ONLINE APPENDIX

Firms, Destinations, and Aggregate Fluctuations

Julian di Giovanni Andrei A. Levchenko Isabelle Méjean

15 April, 2014

Appendix A Intensive and Extensive Margins

This appendix decomposes the growth rate of aggregate sales into the intensive and extensive components, and shows that the bulk of the aggregate sales volatility is driven by the intensive margin. The intensive component at date t is defined as the growth rate of sales of firm-destination pairs that had positive sales in both year t and year $t-1$. The extensive margin is defined as the contribution to total sales of the appearance and disappearance of firm-destination-specific sales. The log-difference growth rate of total sales can be manipulated to obtain an (exact) decomposition into intensive and extensive components:

$$\begin{aligned}
\tilde{\gamma}_{At} &\equiv \ln \sum_{f,n\in I_t} x_{fnt} - \ln \sum_{f,n\in I_{t-1}} x_{fnt-1} \\
&= \ln \frac{\sum_{f,n\in I_{t/t-1}} x_{fnt}}{\sum_{f,n\in I_{t/t-1}} x_{fnt-1}} - \left(\ln \frac{\sum_{f,n\in I_{t/t-1}} x_{fnt}}{\sum_{f,n\in I_t} x_{fnt}} - \ln \frac{\sum_{f,n\in I_{t/t-1}} x_{fnt-1}}{\sum_{f,n\in I_{t-1}} x_{fnt-1}} \right) \\
&= \underbrace{\gamma_{At}}_{Intensive\ margin} - \underbrace{\ln \frac{\pi_{t,t}}{\pi_{t,t-1}}}_{Extensive\ margin},
\end{aligned} \tag{A.1}$$

where $I_{t/t-1}$ is the set of firm-destination pairs active in both t and $t-1$ (the intensive sub-sample of firms×destinations in year t) and $\pi_{t,t}$ ($\pi_{t,t-1}$) is the share of output produced by this intensive sub-sample of firms in period t ($t-1$). Thus, the extensive margin calculation treats symmetrically entry into domestic production (a new firm appearing) and entry into exporting (an existing firm beginning to export to a particular destination n). Entrants have a positive impact on growth while exiters push the growth rate down, and the net impact is proportional to the share of entrants'/exiters' sales in aggregate sales.[1] Meanwhile, an observation only belongs to the intensive margin if an individual firm serves an individual destination in both periods.

Using equation (A.1), the impact of the intensive and extensive margins on aggregate volatility then can be written as:

$$\tilde{\sigma}_A^2 = \sigma_A^2 + \sigma_\pi^2 - 2\text{Cov}(\gamma_{At}, g_{\pi t}), \tag{A.2}$$

where $g_{\pi t} \equiv \ln \pi_{t,t}/\pi_{t,t-1}$ is the extensive margin component of equation (A.1), σ_π^2 is its variance, σ_A^2 is the variance of the intensive margin growth rate γ_{At}, and $\text{Cov}(\gamma_{At}, g_{\pi t})$ is the covariance between the two.

Inclusive of entry and exit, the volatility of total sales $\tilde{\sigma}_A^2$ is the sum of three components: i) the volatility of output produced by incumbent firms – the intensive margin, ii) the

[1] This decomposition follows the same logic as the decomposition of price indices proposed by Feenstra (1994).

volatility of entries and exits during the sample period – the extensive margin and iii) the (potential) covariance of those two terms. A convenient feature of this decomposition is that it accounts for the impact of extensive margin adjustments on aggregate volatility in a very simple way.

Though we do our best to estimate the extensive margin of firm-destination sales, there are several features of the data that may lead to overestimation of the importance of the extensive margin. First, mergers and acquisitions will appear as exits for the acquired firms, which would incorrectly add to the (negative) extensive margin.[2] Second, we cannot observe a firm's behavior prior to and after our sample period. This censoring will lead to an upward bias of the extensive margin in the first and last year of our sample, and thus we ignore these years in calculating the volatility of the extensive margin. Third, new entrants will be more likely to exhibit high growth rates as they start production and are growing towards their "steady-state" size. If young firms exhibit growth rates above the cutoff in the trimming procedure, we may record short-run entries and exits where only one entry took place. This will again overstate the importance of the extensive margin.[3]

Table OA.1. Intensive and Extensive Margins and Aggregate Volatility

	Whole Economy		*Manufacturing Sector*	
	(1)	(2)	(3)	(4)
	St. Dev.	Relative SD	St. Dev.	Relative SD
Actual	0.0228	1.0000	0.0309	1.0000
Intensive	0.0206	0.9022	0.0260	0.8429
Extensive	0.0083	0.3650	0.0103	0.3322

Notes: This table presents the standard deviations, in absolute and relative terms with respect to the actual, for the two components of aggregate growth: intensive and extensive margins, over 1992–2006.

Table OA.1 presents the standard deviations of the intensive and extensive margins, both in absolute terms and relative to the standard deviation of aggregate sales growth. We restrict attention to the period 1992–2006, because it is not possible to measure the extensive margin in the first and last years of the sample due to sampling issues discussed above. It is clear that the impact of the extensive margin on aggregate volatility is minor. While the

[2] M&A's will also lead to artificially large growth rates for the acquiring firm in the year of the M&A, which will appear in the intensive margin. The data do not record whether an M&A takes place, but our cleaning procedure discussed in Section 3 – i.e., dropping extreme growth rates – should drop the acquiring firm observation because of its large sales growth rate in the first year of acquisition.

