Aggregating Microeconomic Shocks in Production Networks

Jan Schulz[a], Omer Fetai[a], Anna Gebhard[b] and Mishael Milaković[a]

[a]Economics Institute, University of Bamberg*

[b]Mathematics Department, Otto von Guericke University of Magdeburg

September 10, 2026

Abstract

Recent theoretical advances have called the invariance of sales shares with respect to productivity differences into question, arguing that sizable higher-order macroeconomic effects may emerge when sales shares react to productivity shocks. Here we propose a parsimonious model for the adjustment of sales shares in response to productivity shocks, and operationalize it to readily quantify the impact of idiosyncratic shocks on aggregate fluctuations. Using input-output data for the European (EU28) economy, we find that our model significantly outperforms recently suggested specifications that rely either on the granularity of sales shares or the heavily skewed distribution of connections in production networks. While our results confirm earlier findings in the sense that microeconomic shocks are an important driver of macroeconomic fluctuations, we demonstrate that previous approaches substantially underestimate their relative impact because they fail to account for the interaction of granularity and network effects. Our main empirical finding is that idiosyncratic shocks to about three percent of industries already explain more than two thirds of the business cycle once we account for this interaction.

*We thank Simone Alfarano, Herbert Dawid, Sebastian Gechert, Alexander Hempfing, Tim Hagenhoff, Tomasz Makarewicz, Philipp Mundt, Christian Proaño and Frank Westerhoff, as well as participants at the Workshop on Heterogeneous Interacting Agents in London in 2019 and Ancona in 2026, the First Behavioral Macroeconomics Workshop in Bamberg in 2019, the conference of the Forum for Macroeconomics and Macroeconomic Policies in Berlin in 2019, and the International Workshop on Economic Complexity and Macroeconomic Dynamics in Valencia in 2025, for helpful comments and discussions. We gratefully acknowledge financial support from the Hans-Böckler-Foundation through grant PK045.

1 Introduction

We propose a parsimonious model for the perpetual adjustment of sales shares in response to productivity shocks between discrete periods that is readily estimated from available data. We demonstrate that previous empirical approaches substantially underestimate the aggregate impact of idiosyncratic shocks because they fail to account for the interaction of granularity and network effects. Using input-output data for the European (EU28) economy on the industry level, we find that shocks to less than three percent of industries already explain about two thirds of the business cycle when we account for this interaction.

Earlier theoretical advances have emphasized the importance of heavily skewed and fat-tailed distributions in firm size (Gabaix, 2011) or in the connectivity of production networks (Acemoğlu et al., 2012), arguing that idiosyncratic shocks on the microeconomic level can have considerable macroeconomic effects. These approaches essentially rest on the foundational theorem for efficient economies put forward by Hulten (1978), and mark an important deviation from the diversification argument of Lucas (1977), which claims that microeconomic shocks will cancel out in the aggregation process. In light of Hulten's theorem, whereby the aggregate productivity of an economy can be expressed as a sales share-weighted sum of idiosyncratic productivity levels, empirical studies have typically assumed sales shares to be constant between periods of observation. Baqaee and Farhi (2019) call this invariance of sales shares with respect to productivity differences into question and show that sizable higher-order effects can emerge when sales shares do respond to productivity shocks, and that the magnitude of such propagation effects will depend on the structure of the underlying production network. From this perspective, our contribution can also be viewed as a parsimonious operationalization of Baqaee and Farhi's theoretical considerations that explicitly accounts for the dynamical adjustment of sales shares. Their rich framework considers a very general application up to second-order macroeconomic impacts of productivity shocks. Empirical operationalizations are hindered, however, by the dimensionality of the problem that requires to specify mico-level production functions for all producers. The problem is compounded by the fact that the elasticities they employ are equilibrium objects which cannot be empirically estimated in isolation. In contrast, our approach treats all micro and macro elasticities as well as input-linkages as homogeneous, yet permits precise aggregation up to arbitrary order and is fully characterized by observables. This allows for straightforward empirical application, and also for an economic interpretation of several commonly employed estimation strategies in the literature.

The aggregation from microeconomic shocks to macroeconomic behavior is conceptually not trivial because the interactions among economic agents can amplify or attenuate initial shocks. Consequently, we highlight the interaction of heterogeneous industries at an intermediate layer of aggregation within production networks. Those industries are still subject to microeconomic forces, in our case the microeconomic elasticities of substitution, yet the properties emerging from the interaction of industries with each other are irreducible to their individual micro behavior. This irreducibility stems from the fact that relative sizes and network structures are not individual properties of the considered micro entities but can only be defined in relation to all other micro entities at once. While our model is obviously not the only conceivable approach to aggregation, its empirical performance merits consideration and avoids many of the problems that direct aggregation from the micro to the macro level brings about.[1] Intuitively, shocks to firms or industries that are both large in size and central in the production network will have a substantial macroeconomic impact whereas shocks to small and peripheral entities will not. The interaction of granularity and network effects is less intuitive when the two effects work in opposite directions. To gauge the relative importance of the two effects, we also derive the polar case of an economy that is homogeneous in its network structure yet heterogeneous in its size distribution (the purely granular case). The commonly used approximation by Riemann sums that leads to the purely granular case turns out to be a special case of our model when the network is regular. More importantly, though, the interaction of granularity and network structure substantially increases explanatory power compared to the purely granular special case for our European sample.

Our paper relates to at least three strands of previous research. Hulten's theorem and

[1] Kirman (1992) argues that traditional aggregation rules have typically imposed an unrealistically large degree of homogeneity on economic agents, for instance the Gorman (1961) polar form, resulting in the inability of such models to explain the moments of aggregate output time series (Ascari et al., 2015).

Domar (1961) aggregation of idiosyncratic productivity shocks are the basis of modern growth accounting and its empirical applications, suggesting a linear impact of productivity shocks on aggregate fluctuations via sales shares (or Domar weights). The contributions by Gabaix (2011) and Acemoğlu et al. (2012) are important because they demonstrate that even in a world of linear shock responses idiosyncratic shocks will have a significant impact on the business cycle if Domar weights are fat-tailed. Viewed from this perspective, their approaches are two sides of the same coin (Gabaix, 2016), showing that idiosyncratic shocks will generally not cancel out in the process of aggregation. Second, from a microeconomic point of view, it is only the network perspective that sheds light on comovement and shock transmission among producers of intermediate goods, and microeconomic studies indeed find a large degree of sectoral comovement (Saint-Paul, 1993; Foerster et al., 2011; Di Giovanni et al., 2014; Stella, 2015), a property that we can also reproduce in our model. Joya and Rougier (2019) also using a European sample provide causal evidence for contagion effects generating this comovement. Third, Baqaee and Farhi (2019) challenge and qualify Hulten's aggregation rule, extending it beyond first-order terms to allow for nonlinearities that can emerge when sales shares vary with productivity shocks. They argue that Hulten's aggregation rule is an artifact of the constant returns to scale (CRS) property assumed for every micro production function, and recent sectoral level studies estimate elasticities of substitution that differ from the CRS benchmark (Atalay, 2017; Boehm et al., 2019; Barrot and Sauvagnat, 2016). This is also what casual empiricism suggests, since relative sales shares vary on all levels of aggregation from plants to firms to industries or sectors.

To paraphrase Sims (1980), we face a "wilderness" of potential alternatives for the aggregation of production once we depart from the knife-edge Cobb-Douglas case, apparently preventing the estimation of the correct structural model. To circumvent this problem, we opt for a small set of nonparametric assumptions that imply restrictions on the functional form of the considered production function, resulting in a reduced-form model based on the elasticity of substitution among supplier relationships. Aggregate shocks are then measured by the weighted sum of idiosyncratic productivity shocks, allowing us to decompose the aggregate shock into its constituent parts and to isolate the impact of each industry on aggregate fluctuations. Our proposed measure outperforms previous approximations that assume an invariance of sales shares between periods. It is worthwhile to point out that we achieve this performance using raw data only, that is without resorting to any arbitrary manipulation of the productivity growth rate distributions such as *winsorizing*, which considerably distorts the performance of granular approaches (Dosi et al., 2019).[2] Finally, our approach can also be thought of as a short-run complement to the complexity index proposed by Hidalgo and Hausmann (2009), who argue that the relative degree of complexity of national production structures is a useful predictor for national levels of production in the long run. Notably, though, our short-run perspective is transnational and emphasizes instead that aggregate short-run fluctuations depend on industry level characteristics that are not confined to national boundaries.

The next section presents our reduced-form model. Section 3 deals with the econometric identification strategy for testing our model predictions. Our main empirical results are presented in section 4, including several robustness checks. We discuss the implications and limitations of our results in section 5.

2 Idiosyncratic Shocks and Aggregation

Hulten's theorem states that for any efficient economy in general equilibrium with N producers (firms or industries) indexed by i

$$d\,log\,TFP = \sum_{i=1}^{N} \frac{S_i}{Y} d\,log\,TFP_i, \tag{1}$$

where TFP denotes the total factor productivity and Y the GDP of the economy, while S_i are the sales and TFP_i the productivity levels of entity i. Following Gabaix (2011) and assuming that GDP and TFP growth rates are related through a proportionality constant

[2] Since purely granular approaches to aggregate fluctuations are very sensitive to winsorizing, we quantify its impact for productivity shocks that have a normal or a double-exponential distribution in the Online Appendix C. We show analytically that in the empirically relevant leptokurtic case the bias from winsorizing is substantial.

