Combining the Granular and Network Origins of Aggregate Fluctuations

A. Hempfing
J. Schulz
M. Milaković

September 21, 2018

Working Paper

Abstract

The desire to understand business cycle fluctuations, trade interdependencies and co-movement has a long tradition in economic thinking. Recent evidence in both theoretical and empirical studies has shown that a small number of entities such as firms or industries can have sizeable effects on aggregate economic fluctuations. One strand of literature focuses on the size of these "granular" entities, while another points out the relevance of cascade effects due to significant asymmetries in the roles of industries in the economic network. This study combines both approaches and empirically examines the contribution of industry-level idiosyncratic shocks to aggregate fluctuations in Europe. To perform our empirical analysis, we use the World Input-Output Database (WIOD). The WIOD consists of annual time series of World Input-Output Tables (WIOTs) covering, among others, all 28 EU countries for the period from 2000 to 2014. The analysis shows that shocks to large industries are of relevance, explaining roughly 20% of output fluctuations in the EU. This explanatory power increases considerably when controlling for sector-specific shocks. More importantly, when including the network perspective, the explanatory power increases to almost 80% of explained variance in GDP growth rates. We find that the European economy is indeed granular in industries and that combining the granular and network approaches significantly improves the explanatory power for the origins of aggregate fluctuations. Hence, aggregate fluctuations in the European economy seem not only to be dependent on the size of an industry but also on its position within the European production network.

1 Introduction

This paper proposes a simple model unifying the granular and network origins of aggregate fluctuations. We combine both methods to show that weighted idiosyncratic shocks to individual industries significantly improve the explanatory power of fluctuations in the aggregate. The desire to understand business cycle fluctuations building from proper microfoundations at the consumer- or firm-level has characterized macroeconomics for decades. Conventional macroeconomic thinking with its emphasis on representative agents has, however, largely discarded the view of an economy consisting of interconnected and interacting heterogeneous entities such as firms and industries. In contrast, the industrial dynamics literature building on this vision has left the macroeconomic implications of the observed statistical regularities at the micro-level largely unexplored.

It was the two seminal studies by Gabaix (2011) and Acemoğlu et al. (2012) that first brought these two perspectives together. They showed that the skewed and fat-tailed distributions of size- and network-based metrics at the firm- and industry-level imply that a large fraction of business cycle fluctuations arise from idiosyncratic shocks to large and strongly interconnected entities in blatant contradiction to the implications of conventional representative agent-frameworks. In this study, we show that the interaction of both channels yields results whose explanatory power is superior to both approaches on their own. Intriguingly, both the irrelevance result for idiosyncratic shocks to aggregate fluctuations following Lucas (1977) as well as the studies by Gabaix (2011) and Acemoğlu et al. (2012) arguing precisely the opposite are based on the very same foundational theorem. Hulten (1978) showed with minimal assumptions for economies in general equilibrium, that the first-order impact on output of a TFP shock to a firm or an industry is equal to that industry's or firm's sales as a share of output, that is, its respective Domar weight (Domar, 1961). This implies that we can, as a first-order approximation, write the aggregate TFP shock as the Domar weighted average of idiosyncratic TFP shocks to the individual firms or industries. Based on Hulten's theorem, Lucas (1977) can now argue that as economies consist of millions of firms, the law of large numbers implies with near-certainty that the Domar weighted idiosyncratic shocks should average out to zero, if their respective Domar weights are also close to zero.

Gabaix (2011) and Acemoğlu et al. (2012) challenge precisely this idea that sales shares are undoubtedly close to zero. Gabaix (2011) argues that shocks to very large or, in his jargon, "granular" firms will not average out with shocks to much smaller firms and thus might result in aggregate fluctuations. Combining results from Acemoğlu et al. (2010) and Carvalho (2010), Acemoğlu et al. (2012) use a Cobb-Douglas model in the spirit of Long Jr and Plosser (1983) to show that the equilibrium size of sectors is shaped by their associated input-output matrix. Industries are weighted proportionally to their centrality as a supplier in the network and thus, shocks to such central players do not cancel out in the aggregate. Thus, this notion of linkages across sectors can be central to economic performance and dates back at least to Leontief (1936). However, for both the "granular" and network approach to aggregate fluctuations, the Domar weights emerge as the correct weights for idiosyncratic shocks in line with Hulten's theorem. This has led Gabaix (2016) to conclude in a recent survey article that “networks are a particular case of granularity rather than an alternative to it.”

Here, our approach differs. We build on recent theoretical results that the the first-order approximation by Hulten’s theorem is only exact under a Cobb-Douglas production function but higher-order terms become non-negligible for different functional forms, in particular in the presence of complementarities (Baqaee and Farhi, 2017). If complementarities are indeed present in supply chains, as both casual empiricism as well as empirical studies suggest (Atalay, 2017; Boehm et al., 2018), idiosyncratic shocks to central firms or industries in the production network should be weighted more strongly. Such magnifying effects by the propagation of shocks through the production network are analyzed, amongst others, in Di Giovanni et al. (2014), Foerster et al. (2011) and Stella (2015). In this study, we analyze the GDP growth rates of the EU28 countries with their underlying industries. The EU28 seem to be an ideal candidate for the exploration of these theoretical results, as it represents a large and strongly interconnected economic zone through the European single market that was, considering the recent crises, characterized by large variation in GDP growth rates.

We indeed find that the combination of both size- and network-based approaches result in the highest explanatory power of about 80% of explained variance through weighted idiosyncratic productivity shocks to a very small number of large and strongly interconnected industries. We achieve this result without resorting to-questionable procedures applied to the distribution of productivity growth rates such a "winsorizing" the data, like, as we argue, Gabaix (2011) does. We find that the European economy is granular in industries and that combining granular and network approaches significantly improves the explanatory power of weighted idiosyncratic shocks. Hence, aggregate fluctuations of the European economy seem not only to be dependent on the size of an industry but also on its position within the European production network. We speculate that this result might be of great interest for predictive purposes, as these interaction of network- and size-based metrics allows to identify the small number of industries driving aggregate fluctuations. Also, we hope that this study in proposing a new channel for aggregate fluctuations and a further step into the direction for proper microfoundations for macroeconomic modeling away from the representative-agent framework that has found to fail to explain the macro-moments of output time series (Ascari et al., 2015).

The remainder of this paper is organized as follows: Section 2 will provide the basic model we try to empirically assess, our hypotheses and the econometric identification strategy. Section 3 introduces and discusses the database we use, while section 4 presents our generated results. Section 5 concludes and gives a short outlook on possible directions for further research.

2 Model

2.1 Measuring Aggregate Fluctuations

In line with Gabaix (2011), we take Hulten's theorem as the starting point for our model and econometric specification. Hulten (1978) shows that in an economy with linkages, the TFP growth rate of the total economy in general equilibrium is given by the Domar weighted sum of idiosyncratic productivity shocks to N firms or industries as

$$\frac{dTFP_t}{TFP_t} = \sum_{i=1}^{N} \frac{S_{i,t}}{Y_t} d\pi_{i,t}, \tag{1}$$

where $\pi_{i,t}$ is the productivity of firm or industry i in period t, $S_{i,t}$ its sales at t and TFP_t is the Total Factor Productivity in period t. Given that the productivity shocks are Hicks-neutral, these cumulative idiosyncratic shocks multiplied by a factor μ reflecting factor usage can explain GDP growth by

$$\frac{dY_t}{Y_t} = \mu \cdot \sum_{i=1}^{N} \frac{S_{i,t}}{Y_t} d\pi_{i,t}. \tag{2}$$

Gabaix (2011) shows that for sufficiently heterogeneous Domar weights, in particular a firm-size distribution with an upper power-law tail with tail exponent α between unity and two, idionsyncratic shocks to the largest firms create non-vanishing aggregate fluctuations.[1] Assuming that Hulten's theorem indeed holds exactly for the economy in question, we introduce Gabaix' "granular residual" $\mathcal{G}$ as a measure for the impact of these largest industries based purely on observables by

$$\mathcal{G}_t = \sum_{i=1}^{K} \frac{S_{i,t-1}}{Y_{t-1}} (g_{i,t} - \bar{g}_t), \tag{3}$$

where $S_{i,t-1}$ in the gross output of sector i in year $t-1$, Y_{t-1} is the real GDP of the European Union in year $t-1$ and the term $g_{it} - \bar{g}_t$ the median-subtracted labour productivity growth rate. The labour producity growth of a sector i in year t is defined as $g_{i,t} = \ln(\frac{S_{i,t}}{E_{i,t}}) - \ln(\frac{S_{i,t-1}}{E_{i,t-1}})$, where $E_{i,t}$ is the number of employees employed in a specific sector i in year t. $\bar{g}_t$ is the cross-sectional median[2] of the top Q sectors of $g_{i,t}$ with $Q = K$. How to calibrate K and Q will be discussed in the section below.

