benchmark estimation, we assume that French sellers account for a small fraction of foreign buyers' input bundles. This makes each single input supplier atomistic $w_s^b \approx 0$ in which case mark-ups are constant ($d \ln \mu_{sb,t} = 0$). Moreover, the diversification of buyers' input bundle implies that shocks to individual prices compensate so that the price index is roughly constant ($d \ln P_{b,t} \approx 0$). Such assumption is plausible as long as French inputs account for a small fraction of foreign firms' total

inputs, either because foreign buyers mostly rely on domestic inputs or because they import from a large array of different locations.

Under this assumption, eq. (1) simplifies into a linear combination of five orthogonal components which represent the five sources of shocks in the model of Section 2.1¹¹:

$$\begin{aligned} g_{sb,t} &= (1-\sigma)d \ln \frac{\omega_t}{Z_t} + d \ln A_{i,t} \\ &+ (\sigma-1)d \ln z_{s,t} + (\eta-1)d \ln a_{b,t} + (\sigma-1)d \ln a_{sb,t} \\ &\equiv \varepsilon_{i,t} + \varepsilon_{s,t} + \varepsilon_{b,t} + \varepsilon_{sb,t} \end{aligned} \quad (2)$$

In the following, we show how to use high-dimensional fixed effects estimators to recover the elements of eq. (2), in the cross-section of firm-to-firm growth rates. The empirical strategy takes inspiration from the labor literature, which uses the same type of log-linear decomposition to disentangle the worker-effect from the firm-effect at the root of the wage dispersion (Abowd et al., 1999). The model in eq. (2) has a similar structure except that the underlying bipartite graph is made of sellers and buyers instead of workers and firms. Two differences must be noted. First, whereas the labor literature exploits the mobility of workers across firms to separately identify the two types of effects, the structure of the network where a seller can sell to multiple buyers and a buyer can buy from multiple sellers allows separate identification of seller-effects from buyer-effects even in the cross-section of growth rates as long as all buyers and sellers are connected. This implies that seller and buyer effects can be estimated for every (overlapping) couple of years.¹² Second, the labor literature, in which the decomposition is essentially statistical, does not impose orthogonality between the worker and the firm-effects. In our framework, by definition, the different shocks are orthogonal to each other which forces us to augment the high-dimensional fixed effects estimator used in the labor literature with additional moment conditions, as presented just below.

Eq. (2) together with the orthogonality assumptions between the buyer, seller, and match effects can be expressed using the following moment conditions:

$$\left. \begin{aligned} E[\varepsilon_{sb,t} | b, s] &= 0 \\ E[\varepsilon_{b,t} | s] &= 0 \\ E[\varepsilon_{s,t} | b] &= 0 \end{aligned} \right\} \forall t \quad (3)$$

The first moment condition is the equivalent of the exogeneity condition in Abowd et al. (1999) adapted to our network environment. Conditioning on b and s in this setting means that we take as given the position of a given node in the network, which is consistent with the assumption of exogenous sourcing in Section 2.1. In practice, this amounts to conditioning on the matrix of dummy variables for sellers and buyers (see details in Appendix A.1). The next two moment conditions are specific to our structural model as just discussed when presenting (2). Together, these three moment conditions reflect the model's structural assumptions, in which the drivers of the seller, buyer, and seller-buyer components are orthogonal to each other and induced by zero-mean shocks. Would we restrict our attention to the first set of moment condition, the model could be estimated using the algorithm proposed by Abowd et al. (2002). The OLS solution to this problem allows identifying all components in eqs. (2), up to a normalization constraint, as long as the graph is connected (see details in Appendix A.2). The Abowd et al. (2002) estimator does not however impose that the seller and buyer components are orthogonal to each others, i.e. is not consistent with the last two conditions in (3). Instead,

¹¹ In the following, we will not try to separate the supply and demand shocks entering the aggregate component $\varepsilon_{i,t}$ of this equation. The reason is that i) aggregate shocks account for a relatively small share of volatility in our results and ii) the decomposition of $\varepsilon_{i,t}$ into its supply $\left((1-\sigma)d \ln \frac{\omega_t}{Z_t} \right)$ and demand $(d \ln A_{i,t})$ components is no longer possible under the alternative specifications used to assess the robustness of the benchmark model.

¹² See Bernard et al. (2018) for a similar application of insights from the labor literature to firm-to-firm production networks.

we thus rely on a three-step procedure in which firm-to-firm growth is first regressed on a set of buyer effects, the residual of which is then orthogonalized in the seller dimension, before the three recovered components are orthogonalized in the sector-country dimension to recover zero-mean growth components. The aggregate component is given by the difference between the observed growth rate and the estimated individual components. See details in Appendix A.2.

2.2.2. Sources of endogeneity and alternative models

The benchmark model neglects any potential feedback effect of shocks through the buyer's price index or buyer-specific mark-ups. This implies a threat to identification due to potential endogeneity of the buyer- and seller-specific components. Appendices A.3–A.4 present two alternative models that exploit slightly different sets of moment conditions to account for these potential sources of endogeneity. Results based on these alternative models help assess the robustness of estimates recovered from the benchmark case in (2). While the details are left for the Appendix, we will now discuss the economic sources of such endogeneity and the way we control for it.

The first threat to the identification of shocks in eq. (2) is due to the buyer-specific component absorbing the impact of the price index entering eq. (1). Whereas the benchmark model assumes this impact to be negligible, the correlation of this price index with some of its components might be substantial at some points of the network. To see why, notice that the growth of the buyer-specific input price index can be approximated as a weighted average of seller-specific (adjusted) price changes:

$$d \ln P_{b,t} \approx \sum_s w_{s,t-1}^b \left(d \ln \frac{\omega_t}{Z_t} - d \ln z_{s,t} - d \ln a_{sb,t} \right)$$

where $w_{s,t-1}^b$ is the share of seller s in buyer b input purchases and we continue to assume constant mark-ups for simplicity.

If the contribution of French sellers to a buyer's input purchases is small enough (i.e. if the sum of $w_{s,t-1}^b$ across all sellers located in France is close to 0), the price index growth is orthogonal to the different terms entering eq. (2) and neglecting its impact does not constitute a threat to identification. At most, it affects the interpretation of the estimated $\varepsilon_{b,t}$ component, which in part depends on supply shocks affecting the buyer's input suppliers located outside of France. This is no longer true if one or several French sellers are non-atomistic from the point of view of a buyer b . This is the case if French sellers are important input suppliers (e.g. if the sum of $w_{s,t-1}^b$ across all sellers located in France is close to 1) and the distribution of market shares across French sellers is highly concentrated. Under these circumstances, shocks affecting the buyer's main French sellers have a non-negligible impact on its input price index, which implies that $d \ln P_{b,t}$ is correlated with $d \ln z_{s,t}$ and $d \ln a_{sb,t}$.

Unfortunately, we do not observe in our data the whole distribution of market shares $w_{s,t-1}^b$ but only the distribution of sales across French sellers. When a seller accounts for a substantial share of a buyer's input purchases from France, we cannot rule out the possibility that it also represents a non-negligible share in its overall cost, thus inducing endogeneity. The "Monopolistic Competition" model detailed in Appendix A.3 thus takes the extreme view that these two shares are just equal, which amounts to assuming that foreign buyers solely rely on French inputs to produce. Based on this assumption, we derive an alternative set of moment conditions that structurally control for this source of endogeneity. In short, the alternative model is based on a transformation of the growth of sales, that allows to get rid of the price index, and the associated endogeneity issue. The estimation of the transformed model allows to recover an alternative set of shocks, which can then be plugged into eq. (1) to compute the buyer price index and measure the contribution of each type of shocks when the buyer price index is endogenous.

Once the possibility of sellers being non-atomistic in buyers' input purchases is taken into account, it becomes natural to let sellers adjust their mark-up accordingly, as in Atkeson and Burstein (2008). In the "Oligopolistic Model" presented in Appendix A.4, both the buyer price index ($d \ln P_{b,t}$) and markups ($d \ln \mu_{sb,t}$) are thus allowed to respond to the structural shocks. Strategic pricing behavior adds another source of endogeneity because shocks affecting a given seller induce endogenous mark-ups adjustments that are heterogeneous across the firm's buyers. In the benchmark model (2), these endogenous mark-up adjustments are absorbed into the residual, which thus becomes correlated with the seller-specific component. In the "Oligopolistic Competition" model, we control for endogenous changes in match-specific mark-ups in a non-parametric way. Using the definition of mark-ups in Atkeson and Burstein (2008), we can write¹³:

$$d \ln \mu_{sb,t} = f(w_{s,t-1}^b) \left[d \ln \frac{\omega_t}{Z_t} - d \ln z_{s,t} - d \ln a_{sb,t} \right]$$

$$\text{with } f(w_{s,t-1}^b) \equiv - \frac{\mu_{sb,t-1} w_{s,t-1}^b \left[\frac{1}{\eta} - \frac{1}{\sigma} \right] (\sigma - 1)}{1 + \mu_{sb,t-1} w_{s,t-1}^b \left[\frac{1}{\eta} - \frac{1}{\sigma} \right] (\sigma - 1)} < 0,$$

$$\frac{d | f(w_{s,t-1}^b) |}{d w_{s,t-1}^b} > 0$$

Hence, there is incomplete pass-through of cost shocks onto consumer prices, with a pass-through rate that is decreasing in the firm's share of the buyer's input purchases. This heterogeneity is taken into account in the "Oligopolistic Model" by interacting the structural shocks with a step function of $w_{s,t-1}^b$, as detailed in Appendix A.4.

Before concluding, note that there is another source of potential endogeneity that the model in Section 2.1 rules out implicitly. Non-constant returns to scale would indeed imply that there is an additional term entering eq. (1) which would depend on the firm's overall output. Under non-constant returns to scale, the firm's marginal cost, thus its prices, is indeed a function of aggregate output ($y_s = \sum_b \tau_j q_{sb}$), creating a source of endogeneity that is very similar to the case treated in the "Monopolistic Competition" model but involving a correlation between the seller-specific component and other structural shocks to buyer demands. We chose to avoid treating this case explicitly because we have good reasons to think that this source of endogeneity is negligible. Indeed, most foreign buyers are atomistic from the point of view of French sellers.¹⁴ This implies that the correlation between a seller's marginal cost and shocks hitting its foreign clients is close to zero. Since the strength of the endogeneity bias is increasing in a buyer's share of the seller's overall output, the low market shares observed in the data are unlikely to create an important threat to identification.

In the rest of the analysis, we will base our discussion on the estimation of shocks obtained from the Benchmark model. We however checked that our results are robust across alternative models, as detailed in Appendix B.1. Since these alternative models relax one or several assumptions regarding the market structure imposed by the benchmark model but introduce additional constraints to reach identification, we see results based on these three models as complementary. Their comparison is thus informative of the robustness of our Benchmark model findings.

¹³ Note that this is assuming that the firm adjust its mark-up taking the sectoral price index as given.

¹⁴ Using external data from the INSEE-Ficus database of firms' balance-sheet, we are indeed able to recover the share of each foreign buyer in French sellers' overall sales, on top of its share in the seller's European exports which we observe directly in our data. In the sub-sample for which the information is available, the average share of a buyer in total sales is as low as 1.3%. The reason is that firms on average sell most of their production on their domestic market (Eaton et al., 2011; di Giovanni et al., 2019).