[3] To reduce the impact of this effect on the baseline results carried out on the intensive margin, we aggregate the data over three-year periods, and the results are robust (see Section 4.4).

intensive margin aggregate volatility accounts for 90% and 84% of the overall sales volatility in the whole economy and the manufacturing sectors, respectively, the extensive margin accounts for only 37% and 33%. The results are robust to estimation of the extensive margin at three-year intervals, as well as five-year intervals, though there are fewer observations to calculate the variance for the latter, given the length of our sample period.[4]

[4]These results are available upon request.

Appendix B Relationship of $\sigma^2_{A\tau}$ to Aggregate Growth Volatility

Denote by σ^2_A the variance of γ_{At}, the growth rate of aggregate sales. Taking the variance of the right-hand side of (6), σ^2_A can be *exactly* written as the sum of the variances and covariances of the aggregated shocks:

$$\sigma^2_A = \sigma^2_{JN} + \sigma^2_F + COV, \tag{B.1}$$

where $\sigma^2_{JN} = \text{Var}\left(\sum_{j,n} w_{jnt-1}\delta_{jnt}\right)$ is the contribution of the sector-destination-specific shocks to aggregate volatility; $\sigma^2_F = \text{Var}\left(\sum_{f,n} w_{fnt-1}\varepsilon_{fnt}\right)$ is the contribution of firm-specific shocks to aggregate volatility, and $COV = \text{Cov}\left(\sum_{j,n} w_{jnt-1}\delta_{jnt}, \sum_{f,n} w_{fnt-1}\varepsilon_{fnt}\right)$ is the covariance between the shocks from different levels of aggregation.

While equation (B.1) represents an exact decomposition of the variance of γ_{At}, it is inconvenient for our purposes because it conflates the variances of shocks δ_{jnt} and ε_{fnt} with movements of the shares w_{jnt-1} and w_{fnt-1} over time. As a result, the properties of the stochastic processes $\sum_{j,n} w_{jnt-1}\delta_{jnt}$ and $\sum_{f,n} w_{fnt-1}\varepsilon_{fnt}$ are difficult to establish and relate to the properties of the primitive shocks δ_{jnt} and ε_{fnt}.

If the shares were constant over time, and the sample of firms did not change, then the aggregate variance would simply reflect the influence of the volatility of the different shocks, and (B.1) and (8) would coincide. However, this is not the case in our data: the shares and the firm-specific shocks are actually negatively correlated over time. This in turn mechanically reduces the volatility of the aggregated firm-specific shocks. To understand why this would happen, imagine a firm that either has low sales or high sales. When switching from low sales to high sales between $t-1$ and t, the firm's growth rate is large but it is weighted by the sales in $t-1$, which are low, when calculating the aggregated firm-specific component. On the other hand, when switching from high to low, the growth rate is low but this is weighted by lagged sales that are high. A negative covariance between the shocks and weights is then created when computing the contribution of this firm to the aggregate variance.

However, this does not appear to be a large force in practice. While we cannot make precise statements about the stochastic processes governing γ_{At}, $\sum_{f,n} w_{fnt-1}\varepsilon_{fnt}$, and $\sum_{j,n} w_{jnt-1}\delta_{jnt}$, we can use observed δ_{jnt}, ε_{fnt}, and w_{fnt-1}'s to calculate sample variances, which we could think of as estimators of σ^2_A, σ^2_{JN}, and σ^2_F in (B.1). Overall, these match up both qualitatively and quantitatively with the time averages of $\sigma^2_{A\tau}$, $\sigma^2_{JN\tau}$, and $\sigma^2_{F\tau}$ reported in the main text (Table 5): $\sigma_A = 0.021$ and 0.026, and $\sigma_F = 0.009$ and 0.012 for the whole economy and the manufacturing sector, respectively. The firm-specific contribution is somewhat smaller using the definition (B.1): the relative standard deviations, $\frac{\sigma_F}{\sigma_A} = 0.45$ and 0.46 for the whole economy and the manufacturing sector, respectively.

Appendix C Properties of the Estimators for $\sigma^2_{A\tau}$, $\sigma^2_{JN\tau}$, and $\sigma^2_{F\tau}$

C.1 Obtaining δ_{jnt}'s and ε_{fnt}'s

Our approach to computing δ_{jnt}'s and ε_{fnt}'s has a natural fixed effects regression interpretation. It amounts to regressing the cross-section of γ_{fnt}'s in a given year t and market n on the set of sector fixed effects, and retaining the residual as the firm-specific shock. This way of looking at things also makes it clear why we cannot isolate the macroeconomic shocks δ_{nt}. For any given market n at time t the full set of sector effects will span the country effect. Therefore, to include a constant term, a sector effect would have to be dropped, and the constant term would then capture a conflation of the aggregate shock and a shock to a "reference" sector. In turn, sector effects would then pick up sectoral shocks relative to the reference sector shock. Changing this reference sector can affect the values of δ_{nt} and δ_{jnt} as well as their variance. The combined overall impact of the macro and sectoral components remains the same regardless of the choice of the reference sector, and thus does not affect our computed values of firm-specific shocks, or their impact on the aggregate economy. The extended model (12) is implemented by fitting a linear regression on the cross-section of γ_{fnt} for each t and n, in which sector effects are interacted with the observable firm characteristics.