[1]The intuition for these boundary values for the power-law tail exponent is quite clear. The aggregation of values drawn from a distribution with finite variance implies, by the Central Limit Theorem (CLT) converges asymptotically to a Gaussian distribution. A power-law with a tail exponent above 2 has finite variance. Thus, by this CLT argument and as Gabaix (2011) shows, this condition implies vanishing fluctuations at the aggregate level. Conversely, for power-laws with tail exponents between unity and two and thus infinite variance, there exist non-vanishing aggregate fluctuations even for a very large number of firms.

[2]Notice that this is a deviation from Gabaix' proposed specification. He is substracting the *mean*, while we take the *median* of productivity growth rates. This is due to the fact that productivity growth rate distributions are likely to be leptokurtic and the median being less sensitive to outliers is in this sense more representative of an aggregate shock. We discuss this issue in more detail in the subsequent section.

In line with Gabaix (2011) we introduce another specification to control for the median growth $\bar{g}_{I,t}$, the equal-weighted median productivity growth rate among sectors that are in the same industry as i and among the top Q sectors therein. This specification now reads:

$$\tilde{\mathcal{G}}_t = \sum_{i=1}^{K} \frac{S_{i,t-1}}{Y_{t-1}} (g_{i,t} - \bar{g}_{I,t}). \tag{4}$$

Introducing the term $g_{i,t} - \bar{g}_{I,t}$ may control better for industry-specific disturbance, as it is the case for $g_{i,t} - \bar{g}_t$, capturing e.g. industry-wide real price movements. This specification $\tilde{\mathcal{G}}$ is again taken from Gabaix (2011) who shows in the associated online appendix that even for a limit-case, where $\mathcal{G}$ is pure noise, $\tilde{\mathcal{G}}$ still contains information on the *ideal* granular residual that would indeed measure the correct idiosyncratic shocks.

Now we are able to introduce our simple model unifying the granular and network origins of aggregate fluctuations. In the next two specifications, $\mathcal{E}$ and $\tilde{\mathcal{E}}$, we explicitly take the network structure of the given input-output matrix into account and combine it with the approach above. In this, we deviate from both Acemoğlu et al. (2012) and Gabaix (2011) which state that either of both approaches would be sufficient. In contrast to their results, we argue that not only the size of an industry is important but also its position within the respective production network. To capture this notion, the eigenvector centrality is used to weighten the sector-specific granular residuals. The two specifications analagous to specifications (3) and (4) without and with controlling for industry-specific shocks therefore read

$$\mathcal{E}_t = \sum_{i=1}^{K} \left[e_{i,t} \cdot \left(\frac{S_{i,t-1}}{Y_{t-1}} (g_{i,t} - \bar{g}_t) \right) \right] \tag{5}$$

and

$$\tilde{\mathcal{E}}_t = \sum_{i=1}^{K} \left[e_{i,t} \cdot \left(\frac{S_{i,t-1}}{Y_{t-1}} (g_{i,t} - \bar{g}_{I,t}) \right) \right] \tag{6}$$

where $e_{i,t}$ is the sector-specific eigenvector centrality value in time t.

The eigenvector centrality is based on the premise that the importance of a node is proportional to the sum of the influence of its neighbours, that is $\lambda x_i = A_{1i}x_1 + A_{2i}x_2 + ... + A_{ni}x_n$. After a small rearrangement, this can be written in matrix form as the eigenvector equation $\lambda x = Ax$. This means that the centrality vector x is an eigenvector of A, with λ as the corresponding eigenvalue. Since there are commonly multiple eigenvectors, the centrality is taken to be the one corresponding to the largest value of λ.[3] Similar are also Katz centrality and Bonacich centrality. A combination of the latter two has been

[3]This is a result of the additional requirement that all entries in the eigenvector are non-negative and non-zero which implies by the Perron-Frobenius theorem that we have to take the eigenvector corresponding to the largest eigenvalue.

employed by Carvalho (2014) to identify central sectors in his dataset.

How is this specification reconcilable with Hulten's theorem which states that only the sales-to-GDP vector is relevant in weighting the idiosyncratic productivity shocks (Hulten, 1978)? Indeed, as Baqaee and Farhi (2017) show, the first-order approximation of Hulten's theorem for which the change in aggregate GDP levels corresponds to the sum of productivity shocks weighted by the sales-to-output vector only holds exactly for the knife-edge case of a Cobb-Douglas productivity function with constant returns to scale. In particular, the higher-order terms are non-zero in the presence of complementarities in supply chains. Intuitively, these terms magnify negative shocks and attenuate positive shocks. Recent evidence indeed points to the relevdance of complementarities in supply chains and thus, cast doubt on such a Constant Returns to Scale (CRS) Cobb-Douglas specification (Atalay, 2017; Boehm et al., 2018). Furthermore, indirect evidence for the relevance of these higher-order terms can be deduced from an important corollary of the results by Baqaee and Farhi (2017). They show that complementarities in supply-chains are necessary to generate fat-tailed output distributions. This non-Gaussianity of GDP growth distributions is a well established empirical regularity (Ascari et al., 2015; Fagiolo et al., 2007, 2008). This leads us to conclude that indeed network linkages and in particular complementarities in these supply chains need to be considered.

Hence, in the presence of linkages and complementarities, the appropriate additional weight which needs to be placed on each idiosyncratic shock is the eigenvector centrality of each industry in the network. Intuitively, a large eigenvector centrality is assigned to a industry that is relatively strongly connected to industries that are themselves relatively strongly connected. Thus, if complementarities and linkages exist and if thus the gross output of a given industry is positively correlated with the output of its neighbours, a given shock in percentage terms is amplified according to this eigenvector centrality measure. Acemoglu et al. (2015) develop a more rigorous presentation of this argument and show, in particular, that if the interactions between nodes in a given network is linear (as complementarities in supply chains would imply), the macro aggregate should move in line with the shocks to the individual nodes weighted by their eigenvector centrality.

To assess the explanatory power (R^2) of all four specifications (3) to (6), we run a battery of in total eight regressions, equations I to IV, to assess the extent to which idiosyncratic shocks account for aggregate fluctuations

$$I: \ (a)\ g_t^Y = \beta_0 + \beta_1 \mathcal{G}_t + \epsilon_t \qquad (b)\ g_t^Y = \beta_0 + \beta_1 \mathcal{G}_t + \beta_2 \mathcal{G}_{t-1} + \epsilon_t \tag{7}$$

$$II: \ (a)\ g_t^Y = \beta_0 + \beta_1 \tilde{\mathcal{G}}_t + \epsilon_t \qquad (b)\ g_t^Y = \beta_0 + \beta_1 \tilde{\mathcal{G}}_t + \beta_2 \tilde{\mathcal{G}}_{t-1} + \epsilon_t \tag{8}$$

$$III: \ (a)\ g_t^Y = \beta_0 + \beta_1 \mathcal{E}_t + \epsilon_t \qquad (b)\ g_t^Y = \beta_0 + \beta_1 \mathcal{E}_t + \beta_2 \mathcal{E}_{t-1} + \epsilon_t \tag{9}$$

$$IV: \ (a)\ g_t^Y = \beta_0 + \beta_1 \tilde{\mathcal{E}}_t + \epsilon_t \qquad (b)\ g_t^Y = \beta_0 + \beta_1 \tilde{\mathcal{E}}_t + \beta_2 \tilde{\mathcal{E}}_{t-1} + \epsilon_t \tag{10}$$

where g_t^Y is the real GDP per capita growth rate.

2.2 Determining Productivity Shocks

As can easily be seen from equation (2), the standard deviation of GDP growth rates would under the assumption of independence between size and growth rate (Gibrat's law) scale according to

$$\sigma_{GDP} = \mu \cdot h \cdot \sigma_\pi, \tag{11}$$

where σ_{GDP} is the standard deviation of GDP growth rates, h is the square root of the sales Herfindahl, σ_π the cross-sectional standard deviation of productivity growth rates and μ some coefficient reflecting factor usage. As can be seen, the variation of GDP growth rates scales linearly with both the standard deviation of productivity growth rates and the square-rooted Herfindahl index.

However, while most of Gabaix' argument is based on the result that the firm-size distribution has to be sufficiently heavy-tailed for h being large enough in equation (11) to imply plausible values of μ, he pays considerably less attention to σ_π and the distribution implying it. This strikes us as rather unfortunate, since the regression results are heavily dependent on a winsorizing procedure for the productivity shocks potentially strongly affecting σ_π.

In particular, Gabaix winsorizes the productivity growth rates at an absolute value of 0.2, that is, all values below -0.2 and above 0.2 are replaced by -0.2 and 0.2, respectively. He justifies this procedure by referring to extraordinary events such as mergers generating outliers. However, apart from such extraordinary events, it is a well documented empirical regularity that productivity growth rates are heavy tailed. His procedure makes the productivity growth rate distribution Gaussian or even more platykurtic contradicting this well-established strand of industrial dynamics literature (Dosi et al., 2018).