2.3. From shocks to volatility

In this section, we explain how the estimated shocks, together with the structure of trade networks, shape the volatility of trade. We first describe the network structure of the data.

2.3.1. Network structure

A seller s may have multiple buyers in a country i . Let us denote $w_{b,t}^{s,i}$ the fraction of its sales in country i that goes to buyer b . Hence,

$$w_{b,t}^{s,i} = \frac{X_{sb,t}}{\sum_{b'|c(b')=i} X_{sb',t}}$$

where $X_{sb,t}$ is the value of sales to buyer b , at date t . This set of weights for seller s denotes its *trade network* in country i , which we take as given. A similar set of weights, $w_{b,t}^s$, can be computed for seller s as a fraction of each buyer within s 's overall export sales:

$$w_{b,t}^s = \frac{X_{sb,t}}{\sum_i \sum_{b'|c(b')=i} X_{sb',t}} = w_{b,t}^{s,i} w_{i,t}^s$$

where $w_{i,t}^s$ is the share of market i in seller s sales.

The set of all such weights at a given date t constitutes the *Trade Network* of French exporters (all firms s from France) at date t in a country as well as across countries. It is a network, more precisely a bipartite graph, with two types of nodes: the seller-nodes and the buyer-nodes.

2.3.2. Determinants of trade volatility

The volatility of firms' export sales, our object of interest, can be defined as follows:

$$\text{Var}(g_{st}) = \frac{1}{T} \sum_t (g_{st} - \bar{g}_s)^2$$

where g_{st} denotes the growth rate of seller s exports between $t - 1$ and t , \bar{g}_s its mean with both quantities computed across clients and while volatility is over time.

For any given seller s , the volatility of its export sales is a weighted average of the variances and covariances of its firm-to-firm growth rates, observed in the sub-sample of trade flows involving s ' trade network at $\{t - 1, t\}$. Using the decomposition in (2)¹⁵:

$$\begin{aligned} \text{Var}(g_{st}) &= \text{Var}(\varepsilon_{s,t}) + \text{Var}\left(\sum_i w_{i,t-1}^s \varepsilon_{i,t}\right) \\ &+ \text{Var}\left(\sum_i \sum_{b|c(b)=i} w_{b,t-1}^s (\varepsilon_{b,t} + \varepsilon_{sb,t})\right) + \text{Cov} \end{aligned} \quad (4)$$

The *Cov* component represents a sum of covariance terms across the different shocks.

Hence, in the presence of multiple sources of volatility, the seller's variance can be thought of as the sum of multiple variance and covariance terms, each depending on one specific source of volatility. Namely, the first term in eq. (4) measures the micro-level volatility induced by shocks that are specific to the seller. Because it can be of heterogeneous magnitude across sellers, it contributes to the cross-sectional dispersion in firm-level volatilities (Gabaix, 2011, section 2.5 for instance). The second term can be interpreted as the aggregate component of sellers' volatility. Selling to a broader set of markets is a way for the firm to hedge against such country-specific shocks. This possibility is at the root of the argument in Caselli et al. (2015), mostly viewed from a macro-

¹⁵ Here as in the rest of the analysis, we focus on the growth of a seller's sales at the intensive margin, i.e. driven by buyers that belongs to s 's trade network in $t - 1$ and t . This is justified by the empirical strategy which estimates shocks affecting firm-to-firm growth rates. Appendix B in the Online Appendix discusses how extensive adjustments affect our measure of volatility.

economic perspective. Finally, the third term captures the impact of buyer- and seller-buyer shocks. We grouped these terms since their impact depends on the structure of the firm's portfolio of clients, its *trade network*. A less concentrated trade network mechanically reduces the firm's exposure to such shocks. Of course, having a less concentrated portfolio is not the only source of risk hedging in this set-up. A firm might have little exposure to buyer-related shocks if it interacts with less volatile partners and/or if its portfolio is made of partners that co-move negatively. This is the reason why both the structure of portfolios (as measured by the weighting parameters) and the individual shocks jointly determine the size of this component.¹⁶

An equivalent formula can also be obtained at the macro-level by aggregating across sellers. In words, the volatility of aggregate export growth can be decomposed into the sum of the economy's residual exposure to various sources of risk. The residual exposure to seller-specific shocks naturally depends on the extent of diversification in sales across exporters; an argument at the root of the "granular hypothesis" in Gabaix (2011). The aggregate impact of country \times sector-specific shocks depends on the extent of geographic diversification (Caselli et al., 2015). Finally, how much buyer and buyer-seller shocks contribute to aggregate volatility depends on the diversification of import purchases, across buyers and buyer-seller pairs, respectively.

Having discussed how the interaction between the *Trade Network* and the different shocks hitting its nodes determines the extent of volatility in trade markets, we now describe the data used to implement the strategy we have just been describing.

3. Data and stylized facts

3.1. Data

The empirical analysis is conducted using detailed export data covering the universe of French firms. The data are provided by the French Customs. The full data set covers all transactions that involve a French exporter and an importing firm located in the European Union, over the 1995–2007 period.

For each transaction, the data set records the identity of the exporting firm (its SIREN identifier), the identification number of the importer (an anonymized version of its VAT code), the date of the transaction (month and year), the product category (at the 8-digit level of the combined nomenclature) and the value of the shipment.¹⁷ In the analysis, data are aggregated across transactions within a year, for each exporter-importer pair.¹⁸

Statistics on the dimensionality of the *Trade Network*, as observed in 2007, are presented in Table 1. For this particular year, the data set includes 42,888 French firms exporting to 334,905 individual buyers located in 11 countries of the European Union. Total exports by these firms amount to 207 billions euros. Whereas large destination markets naturally involve more firms on both sides of the border, the density of trade networks, as measured by the number of active pairs divided by the potential number of relationships, is actually lower in countries like Germany or Belgium.

¹⁶ Note that the decomposition in equation (4) is exact, in the sense that we do not condition on a constant distribution of weights as in Carvalho and Gabaix (2013) or di Giovanni et al. (2014) but instead take into account the variability in the structure of the Trade Network as a source of volatility in the small. More specifically, we use the following formula for

firm-level volatility: $\text{Var}(g_{st}) = \frac{1}{t-1} \sum_t \left(\sum_b w_{b,t-1}^s g_{sb,t} - \frac{1}{t} \sum_t \sum_b w_{b,t-1}^s g_{sb,t} \right)^2$.

¹⁷ French Customs data track trade between French firms and foreign buyers identified through their VAT number. The data do not capture direct exports to individuals.

¹⁸ Notice that, even though we track each sale a seller makes to each country, we cannot do the same for buyers. More precisely, we cannot know if the same buyer buys from two foreign sellers from two different countries. More generally, since we do not have additional information on the buyer, we cannot say whether it is an affiliate of the same (multinational) firm as the seller or indeed if two buyers in our data are connected through multinational linkages. Such link could affect the volatility of a firm-to-firm relationship, which is left unconstrained in the empirical analysis.

Table 1
Summary statistics on the trade network.

	Value of exports (bil.€)	# French sellers	# foreign buyers	# pairs of buyer–seller
	(1)	(2)	(3)	(4)
Belgium	26.6	29,941	74,427	225,823
Denmark	2.8	8567	9248	22,008
Finland	1.85	5420	5379	12,243
Germany	50.2	25,078	122,568	249,197
Ireland	2.54	6508	6857	16,804
Italy	32.0	20,565	100,115	192,628
Netherlands	15.5	16,851	35,080	73,568
Portugal	4.59	11,980	20,331	44,957
Spain	35.5	22,038	80,178	166,738
Sweden	5.08	7896	10,757	21,832
United Kingdom	30.6	19,289	52,596	115,992
EU11	207	42,888	334,905	1,141,326

Notes: Summary statistics computed on 2007 data describing French bilateral exports. The last line corresponds to the 11 members of the European Union pooled together. The table does not include the transactions for which the CN8 product code is not reported (19,803 sellers accounting for less than 0.05% of exports). Column (1) reports the value of the aggregate trade flow, in billions euros. Columns (2)–(4) respectively report the number of sellers, buyers, and seller-buyer pairs involved in this aggregate trade flow.

Two types of information are not reported in the data which forces us to work with a flexible estimation procedure. First, we do not know whether transactions are arm's length or involve related parties. In 1999 -the only year for which the information exists in French data-40% of French exports was conducted within multinational companies. This represents 10% of seller-product-destination triplets (Davies et al., 2018). While the model does not explicitly take multinationals into account, it might be that intra-firm trade is hit by specific sources of shocks. In our empirical model, such shocks would be absorbed into the residual and contribute to the volatility of seller-buyer growth components.

Second, the industry of foreign importers is not reported in the data.¹⁹ However, a large share of France's exports is likely to be intermediated through foreign wholesalers.²⁰ The technology and demand of wholesalers may differ from the ones of manufacturing buyers, as well as their intrinsic (buyer-specific) volatility. This does not constitute a threat to identification though, as we do not impose any restriction on this source of heterogeneity across buyers.

We impose different restrictions to the data. While the data are exhaustive for transactions within the EU, small exporters are allowed to fill a “simplified” form that does not require the product category of exported goods. This is problematic whenever the empirical strategy controls for sector-specific determinants of the outcome variable since the corresponding transactions cannot be included in the data set. The “simplified” regime concerns firms with total exports in the European Union in a given year below 100,000 euros (150,000 euros since 2006), which are thus dropped from the estimation sample.

In the panel dimension, the different shocks are identified for: i) transactions active for at least two consecutive years - to be able to compute firm-to-firm growth rates; ii) transactions in the main connected group of the *Trade Network*. To minimize the effect of outliers on our

¹⁹ The exception is French exports to Belgium in 2005. For this year, we know the main industry of Belgian buyers for 12% of seller buyer pairs accounting for 65% of Belgian imports from France. Belgian manufacturing firms made about 32% of Belgian imports from France; wholesalers 26%; buyers in other sectors 6% (we do not know the sector for the remaining 35%). Descriptive statistics suggest that identified Belgian wholesalers do not differ markedly from manufacturing firms, in terms of their trade network with France. On average they interact with 6.8 French suppliers, when manufacturing firms interact with 8.7 exporters. French sellers interacting with Belgian wholesalers do not differ significantly from those interacting with Belgian manufacturing firms either.

²⁰ Wholesale trade for instance accounts for 41% of the value of French imports (di Giovanni et al., 2019).

Table 2
Data coverage: exports versus overall sales and overall exports versus estimation sample.