Note that we assume the realizations δ_{jnt} and ε_{fnt} to be observed perfectly, rather than themselves estimated. We can justify this by appealing to the fact that we are working with the universe of French firms, rather than a sample. This assumption is imposed for technical reasons. In order to establish the properties of the sample variances of a set of observed realizations of $\gamma_{At|\tau}$ and its constituent parts as estimators of their variances, as well as state the conditions on the primitives (i.e. properties of δ_{jnt} and ε_{fnt}) under which we can prove results about the properties of this estimator, we rely on the assumptions that (i) there is a well-defined and fixed set of firm-destinations, and (ii) the weights $w_{fn\tau-1}$ are fixed and known for all f, n. If we had instead assumed that we only observe estimates $\hat{\delta}_{jnt}$ and $\hat{\varepsilon}_{fnt}$ of δ_{jnt} and ε_{fnt}, asymptotics would involve proving consistency of $\hat{\delta}_{jnt}$ and $\hat{\varepsilon}_{fnt}$ as estimators of δ_{jnt} and ε_{fnt} as the sample size of firm-destinations goes to infinity. This, however, would not be logically consistent with keeping a fixed set of firm-destinations comprising the summation in $\gamma_{At|\tau}$, or with the assumption of fixed weights $w_{fn\tau-1}$.

C.2 Consistency and Asymptotic Normality

The proof follows the same steps to establish the properties of our estimators for $\sigma^2_{A\tau}$, $\sigma^2_{JN\tau}$, and $\sigma^2_{F\tau}$. Consider a vector-valued random variable $\boldsymbol{\psi_t} = \begin{pmatrix}\psi_{1t} & \psi_{2t} & \cdots & \psi_{Ft}\end{pmatrix}'$, $\boldsymbol{\psi_t} \in \mathbb{R}^F$, and a set of time-invariant weights $\mathbf{w} = \begin{pmatrix}w_1 & w_2 & \cdots & w_F\end{pmatrix}'$, $\mathbf{w} \in \mathbb{R}^F_+$. Denote by $Z_t = \mathbf{w}'\boldsymbol{\psi_t} = \sum_{i=1}^F w_i \psi_{it} \in \mathbb{R}$ a scalar-valued random variable that is a weighted sum of ψ_{it}'s. Assume we observe a stochastic process $\{\boldsymbol{\psi_t} : t = 1, \dots T\}$, and consequently a stochastic process $\{Z_t : t = 1, \dots T\}$.

In specific cases relevant for us, when $\boldsymbol{\psi_t} = \begin{pmatrix}\cdots & \varepsilon_{fnt} & \cdots\end{pmatrix}'$ is the vector of ε_{fnt} and $\mathbf{w} = \begin{pmatrix}\cdots & w_{fn\tau-1} & \cdots\end{pmatrix}'$ is the vector of firm weights at time $\tau-1$, then $Z_t = \sum_{f,n} w_{fn\tau-1}\varepsilon_{fnt}$ is the contribution of firm-specific shocks to $\gamma_{At|\tau}$ (the "granular residual"), and its variance $\sigma^2_{F\tau}$ is what we are interested in estimating. Similarly, when $\boldsymbol{\psi_t} = \begin{pmatrix}\cdots & \delta_{jnt} & \cdots\end{pmatrix}'$ and $\mathbf{w} = \begin{pmatrix}\cdots & w_{jn\tau-1} & \cdots\end{pmatrix}'$, then $Z_t = \sum_{j,n} w_{jn\tau-1}\delta_{jnt}$ with variance $\sigma^2_{JN\tau}$. Finally, when $\boldsymbol{\psi_t} = \begin{pmatrix}\cdots & \varepsilon_{fnt} & \cdots & \cdots & \delta_{jnt} & \cdots\end{pmatrix}'$ is the stacked vector of δ_{jnt} and ε_{fnt} and $\mathbf{w} = \begin{pmatrix}\cdots & w_{fn\tau-1} & \cdots & \cdots & w_{jn\tau-1} & \cdots\end{pmatrix}'$ is the stacked vector of $w_{fn\tau-1}$'s and $w_{jn\tau-1}$'s, then $Z_t = \gamma_{At|\tau}$.

This appendix states a set of sufficient conditions on the properties of the vector-valued stochastic process $\boldsymbol{\psi_t}$ and the weights vector $\mathbf{w}$ such that Z_t is stationary and the sample variance of Z_t for $t = 1, ..., T$ is a well-behaved estimator of the true variance of Z_t. Applying these conditions to the three cases above separately yields a statement of the sufficient conditions under which the sample variance of T realizations of $\gamma_{At|\tau}$ is a well-behaved estimator of $\sigma^2_{A\tau}$, and similarly for the estimators of $\sigma^2_{F\tau}$ and $\sigma^2_{JN\tau}$.

Definition 1 *A sequence $(Z_t)_{t\in\mathbb{N}}$ of random variables is called **$\boldsymbol{\alpha}$-mixing** if*

$$\alpha(m) = sup\left\{\alpha\left((Z_1, ..., Z_k), (Z_j)_{j\geq k+m}\right)|k \in \mathbb{N}\right\} \xrightarrow{m\to\infty} 0 \tag{C.1}$$

where α is the strong mixing coefficient defined as

$$\alpha(Z, X) = \sup_{\substack{A\in\sigma(Z)\\ B\in\sigma(X)}} |P(A\cap B) - P(A)P(B)|, \tag{C.2}$$

where $\sigma(Z)$ is the σ-field defined by Z.