μ reflecting factor usage, we have

$$d\,log\,Y = \mu \sum_{i=1}^{N} \frac{S_i}{Y} d\,log\,TFP_i. \tag{2}$$

This implies that for any given time period the rate of change of GDP is determined by the Domar weighted sum of idiosyncratic productivity shocks. Hulten's theorem holds for any efficient economy, irrespective of the existence or particular shape of input-output linkages in the economy, but merely up to a first-order approximation, or equivalently when the Domar weights S_i/Y are time invariant with respect to idiosyncratic productivity TFP_i. Whenever sales shares do respond to productivity shocks, however, potentially sizeable higher-order effects can emerge from these changes in sales shares. Baqaee and Farhi (2019) consider a second-order approximation and show that the impact of time-varying sales shares on GDP growth is uniquely determined by two effects: the derivatives of sales shares and of the input-output multiplier with respect to productivity shocks. While the former describes the change of sales shares with respect to productivity differences, the latter describes how these changes affect the ratio of total sales to GDP. It turns out that all higher order terms vanish only in the Cobb-Douglas CRS case. Whenever the production functions of some producers deviate from this knife-edge scenario, second-order terms can become quantitatively important, depending on the degree of factor reallocation that is assumed.

The operationalization of these theoretical insights is complicated by at least two concerns. First, the estimation of a complete structural model depends on unobservable general equilibrium elasticities of substitution. Second, since sales shares and productivity shocks are observable, we could construct a measure of the total change in TFP using (2) by integrating over a discrete time interval, where now both productivity levels TFP_i and sales shares S_i/Y can vary over time,

$$\frac{\Delta Y_t}{Y_{t-1}} = \mu \sum_{i=1}^{N} \int_{t-1}^{t} \frac{S_i}{Y}(t) d\,log\,TFP_i(t), \tag{3}$$

yet a good approximation of the integral in continuous time necessitates reasonably high observational frequencies. The discrete approximation by (left) Riemann sums that is used by both Gabaix (2011) and Baqaee and Farhi (2019) reads

$$\frac{\Delta Y_t}{Y_{t-1}} \approx \mu \sum_{i=1}^{N} \frac{S_{i,t-1}}{Y_{t-1}} \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}, \tag{4}$$

and becomes more accurate the higher the frequency of observation. This approximation, however, presupposes that sales shares stay constant between $t-1$ and t and jump to their new level precisely at time t. Given quarterly or annual frequencies, it seems rather implausible for sales shares to stay (even approximately) constant between periods.[3]

Hence we want to improve and generalize approximation (4) in order to account for the adjustment in sales shares within a discrete time interval using observations only at $t-1$ and t. To do so, we employ the following parsimonious set of assumptions that recovers the approximation by Riemann sums in (4) as a special case. Since our data are on the industry level, the exposition will refer to industries, yet in principle it applies equally to firms. One assumption per level of aggregation fully determines our model, providing a convenient modular structure:

1. On the *micro* level, we follow standard procedure and assume that idiosyncratic productivity shocks are exogenous and that each industry's revenue and TFP growth rates are related by a finite proportionality constant $\eta \geq 1$. This condition is very general and is fulfilled by any conventional neoclassical production function on the micro level (Barro, 1999).

2. On the *intermediate* level, we assume a constant and finite propagation intensity $\omega > 1$ between supplying and supplied industries, implying that intermediate inputs are *micro substitutes* in the terminology of Baqaee and Farhi (2019).[4] If two industries do not have a supplier-recipient relationship, there is no direct effect from the growth of revenue in one to the other, but there will generally be an indirect effect that

[3]While the input-output multiplier in our European data is remarkably stable over time, with an average value of around 0.5 and a coefficient of variation (CoV) of 0.01, the average CoV for Domar weights is more than one order magnitude larger at 0.2.

[4]We deliberately refrain from decomposing this substitutability in its partial price and quantity effect, since this would require us to define the degree of factor mobility, in particular labor mobility, which depends on the considered time-horizon (Baqaee and Farhi, 2019).

is mediated by the elasticity parameter.[5] Joya and Rougier (2019) provide causal evidence for the existence of this propagation effect. Apart from this condition, we also impose Gibrat's law in industry sizes and postulate that the size growth rate is uncorrelated to industry size. The rich extant literature on the firm level typically indicates that Gibrat's law is a reasonable approximation for the relationship of the mean growth rate and firm size but fails at higher moments (Sutton, 1997, as the canonical reference). We confirm this finding for the industry level aggregation of our dataset, indicating that Gibrat's law for the first moment is indeed an innocuous assumption.[6]

3. On the *macro* level, we aggregate from end-of-period sales levels to GDP via a constant input-output multiplier $\phi \in (0, 1]$ that measures the ratio of GDP to total sales, so $\phi < 1$ implies the existence of intermediate goods, whereas $\phi = 1$ corresponds to an islands economy without intermediate consumption.

Since we assume linear relations both at the micro and macro level, nonlinear effects have to originate at the intermediate level and it turns out that end-of-period sales levels are determined by the *interaction* of the size distribution with the topology of the production network, mediated by the propagation intensity.

Heuristically, the idea is to trace how a productivity shock to one industry propagates through the production network, and then to aggregate this effect over all industries experiencing simultaneous shocks. This will lead to an expression that resembles the centrality measure first proposed by Katz (1953), and it is also instrumental in deriving the purely granular and the purely networked economy as special cases. The linkages in the production network (or graph) G are described by the binary adjacency matrix A, where $A_{ij} = 1$ if j is a supplier to i and $A_{ij} = 0$ otherwise.[7] Let d denote the graph distance between two industries and consider what we term the *order effect* of a productivity shock to an industry j as a function of the graph distance to an industry i that is (possibly but

[5] As a microfoundation for this incomplete pass-through, Liu (2020) consider (iceberg) adjustment costs. Coincidentally, with a similar but distinct model, they arrive at a very similar aggregation rule.

[6] Material available upon request.

[7] Note that we are deviating from convention here, with i typically being the supplier in standard IO tables. Our adjacency matrix is thus simply the transpose of the standard version.

not necessarily) different from j. Denote the rate of change in the sales of an industry j that has a distance d to the initially affected industry i by $\Delta S^d_{j,t}/S_{j,t-1}$. From our micro level assumption, the initial effect of a productivity shock on the sales of industry i is the zeroth-order effect denoted by $\Delta S^0_{i,t}/S_{i,t-1}$, and is proportional to the initial productivity shock[8]

$$\frac{\Delta S^0_{i,t}}{S_{i,t-1}} = \eta \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}. \tag{5}$$

Suppose that each industry i takes goods produced by some industry j as inputs; given our assumption of a constant propagation intensity $\omega > 1$, the first-order effect of a change in industry i's sales on the sales growth of a supplier j to i is then given by

$$\frac{\Delta S^0_{i,t}}{S_{i,t-1}} = \omega \frac{\Delta S^1_{j,t}}{S_{j,t-1}}, \tag{6}$$

implying that the first-order effect on the sales growth of a supplier industry j stemming from a single productivity shock to industry i is[9]

$$\frac{\Delta S^1_{j,t}}{S_{j,t-1}} = \eta\omega^{-1} \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}. \tag{7}$$

Next consider an industry k that is a supplier to industry j; given a direct path from i to k through j, the second-order effect on k from the initial shock to i is therefore

$$\frac{\Delta S^2_{k,t}}{S_{k,t-1}} = \eta\omega^{-2} \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}, \tag{8}$$

and so forth. It is convenient to rewrite equations (7) and (8) in terms of their order effect

[8] This is a conventional assumption in real business cycle models (see Gabaix, 2011, for a more detailed discussion).

[9] Note that ω might vary between periods and could thus be indexed by $t-1$. We allow for time-varying weights in our empirical application, yet ω is very stable in practice which is why we drop time indices not to clutter notation too much.

for an arbitrary industry j that might be but is not necessarily connected to the initially affected industry i by a distance d. This allows to generalize for an arbitrary number of paths between industries at a distance d by exploiting a property of the d-th power of the binary adjacency matrix A. For this, we note that A^d_{ij} denotes the number of paths between industry j and i through $d-1$ intermediaries. Thus A^d_{ij} gives the number of ways through which industry j is affected by a change in the productivity of industry i in the d-th round of adjustments, mediated by the constant propagation intensity.