	Value of sales	# of obs	Median volatility	# of firms	Corr. vol.
	(1)	(2)	(3)	(4)	(5)
<i>Panel A: Overall versus export sales</i>					
Total sales	33.99	8,841,503	0.20	697,653	1.00
Export sales	5.71	2,152,793	0.50	111,045	0.23
EU 28 exports	2.68	745,173	0.52	40,945	0.17
<i>Panel B: Total export versus estimation sample sales</i>					
EU 15 exports	2.32	15,366,251	0.50	43,938	1.00
EU 11 exports	2.18	14,069,787	0.50	42,986	0.94
Estimation sample	1.23	3,834,655	0.45	22,084	0.34

Notes: This table gives the coverage of the sample in comparison with firms' overall sales (Top Panel, Source: INSEE-Ficus and Customs, Firm-to-destination export data set) and firms' intra-EU exports (Bottom Panel, Source: Customs, Firm-to-firm export data set). The first column reports the value of sales (in trillion euros). The second column presents the number of observations (seller-year in Panel A, seller-buyer-year in Panel B). The third column presents the volatility of the median firm in the sample, where the volatility is defined as the standard deviation of annual sales growth, restricted to second-order moments computed on at least four points and excluding outlier growth rates. The fourth column reports the number of firms used to compute the median volatility. Note that the number of such firms is slightly larger for “EU 15 exports” than “EU 28 exports” due to slightly different coverages across both customs data sets, the firm-to-firm export data set covering export flows that are not integrated in the firm-to-destination data set. The last column reports correlation coefficients across measures of firm-level volatility. All correlation coefficients are significant at the 1% level.

measures of volatility, we further base our estimates on observations for which the seller (log-) growth rate lies in the interval $[-0.8; 4]$. This leaves us with a data set comprising more than 3.8 millions observations (firm-to-firm growth rates) that represent 53% of EU15 sales.

To summarize, our sample is the outcome of two types of restriction. First, we have a geographic restriction as the seller-buyer data are only available for French exports to EU countries. Second, we have restrictions linked to our estimation strategy. Table 2 compares our final sample with different samples relaxing these constraints.²¹ The first five lines compare the size and features of sub-samples of firms which sales are gradually restricted to the geographic coverage of the estimation sample, namely export sales to 11 EU countries. At the level of individual firms, we find a strong correlation between the volatility of sales to the 11 countries that compose the estimation sample and overall export sales (see the last column in lines 2–5 in Table 2). The correlation of the volatility of overall and export sales computed based on the first two lines is however smaller, at 23%, although still significant. This suggests that our estimation sample is representative of firms' export sales volatility and somewhat correlated with the volatility of their overall sales. More costly are the various restrictions imposed by the estimation strategy, that force us to cover only half of French firms' EU exports (see the last line in Table 2) and 50% of the overall population of exporters. In this sub-population, the correlation between the firms' volatility in the estimation sample and the volatility of their EU11 export sales is equal to 34%. The reason is that the estimation strategy imposes a number of constraints that amount to neglecting many small and potentially volatile export flows, notably affecting the extensive margin of exports. This explains that the median volatility of a firm's exports in this sub-sample

²¹ Here and in the rest of the paper, the volatility of a firm is defined as the variance of its annual export growth. We restrict our attention to the subset of second-order moments computed on at least four points. Whereas four points may seem a small number for computing measures of volatility, this restriction reveals itself rather constraining once one realizes that it relies on the number of years a given firm serves a specific destination or even the number of years a given seller-buyer pair is active. In our data, about 75% of seller-buyer relationships last less than 4 years. These are tiny transaction accounting for less than 15% of exports. The low average duration of relationships also explains why we do not compute time-varying measures of volatility based on sub-periods, as is often done in the macroeconomic literature.

is somewhat reduced, at 0.45. In Section B of the Online Appendix, we discuss into more details the impact of extensive adjustments on firm-level and aggregate volatility.

3.2. Diversification in trade networks

To conclude this data section, let us discuss a number of stylized facts regarding the amount of diversification implied by the data. As discussed in Section 2.3, the structure of exporters' portfolio interacts with the shocks to determine firms and countries' residual volatility. Section A of the Online Appendix provides a detailed analysis of the structure of the *Trade Network*. The main result which emerges from the analysis is the low degree of diversification in individual firms' export portfolios, both across and within countries. Using the notations of eq. (4), the distribution of the $w_{i,t}^s$ and $w_{b,t}^s$ weights is skewed within a firm and in the population of French exporters.

This (lack of) diversification is illustrated in Fig. 1, which shows the cumulated distribution of French exporters, ordered by the number of EU destinations they serve (left panel), and the number of buyers they are connected to, within a destination (right panel).²² In our sample, 25% of French exporters serve a single destination, within the EU, i.e. $w_{i,t}^s = 0 \forall i$ but one country. At the other side of the distribution, less than 20% of firms serve more than 6 EU countries.²³ Within a destination, the number of buyers served is also relatively small on average, with 43% of exporters (times their export destination) serving a single partner there ($w_{b,t}^s = w_{i,t}^s$) and only 12% of firms having more than 10 partners.

Participation into export markets as well as the number of partners served within a destination is known to be the result of exporters' self-selecting into exports (Mayer and Ottaviano, 2008; Bernard et al., 2018). This is confirmed by the statistics in Section A of the Online Appendix. The 20% of firms that serve more than 6 destinations thus account for almost 70% of the value of French exports while the 12% of exporters with more than 10 partners in a destination represent 40% of the aggregate flow. Large exporters tend to serve more destinations and more buyers within a destination, and should thus end up better diversified across and within countries.

While this is indeed the case in our data, the additional lines in Fig. 1 show that this does not imply that large firms end up fully diversified. The reason is that the within-firm distribution of export sales tends to be skewed towards one or two main partners, even when the number of such partners is relatively large. To illustrate this, we reproduced the cumulative distributions, restricting the analysis to each firm's top destinations and partners, that cumulate at least 90%, 50% and 10% of the firm's exports (see the diamond, triangle and square lines in Fig. 1, respectively). This shifts the distributions up, with the last two distributions being very close to a horizontal line. The shift is a direct consequence of the distribution of sales being skewed, across partners within an exporter's portfolio. To see why, consider the distribution of the number of buyers per seller based on total and the top 90% sales. If sales were equally distributed across partners, restricting the sample to the top 90% of a firm's sales would not affect the distribution since no buyer importing from an exporter serving at least two buyers would represent less than 10% of this exporter's sales. Likewise, focusing

²² The circles line in the left panel thus shows the share of French exporters serving x destinations or less, which is naturally equal to 100% when we reach 11 countries, our sample. Likewise, the circles line in the right panel measures the share of exporters serving x buyers or less in a given destination. The other three lines in each panel depict the cumulated distributions obtained from firms' sales, when the smallest destinations / partners are excluded from the analysis. These are meant to show that even firms serving many destinations and partners usually have quite concentrated export portfolios with most of their sales going to their main destination / partner.

²³ The limited number of destinations served by the typical seller in these data is not due to the data being restricted to EU exports. When the equivalent distribution is calculated on the whole sample of firm-level export flows, which neglects the buyer dimension but is not confined to EU exports, results are even sparser. In these data, 44% of French exporters serve a single destination while 22% serve 7 destinations or more.

on the firm's top 50% sales would hardly change the distribution, for firms serving three partners or more. Instead, what the figure shows is that the distributions are very different. Indeed, for the vast majority of French exporters a single partner in a single destination is responsible for at least 50% of export sales. This explains that the distributions restricted to the Top 50 and Top 10% of a firm's sales are close to horizontal and that even the distribution based on top 90% sales is relatively far from the overall distribution.

This lack of diversification has consequences in the aggregate, as illustrated in Table 3. Our statistics presented in Column (1) of Table 3 confirm results in Gabaix (2011) and Mayer and Ottaviano (2008): the distribution of firms' size is extremely skewed and the concentration of sales is even stronger among exporters. Indeed, the Herfindahl index of French exports is more than 200 times what it would be shall exports be split evenly across individuals. The top 10% firms are thus responsible for 90% of aggregate trade. Less documented is the fact that the concentration of trade is even more pronounced across importers (column (2)) and across seller-buyer pairs (column (3)). Such levels of export concentration imply that the overall economy is strongly exposed to individual shocks hitting the largest firms at both sides of the border.

4. Results

4.1. Summary statistics on estimated shocks

Summary statistics on the estimated shocks are provided in Table 4. All three individual components are identified for 3.8 millions (growth) observations. There are 12 years, 11 countries and 35 2-digit industries, hence more than 4300 "aggregate" (i.e. sector \times country \times year) shocks. Finally, we are in position to identify more than 200,000 seller (time) effects, using an average of 13 observations per effect and 930,000 buyer (time) effects, using on average 4 observations per effect.

Without much surprise, the residual match-specific component is found to have the largest variance. But the other two individual components also display substantial variability. The dispersion of the macro shocks is instead an order of magnitude lower. The robustness of these findings can be assessed based on Table A1 which reproduces the same statistics obtained from the two alternative models, namely the "Monopolistic competition" and "Oligopolistic competition" models. The estimated variance of the seller-specific shocks is somewhat larger in these estimations that control for feedback effects of such shocks onto the buyer input price index. Likewise, the estimated variance of aggregate shocks is inflated in the "Oligopolistic Competition" model, when these shocks are allowed to be (heterogeneously) passed-through into buyer-specific prices.²⁴

Of course, these standard errors hide a huge amount of heterogeneity. Some of this heterogeneity is due to variations in the precision of estimates, which can be relatively low for some of these fixed effects due to the sparsity of the underlying graph (Jochmans and Weidner, 2017). Section B.2 of the Appendix discusses how the main results vary when we exclude shocks identified over a small number of observations. But the heterogeneity also reflects different volatility patterns for shocks, across individuals and countries. In the rest of the analysis, we will take this heterogeneity as given and study what it implies for the residual exposure of firms and countries to the shocks.

²⁴ One may be concerned that estimates are biased by the selection of firms into the sample. Since the fixed effects estimator is applied to the cross-section of growth rates, extensive margin adjustments are not taken into account in the estimation. If firms' exit is endogenous to (negative) shocks, this might have a significant impact on the distribution of shocks. We did try to account for attrition after a negative shock. Namely, we estimated the model again using a weighted estimator that gave more weight to observations that were more likely to exit the sample. Exit probabilities were recovered in a first step from a selection equation using the structure of the firm's portfolio as exogenous regressor. Results available upon request were virtually unchanged.




Fig. 1. Diversification of exporters, across destinations and buyers. Notes: The left panel displays the proportion of sellers that serve x destination markets or less, in 2007. The right panel corresponds to the share of seller that serve x buyers or less in a given destination, also in 2007. The “Total Sales” distributions correspond to total exports. The distributions labeled “Top $X\%$ Sales” are computed restricting the amount of each firm’s sales to the X first percentiles of the distribution of sales when transactions are ordered by their decreasing share in the firm’s total sales. In the left panel, the grey diamonds for instance are to be interpreted as follows: If, for each exporter, we neglect the set of the smallest markets contributing to the last 10% of the exporter’s sales, more than 60% of exporters have a degree of one market while less than 1% serve 6 countries or more.

4.2. Firms’ diversification and exposure to micro-demand shocks

To establish the link between the within-firm distribution of export sales and firms’ exposure to microeconomic demand risks, we now combine the estimated shocks with the observed structure of trade networks to compute the determinants of a firm’s volatility as in eq. (4). Results are summarized in Table 5. The first line displays the volatility of the average firm’s export growth in the estimation sample. The subsequent lines report “counterfactual” volatilities whereby one or several determinants of the firm’s growth are muted. The percentage impact on firms’ volatility reported in the second column is indicative of the relative contribution of this particular shock to fluctuations in export growth. In practice, we shut down all the variance and covariance components of eq. (4) that involve a particular shock. As should be clear from eq. (4), the impact of muting a shock depends on: i) the variance of the shock for the particular firm under study (e.g. $Var(\epsilon_{s,t})$ when the seller-specific shock is muted); ii) the extent to which this shock is diversified along the firm’s portfolio (which depends on the firm’s trade network as established in Section 2.3); and iii) the empirical covariance of this shock with other growth components.²⁵ While the identification strategy imposes orthogonality across the shocks on average, empirical covariances within a particular exporter can be a significant source of volatility, which explains that the various counterfactuals do not exactly sum up.