Lemma 1 *Let $\boldsymbol{\psi_t}$ be a vector-valued jointly stationary and α-mixing stochastic process of dimension $F \times 1$ with mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\Omega}$. Denote by $Z_t \equiv \mathbf{w}'\boldsymbol{\psi_t}$ the scalar-valued process that corresponds to the weighted sum of the individual elements of $\boldsymbol{\psi_t}$. Then,*

1. *Z_t is a stationary, α-mixing process with mean μ_Z and variance σ_Z^2.*

2. *If Z_t satisfies $E|Z_t|^8 < \infty$ and $\alpha(T) = O(T^{-3})$,[5] then the sample variance*

$$s_Z^2 = \frac{1}{T-1}\sum_{t=1}^{T}(Z_t - \bar{Z}_t)^2, \tag{C.3}$$

where $\bar{Z}_t = \frac{1}{T}\sum_{t=1}^{T} Z_t$ is the sample mean, is a consistent estimator of the variance σ_Z^2 of Z_t, with a limiting distribution characterized by

$$\sqrt{T}(s_Z^2 - \sigma_Z^2) \xrightarrow{d} N(0, \xi^2), \tag{C.4}$$

where

$$\xi^2 = Var\left[(Z_t - \mu_Z)^2\right] + 2\sum_{k=1}^{\infty} Cov\left[(Z_t - \mu_Z)^2, (Z_{t+k} - \mu_Z)^2\right]. \tag{C.5}$$

Proof: The function $Z(\mathbf{x}) = \mathbf{w}'\mathbf{x}$ is measurable since $\mathbf{w}$ is known and not time-varying. Theorem 1.1 in Durrett (2005, p. 333) combined with joint stationarity of $\boldsymbol{\psi_t}$ delivers the result that $Z_t = Z(\boldsymbol{\psi_t}) = \mathbf{w}'\boldsymbol{\psi_t}$ is stationary (a measurable function of a stationary process is itself stationary). Similarly, Theorem 3.49 in White (2001, p. 50) combined with the assumption that $\boldsymbol{\psi_t}$ is α-mixing of size $-a$ delivers the result that Z_t is also α-mixing of size $-a$ (a measurable function of an α-mixing process is itself α-mixing). This proves the first claim. The second claim follows directly from Theorem 1.8 of Dehling and Wendler (2010, p. 128), since Z_t satisfies all the conditions required in that theorem and it is easily verified that the $U-$statistic corresponding to the sample variance satisfies the moment and continuity conditions of that theorem. ■

C.3 Standard Errors

As is customary, in our empirical implementation we will compute the confidence intervals based on the empirical counterpart of (C.5):

$$\hat{\xi}^2 = \frac{1}{T-1}\sum_{t=1}^{T}\left[(Z_t - \bar{Z}_t)^2 - s_Z^2\right]^2 + 2\sum_{k=1}^{T-2}\frac{1}{T-k-1}\sum_{t=1}^{T-k}\left[(Z_t - \bar{Z}_t)^2 - s_Z^2\right]\left[(Z_{t+k} - \bar{Z}_t)^2 - s_Z^2\right] \tag{C.6}$$

[5] The statement of these conditions can be made more general. Namely, the proposition holds if $\exists\ \nu > 0$ and $\exists\ \phi > 0$ such that $E|Z_t|^{\max\{\phi, 2(2+\nu)\}} < \infty$ and $\alpha(T) = O(T^{-\rho})$ for $\rho > \frac{3\phi\nu+\nu+5\phi+2}{2\phi\nu}$. This more general statement of the conditions captures the tradeoff between the number of finite moments and the degree of time dependence: one can allow for more time dependence (lower ρ) if one assumes existence of higher order finite moments, and vice versa.

For large k, the object $\sum_{t=1}^{T-k}\left[(Z_t-\bar{Z}_t)^2-s_Z^2\right]\left[(Z_{t+k}-\bar{Z}_t)^2-s_Z^2\right]$ cannot be precisely estimated. Thus, we cut the number of maximum allowable lags to $q \ll T-2$ and use the HAC estimator that downweights more distant covariances (Newey and West, 1987):

$$\begin{aligned}\hat{\xi}_{HAC}^2 \quad = \quad & \frac{1}{T-1}\sum_{t=1}^{T}\left[(Z_t-\bar{Z}_t)^2-s_Z^2\right]^2 + \\ & 2\sum_{k=1}^{q}\left[1-\frac{k}{q+1}\right]\frac{1}{T-k-1}\sum_{t=1}^{T-k}\left[(Z_t-\bar{Z}_t)^2-s_Z^2\right]\left[(Z_{t+k}-\bar{Z}_t)^2-s_Z^2\right]\end{aligned} \tag{C.7}$$

Following Andrews (1991), we choose q as a function of sample size according to the following rule of thumb:

$$q+1 \approx 0.75T^{1/3}$$

For us, with $T = 17$, this amounts to $q + 1 \approx 2$, so we only use the covariance for one lag. We are interested in the standard error of s_Z^2, which is obtained by dividing $\hat{\xi}_{HAC}^2$ by $\sqrt{T}$. Finally, the figures and tables in the main text report the results expressed in terms of the standard deviation s_Z. We use the delta method to obtain the standard error of the standard deviation.