If we assume that the network topology does not change during a discrete time interval (but might very well change from one period to the next), the n-th order effect on an arbitrary industry j is then given by

$$\frac{\Delta S^n_{j,t}}{S_{j,t-1}} = A^n_{ij}\eta\omega^{-n}\frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}. \tag{9}$$

Equation (9) and $\omega > 1$ imply that the correlation between the sales growth rates of different industries has to be positive and decreasing in the network distance between industries, and this is exactly what Carvalho (2014) finds for his dataset of US industries, lending empirical support to the plausibility of this requirement. To get the total effect from the shock on i to an industry j from all orders, we sum over n

$$\frac{\Delta S_{j,t}}{S_{j,t-1}} = \eta\sum_{n=0}^{\infty} A^n_{ij}\omega^{-n}\frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}, \tag{10}$$

and $\omega > 1$ is a sufficient condition for this sum to converge. In the next step, we consider the total effect of simultaneous TFP shocks to all N industries apart from the singular one we considered so far and on their sales growth rates by summing both sides of the equation over the set of all industries by

$$\sum_{i=1}^{N}\frac{\Delta S_{i,t}}{S_{i,t-1}} = \eta\sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{n=0}^{\infty} A^n_{ij}\omega^{-n}\frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}. \tag{11}$$

The expression in (11) closely resembles the Katz centrality vectors frequently employed in the network literature to measure the centrality of nodes when higher-order effects are relevant. To see this, consider the vector of Katz centrality scores k for the graph G with a binary but not necessarily symmetric adjacency matrix A depending on some attenuation parameter $\beta > 1$ and given by

$$k(\beta) = ((\mathcal{I}_{N\times N} - \beta A)^{-1} - \mathcal{I}_{N\times N})\mathcal{I}_{N\times 1} \tag{12}$$

$$= ((\mathcal{I}_{N\times N} + \beta^{-1} A + \beta^{-2} A^2 + ...) - \mathcal{I}_{N\times N})\mathcal{I}_{N\times 1} \tag{13}$$

$$= \left(\sum_{j=1}^{N} \sum_{n=0}^{\infty} \beta^{-n} A_{ij}^{n} - 1 \right)_{i=1}^{N}, \tag{14}$$

where $\mathcal{I}_{N\times N}$ is the identity matrix with dimension $N \times N$ and $\mathcal{I}_{N\times 1}$ is a unit vector with length N. Bonacich (1991) shows that if β approaches the dominant (Perron) eigenvalue λ_{max} from above, the vector of Katz centrality scores approaches the vector of eigenvector centrality scores denoted e. For $\beta = \omega$ we have from equation (14) that $k(\omega) = (\sum_{j=1}^{N} \sum_{n=0}^{\infty} A_{ij}^{n} \omega^{-n} - 1)_{i=1}^{N}$, so this term corresponds to the vector of Katz centrality scores. Assuming that $\omega = \lambda_{max}$, that is the coefficient ω corresponds to the Perron eigenvalue, this implies that $k(\lambda_{max}) = e$, so the Katz centrality vector is equal to the vector of eigenvector centrality scores, and we have $\mathcal{I}_{N\times 1} + e = (\sum_{j=1}^{N} \sum_{n=0}^{\infty} A_{ij}^{n} \omega^{-n})_{i=1}^{N}$.

Let $\theta_{i,t-1} = 1 + e_{i,t-1}$ denote the i-th entry of the *shock transmission vector* θ_{t-1}, capturing both the initial impact of the productivity shock as well as all higher-order propagations between t and $t-1$. Expressing equation (11) using $\theta_{i,t-1}$, it follows that

$$\sum_{i=1}^{N} \frac{\Delta S_{i,t}}{S_{i,t-1}} = \eta \sum_{i=1}^{N} \theta_{i,t-1} \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}. \tag{15}$$

By imposing the restriction $\omega = \lambda_{max}$, the total rate of change resulting from the sum of arbitrary order-effects between t and $t-1$ will correspond to the centrality weighted sum of all productivity shocks. Hence we have characterized the end-of-period state of all industries in the economy, which completes the aggregation at the intermediate level. To close the system and to aggregate to the macro level, we recall from our macro assumption that

$$\phi \sum_{i=1}^{N} \Delta S_{i,t} = \Delta Y_t, \tag{16}$$

where ΔY_t denotes the change in GDP between $t-1$ and t, and the input-output multiplier $\phi < 1$ indicates the presence of intermediate goods. From equation (16) we have

$$\Delta Y_t = \phi \sum_{i=1}^{N} \Delta S_{i,t} = \phi \sum_{i=1}^{N} \frac{\Delta S_{i,t}}{S_{i,t-1}} S_{i,t-1}, \tag{17}$$

and therefore

$$\frac{\Delta Y_t}{Y_{t-1}} = \phi \sum_{i=1}^{N} \frac{\Delta S_{i,t}}{S_{i,t-1}} \frac{S_{i,t-1}}{Y_{t-1}}. \tag{18}$$

To combine the condition in equation (15) with the aggregation condition in equation (18), we need to impose the statistical regularity of Gibrat's law, i.e., industry size measures $\left(S_{i,t-1}/Y_{t-1}\right)_{i=1}^{N}$ and size growth rates $\left(\Delta S_{i,t}/S_{i,t-1}\right)_{i=1}^{N}$ being uncorrelated. It proves instructive to consider both vectors as sequences of realizations of random variables with finite first moment to rewrite equation (18) in terms of expectation values by

$$\phi \sum_{i=1}^{N} \frac{\Delta S_{i,t}}{S_{i,t-1}} \frac{S_{i,t-1}}{Y_{t-1}} = \phi N \cdot E\Big[\frac{\Delta S_{i,t}}{S_{i,t-1}} \cdot \frac{S_{i,t-1}}{Y_{t-1}}\Big] = \phi N \cdot E\Big[\frac{\Delta S_{i,t}}{S_{i,t-1}}\Big] \cdot E\Big[\frac{S_{i,t-1}}{Y_{t-1}}\Big], \tag{19}$$

with the last step exploiting Gibrat's law. This leads to

$$\phi N \cdot E\Big[\frac{\Delta S_{i,t}}{S_{i,t-1}}\Big] \cdot E\Big[\frac{S_{i,t-1}}{Y_{t-1}}\Big] = \phi \sum_{i=1}^{N} \frac{\Delta S_{i,t}}{S_{i,t-1}} \cdot E\Big[\frac{S_{i,t-1}}{Y_{t-1}}\Big], \tag{20}$$

and by equation (15)

$$\phi\eta \sum_{i=1}^{N} \theta_{i,t-1} \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}} \cdot E\Big[\frac{S_{i,t-1}}{Y_{t-1}}\Big] = \phi\eta \sum_{i=1}^{N} \theta_{i,t-1} \frac{S_{i,t-1}}{Y_{t-1}} \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}, \tag{21}$$

where the last equality assumes independence of Domar weights and weighted TFP shocks.[10] This allows us finally to restate the GDP growth rate as a function of industry

[10]Note that this assumption is latter independence assumption is approximately true for our dataset,

level weights and idiosyncratic TFP shocks by

$$\frac{\Delta Y_t}{Y_{t-1}} = \phi\eta \sum_{i=1}^{N} \theta_{i,t-1} \frac{S_{i,t-1}}{Y_{t-1}} \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}. \tag{22}$$

Equation (22) is our central result and relates the macro state of the economy to micro productivity shocks through the interaction of network structure and Domar weights in *interaction* weights $\theta S/Y$. We obtain equation (22) under the crucial assumption of micro substitutability ($\omega > 1$) and the technical requirement that the parameter ω is equal to the Perron eigenvalue λ_{max}. While the former is theoretically intuitive if we want to exclude explosive growth, the latter is certainly plausible from an empirical point of view as the Perron eigenvalue is remarkably stable over time and close to twenty, implying that about five percent of productivity shocks are transmitted to suppliers. We also invoke the statistical requirement of Gibrat's law that we can empirically confirm for our dataset.

Note that the Riemann sum approximation (4) is a special case of our model specification. We can interpret this case economically as the correct specification for homogeneous network topologies or, more precisely, k-regular graphs (not merely as a discrete approximation with desirable mathematical properties). For k-regular graphs, all industries exhibit equal (or representative) degree $k \geq 0$. For this condition, $\theta_{i,t-1} = 1 + \bar{e} = \bar{\theta} \geq 1$ for all $i = 1, \ldots, N$. Letting $\mu = \phi\,\eta\,\bar{\theta}$, this leaves us with the formulation used by Gabaix (2011) and Baqaee and Farhi (2019),

$$\frac{\Delta Y_t}{Y_{t-1}} = \phi\eta\bar{\theta} \sum_{i=1}^{N} \frac{S_{i,t-1}}{Y_{t-1}} \cdot \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}} = \mu \sum_{i=1}^{N} \frac{S_{i,t-1}}{Y_{t-1}} \cdot \frac{\Delta TFP_{i,t}}{TFP_{i,t-1}}. \tag{23}$$

Even though the result in (23) is formally equivalent to (4) used in Gabaix (2011), μ depends in our formulation not only on factor usage but also on the input multiplier as well as the average connectivity of producers. This result might help to explain why empirical studies after Gabaix' seminal contribution tend to find values for μ far below unity (Ebeke and Eklou, 2017; Miranda-Pinto and Shen, 2019; Wagner and Weche, 2020), as do we, in obvious violation to the interpretation by Gabaix (2011) in terms of different growth

with weighted TFP shocks and Domar weights being only weakly connected on average with $\rho \approx 0.1$.

models which unanimously imply $\mu > 1$. The invariance of sales shares here is essentially a network homogeneity result. This alternative derivation is in line with the theoretical result in Baqaee and Farhi (2019) that the invariance of sales shares is a linearity result, since asymmetries in the network topology are the only possible source of nonlinearities here. We include this purely granular specification as a robustness check in the empirical part. Given the modular outline of the model, the purely granular case corresponds to a direct aggregation from the micro to the macro level without any heterogeneity at the intermediate level. In addition to a readily testable aggregation rule, the modular structure of our reduced-form model thus provides a convenient way to quantify the relevance of the intermediate in the aggregation.