Results in Table 5 confirm that volatility at the firm-level is mostly driven by microeconomic shocks. Muting all “macro” growth components hardly affect firms’ volatility. Instead, muting the seller-specific shocks or the two buyer-related ones has a strong impact on firms’ volatility. The volatility under these two counterfactual experiments is almost halved. The relative contribution of shocks is comparable in the alternative “Monopolistic competition” and “Oligopolistic competition” models, as shown in Table A2.

²⁵ It has to be noted that these “counterfactuals” are computed under the assumption that the distribution of weights, the firm’s trade network using the wording of Section 2.3, is left unaffected. A full fledge counterfactual exercise would need to take into account that the observed distribution of weights is conditional on the realization of past shocks, and is potentially different from the counterfactual one. Since the counterfactual experiments are used as a way to measure the contribution of each shock to the volatility of exports, we decided to neglect this indirect effect and focus on the impact of muting one source of shocks conditional on the structure of the network.

Table 3
Concentration of aggregate trade flows.

	Concentration across		
	Sellers (1)	Buyers (2)	Seller-buyer pairs (3)
Herfindahl Index			
Absolute value	0.005	0.001	0.001
Relative to symmetry	234	497	755
Share in aggregate of			
Top 10%	90%	94%	94%
10 largest	15%	7%	6%

Notes: Summary statistics computed on 2007 data describing French multilateral exports. This table reports statistics regarding the concentration of exports across French exporters (column (1)), foreign importers (column (2)) and the pairs they form (column (3)). Concentration of export sales is measured by i) the Herfindahl index, either expressed in absolute value or in relative terms with the value one would obtain would existing individuals be symmetric in size (computed as the Herfindahl times the number of individuals), ii) the share in aggregate exports of the top decile of the distribution or the 10 largest individuals.

Table 4
Summary statistics on the estimated shocks to firm-to-firm growth.

	Mean	Std. dev.	Count
Firm-to-firm growth	-0.013	0.689	3,834,655
Aggregate component	-0.013	0.045	4310
Seller-specific component	0.000	0.197	283,032
Buyer-specific component	0.000	0.358	933,888
Match-specific component	0.000	0.552	3,834,655

Notes: This Table gives summary statistics on the estimated shocks in the benchmark model. The first column reports the mean of the different components, the second column provides its standard deviation and the third column the number of individuals used to compute these statistics. The first line corresponds to summary statistics on firm-to-firm growth and the next four lines detail estimates on the estimated shocks, the aggregate component ($\epsilon_{i,t}$), the seller-specific shock ($\epsilon_{s,t}$), the buyer-specific shock ($\epsilon_{b,t}$) and the match-specific component ($\epsilon_{sb,t}$).

Using the insight of eq. (4) and the results in Table 4, microeconomic seller-specific shocks should have a strong impact on a firm’s volatility: these shocks are volatile and have a one-to-one impact on the variance of a firm’s export growth. The large impact of microeconomic buyer-

Table 5
Summary statistics on the actual and counterfactual distributions of firm-level volatilities.

	Mean	%Change
Actual volatility $Var(g_{s,t})$	0.192	
Volatility when muting		
Aggregate shocks $Var(\cdot \varepsilon_{it}, t = 0)$	0.191	-0.007
Seller-specific shocks $Var(\cdot \varepsilon_{s,t}, t = 0)$	0.106	-0.447
Buyer-related shocks $Var(\cdot \varepsilon_{b,t}, \varepsilon_{sb,t} = 0)$	0.097	-0.495
One buyer-related shock after the other		
Buyer-specific $Var(\cdot \varepsilon_{b,t}, t = 0)$	0.146	-0.238
Match-specific $Var(\cdot \varepsilon_{sb,t}, t = 0)$	0.151	-0.214

Notes: This table gives summary statistics on the actual and counterfactual levels of firm-level volatility, at the mean of the distribution. The counterfactuals are obtained by muting different shocks one after the other. The first column is the level of the variance at the mean, in the actual and counterfactual distributions. The second column reports the percentage difference between the corresponding counterfactual and actual variance levels.

related shocks is instead more surprising as this source of risk naturally diversifies along the firm's portfolio of sales. These shocks have a substantial impact on the firms' volatility because most firms are poorly diversified.

This statement is illustrated in Fig. 2 which shows the median exposure to buyer-related microeconomic shocks (as measured by $Var(g_{s,t} | \varepsilon_{s,t}, \varepsilon_{i,t}, t = 0)$) in each decile of the distribution of firms' sales concentration indices. The clear upward-trending relationship means that residual exposure to microeconomic demand shocks is stronger when firms have more concentrated sales. Firms in the tenth decile of the distribution of concentration indices are almost four times more exposed to these shocks than firms in the first decile. This is true even though firms have the possibility to compensate a high degree of concentration in sales by the interaction with less volatile partners. For instance, one might explain the difference in portfolio concentration by different strategies for serving foreign markets. As discussed in Bernard et al. (2010), some firms directly serve foreign clients abroad while others interact with a single intermediary, which is later in charge of serving the foreign market with the firm's products. Direct trade implies a larger number of partners served, in comparison with intermediated trade which mostly involves a single intermediary. This intermediary itself aggregating the demand of many foreign partners, one should expect the volatility of its demand to be lower than the typical final consumer. While our strategy does not rule out this possibility and leaves buyer-

specific volatilities unconstrained, our results show that such counteracting force is not sufficiently strong to compensate for the mechanical effect that diversification in sales has on firms' exposure to microeconomic demand shocks.

4.3. Exposure to micro demand shocks along the size distribution

The previous sub-section has established that firms in our sample are strongly exposed to microeconomic demand shocks and that such exposure declines with the diversification of export portfolios. Since such diversification in part reflects the self-selection of firms into exporting, one should expect this to contribute to heterogeneity in volatilities along the size distribution. Fig. 3 shows that this is indeed the case; although the correlation is weaker than expected.

Namely, Fig. 3 plots the median volatility of firms in each decile of the firms' size distribution, together with its main components. Consider first the dark bars which correspond to the raw data. As already documented in the literature (di Giovanni et al., 2014), firm-level volatility is strongly decreasing in size, with firms in the first decile of the distribution being more than twice as volatile as the 10% largest exporters. As shown by the light grey bars of Fig. 3, the size-volatility relationship mimics almost perfectly the tendency of large firms to be exposed to less volatile seller-specific shocks. This is consistent with the interpretation found in the literature, that idiosyncratic supply risk is lowered for large firms (see for instance, Gabaix (2011), section 2.5).

The correlation is less pronounced when focusing on the microeconomic demand shocks, which are diversifiable within a firm. The size of the medium grey bars is indeed decreasing slowly along the distribution of size decile up to the eighth decile. The reason is that medium size firms are only slightly better diversified in sales than small firms. Even firms at the top of the distribution have relatively skewed sales, thus being left quite exposed to microeconomic demand risk. Because of this, we shall expect these individual shocks to continue contributing significantly to the volatility "in the large", of aggregate exports.

4.4. Exposure to foreign risk and aggregate volatility

The distribution of firms' size is extremely skewed, as illustrated using our data based on Table 3. This skewness implies that shocks to the largest firms have a substantial impact on aggregate fluctuations (Gabaix, 2011). This is confirmed in Table 6 which decomposes the




Fig. 2. Firms' sales concentration and buyer-driven volatility. Notes: This graph represents the amount of volatility driven by micro buyer-related shocks for firms with different level of sales concentration. Volatility driven by buyer-related shocks is defined as $Var(g_{s,t} | \varepsilon_{i,t}, \varepsilon_{s,t}, t = 0)$. Concentration (diversification) is measured using the year-average Herfindahl index of French exporters computed over their European buyers. Firms are grouped by decile of concentration.




Fig. 3. Volatility, by decile of firms' size. Notes: This figure represents the median volatility of sellers' exports across deciles of sellers' size. Sellers are grouped into size bins based on their initial size, with bin 1 corresponding to the 10% smallest exporters. For each decile, the figure reports the median volatility of sellers ("Volatility seller" defined as $Var(g_{it})$), the median volatility attributable to buyer-related microeconomic shocks ("Buyer-driven vol." defined as $Var(g_{s,t} | \epsilon_{jt,t}, \epsilon_{s,t} = 0)$) and the median volatility induced by seller-specific shocks ("Seller-driven vol." defined as $Var(g_{s,t} | \epsilon_{b,t}, \epsilon_{sb,t} = 0)$).

volatility in the large, i.e. the variance of aggregate (bilateral and multilateral) export growth, into its various components. Here again, we compare aggregate volatilities with a number of "counterfactuals" that mute one or several sources of volatility, leaving the structure of the Trade Network unchanged. This helps deal with empirical covariances across shocks.

Once the analysis is performed at the aggregate level, common aggregate shocks naturally start contributing substantially to the overall volatility. Muting these shocks thus reduces the volatility of exports by 38% in the bilateral case and 62% for multilateral exports. However, this does not mean that the three microeconomic shocks become irrelevant. Instead, their combined contribution to overall volatility is comparable to that of aggregate shocks at the multilateral level and even stronger when volatility is computed country-by-country (see the comparison of the second and third lines in Table 6).²⁶ The reason why individual shocks have a substantial impact in the aggregate is the strong concentration of exports documented in Table 3. To show this, the last line in Table 6 reports the results of another counterfactual experiment in which all shocks are left unaffected but firm-to-firm relationships are assumed symmetric within and across each firm's portfolio (i.e. $w_{sb,t-1} = w_{t-1}, \forall (s,b)$). In such "symmetric" Trade Network, the amount of volatility is as low as it is when all microeconomic shocks are muted. The significant impact of microeconomic shocks is due to the strong concentration of sales in international markets.

The large contribution of microeconomic shocks to aggregate fluctuations confirms insights of the granularity literature, in the context of export sales. di Giovanni et al. (2014) estimate that around 50% of fluctuations in French aggregate sales is attributable to microeconomic growth components. What the structural approach developed in this paper adds to this literature is a more comprehensive decomposition of these "granular" fluctuations with the explicit consideration of micro demand shocks. Namely, the additional counterfactuals in Table 6 allow us to compare the relative contribution of each individual shock, for the volatility in the large. First, notice that the relative contribution of seller-specific components is naturally larger in the multilateral than in the bilateral

²⁶ The reason why the relative contribution of microeconomic shocks is reduced in the multilateral case is that selling goods to several countries mechanically induces a diversification of sales across individual partners. As a consequence, diversification in sales is larger in the multilateral than in the bilateral case.

Table 6
Summary statistics on the actual and counterfactual levels of volatility in the large.