Appendix D Detailed Data Description

The sales data, as well as additional variables, come from the balance sheet information collected from firms' tax forms. The French tax system distinguishes three different regimes, the "normal" regime (called BRN for *Bénéfice Réel Normal*), the "simplified" regime (called RSI for *Régime Simplifié d'Imposition*) that is restricted to smaller firms, and the "micro-BIC" regime for entrepreneurs. The amount of information that has to be provided to the fiscal administration is more limited in the RSI than in the BRN regime, and even more so for "micro-BIC" firms. Under some conditions, firms can choose their tax regime. An individual entrepreneur can thus decide to enroll in the "micro-BIC" regime if its annual sales are below 80,300 euros. Likewise, a firm can choose to participate in the RSI rather than the BRN regime if its annual sales are below 766,000 euros (231,000 euros in services).[6]

Throughout the exercise, "micro-BIC" and "RSI" firms are excluded. We do not have enough information for "micro-BIC" firms. We also exclude "RSI" firms, both because their weight in annual sales is negligible and because it is difficult to harmonize these data with the rest of the sample. In 2007, those firms represented less than 4% of total sales and about 11% of total employment. Thus, our sample represents the bulk of the aggregate French economy.

The BRN dataset contains detailed information on the firms' balance sheets, including total, domestic, and export sales, value added, as well as many cost items including the wage bill, materials expenditures, and so on, as well as NAF sectors in which the firm operates.[7] This represents around 30% of industrial and service firms but more than 90% of aggregate sales.[8] We do not have any information at the plant level, however.

The information collected by the tax authorities is combined with the firm-level ex-

[6] Those thresholds are for 2010. They are adjusted over time, but marginally so.

[7] "NAF", *Nomenclature d'Activités Française*, is the French industrial classification. Our baseline analysis considers the level of aggregation with 60 sectors. This corresponds to the 2-digit ISIC (Revision 3) nomenclature. We merge together some sectors in order for our nomenclature to be consistent with the one used in the input-output tables. Namely, we merge agriculture, forestry and fishing (NAF 1, 2 and 5), all mining and quarrying activities (NAF 10 to 14), tobacco and other food industries (NAF 15 and 16), textile, wearing apparel and leather (NAF 17, 18 and 19), paper products and publishing (NAF 21 and 22), manufacturing n.e.c and recycling (NAF 36 and 37), all activities related to electricity gas and water (NAF 40 and 41), wholesale and retail trade (NAF 50, 51 and 52), transport and storage activities (NAF 60 to 63) and all community, social and personal services (NAF 90 to 93). We also drop NAF sectors 95 (domestic services), and 99 (activities outside France). The NAF nomenclature has been created in 1993, as a replacement for the "NES" (Nomenclature Economique de Synthèse). Data for 1990–92 are converted into the NAF classification using a correspondence table.

[8] We drop the banking sector because of important restructuring at the beginning of the 2000s that artificially adds a large amount of volatility to the dataset. This sector represents less than 4% of total sales in 1990 but more than 25% by the end of the period.

port data for each foreign destination market from the French customs authorities. The datasets can be merged using a unique firm identifier, called SIREN. In merging together the customs and balance sheet data, we had to make a number of adjustments. First, we drop observations for firms that appear in the customs but do not appear in the BRN data (some of these firms may produce farming goods, which are not in the balance sheet data). Second, a number of firms declare positive exports to the tax authorities but are not in the customs files. Since our procedure exploits the bilateral dimension of exports, and the customs data are the most reliable source of exporting information, we assume that those firms are non-exporters. Third, in a small fraction (6.6%) of exporter-year observations present in both the customs and the BRN data, the value of export sales is not the same in the two databases. We thus use the customs data to compute the share of each destination market in total firm exports and apply these shares to export sales provided in the BRN file.

The customs data are quasi-exhaustive. There is a declaration threshold of 1,000 euros for annual exports to any given destination. Below the threshold, the customs declaration is not compulsory. Since 1993, intra-EU trade is no longer liable for any tariff, and as a consequence firms are no longer required to submit the regular customs form. A new form has however been created that tracks intra-EU trade. Unfortunately, the declaration threshold for this kind of trade flows in much higher, around 150,000 euros per year. A number of firms continue declaring intra-EU export flows below the threshold however, either because they don't know ex ante that they will not reach the 150,000 Euro limit in a given fiscal year, because they apply the same customs procedure for all export markets they serve, or because they delegate the customs-related tasks to a third party (e.g., a transport firm) that systematically fills out the customs form. Below-cutoff exports missing from customs data can potentially create two problems (i) some export sales might be counted as domestic, affecting the computation of domestic shocks; and (ii) some export sales that occur in reality (a subset of those below 150,000 euros) are missing from our data, affecting the computation of export shocks. We use the information contained in the tax forms to both deal with this problem and assess its extent. On the tax form, the firms report their total exports. Thus, we can conjecture that firms that do not appear in customs data but report positive exports on their tax forms are those for whom exports (by destination) fall below the customs cutoff. We address problem (i) by calculating the firm's domestic sales as the difference between their reported total sales and their exports reported on the tax form. In this way, we do not "contaminate" domestic sales with erroneously classified exports.

Below, we report our main results for domestic sales only, and they are robust. For problem (ii), this fix is not available. We can judge how many exports we are missing by comparing exports declared on tax forms to exports declared to customs. It appears that the problem is relatively minor. In 10% of firm-year observations, the tax form reports exports but the customs data do not. These observations account for 7% of overall exports. On average, the total exports reported in the tax form but missing from customs (413,000 euros per year) are an order of magnitude smaller than average exports in the whole sample, which are 3,056,000.