More importantly, the interaction of such winsorizing with a fat-tailed initial distribution might heavily distort σ_π as the measured standard deviation for productivity growth rates. To allow for a more general investigation, consider winsorizing not at absolute values but at probabilities. Winsorizing at a probability of p would thus result in a distribution censored at the pth and $1-p$th quantile. Assuming a Laplacian initial distribution as one of the most intensively studied symmetric heavy-tailed distributions, winsorizing at a probability of merely 1% results in a huge reduction of σ_π by more than 12.5%.[4] Given any Gaussian initial distribution, winsorizing at 1% or equivalently censoring the distribution by in total $2p = 2\%$ only leads to a less than proportional decrease in σ_π by about 1.8%. This vast difference between both results highlights the amplifying interaction effect of fat tails. Given the linear scaling relationship in equation (11), the induced volatility in GDP growth would thus also be scaled down by more than 12.5% for any Laplacian initial

[4] Notice that the Laplacian distribution is a reasonable benchmark for the distributions of labour productivity, as, e.g., shown by Dosi et al. (2012) and Yu et al. (2015). In there, they use an Asymmetric Exponential Power Distribution (AEP) to show that the productivity growth rates below the mode are more leptokurtic than the Laplacian would imply and above them more platykurtic. They use, however, a differing definition of labour productivity, that is, Value Added per Employee.

distribution, but only 1.8% for any Gaussian initial distribution. It is hard to reconcile this probably hugely distorting effect for the regression procedure, give the robust empirical evidence for the heavy tailedness of productivity growth distributions with Gabaix' complementary calibration excercise. Also, the fact that Gabaix' procedure necessitates Gaussian or sub-Gaussian productivity growth distributions in obvious contradiction to the super-Gaussian distributions well established in the industrial dynamics literature as an empirical regularity casts doubt on the validity of his procedure (Castaldi and Dosi, 2009; Mundt et al., 2015; Bottazzi et al., 2017).

In contrast to Gabaix' firm level analysis, we do not need any winsorizing for our results to hold and be significant. We speculate that this is strong evidence that the industry-level of aggregation is more appropriate compared to a more fine-grained firm-level disaggregation. By construction, analyses on the firm-level ignore the various economic and administrative linkages between these firm-level entities. This is especially true considering that the boundaries between these entities are often merely resulting from legal and not economic considerations and thus, two legally separate entities might be subject to the same idiosyncratic shock such as a change in administrative policy or a strike. In our higher industry-level aggregation, such intra-industry linkages are captured by construction, as only the net aggregate effect after the propagation of the shock through the intra-industry network is reported. Also, as described in the subsection above, we explicitly take inter-industry linkages into account to capture all intra- and inter-industry linkages.

Gabaix' initial approach still leaves two important open questions. First, he bases his argument on *idiosyncratic shocks* to individual entities as opposed to aggregate shocks hitting all entities of the total economy or of an industry. In his specification, an aggregate shock is defined as the mean productivity growth rate over Q firms that are considered and consequently, an idiosyncractic shock is given by the Q-demeaned productivity growth rate for each firm. This Q, however, is somewhat arbitrarily fixed to $Q = 100$ or $Q = 1,000$. Without any justification from economic theory given, it is therefore hard to see why of all possible Q values, shocks to these 100 or $1,000$ firms should constitute aggregate shocks in comparison to, e.g., the whole sample average.

Secondly, his theorem establishing the significance of idiosyncratic shocks is based on the empirically well-established power-law distribution of firm sizes. Specifically, he shows that if the tail exponent α of a given power-law is between unity and two, aggregate effects emerge from idiosyncratic shocks to these power-law distributed entities. While he concedes that aggregate effects might also emerge from size distributions that are not power-law but sufficiently fat-tailed, he does not give boundary conditions in the same fashion for other distributions. For our data, this poses a problem, as industry-level gross outputs are clearly not power-law distributed (c.f. Appendix A.2).

Our approach to empirically calibrate Q resolves both issues in a unified fashion. One desideratum for Q is for it to encompass the minimally possible set of granular industries, as, in the spirit of granularity, a relatively small number of large entities should drive aggregate fluctuations. We define granular entities as those entities whose measured concentration is between the boundary values given by Gabaix' for a power-law size distribution. Again in line with Gabaix and as can be seen in equation (11), the square-rooted Herfindahl index h is here the correct measure of concentration. Let $h^{PL}(\alpha, N)$ be the square-rooted Herfindahl index given for a perfect power-law with a tail-exponent α and number of entities N.[5] Q_t is thus given by the minimal set of N largest industries whose $h_t^{emp}(N)$ as their square-rooted Herfindahl index is between the Herfindahl indices $h^{PL}(\alpha, N)$ for N entities distributed as a power law with a tail-exponent α of 1 or 2, respectively. We thus take Q_t to be the minimal number of industries N for which $h^{PL}(1, N) \geq h_t^{emp}(N) \geq h^{PL}(2, N)$ holds. We call this condition the *Minimum Granularity Condition* (MGC). Interestingly and conveniently, the empirical h stays in these bounds also for all other higher number of industries N. The results for the WIOT data are given in Table 1.

Year	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013
Q_t^*	219	219	214	207	205	204	204	208	224	205	219	224	221	218
Zipf	0.087	0.087	0.088	0.089	0.089	0.090	0.090	0.089	0.086	0.089	0.087	0.086	0.086	0.087
Rep. Industry	1.284	1.284	1.282	1.279	1.279	1.279	1.278	1.280	1.285	1.279	1.28	1.286	1.285	1.284

Table 1: Empirical HHI table for minimal granular sample sizes compared to the Zipfian and the representative industry benchmark.

Note: Q_t^* gives the minimal granular sample sizes as defined by the MGC, column *Zipf* compares the empirical square-rooted HHI to the implied square-rooted HHI for the same $N = Q_t^*$ and a Zipfian distribution by $\frac{h^{emp}(Q_t^*)}{h^{PL}(1, Q_t^*)}$, while column *Rep. Industry* reports the ratio of the empirical h for Q_t^* to the h for the representative industry benchmark given by $h = \frac{1}{\sqrt{Q_t^*}}$.

Column Q_t^* gives the minimal granular sample sizes for the calibration of our model. Notice that the measured concentration is more than one order of magnitude smaller than the Zipfian benchmark typically characterizing firm-size distributions (row 3), while only being about 28% above the equal weights or representative industry benchmark. This is further evidence that Gabaix is tinkering with the wrong distribution, as the productivity growth rates sustain their fat-tailedness on this arguably more appropriate level of aggregation, while the size-distribution strongly responds to this aggregation of firms to industries. Finally, we need a consistent measure of aggregate shocks or, equivalently, a Q^* for the whole sample period. We define thus a *Longitudinal Minimum Granularity Condition* (LMGC) for the minimum Q^* for which for all years, the considered Q industries are within the boundaries defined by Gabaix. This is just the maximum of all observed values, as the empirical h^{emp} stays in the boundaries defined by Gabaix for all number of industries N larger than Q_t for all considered periods. Q^* is thus $Q^* = 224$ which we will use in our regression model.

[5] The derivation as well as a definition for $h^{PL}(\alpha, N)$ is given in Appendix A.3.

3 Data

In order to perform our empirical analysis, we use the World Input-Output Database (WIOD). More specifically, the Input-Output Tables for output and the Social Economic Accounts for employment data. Macroeconomic data (GDP, GDP per capita growth rate) are taken from the World Bank's Development Indicators database (World Bank, 2018). Central to the WIOD is that it consists of annual time series of WIOTs. These tables can be regarded as a set of national input-output tables that are connected with each other by bilateral international trade flows. The second release from 2016 which is used for this paper covers all 28 EU countries and 15 other major countries in the world for the period from 2000 to 2014 (World Input-Output Database, 2017). Together, the countries included cover more than 85% of world GDP in 2008 (at current exchange rates). Moreover, a model for the Rest-Of-World (ROW) has been estimated to cover other countries not explicitly stated.[6]

However, we focus on the EU28 as it seems to be an ideal candidate for the exploration of our model, as the European Union represents a large and strongly interconnected economic zone that was, considering the recent crises, characterized by large variation in GDP growth rates. Moreover, the European Single Market seeks to guarantee the free movement of goods, services, capital and labour (European Commission, 2018). This is a major advantage vis-à-vis a more unassociated global economy.

The WIOTs have an industry-by-industry format reflecting the economic linkages across industries. They provide details for 56 industries mostly at the two-digit ISIC rev. 4 level or groups thereof. These tables can be regarded as a set of national input-output tables that are connected with each other by bilateral international trade flows. Thereby, the columns in the WIOT contain information on production processes, while in rows the distribution of the output of industries over user categories is indicated. Products can be used as intermediates by other sectors or as final products by households, governments or firms. A critical accounting identity of the WIOTs is that gross output of each industry (last element of each column) is equal to the sum of all users of the output from that industry (last element of each row). Further, imports are broken down according to the country and industry origin in a WIOT (Timmer et al., 2015).