	Bilateral		Multilateral	
	Variance	%Change	Variance	%Change
	(1)	(2)	(3)	(4)
Actual variance $Var(\cdot)$	0.0042		0.0015	
Mute				
Aggregate shocks $Var(\cdot \epsilon_{jt,t} = 0)$	0.0026	-0.382	0.0006	-0.620
Micro. shocks $Var(\cdot \epsilon_{s,t}, \epsilon_{b,t}, \epsilon_{sb,t}, \epsilon_{t} = 0)$	0.0008	-0.803	0.0004	-0.741
Seller-specific shocks $Var(\cdot \epsilon_{s,t} = 0)$	0.0035	-0.140	0.0010	-0.336
Buyer-specific $Var(\cdot \epsilon_{b,t} = 0)$	0.0044	-0.110	0.0012	-0.173
Match-specific $Var(\cdot \epsilon_{sb,t} = 0)$	0.0030	-0.177	0.0013	-0.161
Granularity $Var(\cdot w_{sb,t} = w_t)$	0.0008	-0.809	0.0004	-0.741

Notes: Column (1) reports the median variance of bilateral export growth volatility, in the data and the counterfactual experiments. Column (2) is the median percentage change in the volatility of bilateral exports, in the counterfactual in comparison with the actual variance. Column (3) corresponds to the actual and counterfactual variances of multilateral exports. The same results expressed in percentage change from the actual multilateral variance are reported in Column (4).

case. This is because this source of volatility is not diversified across destinations while the other two are. In bilateral data, all three microeconomic shocks instead contribute almost equally to the variance of aggregate exports.²⁷

Together, these results indicate that both macro and micro shocks contribute substantially to explaining the level of volatility in export data. We conclude this section by assessing their contribution to the heterogeneity across destination markets. Namely, we decompose the

²⁷ This result is somewhat sensitive to market structure assumptions as shown by the comparison of results based on the three alternative models in Table A3. The relative impact of removing buyer-specific components is thus found larger in the Monopolistic and Oligopolistic competition cases than in the benchmark model. Under these two market structure assumptions, the direct impact of the seller and seller-buyer shocks is indeed attenuated by feedback effects on the buyer's price index. The relative contribution of seller-specific and seller-buyer shocks is further reduced in the Oligopolistic case, because markup adjustments induce an incomplete pass-through of such shocks onto the price paid by buyers.

cross-sectional variation in destination-specific growth volatilities into its different drivers using the following identity:

$$\begin{aligned} \text{Var}(g_{it}) = & \text{Cov}(g_{it}, \varepsilon_{i,t}) + \text{Cov}\left(g_{it}, \sum_s w_{s,t-1}^i \varepsilon_{s,t}\right) \\ & + \text{Cov}\left(g_{it}, \sum_{b|c(b)=i} w_{b,t-1}^i \varepsilon_{b,t}\right) + \text{Cov}\left(g_{it}, \sum_s \sum_{b|c(b)=i} w_{sb,t-1}^i \varepsilon_{sb,t}\right) \end{aligned} \quad (5)$$

where $w_{s,t-1}^i$, $w_{b,t-1}^i$ and $w_{sb,t-1}^i$ respectively denote the share of seller s , buyer b , and the seller-buyer pair (s, b) in bilateral exports, at $t - 1$. By regressing each term on the right-hand side on the overall variance of destination-specific export growth, one obtains a measure of how much of the variation in bilateral volatilities can be attributed to each component. The results of the regressions under various market structure assumptions can be found in Table A4 of the Appendix. The conclusions presented below are robust to these alternatives.

Results confirm the overriding role of microeconomic shocks, most notably buyer-specific ones as a driver of volatility. Indeed, 50% of the cross-country heterogeneity in export growth volatilities is attributable to the heterogeneity in the preponderance of these shocks, while seller-specific, aggregate and seller-buyer shocks respectively account for 24, 18 and 9% of the cross-sectional variance. The large effect of microeconomic shocks is due to the structure of the *Trade Network* being quite heterogeneous across destinations, especially on the buyer side. More concentrated and more connected destinations induce a stronger exposure to buyer-specific shocks, exacerbating the volatility of exports. The contribution of seller-specific shocks, although lower, is also substantial and triggered by a similar heterogeneity in the structure of trade networks. Instead, the degree of exposure to seller-buyer shocks does not seem key to explain the variance of export sales in various destinations, and the corresponding coefficient is not significant. Although these results need to be taken with precaution given the small number of destinations in the sample and their specificities, they confirm that the

Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jinteco.2019.103276>.

Appendix A. Details on the estimation strategy

A.1. General structure of the problem

The three estimated equations detailed below share a common structure, namely a decomposition of a firm-to-firm growth term into various components:

The three estimated equations detailed below share a common structure, namely a decomposition of a firm-to-firm growth term into various components:

$$g_{sb,t} = \varepsilon_{i,t} + \varepsilon_{s,t} + \varepsilon_{b,t} + \varepsilon_{sb,t}$$

or, in matrix format:

$$G_t = \alpha_t \Xi_t^C + \chi_t \Xi_t^S + \beta_t \Xi_t^B + \nu_t$$

where G_t is the vector containing the growth terms (which is $N_t \times 1$, N_t designating the number of observations for year t), α_t is the design matrix for the year- t country-sector effects ($N_t \times N_t^C$ with N_t^C the number of country-sector for year t),²⁸ χ_t is the design matrix for the year- t seller effects ($N_t \times N_t^S$, where N_t^S is the number of sellers for year t), β_t is the design matrix for the year- t buyer effects ($N_t \times N_t^B$, where N_t^B is the number of buyers b , at date t), and ν_t is the vector of residuals ($N_t \times 1$).

Such structure is comparable to the equation estimated in Abowd et al. (1999), where identification is achieved assuming:

$$E[\varepsilon_{sb,t}|b, s, i] = 0 \quad \forall t \Leftrightarrow E(\nu_t|\alpha_t, \chi_t, \beta_t) = 0$$

In the following, we will explain how the estimation is performed when this orthogonality assumption is complemented with additional restrictions on the correlation between the other components. Before this, let us summarize how the estimation works in the non-structural case.

²⁸ Because some of our firms sell multiple products, the definition of the firm's "industry" is not necessarily straightforward. We chose to affect each seller-buyer pair to the "industry" that corresponds to the most important product constituting the corresponding trade flow. Industries are defined by the 2-digit level of the HS nomenclature.

examination of the micro-structure of trade networks and the nature of shocks hitting individual firms is useful to understand the determinants of macroeconomic volatility. Individual shocks hitting the foreign clients together with the diversification of sellers and the connectedness of these clients turn out to be an important driver of volatility. These should not be neglected in future work on trade openness and volatility.

5. Conclusion

In this paper, we provide an integrated account of the sources of individual and aggregate export volatility. The novelty is to account for the possibility that individual shocks hitting exporters' foreign clientele drive part of this volatility. We first propose a structural method for identifying different sources of fluctuations in seller-buyer relationships. This method allows us to analyze export sales volatility in terms of i) microeconomic sources of fluctuations; ii) the microeconomic structure of trade networks. Our emphasis on the firms' portfolio of clients, and the identification of buyer-related shocks as a key driver of fluctuations are, we believe, novel contributions. In presence of buyer-related shocks, differences in the diversification of individual exporters remain a key driver of these firms' heterogeneous volatilities. This is true even though entering foreign markets (almost mechanically) reduces the volatility of individual exports and therefore allows firms to diversify this buyer-related source of risk. Furthermore, even the largest exporters are not very diversified and end up being exposed to microeconomic demand risks. In turn, the limited diversification of large exporters and the connectedness of the largest buyers bring a large amount of "granular" risk to the overall economy, through their exposure to microeconomic supply and demand shocks. Since international trade tends to inflate the importance of these large firms in the aggregate, these combined mechanisms increase the amount of macroeconomic volatility and explain why individual-level foreign demand shocks are not only an important source of firm-level volatility but also of aggregate fluctuations.

As in [Abowd et al. \(1999\)](#), identification of the seller and buyer components requires that the buyers and sellers must be connected in the sense of belonging to a connected group. For each connected group, all the buyer and seller effects are identifiable up to a normalization constraint. To have comparable effects, we focus our analysis on the largest component. Since trade networks are extremely well connected, this restriction does not affect our conclusions since the largest component comprises more than 95% of all observations.²⁹

Note that the above equation could be estimated in the panel dimension. Since the ultimate objective is to use the estimated effects to discuss the sources of *volatility* in the data, we decided to estimate the model year-by-year, relying exclusively on the cross-sectional dimension to identify the estimated effects. This strategy allows us to avoid imposing undue structure on the correlation of growth components through time. Whereas in the model of [section 2.1](#), shocks are implicitly not autocorrelated, estimating the model year-by-year does not impose any restriction on the correlation over time of the various components; the only constraints are imposed on the cross-sectional dimension of the growth components.

Let us now give a concrete example on how identification is achieved in these high-dimensional fixed effects regressions. Suppose that the largest connected group for a given year is made of two sellers s_1 and s_2 and two buyers b_1 and b_2 . s_1 exports to b_1 and b_2 while s_2 only serves b_2 . To simplify, neglect for now aggregate shocks, by assuming that the left-hand side variables have already been purged from their common component. Hence:

$$\begin{aligned} g_{s_1 b_1} &= \varepsilon_{s_1} + \varepsilon_{b_1} + \varepsilon_{s_1 b_1} \\ g_{s_1 b_2} &= \varepsilon_{s_1} + \varepsilon_{b_2} + \varepsilon_{s_1 b_2} \\ g_{s_2 b_2} &= \varepsilon_{s_2} + \varepsilon_{b_2} + \varepsilon_{s_2 b_2} \end{aligned}$$

The OLS solution to estimating the fixed effects ε_{s_1} , ε_{s_2} , ε_{b_1} and ε_{b_2} implies minimizing the sum of squared residuals ie

$$\text{Min} \left[(g_{s_1 b_1} - \varepsilon_{s_1} - \varepsilon_{b_1})^2 + (g_{s_1 b_2} - \varepsilon_{s_1} - \varepsilon_{b_2})^2 + (g_{s_2 b_2} - \varepsilon_{s_2} - \varepsilon_{b_2})^2 \right]$$

with respect to ε_{s_1} , ε_{s_2} , ε_{b_1} and ε_{b_2} . Moreover, the identification conditions for the buyer and seller fixed effects (i.e. their weighted sum being equal to zero, as described in [Abowd et al. \(2002\)](#)) translate in our case into

$$\begin{aligned} \varepsilon_{b_1} + 2\varepsilon_{b_2} &= 0 \\ \varepsilon_{s_1} + 2\varepsilon_{s_2} &= 0 \end{aligned}$$

Combining these identification constraints with the first order conditions of the minimizing problem finally yields

$$\begin{aligned} \varepsilon_{b_1} &= \frac{2}{3}(g_{s_1 b_1} - g_{s_1 b_2}), \quad \varepsilon_{b_2} = \frac{1}{3}(g_{s_1 b_2} - g_{s_1 b_1}) \\ \varepsilon_{s_1} &= \frac{1}{3}(g_{s_1 b_2} - g_{s_2 b_2}), \quad \varepsilon_{s_2} = \frac{2}{3}(g_{s_2 b_2} - g_{s_1 b_2}) \end{aligned}$$

Thanks to the connectedness in this simple network, all fixed effects can be estimated, including those that correspond to poorly connected nodes (b_1 and s_2 in the example). A recent paper by [Jochmans and Weidner \(2017\)](#) however suggests that some of these fixed effects might be poorly estimated because the network under study is relatively sparse. To gauge the seriousness of the problem, they derive bounds on the variance of the fixed-effect estimator that uncover the importance of the structure of the network. We apply the proposed strategy to our bipartite graph. This involves first forming the normalized Laplacian of the network (based on the adjacency matrix representing the network) and then computing its eigenvalues. The second largest tells how dense is the network and at the same time intervenes when setting bounds on the standard errors of the estimated fixed effects in bipartite graphs. In our data, the corresponding eigenvalue is equal to 0.002. This is a low value, showing that our bipartite is not strongly connected and that some fixed effects may be imprecisely estimated. The consequence of this previous result is that we must assess the robustness of our conclusions when some, potentially many, effects are poorly estimated. We explain in [Section B.2](#) how we proceed.