3 Data and Estimation

We use the second release of the World input-output database (2017), henceforth WIOD, that presently covers the period from 2000 to 2014. A detailed breakdown of the WIOD and its various data sources, for the most part based on national accounts, is provided by Timmer et al. (2015). Macroeconomic data for GDP and GDP per capita series as well as several other macro variables used for robustness checks are taken from Eurostat and the World Bank's Development Indicators database. We focus on the EU28, mainly to exploit the large degree of cross-country connectivity in an integrated free trade area with a large ensemble of industries. While the sample period is rather short, it contains at least two important crises periods (the global financial and economic crisis and the European sovereign debt crisis) with substantial variation in GDP growth rates. Thus the sample period provides a reasonably challenging testing ground for the empirical modeling of aggregate fluctuations in general, and it should also help in discriminating between two different cases we have derived in the previous section.

3.1 Data

The WIOD provides annual series for 56 industries per country, leaving us with 1568 distinct industries in the EU28. The size-based metrics are taken from the Socio-Economic

Accounts (SEA), in particular the annual volume of industry sales (gross output) and the number of employees used to calculate Domar weights and productivity shocks, respectively. The complementary cumulative distribution functions (CDFs) of industry sales, reported for each sample year in Online Appendix D, reveal a large degree of heterogeneity in industry size, but they do not exhibit power law-like behavior. Thus aggregation from the firm to the industry level apparently prevents the straightforward application of 'pure' granular theory that rests on the power law property of firm size distributions.[11]

In addition to size-based metrics from the SEA, our model also requires the specification of adjacency matrices in order to account for the dynamic adjustment of sales shares through supplier relations. To construct these matrices, we binarize the World Input-Output Tables (WIOTs) from the WIOD that describe inter-industry flows for all industry pairs, leading to a 1568×1568 matrix for each year. Following Carvalho (2014), we apply a standard threshold of one percent of an industry's total input deliveries for binarization. The resulting networks are extremely sparse and visualized in Figure 1 for the year 2000, where we observe 26 924 non-zero edges out of a possible $1568^2 \approx 2.5 \times 10^6$, resulting in a low network density of around 0.01, with an average degree of 17 in that year. Since almost all industries are connected through a (weak) input linkage, this leads to a massive reduction in average connectivity by almost two orders of magnitude, from 90 % of non-zero edges in the initial network to the 1% after binarization. This discrepancy leaves us reasonably confident that we effectively separated noise from signal and that the remaining links represent economically meaningful linkages. The size of each node in the figure is proportional to its weight score,[12] and the network looks very similar in other years as well. After all, the binarization threshold does not artificially inflate connectivity and should consequently not have a spurious impact on aggregate fluctuations. Yet, our aggregation perspective emphasizes the importance of higher-order effects in the production network, and the figure shows that all clusters are still connected to each

[11]By contrast, the distribution of our construct of 'interaction weights' summarized in Online Appendix D is much closer to the power law benchmark for all years, indicating that size-based measures might not be able to capture the full extent of an industry's impact on aggregate fluctuations.

[12]In the empirical part the transmission vector corresponds to eigenvector centralities, which are much more efficiently computed by a standard power iteration method than Katz centralities. This does not bias the estimates in our dataset because the initial impact of productivity shocks rapidly vanishes in relation to the higher-order terms in equation (11), even for a finite-sum approximation.

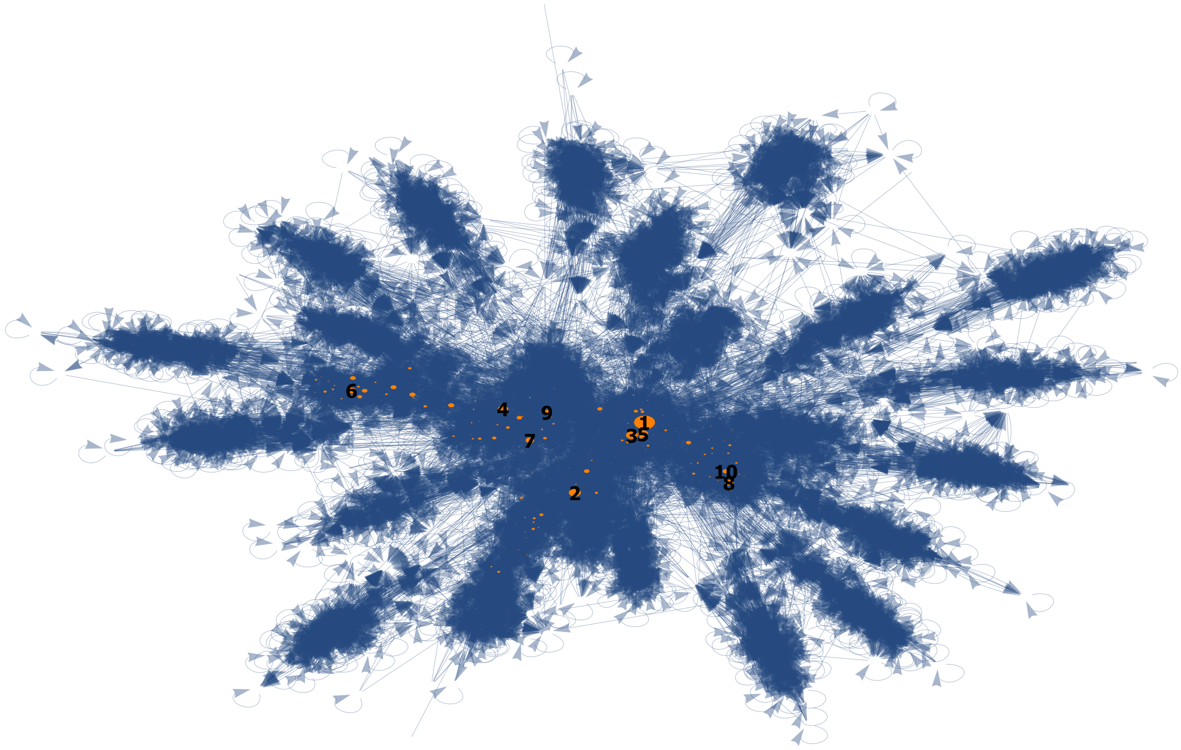

1	German manufacture of motor vehicles, trailers and semi-trailers	6	Swedisch real estate
2	French construction	7	British human health and social work activities
3	German construction	8	Italian construction
4	British construction	9	British Public administration and defence; compulsory social security
5	German manufacture of machinery and equipment n.e.c.	10	Italian wholesale trade

Figure 1: The EU28 production network exhibits substantial heterogeneity in weights and connectivity, along with clusters that mostly coincide with national borders. The ten industries with highest weights are listed in descending order and node sizes are proportional to interaction weights.

other via some intermediaries, corroborating the importance of shock transmission through higher-order connections in spite of the network's low density and average degree. The industry with highest impact on aggregate flucations is the German automotive sector, in line with the intuition that industries producing highly complex final products rely on long and diversified supply chains and are thus also of systemic importance.

3.2 Estimation

To quantify the impact of industry level shocks on aggregate fluctuations we essentially adhere to the procedure of Gabaix (2011). First we construct *residuals* that represent the sum of idiosyncratic shocks to the $K < N$ largest industries, here weighted according to the theoretically motivated specifications in equations (22) and (23).[13] Then we run linear regressions on the various residuals and consider the coefficients of determination, in other words their respective (adjusted) R^2, in order to compare how much of the variation in GDP growth can be explained by each of the two model specifications. To do so, we need to measure productivity shocks and here we will use labor productivity instead of TFP.[14] Let $g_{i,t} - \bar{g}_t$ denote the productivity shock to industry i at time t, where $g_{i,t} = \ln(S_{i,t}/E_{i,t}) - \ln(S_{i,t-1}/E_{i,t-1})$ is the observed change in i's labor productivity, with $E_{i,t}$ being the number of employees in industry i at time t, while $\bar{g}_t$ measures the aggregate productivity shock. This conventional formulation assumes that the observed productivity signal $g_{i,t}$ is additively separable into an idiosyncratic and an aggregate component. We calculate the aggregate shock following Gabaix (2011) by taking the median over the K^* industries with highest weight in each period. The specification of $\bar{g}_t$ thus leaves one additional degree of freedom - the number of industries K^* with highest weight over which we calculate the aggregate shock. Our results remain remarkably robust for several different K^* that we use for our calibration.