The WIOTs can be mainly divided into four parts: a supply-and-use matrix, a final demand matrix as well as a value added and total/gross output vector. For our purpose, we are interested in the annual volume of gross output and the corresponding number of employees which we are obtaining from the Socio-Economic Accounts also available from the WIOD, as well as the activity carried out by each industry, captured in the 56 industry codes.

[6] The specific procedure is described in Timmer et al. (2015)

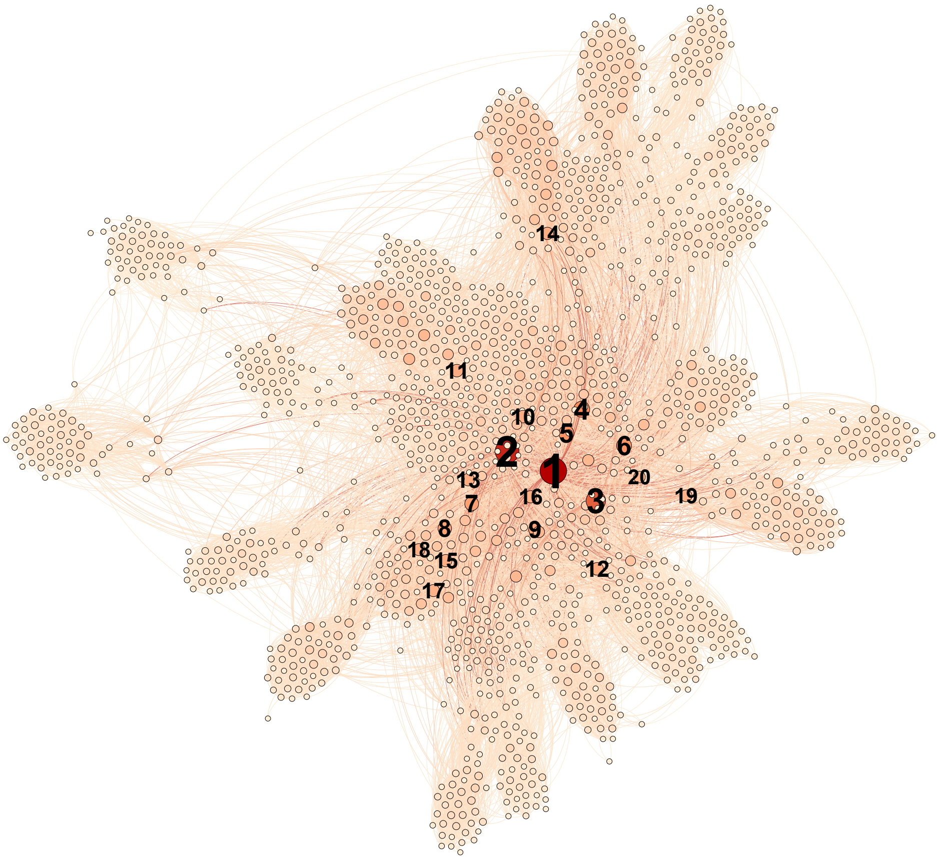

Figure 1: The European Production Network Corresponding to the WIOT Input-Output Data in 2014.

Thus, we obtain a 1568×1568 matrix. This input-output matrix can be translated into the network visualized above. Each nonzero (i, j) entry is a directed edge within the network. We only use the binary information contained in the input-output data, considering a link to be present if the associated input transaction is above 1% of a sector's total input purchases used for total output. This threshold has also been applied by Carvalho (2014). Even though 1% might be seen considerable, roughly 80% of trade information between industries was able to be preserved in the median.

Figure 1 provides the production network representation corresponding to the WIOT input-output data for the EU28 in 2014. A first-order characterization of the network is its low density: there are only 25,825 nonzero edges out of a possible 1509^2, causing a network density of 0.01. Also, the average degree of the network is with 17, relative low. Thus, on average, an industry is connected to 17 other industries. These network-averages could suggest that the network is relatively weakly interconnected and thus, network linkages play a negligible role for the transmission of idiosyncratic shocks. The size of each node is proportional to its respective eigenvector centrality for which the largest 20 are labeled. These are, among others, the German automotive (1) and machine manufacturing (2) industry, the German and French construction sector (3,4), the Italian machine manufacturing industry (8), the British health-care sector (11) or the Italian wholesale trade industry (18). As is easily observable, the EU28 production network is characterized by a large degree of heterogeneity in terms of their eigenvector centrality. By the discussion in the subsection above, their ability to transmit shocks is thus also highly heterogeneous.

Following the arguments of Timmer et al. (2015), the WIOD has some significant advantages. First and foremost, the WIOTs have been designed to trace developments over time. Hence, they allow the benchmarking of output, value added, trade and consumption from national accounts statistics as time series. Secondly, the WIOD is based on official and publicly available data and constructed within the framework of the International System of National Accounts, adhering to its concepts and accounting identities. By checking the data quality in preparation of the data analysis, it was possible to confirm this statement. For example, the deviation from national value added statistics of several European countries was less than 0.5% of the WIOT value added at basic prices for 2014.[7] Thirdly, the WIOTs have been constructed on the basis of sets of national supply and use tables (SUTs). These bear significant advantages. For example, the possibility to combine them with trade statistics that are product-based and employment statistics that are industry based. This allows one to take the multi-product nature of firms into account (Timmer et al., 2015).

Yet, according to Timmer et al. (2015), some limitations need to be borne in mind. The first limitation is technology heterogeneity. This is a limiting assumption in any input–output table, which assumes homogeneity within industries. Thereby, a column in an input-output table only provides the average production structure across all firms in a particular industry. However, these structures might substantially differ for various types of businesses. A second limitation is the ROW model which is necessary to have a complete description of all flows in the global economy. Thirdly, the trade in service and intangibles provides a bottleneck because services trade data has not been gathered at the same level of detail and accuracy as goods trade data (Timmer et al., 2015).

Despite the limitations mentioned above, the data at hand provides a substantial possibility to analyse the European economy and its production networks. The neglible deviations between the used datasources show the consistency of our database. Also, the high degree of observed heterogeneity in terms of the eigenvector centrality reveals the WIOD to be an ideal testing ground for our outlined hypotheses and the relevance of second-order effects contrary to Hulten's theorem. The WIOTs make it possible to investigate economic networks on a global scale, while examining aggregate fluctuations as well as input-output linkages on a national or sectoral level over time, too. Therefore, in the upcoming section, we present our results.

[7] Material available upon request.

4 Results

4.1 Cross-Sectional Calibration

By computing the explanatory power of the granular residual of a given number of large sectors, we check whether the European economy is granular or not following the specification of equation 3. Table 2 shows the result of the estimated coefficients β_0, and β_1 respectively, of the OLS regression in specification (7).

	I	
	(a)	(b)
$\mathcal{G}_t$	0.957	0.921*
	(0.576)	(0.470)
$\mathcal{G}_{t-1}$		1.549**
		(0.528)
Intercept	0.009*	0.011**
	(0.005)	(0.004)
T	14	13
R^2	0.187	0.555
Adjusted R^2	0.119	0.466
F Statistic	2.761	6.240**

Table 2: g_t^Y (GDP per capita growth rate) is regressed on the granular residual $\mathcal{G}_t$ without a and with one lag, when $Q = K = 224$. Standard errors are given in parentheses. Significance at the *** 1%, ** 5%, * 10% level.

In the first specification, where we subtract the median cross-sectional growth rates from the sector specific growth rate, one might anticipate that our results indicate that the European economy is granular. At least when running the regression with no lag, the results are significant and the granular residual accounts for around 20% of variations of GDP growth according to the coefficient of determination.

However, the left panel of Figure 2 shows that our specification is still less than satisfactory. As, by Hulten's theorem, aggregate fluctuations should result from the (Domar) weighted idiosyncratic shocks to individual industries, the empirical R^2 plot should increase (almost) monotonically. This is the case as a larger number of K largest industries considered should explain a larger proportion of the composite shock emerging from these single idiosyncratic shocks. Given the fact that the industry size distribution exhibits a large degree of heterogeneity and that thus the idiosyncratic shocks large industries are associated by a disproportionately large (Domar) weight, the R^2 plot should also show a huge jump for small values of K and settle down on a plateau rather quickly (for now relatively small industries).