A.2. Benchmark model

In the baseline model, the estimated equation takes the following form:

$$g_{sb,t} = \varepsilon_{i,t} + \varepsilon_{s,t} + \varepsilon_{b,t} + \varepsilon_{sb,t}$$

with the moment conditions:

$$\left. \begin{aligned} 0 &= E[\varepsilon_{sb,t} | b, s] \\ 0 &= E[\varepsilon_{b,t} | s] \\ 0 &= E[\varepsilon_{s,t} | b] \end{aligned} \right\} \forall t$$

While the first moment condition is the one used in [Abowd et al. \(1999\)](#), the other two conditions preclude the use of the algorithm proposed by [Abowd et al. \(2002\)](#). Instead, we decided to use a three-step strategy in which growth is first regressed on buyer effects, the residuals of which we regress onto seller effects. This separates growth into a buyer term, a seller term and a match component. Then, each component is orthogonalized in the sector-country dimension to recover zero-mean growth components. Finally, the aggregate shock is computed as the difference between

²⁹ In a previous version of this paper, we adopted a slightly different version of the model, in which seller effects were country-specific. Therefore, the equation above could be estimated country by country. In this version, a seller is endowed with a unique seller-effect in the year common to all its destinations. Obviously, buyers are all country-specific since there is no common identifier. The benefits of this new strategy are clear on at least two grounds. First, the network is denser and many more observations are now "connected"; hence with identified effects. Second, because we have more available observations to estimate each of the seller effects (at least for those sellers that export to at least two countries), the precision of the estimated seller effects is increased.

overall growth and the sum of the three individual components: $\hat{\varepsilon}_{i,t} = g_{sb,t} - \hat{\varepsilon}_{s,t} - \hat{\varepsilon}_{b,t} - \hat{\varepsilon}_{sb,t}$. This strategy implies a set of estimates that satisfies the above moment conditions.³⁰

A.3. Monopolistic competition

The Monopolistic Competition model helps underline a potential endogeneity issue in the benchmark case. Whenever some sellers are not atomistic in the buyer's input bundle, the buyer-specific input price becomes endogenous to supply and match-specific taste shocks, which invalidates the orthogonality conditions surrounding the system in (3). More formally,

$$d \ln P_{b,t} = \sum_s w_{s,t-1}^b \left[d \ln \frac{\omega_t}{Z_t} - d \ln z_{s,t} - d \ln a_{sb,t} \right]$$

Seller-specific productivity shocks and match-specific preference shocks can have a substantial negative impact on the buyer-specific price index whenever $w_{s,t-1}^b$ is sufficiently large. This introduces a correlation between the buyer component ($\varepsilon_{b,t} + (\sigma - \eta)d \ln P_{b,t}$) and the corresponding seller's $\varepsilon_{s,t}$ and $\varepsilon_{sb,t}$ terms.

Controlling for shocks affecting the buyer-specific input price index makes the growth equation a more complicated function of the fundamental shocks. To solve this problem, we estimate a transformation of the model that allows to get rid of the input price index, in the special case when all sellers entering $d \ln P_{b,t}$ are located in France. The estimated equation takes the following form:

$$\begin{aligned} \tilde{g}_{sb,t} &= (1 + \lambda)(1 - \eta)d \ln \frac{\omega_t}{Z_t} + (1 + \lambda)d \ln A_{i,t} + (\sigma - 1)d \ln z_{s,t} + (1 + \lambda)(\eta - 1)d \ln a_{b,t} + (\sigma - 1)d \ln a_{sb,t} \\ &\equiv \tilde{\varepsilon}_{i,t} + \varepsilon_{s,t} + \tilde{\varepsilon}_{b,t} + \varepsilon_{sb,t} \end{aligned} \quad (\text{A.1})$$

where

$$\tilde{g}_{sb,t} = g_{sb,t} + \lambda \sum_s w_{s,t-1}^b g_{sb,t} \quad \text{and} \quad \lambda = \frac{\sigma - \eta}{\eta - 1}$$

Intuitively, adding to the LHS variable a weighted average of the buyer's growth rates allows to control for the seller-specific shocks entering the price index in the right-hand side of eq. (1). The choice of the loading factor λ is such that the two cancel out exactly.

Given the structure of the model, the components in eq. (A.1) satisfy the following moment conditions:

$$\left. \begin{aligned} 0 &= E[\varepsilon_{sb,t} | b, s] \\ 0 &= E[\tilde{\varepsilon}_{b,t} | s] \\ 0 &= E[\varepsilon_{s,t} | b] \end{aligned} \right\} \forall t$$

Given a value for λ , the components of the equation can be identified using [Abowd et al. \(2002\)](#) in the cross-section of year-specific growth rates under the first moment condition in the above system. We identify the λ parameter using the other two conditions that imply orthogonality between the seller and buyer effects. Since the model is linear, conditional on λ , the relationship between λ and the magnitude of the correlation between the seller and the buyer fixed effects is monotonic ([Blundell and Robin, 1999](#)). As they suggested, we implement a grid-search algorithm on all the possible values of λ and pick the value which best satisfies the model-implied orthogonality condition. Here as well, the aggregate shocks are recovered by orthogonalizing each component in the sector \times country dimension.

In models with two-way effects, even when data are simulated with no correlation between the individuals at each side of the graph (here, between buyers and sellers), estimated effects can end up being negatively correlated ([Abowd et al., 2018](#)). The intuition for this result is quite straightforward. In such additive models, when an estimation error is made on one effect, there is a corresponding estimation error of the opposite sign on the other effect. Because the standard error of these effects decreases as the number of observations used to estimate them increases, the larger the number of buyers connected to a seller, or conversely the number of sellers connected to a buyer, the more precise these effects become (see [Andrews et al., 2008](#), for a more systematic analysis of this "limited connectivity bias").

Based on these results, we argue that the structure of the network by itself may induce a bias in the estimated seller and buyer effects. To quantify the magnitude of this bias, we generate uncorrelated seller and buyer effects from a normal distribution with fixed, known variance for each node of the network, as well as a residual, also drawn in a normal distribution.³¹ Adding these effects, we generate simulated growth rates. These growth rates are used to estimate the seller and buyer effects using the [Abowd et al. \(2002\)](#) procedure and, then, compute the associated correlation between the two. This procedure is repeated 100 times. This yields a distribution of the bias using our simulated effects and the realized structure of the network since, by construction, the true correlation between these effects is equal to zero. We select the mean of this distribution as our target bias, which is -0.0670 in our data. We then take into account the limited connectivity bias by targeting this value for the correlation between the buyer and seller effects instead of the strict orthogonality condition implied by the model.³²

Using this strategy, we recover a value for λ of 0.77, which is consistent with the model's assumptions (namely, $1 < \eta < \sigma$). For instance, $\lambda = .77$ is consistent with an upper-level elasticity of $\eta = 3$ and a lower-level one at $\sigma = 4.5$. For this value of λ , the moment conditions in the above system are satisfied and we can thus recover orthogonal individual components in eq. (A.1). Based on these, one can finally go back to the original

³⁰ Note that the alternative strategy in which growth is first regressed on seller effects before residuals are orthogonalized in the buyer dimension also satisfies the moment conditions. We checked that both strategies deliver highly correlated, although not strictly equivalent, shocks.

³¹ For all three components, the variance of the underlying normal distribution is calibrated using the mean variance estimated in the baseline model, that corresponds to $\lambda = 0$.

³² Note that this procedure is an attempt to control for the limited connectivity bias but is not exact as the Monte Carlo in which the bias is calibrated is not the same environment as the estimation itself. In principle, one would like to adopt a more sophisticated method of simulated moments to estimate λ . Unfortunately, there is no clear theoretical support for setting up a sampling procedure necessary to perform a simulated method of moments in our context. The reason is that the bipartite graph under consideration is sparse and, to our knowledge, there exists no theory of bootstrap in bipartite networks in a sparse setting.

decomposition:

$$g_{sb,t} = (1-\eta)d \ln \frac{\omega_t}{Z_t} + d \ln A_{i,t} + (\sigma-1)d \ln z_{s,t} + (\eta-1)d \ln a_{b,t} + (\sigma-1)d \ln a_{sb,t} + (\eta-\sigma) \sum_s w_{st-1}^b [d \ln z_{s,t} + d \ln a_{sb,t}]$$

$$\equiv \frac{1}{1+\lambda} \tilde{\varepsilon}_{i,t} + \varepsilon_{s,t} + \frac{1}{1+\lambda} \tilde{\varepsilon}_{b,t} + \varepsilon_{sb,t} + \frac{\eta-\sigma}{\sigma-1} \sum_s w_{st-1}^b [\varepsilon_{s,t} + \varepsilon_{sb,t}]$$

A.4. Oligopolistic competition

Let us finally consider an extension of the Monopolistic Competition model in which those large sellers than are non-atomistic in foreign buyers' input purchases can use their market power to price strategically and set mark-ups that are heterogeneous across buyers. In this model, the growth equation becomes:

$$g_{sb,t} = (1-\eta)d \ln \frac{\omega_t}{Z_t} + d \ln A_{i,t} + (\sigma-1)d \ln z_{s,t}$$

$$+ (\eta-1)d \ln a_{b,t} + (\sigma-1)d \ln a_{sb,t} + (1-\sigma)d \ln \mu_{sb,t}$$

$$+ (\eta-\sigma) \sum_s w_{st-1}^b [d \ln z_{s,t} + d \ln a_{sb,t} - d \ln \mu_{sb,t}]$$

with

$$d \ln \mu_{sb,t} = \left(\frac{1}{\eta} - \frac{1}{\sigma}\right) \mu_{sb,t-1} w_{st-1}^b d \ln w_{st}^b$$

$$d \ln w_{st}^b = (1-\sigma)d \ln \frac{p_{sb,t}}{a_{sb,t} p_{b,t}}$$

If we assume that firms neglect the impact of their markup adjustments on the aggregate price index, this simplifies into:

$$d \ln \mu_{sb,t} = f(w_{st-1}^b) \left[d \ln \frac{\omega_t}{Z_t} - d \ln z_{st} - d \ln a_{sb,t} \right]$$

with

$$f(w_{st-1}^b) = \frac{(1-\sigma) \left(\frac{1}{\eta} - \frac{1}{\sigma}\right) \mu_{sb,t-1} w_{st-1}^b}{1 - \left(\frac{1}{\eta} - \frac{1}{\sigma}\right) (1-\sigma) \mu_{sb,t-1} w_{st-1}^b}$$