Our approach involves working with the sales growth rates of firms to individual markets. One concern with these data is that firm sales could be measured with error, and thus the volatility of firm-specific shocks we estimate may simply be the variance of the measurement error. As is typical of micro data, there is a great deal of dispersion in the set of individual growth rates we obtain. There are a number of reasons for which the data may contain outliers. For instance, the BRN file does not provide any information on firms undergoing a *controle fiscal* – i.e. a tax audit – during a given year. For these firms, the "sales" variable is either zero or missing, which results in either extreme growth rates or artificial exits and re-entries around those years.[9] Also, firms can change their organizational structure in a given year, grouping activities together in different entities, which can result in a number of large "exits." In a number of cases, firms decided to create new holding companies that pooled together the charges and benefits of all firms comprising the group. The members of those groups, that before filed separate tax forms, would then disappear from the fiscal files.

While measurement error is by construction impossible to rule out, we believe that our results are not unduly driven by it for a number of reasons. First, the French data we are working with are high quality, coming from tax and customs records. These are the data underlying the national accounts for France. Second, in order for extreme observations not to introduce noise in the estimation and aggregation exercise, we apply a trimming procedure. Namely, we drop the individual growth rates in which sales are either double or half their previous year's value. Third, we repeat the analysis on 3-year growth rates instead of annual growth rates as one of the robustness checks, a procedure that should help average out year-to-year measurement error. The fact that 3-year growth rates continue

[9] The audit of the firm's tax statements is over the period going back 3 years, and up to 10 years if fraud is detected. This is relevant for us because this process often lasts for many months, and during the year the company is in *controle fiscal*, there is no "regular" BRN declaration which may result in missing data values for certain firm-years, even for big companies. Unfortunately, we do not have data on which firm-years are under *controle fiscal*.

to produce a significant firm component for aggregate fluctuations suggests that the main results in the paper are not driven by measurement error.

Appendix E A Simple Model of Input-Output Linkages at the Firm Level

This appendix presents a simple extension of the baseline model of Section 2 to illustrate how interconnections between firms can generate positive correlation in the estimated firm-specific shocks. We model the interconnection through input-output linkages.

Suppose that the sales of a firm are given by (3), but the cost of the input bundle is now firm- rather than sector-specific:

$$x_{fnt} = \omega_{fnt} \frac{\varphi_{jnt} Y_{nt}}{(P_{jnt})^{1-\theta}} \left(\frac{\theta}{\theta - 1} \kappa_{jnd} c_{fdt} a_{fdt} \right)^{1-\theta},$$

where

$$c_{fdt} = A h_{dt}^{\lambda_f} \prod_{g \in \Xi_{fdt}} p_{gdt}^{(1-\lambda_f)\rho_{fg}}, \quad \sum_g \rho_{fg} = 1.$$

This specification assumes that the cost of firm f's input bundle c_{fdt} has a Cobb-Douglas form in labor, paid the equilibrium wage h_{dt}, and the set Ξ_{fdt} of inputs bought from the firm's input providers at their equilibrium price p_{gdt}. The parameter λ_f measures the share of labor in the firm's cost function, and ρ_{fg} is the share of spending on inputs produced by firm g in the total intermediate input spending by firm f. Finally, A is a constant that depends on the parameters of the production function.

Productivity shocks to an input provider g have a direct effect on its sales: $d \ln x_{gmt} / d \ln a_{gmt} = 1 - \theta$. Because of input-output linkages, they also transmit to firm f with the following elasticity:

$$\frac{d \ln x_{fnt}}{d \ln a_{gmt}} = (1-\theta)(1-\lambda_f)\rho_{fg}.$$

Intuitively, a positive productivity shock decreases the upstream firm's output price and thus the downstream firm's input cost, positively affecting its sales. This transmission of shocks via the IO linkage implies that the sales growth rates of firms f and g exhibit positive comovement.

In particular, if idiosyncratic firm-specific productivity shocks are the only source of shocks in the economy, the covariance of the firm-specific sales growth components between any two firms f and g is

$$\begin{aligned}\mathrm{Cov}\left(\varepsilon_{fnt},\varepsilon_{gmt}\right) = (1-\theta)^2 \Bigg[& \underbrace{(1-\lambda_g)\rho_{gf}\mathrm{Var}(a_{fdt})}_{\text{Propagation from } f \text{ to } g} + \underbrace{(1-\lambda_f)\rho_{fg}\mathrm{Var}(a_{gdt})}_{\text{Propagation from } g \text{ to } f} \\ & + \underbrace{\sum_{h\in\Xi_{fdt}\cap\Xi_{gdt}} (1-\lambda_f)(1-\lambda_g)\rho_{fh}\rho_{gh}\mathrm{Var}(a_{hdt})}_{\text{Propagation through common input providers}} \Bigg]. \end{aligned} \tag{E.1}$$

Summing over all firms connected to f and assuming that the variance of shocks is homogeneous over firms ($\mathrm{Var}(a_{fnt}) = \sigma^2 \ \forall\ f, n$), one can recover the contribution of a single firm to the overall linkage factor (neglecting the impact of weights):

$$\begin{aligned}\sum_{g,m}\mathrm{Cov}(\varepsilon_{fnt},\varepsilon_{gmt}) = (1-\theta)^2\sigma^2 \Bigg[& \underbrace{\sum_g (1-\lambda_g)\rho_{gf}}_{\text{Weighted out-degree } d_f} + (1-\lambda_f) \\ & + (1-\lambda_f)\underbrace{\sum_{g,m}\sum_{h\in\Xi_{fdt}\cap\Xi_{gdt}} (1-\lambda_g)\rho_{fh}\rho_{gh}}_{\text{Second-order degree } q_f} \Bigg]. \end{aligned} \tag{E.2}$$

As in Acemoglu et al. (2012), the impact of one single firm on the aggregate volatility depends on how connected it is to the rest of the economy. Shocks affecting a firm that provides inputs to a large number of downstream players, i.e., that has a large "weighted out-degree" d_f in the words of Acemoglu et al. (2012), will have a larger impact. This is what the first term of (E.2) captures. The second term accounts for the fact that firms that use more inputs will fluctuate more as a result of productivity shocks affecting their input providers. Finally, the third term captures "second-order connections" as denoted by Acemoglu et al. (2012) – namely the fact that common input suppliers magnify the propagation of shocks across firms.