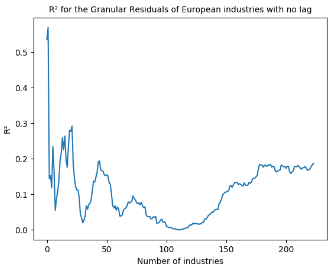

(a) R^2 - No Lag

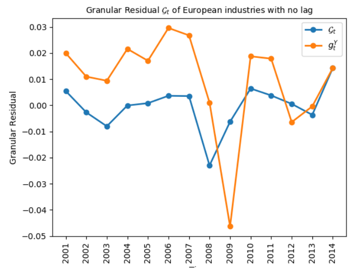

(b) Granular Series - No Lag

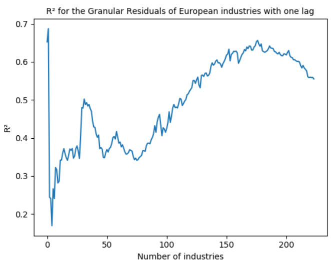

(c) R^2 - One Lag

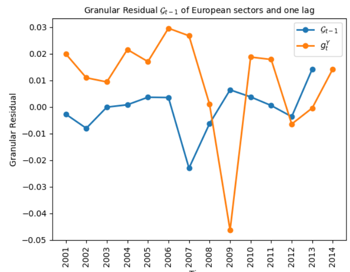

(d) Granular Series - One Lag

Figure 2: R^2 and Granular Series for Europe with no and one lag.

The empirical plot, however, shows no such jump for small K values and sizeable oscillations in the R^2 plot. The model therefore seems to be misspecified both in the definition of idiosyncratic shocks and also in its negligence of effects resulting from the network structure.[8] Figure 2 (b) and (d) show the Granular time series at $K = 224$. It can be seen that the granular series $\mathcal{G}_t$ and $\mathcal{G}_{t-1}$ do not follow the GDP per capita growth rate g_t^Y well. Hence, it is no surprise that the results following out of specification (3) cannot convince.

4.2 Sector-Specific Calibration

By introducing the term $g_{i,t} - \bar{g}_{I,t}$ in the second specification, we may control better for industry-wide (sector-specific) disturbances. Computing the explanatory power of the granular residual of a given number of large sectors, we check whether the European economy is granular or not following the specification of equation (4). Table 3 shows the result of the estimated coefficients β_0, and β_1 respectively, of the OLS regression in equation 8.

[8]This also implies that the results reported by Gabaix might be misleading in another respect. Even though an economy might seem highly granular for a given K = 100, this, as is easily observable in Figure 2 (c), might not hold for higher values of K, where the expected explanatory power should, however, be increasing.

	II	
	(a)	(b)
$\tilde{\mathcal{G}}_t$	1.1186***	1.1046***
	(0.2474)	(0.2308)
$\tilde{\mathcal{G}}_{t-1}$		0.3295
		(0.2370)
Intercept	0.0082**	0.0068*
	(0.0032)	(0.0030)
T	14	13
R^2	0.630	0.738
Adjusted R^2	0.599	0.685
F Statistic	20.43***	14.05***

Table 3: g_t^Y (GDP per capita growth rate) is regressed on the granular residual $\tilde{\mathcal{G}}_t$ with a and without lag, when $Q = K = 224$. Standard errors are given in parentheses. Significance at the *** 1%, ** 5%, * 10% level.

In the second specification where we subtract the median sector-specific growth rates from the sector i's growth rate, our results indicate that the European economy is granular since the granular residual accounts for more than 60% of variations of GDP growth. Even though the difference is marginal, one can see in Figure 3 (a) that the explanatory power (R^2) stabilizes quicker and then gradually increases as K increases towards 224.

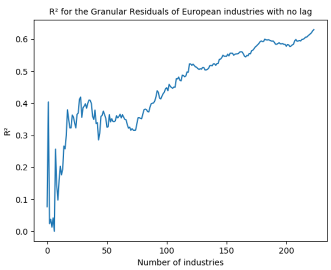

(a) R^2 - No Lag

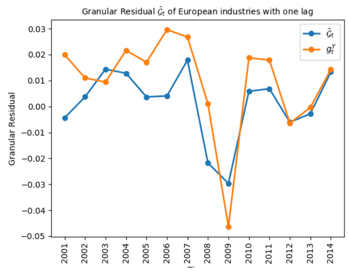

(b) Granular Series - No Lag

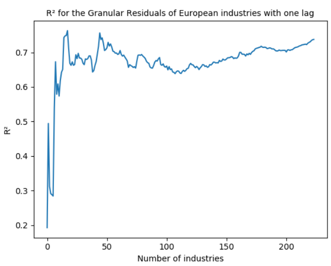

(c) R^2 - One Lag

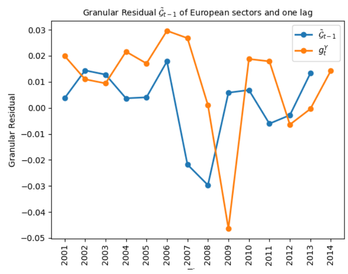

(d) Granular Series - One Lag

Figure 3: R^2 and Granular Series for Europe with no and one lag.

When introducing one lag this effect sets in even quicker, as can be seen in Figure3 (c). However, when looking at the Granular residual series there is just a marginally visible effect in a different movement of the Granular series. Also it seems that the granular residual series just follows the GDP per capita growth series only from 2007 onwards.

Nevertheless, this is much more in line with the desired plot for a correct specification outlined above. It thus seems that the median sector-specific growth rate is a better proxy for an aggregate shock and thus our specification for the idiosyncratic shock is improved compared to the whole-sample median.

4.3 Network-Specific Calibration

By introducing the eigenvector centrality in the third specification we are trying to capture the hypothesis that not only the size of a firm or sector needs to be taken into account when thinking about aggregate fluctuations, but also its position within the network. Hence, by computing the explanatory power of the granular residual of a given number of large sectors again, we check whether the European economy is granular or not following the specification of equation 5. Table 4 shows again the result of the estimated coefficients β_0, and β_1 respectively, of the OLS regression in equation 9.

	III	
	(a)	(b)
$\mathcal{E}_t$	0.4810***	0.4725***
	(0.0927)	(0.0890)
$\mathcal{E}_{t-1}$		0.1766*
		(0.0907)
Intercept	0.0011***	0.0012***
	(0.0002)	(0.0002)
T	14	13
R^2	0.691	0.772
Adjusted R^2	0.666	0.726
F Statistic	26.88***	16.89***

Table 4: g_t^Y (GDP per capita growth rate) is regressed on the granular residual $\mathcal{E}_t$ without a and with lag, when $Q = K = 224$. Standard errors are given in parentheses. Significance at the *** 1%, ** 5%, * 10% level.

In the third specification, where we are weighting the granular residuals by their industry-specific eigenvector centrality. As can be seen, the results indicate that the European economy is granular and significantly dependent on the underlying production network structure, as the granular residual accounts for almost 70% of variations of GDP growth without lag and more than 75% with lag. Moreover, the left panel of Figure 4 shows that our specification is much more satisfactory than the size-based specifications in the previous subsections. As predicted by Hulten's theorem, the empirical R^2 plot increases (almost) monotonically, with some oscillations in the beginning and settles on a plateau when K approaches its limit.

Furthermore, when referring to Figure 4 (b) and (d) it can be easily seen that the granular residual time series follows the GDP per capita growth series quite well, comoving with growth and shrinking periods of the European GDP per capita growth rates.

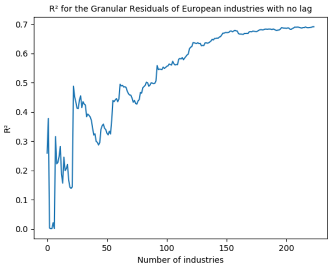

(a) R^2 - No Lag

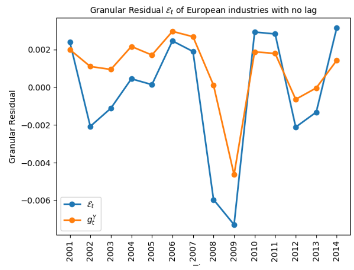

(b) Granular Series - No Lag

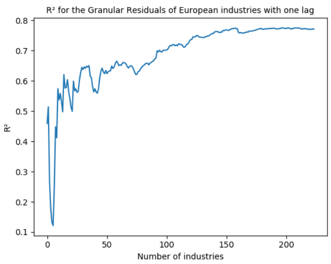

(c) R^2 - One Lag

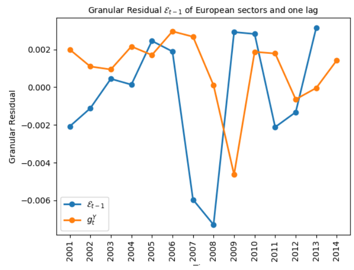

(d) Granular Series - One Lag

Figure 4: R^2 and Granular Series for Europe with no and one lag.

Note: g_t^Y (GDP per capita growth rate) has been divided by ten to level multiplication effects in the calculation.

4.4 Network-and-Sector-Specific Calibration

By introducing the eigenvector centrality and controlling for sector-specific movements in the last specification, we are trying to capture the hypothesis that not only the size of a firm or sector needs to be taken into account when thinking about aggregate fluctuations, but also its position within the network. Moreover, we are controlling for sector-specific disturbances.