Now we are in a position to compute residuals for our respective model specifications, starting with the *interaction residual* $\mathcal{I}$ that aggregates the idiosyncratic productivity shocks to the K largest industries $i = 1, \ldots, K$ in period t using the interaction weight $\theta_{i,t-1} \cdot S_{i,t-1}/Y_{t-1}$,

$$\mathcal{I}_t = \sum_{i=1}^{K} \theta_{i,t-1} \frac{S_{i,t-1}}{Y_{t-1}} (g_{i,t} - \bar{g}_t). \tag{24}$$

For robustness we also consider the case of purely linear shock responses on a k-regular graph economy with the *granular residual* $\mathcal{G}$, using merely Domar weights

$$\mathcal{G}_t = \sum_{i=1}^{K} \frac{S_{i,t-1}}{Y_{t-1}} (g_{i,t} - \bar{g}_t), \tag{25}$$

[13]The granular case studied by Gabaix corresponds to eq. (23), but here for industries instead of firms.

[14]Our model is built on TFP shocks that are Hicks-neutral, while labor productivity shocks generally are not. This is at least one of the reasons why the growth of TFP and labor productivity follow different trajectories, as shown for instance by Syverson (2004) for the manufacturing industry on the plant level. Nevertheless we employ labor productivity for three reasons: from a theoretical viewpoint, the estimation of TFP presupposes precisely the CRS assumption we challenge with our model; from an empirical viewpoint, we simply do not have TFP data in the WIOD; finally, the use of labor productivities facilitates the comparison with the granular results of Gabaix, who also theorizes with TFP but then utilizes labor productivity for his empirical exercise.

As a final robustness check, we consider a (discretized version of a) Divisia index for the granular residual in eq. (25). Following an idea of Törnqvist (1936), this index is intended to capture some of the nonlinearities by averaging weights over two time periods and leads to considerable performance gains in the application by Baqaee and Farhi (2019).

$$\mathcal{D}_t = \sum_{i=1}^{K} \frac{1}{2}\left(\frac{S_{i,t-1}}{Y_{t-1}} + \frac{S_{i,t}}{Y_t}\right)(g_{i,t} - \bar{g}_t), \tag{26}$$

The interaction model, its granular variant and the Divisia index translate to the following regressions,

$$I: \ (a) \ g_t^Y = \beta_0 + \beta_1 \mathcal{I}_t + \epsilon_t \qquad (b) \ g_t^Y = \beta_0 + \beta_1 \mathcal{I}_t + \beta_2 \mathcal{I}_{t-1} + \epsilon_t \tag{27}$$

$$II: \ (a) \ g_t^Y = \beta_0 + \beta_1 \mathcal{G}_t + \epsilon_t \qquad (b) \ g_t^Y = \beta_0 + \beta_1 \mathcal{G}_t + \beta_2 \mathcal{G}_{t-1} + \epsilon_t \tag{28}$$

$$III: \ (a) \ g_t^Y = \beta_0 + \beta_1 \mathcal{D}_t + \epsilon_t \qquad (b) \ g_t^Y = \beta_0 + \beta_1 \mathcal{D}_t + \beta_2 \mathcal{D}_{t-1} + \epsilon_t \tag{29}$$

where aggregate fluctuations are denoted g_t^Y and measured by the growth rates of (per capita real) GDP. The left panel, marked (a), corresponds to equations (22) and (23) with III using the Divisia index in (26), while the right panel, marked (b), includes a lagged term of the respective residual. The lagged term accounts for the possibility that the adjustment of sales shares takes time such that GDP responds to micro productivity shocks with a delay, and controls for reverse causality.

A useful tool for our investigation are plots that show how much of the explained variation R^2 in GDP growth rates can be attributed to the largest K industries.[15] To establish a benchmark we show in Appendix A that the expected R^2 in the presence of idiosyncratic shocks, denoted $E[R^2]$, for K industries with weights $w \in \{\theta\, S/Y, S/Y\}$ is given by the ratio of Hirschman-Herfindahl indices H for the weights of these K industries relative to the weights for all N industries,

$$E[R^2] = \frac{H(K)}{H(N)} = \sum_{i=1}^{K} \frac{(w_i)^2}{\sum_{j=1}^{N} (w_j)^2}. \tag{30}$$

[15] This approach is inspired by Blanco-Arroyo et al. (2018), who investigate purely granular regimes, though. Here we generalize the comparison of empirical R^2 plots to the benchmark of correctly identified idiosyncratic shocks and weights.

Since both weights exhibit considerable leptokurtosis, we should observe a large initial increase in R^2 already for a small number of industries, followed by a roughly monotonous yet moderate increase that settles on a plateau when K approaches its upper bound, where the size of industry weights decreases rapidly. These three qualitative features will guide our comparison and assessment of the different specifications in the next section.

Notably, and in stark contrast to conventional granular approaches (Gabaix, 2011; Ebeke and Eklou, 2017; Blanco-Arroyo et al., 2018), we only use raw data in all our specifications, and do not winsorize the empirical density of productivity growth rates.[16] In Appendix B, we show that winsorizing should actually *decrease* rather than increase explanatory power for a correct model specification, essentially due to a decrease in residual variance, and this effect is amplified by the well-documented leptokurtic nature of empirical productivity growth rate densities (Dosi et al., 2012; Yu et al., 2015; Dosi et al., 2019).[17] Explanatory power thus only increases when the winsorizing procedure generates spurious comovement between industries, inflating regression coefficients by compensating for the initial drop in variance and explanatory power. Viewed from this perspective, the fact that a purely granular approach needs to generate comovement through winsorizing to push the estimates over the brink of significance already indicates the crucial importance of the interaction between size and network effects.

4 Results

The empirical formulation of our main idea corresponds to eq. (27), so we estimate eq. I(a), along with the variant I(b) that includes one lagged term of the interaction residual. We set the number of relevant industries K^* to $K^* = 200$ but our results remain robust for

[16]Winsorizing is typically justified by reference to extraordinary events such as large mergers. Yet Bottazzi et al. (2019) show in their appendix that for each of the idiosyncratic events causing the four minimum and maximum growth rates in the Gabaix (2011) sample, the notion of "extraordinary events" can be reasonably challenged.

[17]We show that symmetric winsorizing at merely one percent already diminishes the explanatory power by more than five percent when productivity growth rates follow a Laplace distribution, as the empirical literature strongly suggests. The loss in explanatory power is considerably smaller for a Gaussian distribution at about 1.8 percent. We discuss the details of these computations in Appendix B and Online Appendix C. Recent studies also recognize the relevance of leptokurtic productivity growth rate distributions (Carvalho and Tahbaz-Salehi, 2019; Daniele and Stüber, 2023).

different K^* and K. Figure 2 shows our main comparative results for the interaction and granular residual for all different specifications. The interaction specifications lead to highly significant contemporaneous coefficient estimates for both considered specifications, while coefficients for the granular special case remain insignificant. The specification accounting for higher order effects leads to considerable performance gains to the purely granular case, lending robust support for the interaction hypothesis. The explanatory power of the interaction specification as measured by the coefficient of determination is sizeable, more than one half or more than two thirds of explained variance for the purely contemporaneous and specification with lagged term, respectively. Our findings do not appear to be driven by misspecification or anomalous behavior either, with all coefficients exhibiting the expected (positive) sign[18] and very similar magnitude (of $\approx 10^{-2}$) and thus not artificially inflating model performance or biasing results in favor of one residual. The time series of both residuals against GDP growth rates suggest that the superior explanatory power of the interaction residual is primarily the result of its performance during the twin crises within our sample period that the interaction residual tracks reasonably well but are not at all reflected in the behavior of the purely granular variant. This invites to speculate that the higher-order effects the interaction specification is designed to capture are particularly relevant to account for cascading failures that are presumedly most impactful in crisis periods.

The incremental gains of explanatory power by adding industries are extremely asymmetric for the interaction specification.[19] The R^2 plot quickly stabilizes after only about 50 industries - or 3% of all with systemic importance and thus follows our theoretical predictions on the functional form of R^2 plots with leptokurtic weight distributions. The R^2 plots furthermore indicate that the 'snapshot' results we report in the right panel of Figure 2 are not misleading and the coefficients of determination are robust against incremental variation of K. Most importantly, the interaction residual outperforms for

[18]Since the eigenvector centrality we employ can, by the nature of eigenvectors, give centrality scores only to up to a common factor, both coefficient estimates are not directly quantitatively interpretable. We can thus gauge the empirical success of the interaction hypothesis only by the sign and significance of the estimated coefficients.

[19]Note that we take the composite K largest industries, e.g. when $K = 10$ we take the ten largest industries in a given year so their identities might change over time.