From this, it comes that variable mark-ups can induce endogeneity since adjustments in seller-buyer-specific mark-ups are systematically correlated with technology shocks. Since the elasticity of markups to marginal costs adjustment is unambiguously negative and increasing (in absolute value) in the seller's market share of buyer *b*'s input purchases, one can interact the aggregate and seller components with a function of the seller's share of the buyer's input purchases to account for heterogeneous reaction of markups to the corresponding shocks along the distribution of a seller's partners. In order to reduce the dimensionality of the problem, we approximate $f(w_{s,t-1}^b)$ using a step function that allows a firm's pricing strategy to be different across two sub-samples of buyers, those on which the firm has low versus high market power. The definition of these sub-samples is performed at the level of each seller using the following procedure. Our measure of a seller's market power is based on its market share in its clients' overall input purchases in France ($w_{s,t-1}^b$). A seller is supposed to have more market power on those buyers that rely heavily on her as a supplier. For each seller, the distribution of market shares is then cut into two sub-samples, delimited by the firm's median market share. Based on this, we define a dummy variable $HighSh_{s,t-1}^b$ which is equal to one for all buyers that end up in the sub-sample above the median. The estimated equation that controls for heterogeneous mark-up adjustments to shocks is then:

$$\tilde{g}_{sb,t} = (1+\lambda)(1-\eta) \left(1 + \alpha HighSh_{s,t-1}^b\right) d \ln \frac{\omega_t}{Z_t} + (1+\lambda)d \ln A_{i,t}$$

$$+ (\sigma-1) \left(1 + \alpha HighSh_{s,t-1}^b\right) d \ln z_{s,t} + (1+\lambda)(\eta-1)d \ln a_{b,t}$$

$$+ (\sigma-1) \left(1 + \alpha HighSh_{s,t-1}^b\right) d \ln a_{sb,t}$$

$$\equiv \tilde{\varepsilon}_{i,t} + \left(1 + \alpha HighSh_{s,t-1}^b\right) \varepsilon_{s,t} + \tilde{\varepsilon}_{b,t} + \tilde{\varepsilon}_{sb,t} \tag{A.2}$$

Eq. (A.2) is estimated under the following moment conditions:

$$\left. \begin{aligned} 0 &= E[\tilde{\varepsilon}_{sb,t} | b, s] \\ 0 &= E[\tilde{\varepsilon}_{b,t} | s] \\ 0 &= E\left[\left(1 + \alpha HighSh_{s,t-1}^b\right) \varepsilon_{s,t} | b\right] \end{aligned} \right\} \forall t$$

where $\tilde{g}_{sb,t} = g_{sb,t} + \lambda \sum_s w_{s,t-1}^b g_{sb,t}$ and λ is calibrated as in the Monopolistic Competition case.

In order to recover the shocks and an estimate for α , we implement an iterative strategy.³³ Namely, the Guimaraes and Portugal (2010) algorithm is first used to recover an estimate for the seller and buyer components. These estimated effects are then passed onto the LHS to compute a measure of growth which is purged from the individual components ($\tilde{g}_{sb,t} - \hat{\varepsilon}_{s,t} - (1 + \lambda)\hat{\varepsilon}_{b,t}$). This transformed variable is then regressed on the interaction between the (just estimated) seller component and the $HighSh_{s,t-1}^b$ dummy. This provides us with a first estimate of α . It is then possible to use this estimate to compute a new set of seller and buyer components, where the seller component is now estimated taking into account price discrimination, as captured by the $(1 + \alpha HighSh_{s,t-1}^b)$ component of the estimated equation. The new set of seller and buyer components is again passed onto the LHS in order to recover an updated estimate for α . The process continues until the mean quadratic errors for seller and buyer effects stabilize. The iterative procedure leaves us with an estimate of α of -0.71 which is significant at the 1% level. The negative coefficient is consistent with the theoretical model, in which shocks to the seller's marginal cost are incompletely passed through on buyers over which the seller has sufficient market power.

Finally, note that we have also tested the same procedure using an alternative definition of $HighSh_{s,t-1}^b$. In the above-described procedure, we impose that each seller discriminates across two groups of partners, of equal size. In the alternative specification, $HighSh_{s,t-1}^b$ is instead defined to be one whenever a seller s sells more than all its competitors to a given buyer b . A consequence of this definition is that not all firms display variability in $HighSh_{s,t-1}^b$, and are thus used to identify α . Some firms always sell relatively small volumes while other always dominate their competitors in terms of the quantity sold. On the other hand, this definition avoids to artificially impose variability in $HighSh_{s,t-1}^b$, between buyers over which a given seller has roughly the same market power, but a market share which is just above and below the median market share used to define the sub-samples at the root of the benchmark $HighSh_{s,t-1}^b$. Since each strategy has its pros and cons, we have tested both of them. In this Appendix, we solely report results based on the first definition of $HighSh_{s,t-1}^b$. Results based on the second definition of $HighSh_{s,t-1}^b$ are qualitatively and quantitatively similar and are available upon request.

Appendix B. Additional results

B.1. Alternative models

In this section, we reproduce the results in Tables 4, 5, 6 and the decomposition in (5) using the shocks estimated under the alternative models described in Sections A.3 and A.4. Results are summarized in Tables A1–A4. The main differences with respect to the baseline are commented in the text.

Table A1

Summary Statistics on the estimated shocks: Comparison across models.

	Mean	Std. dev.	Count
<i>Benchmark Model</i>			
Firm-to-firm growth	-0.013	0.689	3,834,655
Aggregate component	-0.013	0.045	4310
Seller-specific component	0.000	0.197	283,032
Buyer-specific component	0.000	0.358	933,888
Match-specific component	0.000	0.552	3,834,655
<i>MC Model</i>			
Firm-to-firm growth	-0.013	0.689	3,834,655
Aggregate component	-0.052	0.047	4310
Seller-specific component	0.000	0.269	283,032
Buyer-specific component	0.000	0.361	933,888
Match-specific component	0.000	0.542	3,834,655
Buyer inputs cost	0.038	0.141	933,888
<i>OC Model</i>			
Firm-to-firm growth	-0.012	0.688	3,670,753
Aggregate component	-0.049	0.125	4098
Seller-specific component	0.004	0.249	242,018
Buyer-specific component	0.000	0.364	897,460
Match-specific component	0.000	0.522	3,670,753
Buyer inputs cost	0.040	0.153	897,460

Notes: This Table gives summary statistics on the estimated shocks under three market structure assumptions. The first column is the mean volatility / estimated shock, the second column provides its standard deviation and the third column the number of individual used to compute these statistics. For each model, the table first reports summary statistics on the corresponding sample population of firm-to-firm growth rates. It then details estimates on the four growth shocks, the aggregate component ($\varepsilon_{i,t}$), the seller-specific shock ($\varepsilon_{s,t}$), the buyer-specific shock ($\varepsilon_{b,t}$) and the match-specific component ($\varepsilon_{sb,t}$). In the MC and OC models, the "Buyer inputs cost" terms measure the feedback effect of shocks through the buyer's input price index.

³³ Here as before, the aggregate term is recovered ex-post, by orthogonalizing the growth components in the sector \times country dimension.

Table A2
Actual and counterfactual distributions of firm-level volatilities: Comparison across models.

	Bench.	MC	OC
Actual variance $Var(g_{s,t})$	0.192	0.192	0.180
Change in the volatility induced by muting			
Macroeconomic shocks $\Delta Var(. \epsilon_{ji,t} = 0)$	-0.007	-0.007	-0.096
Seller-specific shocks $\Delta Var(. \epsilon_{s,t} = 0)$	-0.447	-0.414	-0.343
Buyer-related shocks $\Delta Var(. \epsilon_{b,t}, \epsilon_{sb,t} = 0)$	-0.495	-0.295	-0.282
One buyer-related shock after the other Buyer-specific $\Delta Var(. \epsilon_{b,t} = 0)$	-0.238	-0.127	-0.090
Match-specific $\Delta Var(. \epsilon_{sb,t} = 0)$	-0.214	-0.185	-0.200

Notes: This table gives summary statistics on the actual and counterfactual dispersions of firm-level volatilities, when the counterfactuals are obtained by muting different shocks one after the other. The counterfactual results are expressed as the percentage change between the average volatility of exports and the average counterfactual volatility.

Table A3
Actual and counterfactual levels of volatility in the large: Comparison across models.

	Change in volatility if muting:						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Volatility	Macro	Micro	Granu.	Seller	Buyer	Match
		$\epsilon_{i,t} = 0$	$\epsilon_{s,t}, \epsilon_{b,t}, \epsilon_{sb,t} = 0$	$w_{sbt-1} = w_{t-1}$	$\epsilon_{s,t} = 0$	$\epsilon_{b,t} = 0$	$\epsilon_{sb,t} = 0$
<i>Bilateral</i>							
Benchmark	0.0042	-0.382	-0.803	-0.809	-0.140	-0.110	-0.177
MC	0.0042	-0.393	-0.754	-0.699	-0.118	-0.404	-0.172
OC	0.0045	-0.408	-0.583	-0.520	-0.050	-0.410	-0.132
<i>Multilateral</i>							
Benchmark	0.0015	-0.620	-0.741	-0.741	-0.336	-0.173	-0.161
MC	0.0015	-0.697	-0.664	-0.769	-0.320	-0.300	-0.113
OC	0.0016	-0.784	-0.535	-0.614	-0.162	-0.360	-0.065

Notes: Column (1) reports the variance of aggregate export growth computed country-by-country or using multilateral sales ("Multilateral" line). Columns (2) and (3) are the counterfactual variances one would observe in the absence of macro-economic shocks and in the absence of all three individual shocks, respectively (i.e. $Var(g_i|\epsilon_{i,t} = 0)$ and $Var(g_i|\epsilon_{s,t}, \epsilon_{b,t}, \epsilon_{sb,t} = 0)$, resp.). Column (4) is the counterfactual variance computed using all four shocks but assuming individual transactions to be symmetric in size (i.e. $w_{sbt-1} = w_{t-1}$, $\forall(s,b)$). Finally, columns (5)–(7) are the counterfactual variations in the absence of seller-specific, buyer-specific and match-specific shocks, respectively ($Var(g_i|\epsilon_{s,t} = 0)$, $Var(g_i|\epsilon_{b,t} = 0)$ and $Var(g_i|\epsilon_{sb,t} = 0)$, resp.). All counterfactuals are expressed in relative terms with respect to the actual variance.

Table A4
Sources of heterogeneity in aggregate volatility: Comparison across models.

	Sources of volatility (shocks):			
	Macro.	Seller	Buyer	Match
<i>Model:</i>				
Benchmark	0.18*** (0.024)	0.24*** (0.027)	0.50*** (0.102)	0.09 (0.101)
Monop. Comp.	0.23*** (0.023)	0.16*** (0.034)	0.61*** (0.069)	0.02 (0.095)
Oligop.	0.12*** (0.028)	0.35*** (0.048)	0.62*** (0.033)	-0.09 (0.052)

Notes: Contribution of each family of shocks to the dispersion in destination-specific export volatilities. Contributions are computed using the decomposition in (5) and regressing each RHS term on the overall volatility. Standard errors in parentheses with ***, **, * and * respectively denoting significance at the 1%, 5% and 10% levels.