	IV	
	(a)	(b)
$\tilde{\mathcal{E}}_t$	0.3975***	0.4120***
	(0.0775)	(0.0687)
$\tilde{\mathcal{E}}_{t-1}$		0.0963
		(0.0692)
Intercept	0.0009**	0.0007**
	(0.0002)	(0.0030)
T	14	13
R^2	0.687	0.794
Adjusted R^2	0.660	0.753
F Statistic	26.29***	19.25***

Table 5: g_t^Y (GDP per capita growth rate) is regressed on the granular residual $\tilde{\mathcal{E}}_t$ with one lag, when $Q = K = 224$. Standard errors are given in parentheses. Significance at the *** 1%, ** 5%, * 10% level.

Therefore, we are following the specification of equation 6. Table 5 shows again the result of the estimated coefficients β_0, and β_1 respectively, of the OLS regression in equation (10).

In the last specification, our results indicate that the European economy is granular since the granular residual accounts for almost 80% of variations of GDP growth. As both the resulting R^2 for $K = Q = 224$ is very large and the R^2 plot is surprisingly close to the ideal shape outlined above, we conclude that this is the correct specification. This can be especially seen in Figure 5 (a) and (c) where the explanatory power (R^2) stabilizes quick and then builds a plateau from roughly the 50s industries towards the end.

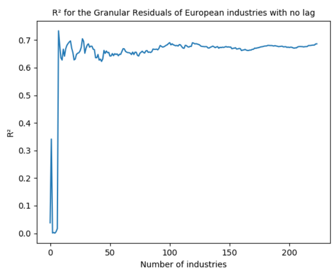

(a) R^2 - No Lag

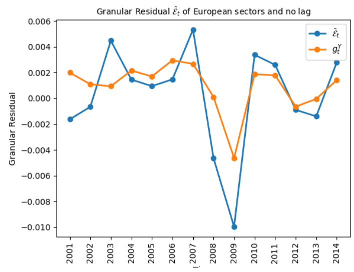

(b) Granular Series - No Lag

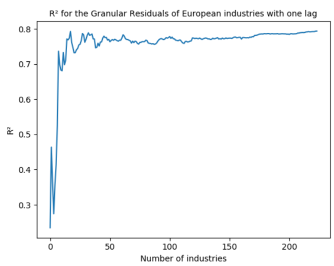

(c) R^2 - One Lag

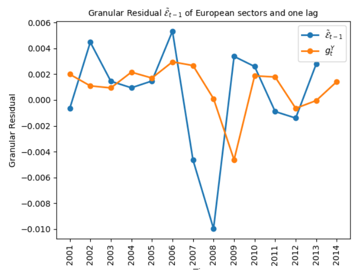

(d) Granular Series - One Lag

Figure 5: R^2 and Granular Series for Europe with no and one lag.
Note: g_t^Y (GDP per capita growth rate) has been divided by ten to level multiplication effects in the calculation.

Moreover, it seems that less than 50 industries shape the aggregate fluctuations of the European economy. In fact, if we sort out the top 50 sectors per year and remove all duplicates we identify 62 country specific industries. Further decomposed these are 22 industries from Germany, Great Britain, France, Italy, Spain and the Netherlands which can be held accountable for the result displayed in Figure 5 (a) and (c). Intriguingly, these are also the six largest economies of the European Union.

5 Discussion and Conclusion

This study has proposed a simple model unifying the granular and network origins of aggregate fluctuations following the seminal work of Gabaix (2011) and Acemoğlu et al. (2012). Moreover, we examined the explanatory power of industry-level idiosyncratic shocks to fluctuations in the aggregate of the European economy.

The analysis shows that shocks to large industries are of relevance, explaining roughly 20% of output fluctuations in the EU. This explanatory power increases considerably when controlling for sector-specific shocks. More importantly, when including network-based metrics, the explanatory power increases to almost 80% of explained variation of GDP growth rates. Furthermore, there are 62 country specific industries that shape the aggregate fluctuations of the European economy. We could thus show that the European economy is indeed granular in industries and that combining the granular and network approaches significantly improves the explanatory power for the origins of aggregate fluctuations. Hence, aggregate fluctuations in the European economy seem not only to be dependent on the size of an industry but also on its position within the European production network. Unlike previous studies, note that these results could be achieved by just using raw data on productivity growth rates, and thus without questionable procedures at odds with the stylized facts from the industrial dynamics literature on productivity growth rate distributions.

Moreover, our findings might have both implications for macroeconomic model-building as well as issues pertaining policy-making. First, as we could find evidence that complementarities in supply chains lead to sizeable second-order effects resulting from input-output linkages, a further investigation of the mechanisms that give rise to specific network structures might prove valuable. For example, it would be interesting to investigate how the transmission of shocks evolves while propagating through the (European) economic network. Second, we have taken the heterogeneity in network- and size-based metrics at a meso-level as our starting point to explain macroeconomic fluctuations. Explaining these phenomena at the meso-level themselves from a more fine-grained firm- or even plant-level processes could thus be a step towards a truly microfounded macroeconomic theory. For this, a deeper understanding of propagation mechanisms of idiosyncratic shocks and their effects on factor reallocation will be indispensable, for which the Input-Output Tables and Social Economic Accounts from the WIOD could also prove useful (Jones, 2011; Timmer et al., 2015). Third, as we identified a mere 22 industries from 6 countries to be the main drivers of aggregate fluctuations in Europe, these could be use as more fine-grained targets for industrial policy to boost aggregate performance more efficiently while protecting other industries getting hit by a combination of size and network effects.

In sum, we anticipate our paper to be a valuable contribution to the current discussion in the economics literature and think that combining the granular and network origins approach may lead to a better understanding of aggregate fluctuations.

References

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608.

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

Acemoğlu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2010). Cascades in networks and aggregate volatility. *NBER Working Paper Series*, 16516.

Ascari, G., Fagiolo, G., and Roventini, A. (2015). Fat-tail distributions and business-cycle models. *Macroeconomic Dynamics*, 19(2):465–476.

Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics*, 9(4):254–80.

Baqaee, D. R. and Farhi, E. (2017). The macroeconomic impact of microeconomic shocks: beyond hulten's theorem. Technical report, National Bureau of Economic Research.

Boehm, C., Flaaen, A., and Pandalai-Nayar, N. (2018). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake. *The Review of Economics and Statistics*, forthcoming.

Bottazzi, G., Li, L., and Secchi, A. (2017). Aggregate fluctuations and the distribution of firm growth rates. *LEM Working Paper Series*, 2017/24.

Carvalho, V. M. (2010). Aggregate fluctuations and the network structure of intersectoral trade (phd thesis). *Retrieved from github.io*.

Carvalho, V. M. (2014). From micro to macro via production networks. *The Journal of Economic Perspectives*, 28(4):23–47.

Castaldi, C. and Dosi, G. (2009). The patterns of output growth of firms and countries: Scale invariances and scale specificities. *Empirical Economics*, 37(3):475–495.

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

Di Giovanni, J., Levchenko, A. A., and Méjean, I. (2014). Firms, destinations, and aggregate fluctuations. *Econometrica*, 82(4):1303–1340.

Domar, E. D. (1961). On the measurement of technological change. *The Economic Journal*, 71(284):709–729.

Dosi, G., Grazzi, M., Tomasi, C., and Zeli, A. (2012). Turbulence underneath the big calm? The micro-evidence behind Italian productivity dynamics. *Small Business Economics*, 39(4):1043–1067.

Dosi, G., Napoletano, M., Roventini, A., Treibich, T., et al. (2018). Debunking the granular origins of aggregate fluctuations: From real business cycles back to Keynes. *LEM Working Paper Series*, 2018/35.

European Commission (2018). The european single market. `http://ec.europa.eu/growth/single-market/`. [Online; accessed 2018-08-30].

Fagiolo, G., Napoletano, M., and Roventini, A. (2007). How do output growth-rate distributions look like? some cross-country, time-series evidence. *The European Physical Journal B*, 57(2):205–211.

Fagiolo, G., Reyes, J., and Schiavo, S. (2008). On the topological properties of the world trade web: A weighted network analysis. *Physica A: Statistical Mechanics and its Applications*, 387(15):3868–3873.

Foerster, A. T., Sarte, P.-D. G., and Watson, M. W. (2011). Sectoral versus aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy*, 119(1):1–38.

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

Gabaix, X. (2016). Power laws in economics: An introduction. *Journal of Economic Perspectives*, 30(1):185–206.

Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, 45(3):511–518.

Jones, C. I. (2011). Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics*, 3(2):1–28.

Leontief, W. W. (1936). Quantitative input and output relations in the economic systems of the united states. *The Review of Economic Statistics*, pages 105–125.

Long Jr, J. B. and Plosser, C. I. (1983). Real business cycles. *Journal of Political Economy*, 91(1):39–69.

Lucas, R. E. (1977). Understanding business cycles. In *Carnegie-Rochester conference series on public policy*, volume 5, pages 7–29. Elsevier.