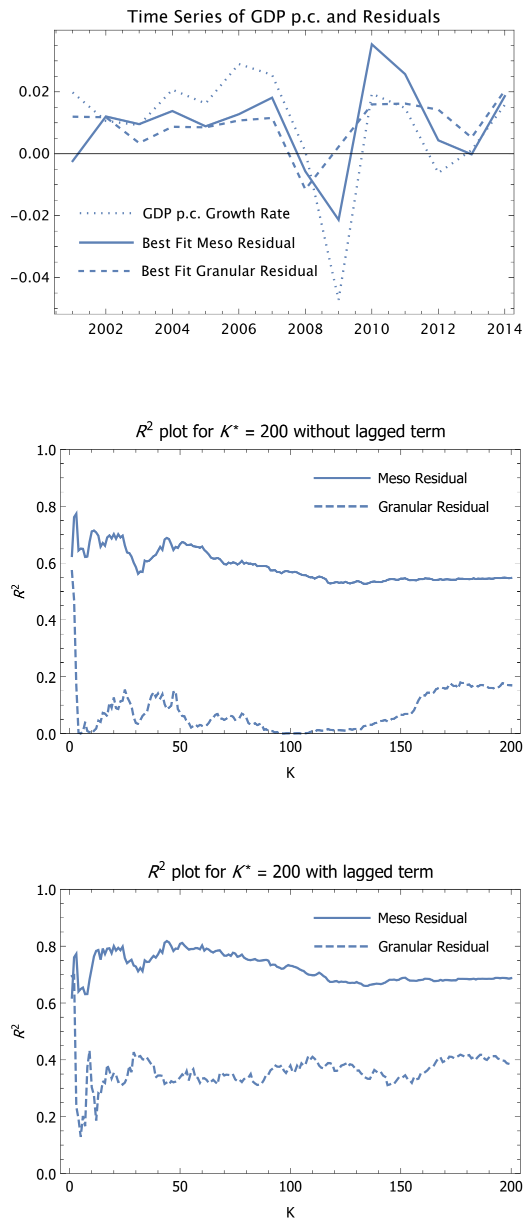

	I	
	(a)	(b)
$\mathcal{I}_t$	5.984***	6.761***
	(1.569)	(1.456)
$\mathcal{I}_{t-1}$		1.376
		(1.441)
Intercept	0.010**	0.009**
	(0.004)	(0.003)
T	14	13
R^2	0.55	0.69
Adjusted R^2	0.51	0.62
F Statistic	14.54***	11.00***
	II	
	(a)	(b)
$\mathcal{G}_t$	0.697	0.608
	(0.445)	(0.418)
$\mathcal{G}_{t-1}$		0.871*
		(0.456)
Intercept	0.009*	0.010*
	(0.005)	(0.005)
T	14	13
R^2	0.17	0.39
Adjusted R^2	0.10	0.26
F Statistic	2.45	3.148*

Figure 2: The interaction residual outperforms the granular special case for all specifications and all considered K^*. The left panel shows the time series of the interaction and granular residuals against g_t^Y as well as the respective R^2 plots without and with lagged term. The right panel reports the corresponding regression results for all specifications. Standard errors are given in parentheses; significance at the *** 1%, ** 5%, * 10% level.

the whole range of $K = 1, \ldots, K^*$, demonstrating the importance of higher-order shock transmission for aggregate fluctuations.[20] The lagged terms change the overall picture only marginally. The granular variant gains relatively more from including a lagged term but is still robustly inferior to the interaction variant. Both contemporaneous coefficients are also rather stable across specifications and remain highly significant. The lagged term in the interaction specification adds little to the explanatory power of the interaction residual and is insignificant at the usual significance levels. In accordance with the interaction hypothesis, it is thus primarily contemporaneous developments in the interaction residual that appears to comove with aggregate flucations. By contrast, Domar-aggregated shocks in the granular special case fail to explain much of the variation in GDP growth rates in both variants. Estimates from conventional Domar aggregation thus appear to strongly underestimate the impact of microeconomic idiosyncratic shocks. Since purely granular economies are a special case of general interaction derivation for symmetric networks, this finding also provides indirect evidence for the relevance of accounting for *asymmetric* topologies, with only a few industries occupying central roles within the production network.

4.1 Robustness: Divisia Index

As a first robustness check, we compare the performance of the interaction residual to a Divisia index for the Domar weights and estimate equation (29). Baqaee and Farhi (2019) show that such an index captures second-order effects of productivity shocks beyond the mere first-order impact captured by the granular residual of Gabaix (2011). In this sense, the Divisia index $\mathcal{D}_t$ is an alternative measure to capture nonlinearities using only size-based measure without accounting for the input-output structure that we include in our interaction residual $\mathcal{I}_t$ for the same goal.

Perhaps unexpectedly, using a Divisia index in place of simple single-period Domar weights does not increase model performance. Instead, the right panel of Figure 3 shows that for $K = 200$, the explanatory power even decreases, with still insignificant

[20]This result is robust for variation in $K^* = 50,\ 100,\ 150,\ 250$. We provide these results upon request. See also section 4.2 for robustness checks with $K = K*$ for various $K = 2, \ldots, 250$.

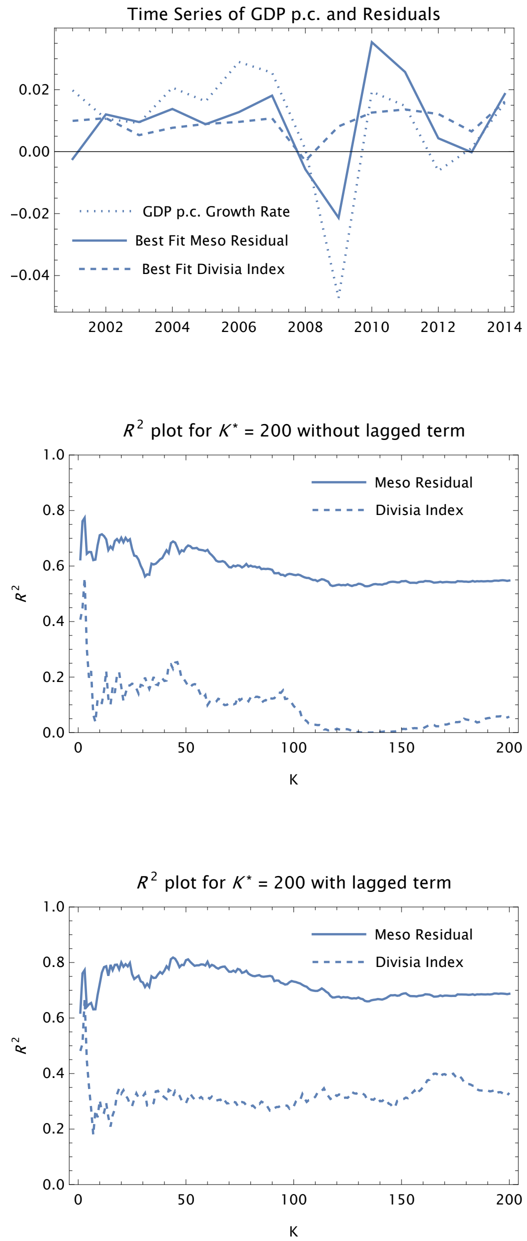

	I	
	(a)	(b)
$\mathcal{I}_t$	5.984***	6.761***
	(1.569)	(1.456)
$\mathcal{I}_{t-1}$		1.376
		(1.441)
Intercept	0.010**	0.009**
	(0.004)	(0.003)
T	14	13
R^2	0.55	0.69
Adjusted R^2	0.51	0.62
F Statistic	14.54***	11.00***
	II	
	(a)	(b)
$\mathcal{D}_t$	0.444	0.459
	(0.521)	(0.476)
$\mathcal{D}_{t-1}$		1.073*
		(0.529)
Intercept	0.008	0.007*
	(0.005)	(0.005)
T	14	13
R^2	0.06	0.33
Adjusted R^2	-0.02	0.20
F Statistic	0.723	2.465*

Figure 3: The interaction residual also outperforms the Divisia index of the granular residual for all considered K^*. The left panel shows the time series of the interaction residual and Divisia index against the GDP per capita growth rate g_t^Y as well as the respective R^2 plots without and with lagged term. The right panel reports the corresponding regression results for all specifications. Standard errors are given in parentheses; significance at the *** 1%, ** 5%, * 10% level.

contemporaneous coefficients in both specifications. For other K, the performance of the Divisia index closely mirrors the one-period granular residual, as is readily visible by comparing the left panels of Figures 2 and 3. Our preferred interpretation of this finding is that adequately accounting for the topology of the production network is necessary to capture the impact of microeconomic shocks - and that purely size-based measures are insufficiently capable of doing so.

4.2 Robustness: Identifying Aggregate Shocks

The validity of our comparative exercises crucially hinges on the identification of idiosyncratic shocks. One worry raised by Dosi et al. (2019) regarding the original results of Gabaix (2011) relates to the definition of the aggregate shock that is subtracted from an industry's growth rate to calculate the idiosyncratic disturbance: As detailed in subsection 3.2, K^* is a free parameter determining the number of industries considered for the calculation of the median as a measure of this aggregate shock. Results might be sensitive to the choice of the number of considered industries K^*. To examine the robustness of our results in this regard, we again employ R^2 plots, where we now jointly vary $K = K^*$ for all considered regression equations (27) to (29).