Ideally, one would like to investigate the role of firm-level linkages in aggregate fluctuations using the insights of (E.1) and (E.2). Using these equations, it is possible to correlate the magnitude of covariances at the firm-level to appropriate measures of linkages. Unfortunately, such firm-level measures of IO linkages are not available for France. Instead, we use sectoral data on IO linkages as a proxy for the intensity of production networks. The implicit assumption is that those sectoral measures of IO linkages are a good proxy for the magnitude of interconnections between firms belonging to those sectors. Since the information is available at the level of each sector pair, we need to correlate them with measures of the $LINK$ term that are also defined by sector pair.

Recall the definition of the $LINK$ term and write it as the sum over all sector pairs in the economy:

$$LINK_{\tau} = \sum_{g \neq f, m \neq n} \sum_{f,n} w_{gm\tau-1} w_{fn\tau-1} \text{Cov}(\varepsilon_{gmt}, \varepsilon_{fnt}) = \sum_{i} \sum_{j} LINK_{ij\tau}, \text{ where}$$
$$LINK_{ij\tau} = \sum_{g,m \in j} \sum_{f,n \in i} w_{gm\tau-1} w_{fn\tau-1} \text{Cov}(\varepsilon_{gmt}, \varepsilon_{fnt}),$$

and $\text{Cov}(\varepsilon_{gmt}, \varepsilon_{fnt})$ is defined by (E.1).

Assume that i) individual volatilities are homogeneous across firms: $\text{Var}(a_{fdt}) = \sigma^2 \ \forall \ f$; ii) the IO coefficients are homogeneous between firms within a sector: $(1 - \lambda_f) = (1 - \lambda_i) \ \forall \ f \in i$ and $\rho_{fg} = \rho_{ij} \ \forall \ f \in i, g \in j$, and iii) $\Xi_{fdt} \cap \Xi_{gdt}$ is homogeneous between firms within a sector pair. Then the $LINK$ term becomes

$$LINK_{ij\tau} = w_{jm\tau-1} w_{in\tau-1} \sigma^2 (1-\theta)^2 \left[\underbrace{(1-\lambda_j)\rho^{ji} + (1-\lambda_i)\rho_{ij}}_{\text{First-order}} + \underbrace{\sum_k (1-\lambda_i)(1-\lambda_j)\rho^{ik}\rho^{jk}}_{\text{Second-order}} \right].$$

This expression thus motivates our approach in Section 4.3.2 of looking for a relationship between the $LINK_{\tau}$ term and the strength of IO linkages between the sectors.

Appendix F Heterogeneous Response to Shocks at the Firm Level

This appendix develops a variant of the model in Section 2 with variable markups. In this more general framework, firms react heterogeneously to common shocks. When this is the case, the firm-specific effect in the baseline estimation would capture not only the impact on firm sales of idiosyncratic shocks but also the heterogeneous response of the firm to sector-destination shocks. The model serves to motivate the alternative empirical model (12), in which sector-destination shocks affect firm sales differently depending on firm characteristics. The main results are robust to this alternative conceptual framework and empirical model.

Consider the model in Section 2 that has Cobb-Douglas preferences over sectors and CES preferences over varieties within a sector. As before, each firm faces the following demand in market n:

$$C_{fnt} = \left(\frac{p_{fnt}}{P_{jnt}}\right)^{-\theta} \omega_{fnt} \frac{\varphi_{jnt} Y_{nt}}{P_{jnt}},$$

where variables are defined in Section 2, and p_{fnt} is the consumer price of firm f's product in market n.

The baseline model assumes the standard "iceberg" multiplicative cost of delivering one unit of the good to market n. Suppose instead, following Berman et al. (2012), that the variable trade cost has two components, one multiplicative and one additive. The consumer price in market n is then

$$p_{fnt} = \tilde{p}_{fnt} \kappa_{jndt} + \eta_{jndt},$$

where $\tilde{p}_{fnt}$ is the producer price, κ_{jndt} the multiplicative variable trade cost, and η_{jndt} the additive variable trade cost.[10] Both κ_{jndt} and η_{jndt} are assumed to be the same for all firms within a sector selling goods to the same destination market.