Mundt, P., Alfarano, S., and Milaković, M. (2015). Gibrat's law redux: think profitability instead of growth. *Industrial and Corporate Change*, 25(4):549–571.

Stella, A. (2015). Firm dynamics and the origins of aggregate fluctuations. *Journal of Economic Dynamics and Control*, 55:71–88.

Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., and Vries, G. J. (2015). An illustrated user guide to the world input–output database: the case of global automotive production. *Review of International Economics*, 23(3):575–605.

World Bank (2018). The world bank data. https://data.worldbank.org/indicator/NY.GDP.PCAP.KD.ZG?locations=SE-FI-EU. [Online; accessed 2018-07-04].

World Input-Output Database (2017). Wiod home. http://www.wiod.org/home. [Online; accessed 2017-04-13].

Yu, X., Dosi, G., Lei, J., and Nuvolari, A. (2015). Institutional change and productivity growth in china's manufacturing: the microeconomics of knowledge accumulation and "creative restructuring". *Industrial and Corporate Change*, 24(3):565–602.

A Appendix

A.1 Measuring Aggregate Fluctuations

Assume a Subbotin initial distribution with a probability density function given by

$$f(x, \kappa, \sigma, \mu) = \frac{1}{2\sigma\kappa^{1/\kappa}\Gamma(1 + 1/\kappa)} \exp\left(-\frac{1}{\kappa}\left|\frac{x-\mu}{\sigma}\right|^{\kappa}\right).$$

with support over the real line and where $\kappa, \sigma \in \mathbb{R}^+, \mu \in \mathbb{R}$ and $\Gamma(\cdot)$ denotes the Gamma function. As can be readily verified, the Subbotin distribution includes the Gaussian for $\kappa = 2$ and the Laplacian for $\kappa = 1$ as special cases. The standard deviation of a p-winsorized distribution is equivalent to the standard deviation of a distribution censored at the pth and $1 - p$th quantile, that is, a distribution where the cumulative weight from outside the pth and $1 - p$th quantile is placed at the ends of the censoring interval. Thus, the impact of winsorizing on the standard deviation of a Laplacian initial distribution can be deduced from the standard deviation of a Subbotin density with $\kappa = 1$ censored at the pth and $1 - p$th quantiles. The uncensored Laplacian is the Subbotin density with $\kappa = 1$, that is,

$$f(x, \mu, \sigma) = \frac{1}{2\sigma} e^{\left|\frac{x-\mu}{\sigma}\right|},$$

where $\sigma \in \mathbb{R}^+$ and $m \in \mathbb{R}$ and with support over the real line. Its pth quantile is given by $Q_L(p) = \mu + \sigma \ln(2p)$ and its $1 - pth$ quantile by $Q_L(1-p) = \mu - \sigma \ln(2p)$ for $p \in (0, 0.5)$.

The standard deviation of this distribution censored at the pth and $1 - p$th quantile is thus given by

$$SD_L(\mu, \sigma, p) = \sqrt{\int_{\mu+\sigma\ln(2p)}^{\mu-\sigma\ln(2p)} (x-\mu)^2 dx + 2 \cdot (\mu + \sigma \ln(2p))^2 \cdot p},$$

as, by symmetry, $Q_L(p) - \mu = -Q_L(1-p) + \mu$. Integrating yields

$$SD_L(\sigma, p) = \frac{\sigma}{\sqrt{2}} \sqrt{2p \cdot (-2 - \ln^2(p) - \ln^2(2) + 2\ln(p) - \ln(4)\ln(p) + \ln(4)) - \Gamma(3, -\log(2p)) + 4},$$

now independent of μ. Define finally $\omega_L(p) = 1 - \frac{SD_L(\sigma,p)}{SD_L(\sigma)}$ as a measure of bias for the standard deviation of a winsorized Laplacian distribution compared to the uncensored initial distribution with $SD_L(\sigma) = \sqrt{2} \cdot \sigma$ as

$$\omega_L(p) = \sqrt{\frac{p\left(2 + \ln^2(2) - \ln(4)\right) + p\ln(p)(\ln(p) - 2 + \ln(4)) - 1}{2p - 1}},$$

that is independent of σ. Thus, we can analytically determine the decrease of the standard deviation for p-winsorizing relative to the initial standard deviation for any given initial Laplacian distribution and the whole permitted range of $p \in (0, 0.5)$. In particular, $\omega_L(0.01) = 0.125929$. This implies than p-winsorizing a Laplacian initial distribution by a cumulative censored probability mass of $2p = 0.02$ or 2% leads to a overproportionate (and large) decrease of the standard deviation by about 12.6%.

To examine the effects of winsorizing on a Gaussian distribution, the same method as for the Laplacian can be applied. Unfortunately, however, the quantile values for the Subbotin density with $\kappa = 2$, however, depend on the inverse generalized Gamma function $\Gamma^{-1}(\cdot)$ that is not analytically tractable.

As a way round, a generic quantile function is considered. Assume without loss of generality that $\mu = 0$, since the standard deviation does not depend on the location of the distribution in question. The quantile function $Q(\kappa, \sigma, p)$ is now defined as $Q(\sigma, \kappa, p) = -\sigma \left(\kappa \Gamma^{-1}\left(\frac{1}{\kappa}, 2p\right)\right)^{\frac{1}{\kappa}}$,where $\Gamma^{-1}(\cdot)$ is the inverse generalized Gamma function.

By $\mu = 0$ and by the symmetry of the distribution,

$$Q(\sigma, \kappa, p) = -Q(\sigma, \kappa, 1-p).$$

For these quantile functions, the standard deviation of the censored Subbotin density $SD(\kappa, \sigma, p)$ at p and $1-p$ is thus given by

$$SD(\kappa, \sigma, p) = \sqrt{\int_{-Q(\sigma,\kappa,p)}^{Q(\sigma,\kappa,p)} x^2 f(x, \kappa, \sigma)\, dx + 2Q(\sigma, \kappa, p)^2 \cdot F(Q(\sigma, \kappa, p))}\,,$$

where $F(\kappa, \sigma)$ is the cumulative distribution function of the (uncensored) Subbotin density.

To simplify notation, let $\Gamma^{-1}\left(\frac{1}{2}, 2p\right) = G(p)$. Setting $\kappa = 2$ for a Gaussian density, subsituting the actual function for the generic quantile function and integrating yields $SD_G(\sigma, p)$ for the standard deviation of a Gaussian distribution censored at at the pth and $1-p$th quantiles as a function of p and σ by

$$SD_G(\sigma, p) = \sigma \cdot \sqrt{\operatorname{erf}\left(\sqrt{G(p)}\right) - \frac{2\exp(-\,G(p))\sqrt{G(p)}}{\sqrt{\pi}} + 2G(p)\,\Gamma\left(\frac{1}{2}, G(p)\right)},$$

where $\operatorname{erf}(\cdot)$ is the error function.

Define finally $\omega_G(p) = 1 - \frac{SD_G(\sigma,p)}{SD_G(\sigma)}$ as the ratio of standard deviation of the p-winsorized Gaussian to the standard deviation of the uncensored Gaussian distribution given by $SD_G(\sigma) = \sigma$ as

$$\omega_G(p) = 1 - \sqrt{\operatorname{erf}\left(\sqrt{G(p)}\right) - \frac{2\exp(-\,G(p))\sqrt{G(p)}}{\sqrt{\pi}} + 2G(p)\,\Gamma\left(\frac{1}{2}, G(p)\right)}.$$

While $\omega_G(p)$ is still dependent on $G(p)$ and thus, on the inverse Gamma function that is not analytically tractable, the function is only dependent on p as its sole argument.

Thus, p can be easily analyzed numerically for its whole domain $p \in (0, 0.5)$. In particular, $\omega_G(0.01) \approx 0.018046$. Thus, censoring a Gaussian distribution by a cumulative probability mass $2p = 0.02$ or 2% leads to a less than proportional decrease of the standard deviation by only about 1.8%.

Figure 6 plots both functions $\omega_G(p)$ and $\omega_L(p)$ in their domain $p \in (0, 0.5)$ as well as the cumulative censored probability mass resulting from p-winsorizing, that is, $2p$. Both functions $\omega_G(p)$ and $\omega_L(p)$ are expectedly bounded between 0 and unity, as with a winsorizing probability of $p = 0$, the winsorized distribution corresponds to the uncensored initial distribution with the same standard deviation and therefore $\lim_{p\to 0} \omega_G(p) = \lim_{p\to 0} \omega_L(p) = 0$. For p approaching 0.5, the distribution approaches the degenerate Dirac-Delta case, where the standard deviation and all higher-order moments are 0. Both functions thus approach unity for p going to 0.5.