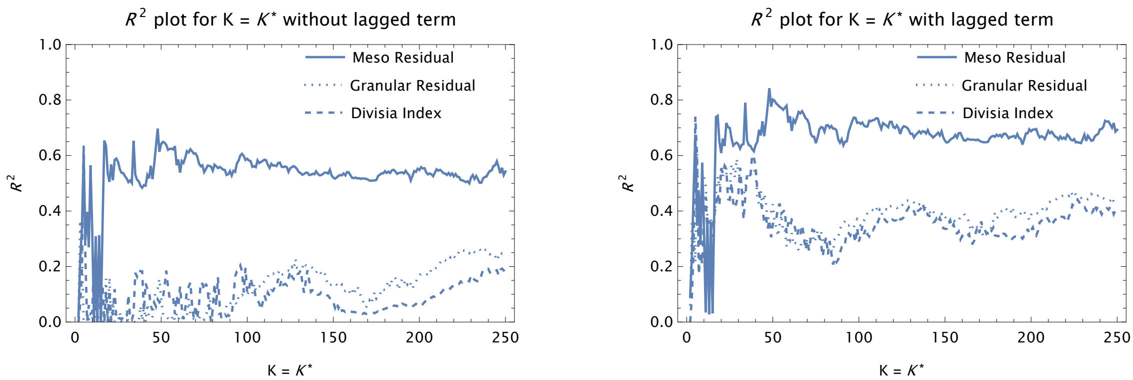

Figure 4: Here we examine the impact of alternative identification strategies for idiosyncratic shocks by varying K^*. The figure shows R^2 plots for $K = K^*$ with $K \in [2, 250]$ for OLS estimates of equations (27), (28) and (29). The interaction residual also robustly outperforms its granular variant and the Divisia index of the granular variant for those alternative measures of idiosyncratic shocks.

In particular, we consider the range of 2 to 250 for $K = K^*$. We start at $K = K^* = 2$, since the three residual variants are identically zero for all t for $K = K^* = 1$. As the results in Figure 4 show, the interaction residual consistently performs better than the other two specifications for these alternative measures of idiosyncratic shocks. After initial large fluctuations,[21] the R^2 quickly stabilizes for both interaction specifications with and without lagged term and hovers around 55 and 70 percent, respectively, confirming the results of our baseline specification. By contrast, the granular residual in both the single-period and Divisia form never reaches the R^2 achieved by the interaction specification and stabilize around 20 and 40 percent. The explanatory power of both size-based measures is roughly the same for all considered $K = K^*$, without one of the two showing superior performance. Overall, we thus unanimously find our initial results robust against varying the number of industries K^* to identify aggregate shocks.

4.3 Robustness: Macro Controls

Our investigation of potential confounding variables starts with two exogenous aggregate disturbances, i.e., the rate of change in oil prices (as oil is close to a universal input) and monetary policy decisions, in order to exclude the possibility that the performance of the interaction model stems from aggregate shocks that might unwittingly be included in our (rather conscientious) construction of idiosyncratic shocks. To preserve a reasonable number of degrees of freedom, we consider only the contemporaneous versions of the control variables in our specifications. We use the following control variables, following Gabaix (2011):

(i) ECB policy rate (denoted ECB, in percent, Euro area) and

(ii) annual growth rate of the real crude oil price (denoted Oil, Brent Europe).

[21]These fluctuations are likely due to the small sample for which the median is calculated that is hardly representing aggregate movements for $K* \leq 20$.

	(Benchmark)	(A)	(B)	(C)
$\mathcal{I}_t$	5.984***	5.999***	5.038**	5.571***
	(1.569)	(1.243)	(1.907)	(1.57)
$\text{ECB}_{control}$		0.006**		0.006**
		(0.002)		(0.002)
$\text{Oil}_{control}$			0.018	0.008
			(0.020)	(0.017)
Intercept	0.010**	-0.006	0.008	-0.006
	(0.004)	(0.006)	(0.004)	(0.006)
R^2	0.55	0.74	0.58	0.75
Adjusted R^2	0.51	0.69	0.50	0.67
AIC	-77.41	-82.78	-76.01	-80.76
BIC	-75.50	-80.22	-73.45	-77.57
F-Statistic	14.54***	15.64***	7.54***	9.78***

Table 1: The table summarizes the various robustness checks to examine the variability and performance of our benchmark jnteraction model against the inclusion of two macro controls (indicated by the subscript "control"). The low variability of effect sizes as well as the fact that the interaction coefficient does not lose significance for any control indicates the robustness of the benchmark specification. The performance only improves when adding a simple proxy for monetary policy that our baseline model cannot account for per construction.

These give rise to the following specifications:

$$(A):\ g_t^Y = \beta_0 + \beta_1\ \mathcal{I}_t + \beta_2\ \text{ECB}_t + \epsilon_t$$

$$(B):\ g_t^Y = \beta_0 + \beta_1\ \mathcal{I}_t + \beta_2\ \text{Oil}_t + \epsilon_t$$

$$(C):\ g_t^Y = \beta_0 + \beta_1\ \mathcal{I}_t + \beta_2\ \text{ECB}_t + \beta_3\ \text{Oil}_t + \epsilon_t$$

Each specification (A) to (C) is estimated by OLS. The results are summarized in Table 1. The $\hat{\beta}_1$ coefficient stays significant across all specifications and, more importantly, effect sizes do not vary much for the inclusion of all of the different controls (A) to (C). The largest deviation occurs for the specification that includes the rate of change in oil prices, but even this drop of less than a fifth in effect size is rather negligible in size. Our previous estimated effect size therefore appears to be robust against the inclusion of standard control variables, indicating that the model performance is not driven by omitted confounders.

The specifications in which we include a simple proxy for monetary policy ((A) and

(C)) would be preferred to our benchmark model according to both information criteria.[22] This result, however, is hardly surprising since the interaction residual is derived in a RBC setting where monetary policy is ineffective in generating real effects (Snowdon et al., 1994). As the New Keynesian literature demonstrates (see, e.g. Galí, 2008), interest rate shocks that are transmitted through the amount of available credit, inflation or asset prices have potentially strong effects on GDP fluctuations under price-stickiness. In this sense, complementing the real-side analysis with simple proxies for monetary policy shocks can be expected to outperform a model built only on real-side effects.[23]

By contrast, sudden surges in oil prices were one of the major historical motivations for supply-side shocks within RBC-type models, since oil price developments can (for economies like the EU28) be reasonably characterized as exogenous disturbances to production (Summers, 1986). Our results leave little room for exogenous aggregate shocks in the form of oil price changes, though, when the microeconomic shocks to individual industries are properly accounted for within the interaction residual. These results therefore not only testify to the general robustness of the interaction channel, they also suggest that conventional aggregate energy price shocks are efficiently identified by microeconomic shocks to a vanishingly small number of industries with disproportionate aggregate influence. Our findings thus corroborate the baseline assumption of RBC models that exogenous productivity shocks are the most important determinant of aggregate fluctuations in the short run but also show that they need to be conceived as operating on a much more granular level.

[22]As Kuha (2004) notes, both criteria select for different types of data-generating processes. The BIC is more suitable for processes with a small number of strong effects, while the AIC selects better for models with a collection of many effects. Model selection without such prior information should therefore be based on both criteria in agreement.

[23]This proxy can, by construction, not account for the (increasingly relevant) unconventional policy measures by the ECB and does only apply to the subset of the 19 Euro zone countries for our EU28 sample. The superior performance from including this imperfect proxy thus points to the relevance of the monetary policy channel and highlights its negligence in RBC-type models even more.

5 Discussion

Since Hulten (1978), growth theorists explicitly or implicitly believed to live in a world of linear shock responses. Our parsimonious model and the empirical results for Europe, however, cast doubt on the conventional wisdom that economies are linear in this sense, and thus on traditional Domar aggregation. We show that it is the strongly non-linear *interaction* of both size- and network-based channels, in the extant literature so far being treated entirely separately, that provides a crucial missing link between idiosyncratic productivity developments and aggregate GDP movement. This interaction channel emerges naturally, whenever we consider the dynamics between observation periods. Our reduced form model explicitly includes the adjustment of sales shares but still allows immediate Domar-type empirical application.

In contrast to much of contemporary macroeconomic theory, we do not counterfactually impose homogeneity in size or independence between industries to enable aggregation. Instead, we explicitly acknowledge industries as strongly heterogeneous and interdependent entities. Conveniently, this also allows us to make full use of the available information in the WIOD, which manifests itself in considerable performance gains. Notably, our aggregation rule also outperforms a prominent nonlinear measure in the form of a Divisia index proposed in the extant literature to capture nonlinearities without accounting for the input-output structure of the production network. As we expect in a model without nominal variables, we find that the performance increases even further by including a simple proxy for monetary policy. Perhaps unexpectedly, however, we find that the relative explanatory power of traditional aggregate shocks in the form of changes in oil prices is rather low compared to the interaction residual.

The robustly superior performance of the proposed aggregation rule based on interaction against the rival granular special case suggests that idiosyncratic shocks are of much greater relevance than previously found and can account for about two thirds of aggregate fluctuations. Notably, we find that the aggregate relevance of shocks is extremely skewed towards those pertaining to merely 50 systemically relevant industries. In the European context of our paper, the proverbial notion of an 'asymmetric shock' therefore lives on in another form, now not with shocks to countries but to specific industries (De Grauwe, 2020). In analogy to established results in complex network theory, we find the system of European industries to be resilient against equiprobable adverse shocks, as the systemically relevant industries are then only very rarely affected (Albert et al., 2000). If they are, however, the effect is disproportionately larger due to cascading failures. By exploiting this asymmetry, the high-dimensional space of macroprudential oversight can thus be reduced to a much more manageable three percent of European industries. We also provide some cautious arguments against policy neutrality in industrial policy. Our results suggest that targeted, highly ranked industries would later generate significant externalities beyond their sales share through the production network as a basic theoretical requirement for welfare-enhancing interventions (Manelici and Pantea, 2021).