A per-unit component of variable trade cost implies that, even under CES preferences, individual markups are not homogeneous across firms. Namely, profit maximization leads to the following producer price:

$$\tilde{p}_{fnt} = \frac{\theta}{\theta - 1} m_{fnt} a_{fdt} c_{jdt},$$

[10] The additive cost η_{jndt} can either be thought of as a distribution cost or a per-unit transportation cost. When thinking of it as a distribution cost, it makes sense to assume this cost is paid using foreign labor. This does not change the main results, but introduces an additional source of sector-destination shocks since the optimal markup then depends on the destination market's wage.

where

$$m_{fnt} \equiv 1 + \frac{\eta_{jndt}}{\theta \kappa_{jndt} a_{fdt} c_{jdt}},$$

is the variable component of markups. Importantly, this component is affected by sectoral cost movements (changes in c_{jdt}) as well as changes in variable trade costs (κ_{jndt} and η_{jndt}). Moreover, the elasticity of m_{fnt} with respect to sector-destination shocks is heterogeneous across firms, and depends on the individual productivity level (a_{fdt}). Identical shocks can thus have different effects on firms sales growth.

Conditional on selling to market n, (f.o.b.) sales by a French firm f (i.e., residing in country d) to market n in period t are thus given by:

$$\begin{aligned} x_{fnt} &= \tilde{p}_{fnt} C_{fnt} \\ &= \omega_{fnt} \frac{\varphi_{jnt} Y_{nt}}{(P_{jnt})^{1-\theta}} \left(\frac{\theta}{\theta - 1} \kappa_{jndt} c_{jdt} a_{fdt} \right)^{1-\theta} \left(\frac{m_{fnt}}{\kappa_{jndt}} \right)^{1-\theta} \left(\frac{p_{fnt}}{\tilde{p}_{fnt}} \right)^{-\theta}. \end{aligned} \tag{F.1}$$

If we were to use (F.1) to write a decomposition of firm sales growth as a function of country, sector-destination and firm-destination shocks as in (4):

$$\gamma_{fnt} = \delta_{nt} + \delta_{jnt} + \varepsilon_{fnt},$$

the firm-specific component would now be

$$\varepsilon_{fnt} = \Delta \log \omega_{fnt} + (1-\theta) \Delta \log a_{fdt} + (1-\theta) \Delta \log m_{fnt} - \theta \Delta \log \left(\frac{\tilde{p}_{fnt}}{p_{fnt}} \right),$$

The first two terms are firm-specific by construction, as before. However, the last two terms, $(1-\theta)\Delta \log m_{fnt} - \theta \Delta \log \left(\frac{\tilde{p}_{fnt}}{p_{fnt}} \right)$, depend on sectoral shocks (and on the macro shocks if the distribution cost is paid in foreign labor). These terms capture firms' heterogeneous response to common shocks.

In particular, the impact of a sectoral cost shock on the firm-level sales is

$$\begin{aligned} \frac{d \ln x_{fnt}}{d \ln c_{jdt}} &= (1-\theta) + (1-\theta) \frac{d \ln m_{fnt}}{d \ln c_{jdt}} - \theta \frac{d \ln \left(\frac{\tilde{p}_{fnt}}{p_{fnt}} \right)}{d \ln c_{jdt}} \\ \text{where} \quad \frac{d \ln m_{fnt}}{d \ln c_{jdt}} &= \frac{-\eta_{jndt}}{\theta \kappa_{jndt} a_{fnt} c_{jdt} + \eta_{jndt}} \in [-1, 0] \\ \text{and} \quad \frac{d \ln \left(\frac{\tilde{p}_{fnt}}{p_{fnt}} \right)}{d \ln c_{jdt}} &= \frac{-\eta_{jndt}}{p_{fnt}} \left(1 + \frac{d \ln m_{fnt}}{d \ln c_{jdt}} \right) < 0 \end{aligned}$$

The first term captures the direct effect of the shock on the firm's marginal cost, which is homogeneous across firms and captured in the δ_{jnt} term of equation (5). The second term, which would be captured in ε_{fnt}, reflects the response of the firm's markup to the shock. When the cost of the input bundle increases, firms reduce their optimal markup, more so the more productive they are. This markup adjustment tends to attenuate the effect of the sectoral shock on sales of the more productive firms. Finally, the third term captures the adjustment in the ratio of the consumer to the producer prices. The combined effect of the cost shock and the markup adjustment on this ratio further attenuates the direct impact of the sectoral shock.

From an econometric point of view, endogenous markup adjustments would induce a negative correlation between the sector-destination fixed effects and the residual term of equation (5). To control for this bias, we thus implement equation (12) that interacts the sector-destination effect with a number of measures, many of which can be thought of as proxies for firm productivity. Following the model laid out in this section, these interaction terms are intended to capture the larger markup adjustment of the more productive firms in response to sector-destination shocks.

Appendix References

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The Network Origins of Aggregate Fluctuations," *Econometrica*, September 2012, *80* (5), 1977–2016.

Andrews, Donald W. K., "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, May 1991, *59* (3), 817–858.

Berman, Nicolas, Philippe Martin, and Thierry Mayer, "How do Different Exporters React to Exchange Rate Changes? Theory, Empirics and Aggregate Implications," *Quarterly Journal of Economics*, 2012, *127* (1), 437–492.

Dehling, Herold and Martin Wendler, "Central Limit Theorem and the Bootstrap for U-Statistics of Strongly Mixing Data," *Journal of Multivariate Analysis*, January 2010, *101* (1), 126–136.

Durrett, Richard, *Probability: Theory and Examples*, Belmont, CA: Brooks/Cole – Thomson Learning, 2005.

Feenstra, Robert, "New Product Varieties and the Measurement of International Prices," *American Economic Review*, March 1994, *84* (1), 157–177.

Newey, Whitney K. and Kenneth D. West, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, May 1987, *55* (3), 703–708.

White, Halbert, *Asymptotic Theory for Econometricians*, San Diego: Academic Press, 2001.