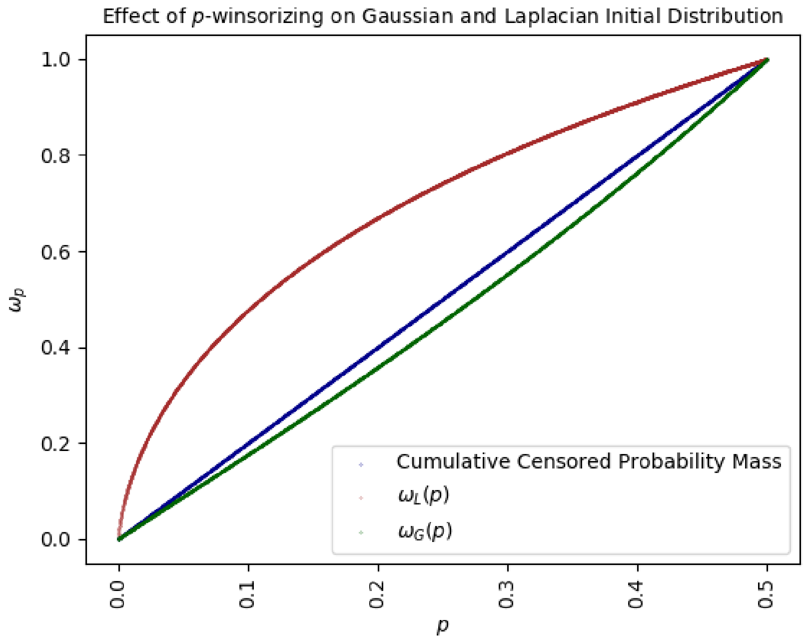

Figure 6: *Relative decrease ω_p of the Standard Deviation for a Gaussian and Laplacian initial distribution from p-winsorizing.*

Apart from these limit cases, however, it can easily be seen that p-winsorizing leads to a overproportionate decrease of the standard deviation for any Laplacian initial distribution that is especially large for small values of p relative to the cumulative censored probability mass. In contrast to that, the decrease for the Gaussian is always less than proportional for the whole domain $p \in (0, 0.5)$. This illustrates the possibly hugely distorting effect of winsorizing distributions with fat tails, such as the Laplacian, while the distortion would be much less problematic for Gaussian distributions.

A.2 Power Law Plots

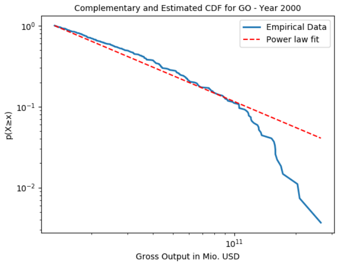

(a) 2000

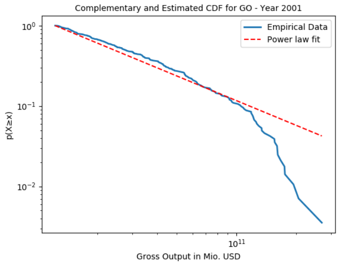

(b) 2001

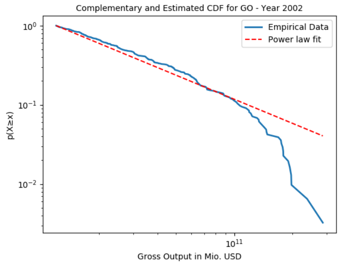

(c) 2002

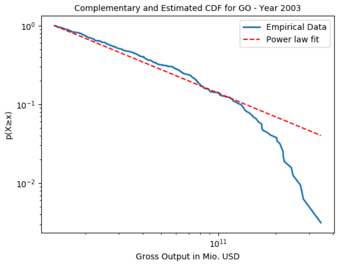

(d) 2003

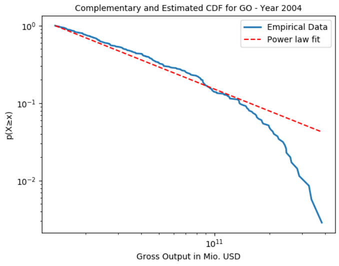

(e) 2004

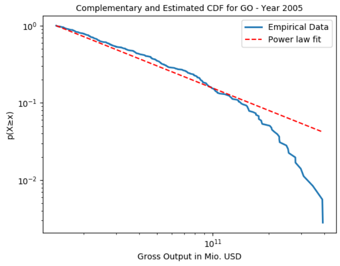

(f) 2005

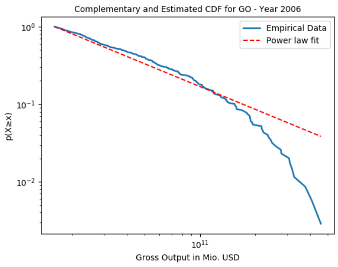

(g) 2006

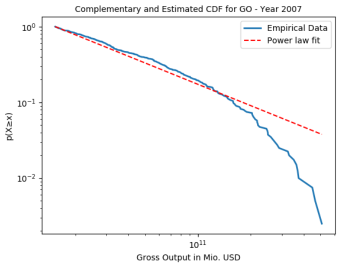

(h) 2007

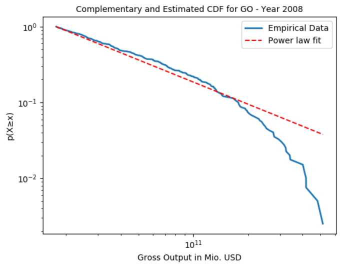

(i) 2008

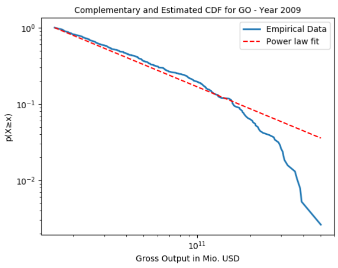

(j) 2009

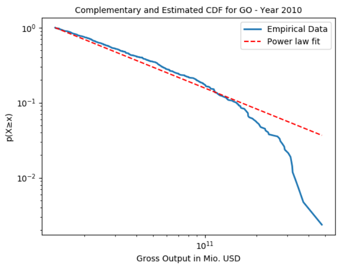

(k) 2010

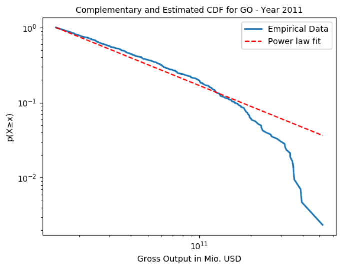

(l) 2011

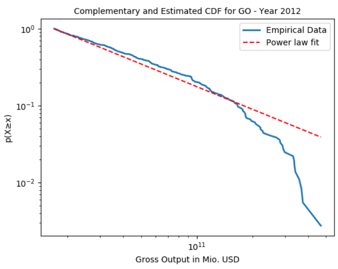

(m) 2012

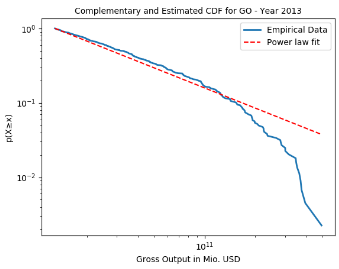

(n) 2013

Figure 7: Complementary and Estimated CDF for Gross Output of all years.

Note: The estimation is carried out following the procedure by Clauset et al. (2009).

A.3 Calibrating Aggregate Shocks

Suppose that Gross Outputs S are perfectly distributed according to a power law with minimum S_{min} and a tail exponent of α. Its values are thus given in the rank-size formulation by

$$S(k) = \frac{S_{min} \cdot N^{\frac{1}{\alpha}}}{k^{\frac{1}{\alpha}}}$$

with $k = 1, 2, ..., N$ as the respective ranks of a given S in a descending order and N as the number of values with $N \in \mathbb{N}^{+}$.

The *Herfindahl-Hirschman-Index* (HHI) is defined as the total sum of N squared market shares. Taking the square root in line with Gabaix, this yields for h as the square root of the HHI

$$h = \sqrt{\sum_{k=1}^{N} \left(\frac{S_k}{\sum_{k=1}^{N} S_k} \right)^2}. \tag{12}$$

With $S_k = S(k)$, this resulting h is given by

$$h = \sqrt{\sum_{k=1}^{N} \left(\frac{\frac{S_{min} \cdot N^{\frac{1}{\alpha}}}{k^{\frac{1}{\alpha}}}}{\sum_{k=1}^{N} \frac{S_{min} \cdot N^{\frac{1}{\alpha}}}{k^{\frac{1}{\alpha}}}} \right)^2} \tag{13}$$

$$= \sqrt{\frac{\zeta_N\left(\frac{2}{\alpha}\right)}{\zeta_N^2\left(\frac{1}{\alpha}\right)}}, \tag{14}$$

where $\zeta_N(\alpha)$ is a truncated Zeta function at N with $\zeta_N(\alpha) = \sum_{k=1}^{N} k^{-\alpha}$. Thus, h under the assumption of a discretized power law is a function of the tail exponent α and the number of Gross Outputs N, that is, $h = h(\alpha, N)$.