B.2. Results based on the most connected nodes

In Section A.1, we discuss potential estimation issues induced by the underlying graph being sparse, which can imply fixed effects that are imprecisely estimated (Jochmans and Weidner, 2017). To account for this possibility, this section tests the robustness of our conclusions when some, potentially many, effects are poorly estimated. To do so, we examine the robustness of our results by focusing on those effects that are precisely estimated.³⁴ Namely, we restrict our attention to the subset of buyers and sellers that are well-connected. Hence we use the distribution of the number of buyers per supplier as well as the number of suppliers per buyer to define the quality of connectedness. Then, we reassess all our results using these restricted networks.

In the benchmark estimation, we restrict to seller-buyer pairs active at least two consecutive years and we trim transactions with extreme sales growth. In this estimation sample, the average node is well connected. The average degree of sellers is 77 and the average degree of buyers is 14.

³⁴ We are thankful to Guillaume Lécué, for his insights on the topic.

However, average degrees hide an important level of dispersion. The median number of buyers per seller drops to 31 while the median number of sellers per buyer is only 5 in our sample.

Table A5

Change in volatility induced by muting shocks: Benchmark and robustness based on sub-samples restricted to the most connected nodes.

	Benchmark (1)	5 sellers per buyer (2)	5 buyers per seller (3)	15 buyers per seller (4)
Mute				
Aggregate shocks $Var(. \varepsilon_{it}, t = 0)$	-0.007	-0.009	-0.016	-0.031
Seller-specific shocks $Var(. \varepsilon_{st}, t = 0)$	-0.447	-0.429	-0.381	-0.347
Buyer-related shocks $Var(. \varepsilon_{bt}, t, \varepsilon_{sb}, t = 0)$	-0.495	-0.680	-0.573	-0.621
One buyer-related shock after the other				
Buyer-specific $Var(. \varepsilon_{bt}, t = 0)$	-0.238	-0.112	-0.216	-0.193
Match-specific $Var(. \varepsilon_{sb}, t = 0)$	-0.214	-0.542	-0.269	-0.295

Notes: This table shows the percentage change in the median firm's volatility of muting various shocks. The same counterfactuals are reproduced using estimated shocks from the benchmark model (column (1)) and various robustness sets obtained from sub-samples restricted to the most connected nodes.

Table A5 replicates the results in Table 5, based on various subsamples restricted to the most connected nodes of the graph, hence those for which the effects are better estimated. The first column is the baseline and the following three columns are the robustness checks. Each line reports the drop in seller-level volatility induced by muting different sources of shocks. In column (2), this counterfactual exercise is performed on a sample restricted to buyers interacting with at least 5 sellers (half of the observations). In this sample, the contribution of the macro and seller shocks is similar to the contribution estimated from the full sample. The contribution of buyer shocks is halved while the one of match-specific shocks is more than doubled. The smaller contribution of buyer shocks is explained by the lower volatility of buyers with at least 5 sellers in the data. The greater role of match-specific shocks is mechanical as the diversification of sellers is lower if one excludes buyers with less than 5 sellers. Overall, the estimation on this sample does not rule out our main finding that buyer-related shocks account for a sizable fraction of individual volatility.

In columns (3) and (4), we reproduce this counterfactual exercise for the sample of sellers interacting with at least 5 and 15 buyers respectively. In these samples, the contribution of seller-specific shocks is slightly lower than in the full sample, and the contribution of buyer and match specific shocks comparable to the baseline sample confirming the role of buyer related shocks for individual volatility.

While Table A5 reports counterfactual exercises for the volatility in the small, we do not report similar exercise for the volatility in the large. In the large, counterfactuals are driven by changes in volatility of the most important nodes (in terms of sales). These nodes are also the most connected nodes. Therefore, excluding poorly connected nodes does not affect counterfactual exercises in the large.

References

- Abowd, John M., Kramarz, Francis, Margolis, David N., March 1999. High wage workers and high wage firms. *Econometrica* 67 (2), 251–334.
- Abowd, John M., Creedy, Robert H., Kramarz, Francis, March 2002. Computing person and firm effects using linked longitudinal employer–employee data. *Longitudinal Employer–Household Dynamics Technical Papers 2002–06*. Center for Economic Studies, U.S. Census Bureau.
- Abowd, John M., Kramarz, Francis, Pérez-Duarte, Sébastien, Schmutte, Ian M., 2018. Sorting between and within industries: a testable model of assortative matching. *Ann. Econ. Stat.* (129), 1–32.
- Acemoglu, Daron, Carvalho, Vasco M., Ozdaglar, Asuman, Tahbaz-Salehi, Alireza, September 2012. The network origins of aggregate fluctuations. *Econometrica* 80 (5), 1977–2016.
- Andrews, M.J., Gill, L., Schank, T., Upward, R., 2008. High wage workers and low wage firms: negative assortative matching or limited mobility bias? *J. R. Stat. Soc. Ser. A* 171 (3), 673–697.
- Antràs, Pol, Fort, Teresa C., Tintelnot, Felix, 2017. The margins of global sourcing: theory and evidence from U.S. firms. *Am. Econ. Rev.* 107 (9), 2514–2564.
- Asker, John, Collard-Wexler, Allan, De Loecker, Jan, 2014. Dynamic inputs and resource (mis)allocation. *J. Polit. Econ.* 122 (5), 1013–1063.
- Atkeson, Andrew, Burstein, Ariel, December 2008. Pricing-to-market, trade costs, and international relative prices. *Am. Econ. Rev.* 98 (5), 1998–2031.
- Barrot, Jean-Noël, Sauvagnat, Julien, 2016. Input specificity and the propagation of idiosyncratic shocks in production networks. *Q. J. Econ.* 131 (3), 1543–1592.
- Bernard, Andrew B., Jensen, J., Bradford, Redding, Stephen J., Schott, Peter K., 2010. Wholesalers and retailers in US trade. *Am. Econ. Rev.* 100 (2), 408–413.
- Bernard, Andrew, Moxnes, Andreas, Ulltveit-Moe, Karen Helene, 2018. Two-sided heterogeneity and trade. *Rev. Econ. Stat.* 10 (3), 424–439.
- Bernard, Andrew B., Moxnes, Andreas, Saito, Yukiko U., 2019. Production networks, geography and firm performance. *J. Polit. Econ.* 127 (2), 639–688.
- Bernard, Andrew, Dhyne, Emmanuel, Magerman, Glenn, Manova, Kalina, Moxnes, Andreas, 2018. The Origins of Firm Heterogeneity: A Production Network Approach.
- Blundell, Richard, Robin, Jean Marc, May–June 1999. Estimation in large and disaggregated demand systems: an estimator for conditionally linear systems. *J. Appl. Econ.* 14 (3), 209–232.
- Boehm, Christoph E., Flaaen, Aaron, Pandalai-Nayar, Nitya, 2019. Input linkages and the transmission of shocks: firm-level evidence from the 2011 Tohoku earthquake. *Rev. Econ. Stat.* 101 (1), 60–75.
- Buch, Claudia M., Döpke, Jörg, Strotmann, Harald, 2009. Does export openness increase firm-level output volatility? *World Econ.* 32 (4), 531–551.
- Campbell, John Y., 2001. Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *J. Finan.* 56 (1), 1–43 02.
- Carballo, Jerónimo, Ottaviano, Gianmarco, Martinus, Christian Volpe, May 2018. The buyer margins of firms' exports. *J. Int. Econ.* 112 (1), 33–49.
- Carvalho, Vasco, Gabaix, Xavier, 2013. The great diversification and its undoing. *Am. Econ. Rev.* 103 (5), 1697–1727.
- Carvalho, V.M., Nirei, M., Saito, Y.U., Alireza Tahbaz-Salehi, A., December 2016. Supply chain disruptions: Evidence from the great East Japan earthquake. *Cambridge Working Papers in Economics 1670*. Faculty of Economics, University of Cambridge.
- Caselli, F., Koren, M., Lisicky, M., Tenreyro, S., 2015. Diversification through trade. *NBER Working Papers 21498*. National Bureau of Economic Research.
- Comin, Diego, Mulani, Sunil, 2006. Diverging trends in aggregate and firm volatility. *Rev. Econ. Stat.* 88 (2), 374–383.
- Comin, Diego A., Philippon, Thomas, March 2006. The rise in firm-level volatility: Causes and consequences. *NBER Macroeconomics Annual 2005, Volume 20*. NBER Chapters. National Bureau of Economic Research, Inc. pp. 167–228.
- Davies, Ronald B., Martin, Julien, Parenti, Mathieu, Toubal, Farid, March 2018. Knocking on tax Haven's door: Multinational firms and transfer pricing. *Rev. Econ. Stat.* 100 (1), 120–134.
- di Giovanni, Julian, Levchenko, Andrei A., August 2009. Trade openness and volatility. *Rev. Econ. Stat.* 91 (3), 558–585.
- di Giovanni, Julian, Levchenko, Andrei A., 2012. Country size, international trade, and aggregate fluctuations in granular economies. *J. Politic. Econ.* 120 (6), 1083–1132.
- di Giovanni, Julian, Levchenko, Andrei A., Mejean, Isabelle, 2014. Firms, destinations, and aggregate fluctuations. *Econometrica* 82 (4), 1303–1340 07.
- di Giovanni, Julian, Levchenko, Andrei A., Mejean, Isabelle, 2019. Foreign Shocks as Granular Fluctuations.
- Eaton, Jonathan, Kortum, Samuel, Kramarz, Francis, 2011. An anatomy of international trade: evidence from French firms. *Econometrica* 79 (5), 1453–1498 09.
- Gabaix, Xavier, 2011. The granular origins of aggregate fluctuations. *Econometrica* 79 (3), 733–772 05.
- Guimaraes, Paulo, Portugal, Pedro, 2010. A simple feasible procedure to fit models with high-dimensional fixed effects. *Stata J.* 10 (4), 628–649.
- Jochmans, Koen, Weidner, Martin, May 2017. Fixed-effect regressions on network data. *CeMMAP working papers CWP26/17*. Centre for Microdata Methods and Practice, Institute for Fiscal Studies.
- Kelly, Bryan, Lustig, Hanno, Van Nieuwerburgh, Stijn, September 2013. Firm volatility in granular networks. *NBER Working Papers 19466*. National Bureau of Economic Research, Inc.
- Koren, Miklós, Tenreyro, Silvana, 2007. Volatility and development. *Quart. J. Econ.* 122 (1), 243–287 02.

- Kurz, Christopher, Senses, Mine Z., 2016. Importing, exporting, and firm-level employment volatility. *J. Int. Econ.* 98 (C), 160–175.
- Magerman, Glenn, De Bruyne, Karolien, Dhyne, Emmanuel, Van Hove, Jan, November 2016. Heterogeneous firms and the micro origins of aggregate fluctuations. Working Papers ECARES ECARES 2016–35. ULB – Universite Libre de Bruxelles.
- Mayer, Thierry, Ottaviano, Gianmarco, May 2008. The happy few: the internationalisation of European firms. *Intereconom. Rev. Eur. Econ. Policy* 43 (3), 135–148.
- Thesmar, David, Thoenig, Mathias, 2011. Contrasting trends in firm volatility. *Am. Econ. J. Macroecon.* 3 (4), 143–180.
- Vannoorenberghe, G., 2012. Firm-level volatility and exports. *J. Int. Econ.* 86 (1), 57–67.
- Vannoorenberghe, Gonzague, Wang, Zheng, Zhihong, Yu, 2016. Volatility and diversification of exports: Firm-level theory and evidence. *Eur. Econ. Rev.* 89 (C), 216–247.