Even though we conscientiously tested the robustness of all our results, our approach faces several limitations, essentially due to data constraints. Within the model, we completely abstracted from the monetary side of the economy, whereas both macroeconomic theory and our empirical evidence strongly hint at the presence and relevance of monetary channels (Ozdagli and Weber, 2017). A possible remedy could be to introduce nominal frictions into the model. Empirically, the WIOD provides only rather short time series, limiting both the generalizability of our results and the number of controls we can reasonably include within the regression specification. Also, while our industry level model is an attempt at disaggregation, industries are of course themselves aggregate abstractions which average over economic decisions conducted at the firm or even plant level (Bernard et al., 2019). We are convinced that datasets without these caveats will pave numerous avenues for further research, not least to establish proper microfoundations for the relationships we only posit in reduced form at the moment. After all, our results underline the importance objects at an intermediate layer of aggregation such as size distributions and contagion networks for proper aggregation from the micro to the macro level already in our present exposition as reduced form relationships.

References

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

Albert, R., Jeong, H., and Barabási, A.-L. (2000). Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382.

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. (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem. *Econometrica*, 87(4):1155–1203.

Barro, R. J. (1999). Notes on growth accounting. *Journal of Economic Growth*, 4(2):119–137.

Barrot, J.-N. and Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics*, 131(3):1543–1592.

Bernard, A. B., Dhyne, E., Magerman, G., Manova, K., and Moxnes, A. (2019). The origins of firm heterogeneity: A production network approach. *NBER Working Paper Series*, 25441.

Blanco-Arroyo, O., Ruiz-Buforn, A., Vidal-Tomás, D., and Alfarano, S. (2018). On the determination of the granular size of the economy. *Economics Letters*, 173:35–38.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake. *Review of Economics and Statistics*, 101(1):60–75.

Bonacich, P. (1991). Simultaneous group and individual centralities. *Social Networks*, 13(2):155–168.

Bottazzi, G., Li, L., and Secchi, A. (2019). Aggregate fluctuations and the distribution of firm growth rates. *Industrial and Corporate Change*, 28(3):635–656.

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

Carvalho, V. M. and Tahbaz-Salehi, A. (2019). Production networks: A primer. *Annual Review of Economics*, 11:635–663.

Daniele, F. and Stüber, H. (2023). The Micro-Origins of Business Cycles: Evidence from German Metropolitan Areas. *Review of Economics and Statistics*, 105(1):70–85.

De Grauwe, P. (2020). *Economics of the monetary union*. Oxford University Press.

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., and Treibich, T. (2019). Debunking the granular origins of aggregate fluctuations: From real business cycles back to Keynes. *Journal of Evolutionary Economics*, 29(1):67–90.

Ebeke, M. C. H. and Eklou, K. M. (2017). The granular origins of macroeconomic fluctuations in Europe. *IMF Working Paper Series*, 17(229).

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.

Galí, J. (2008). Monetary policy, inflation, and the business cycle. *Princeton University Press*.

Gorman, W. M. (1961). On a class of preference fields. *Metroeconomica*, 13(2):53–56.

Hidalgo, C. A. and Hausmann, R. (2009). The building blocks of economic complexity. *Proceedings of the National Academy of Sciences*, 106(26):10570–10575.

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

Joya, O. and Rougier, E. (2019). Do (all) sectoral shocks lead to aggregate volatility? Empirics from a production network perspective. *European Economic Review*, 113:77–107.

Katz, L. (1953). A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43.

Kirman, A. P. (1992). Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2):117–136.

Kuha, J. (2004). AIC and BIC: Comparisons of assumptions and performance. *Sociological Methods & Research*, 33(2):188–229.

Liu, E. (2020). Industrial policies in production networks. *The Quarterly Journal of Economics*, 134(3):1883–1948.

Lucas, R. E. (1977). Understanding business cycles. In *Carnegie-Rochester Conference Series on Public Policy*, volume 5, pages 7–29. Elsevier.

Manelici, I. and Pantea, S. (2021). Industrial policy at work: Evidence from Romania's income tax break for workers in IT. *European Economic Review*, 133:103674.

Miranda-Pinto, J. and Shen, Y. (2019). A granular view of the Australian business cycle. *Economic Record*, 95(311):407–424.

Ozdagli, A. and Weber, M. (2017). Monetary policy through production networks: Evidence from the stock market. Technical report, National Bureau of Economic Research.

Saint-Paul, G. (1993). Productivity growth and the structure of the business cycle. *European Economic Review*, 37(4):861–883.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, 48(1):1–48.

Snowdon, B., Vane, H. R., and Wynarczyk, P. (1994). *A modern guide to macroeconomics: An introduction to competing schools of thought*. Edgar Elgar Publishing.

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

Summers, L. H. (1986). Some skeptical observations on real business cycle theory. *Federal Reserve Bank of Minneapolis Quarterly Review*, Fall:23–27.

Sutton, J. (1997). Gibrat's legacy. *Journal of Economic Literature*, 19:40–59.

Syverson, C. (2004). Market structure and productivity: A concrete example. *Journal of Political Economy*, 112(6):1181–1222.

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.

Törnqvist, L. (1936). The Bank of Finland's consumption price index. *Bank of Finland Monthly Bulletin*, 10:1–18.

Wagner, J. and Weche, J. P. (2020). On the granularity of the German economy–first evidence from the top 100 companies panel database. *Applied Economics Letters*, 27(21):1768–1771.

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 Generic Residuals and Simple Linear Regression

Consider a generic weight w_i for industry i. For our thwo specifications, this weight could be the Domar weight for the granular case, S_i/Y and $S_i/Y \cdot \theta_i$ for the interaction specification. We assume that the TFP growth rate of the total economy is a weighted sum of idiosyncratic productivity shocks to N industries which implies

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

where $\pi_{i,t}$ is the productivity of firm i in period t, $w_{i,t}$ its weight at t and TFP_t is the Total Factor Productivity in period t. We assume that the idiosyncratic productivity shocks $d\pi_i$ are independent of the weights w_i which, for the granular case, would imply Gibrat's law. In the following, we call this condition *Generalized Gibrat's Law.* Assume that that productivity shocks are Hicks-neutral and that GDP growth y_t is proportional to TFP growth. For these assumptions, the cumulative idiosyncratic shocks multiplied by a factor ψ yield the GDP growth rate y_t which is given by

$$y_t = \psi \cdot \sum_{i=1}^{N} w_{i,t} d\pi_{i,t}.$$

It follows that the variance of GDP growth rates is given by

$$var(y) = \psi^2 \cdot \sigma_\pi^2 \cdot H(N),$$

where σ_π^2 is the variance of productivity growth rates and H the Hirschman-Herfindahl index (HHI) of the weights w. For generic weights w, $H(N)$ is therefore just the sum of N squared weights, that is,

$$H(N) = \sum_{i=1}^{N} (w_i)^2.$$

Consider now a generic residual $\mathcal{X}_t$ based purely on observables given by

$$\mathcal{X}_t = \sum_{i=1}^{K} w_{i,t} d\pi_{i,t},$$

where K denotes the K largest industry weights and assuming that we identified the correct productivity differences $d\pi_{i,t}$ for all K firms. If we assume the Generalized Gibrat's law to hold, that is, productivity growth rates to be independent of weights, σ_π has to be independent of $H(K)$ as it is defined by the weight distribution of industries. Also, this implies that $E[\sigma_\pi(K)] = \sigma_\pi(N)$, that is, the variance of productivity growth rates of any subsample of K industries and $K > 1$ is expectationally equivalent to the variance of productivity growth rates for the whole sample. The variance of generic residual $\mathcal{X}$ for the K largest industries in terms of weights is given by

$$var(\mathcal{X}^K) = \sigma_\pi^2(K) \cdot H(K).$$

Taking expectations to define the expected induced variance yields

$$\begin{aligned} E[var(\mathcal{X}^K)] &= E[\sigma_\pi^2(K)] \cdot H(K) \\ &= \sigma_\pi^2 \cdot H(K), \end{aligned}$$

as $H(K)$ is independent of σ_π by the generalized version of Gibrat's law.

Consider a simple linear regression with

$$y_t = \beta_0 + \beta_1 \cdot \mathcal{X}_t + \epsilon_t.$$

An OLS estimation would yield for this specification the estimated $\hat{\beta}_1$

$$\hat{\beta}_1 = \frac{cov(y, \mathcal{X})}{var(\mathcal{X})}.$$

Rewriting the initial specification gives

$$y_t = \psi \cdot \mathcal{X}_t^N = \psi \cdot (\mathcal{X}_t^K + \mathcal{X}_t^{N-K}),$$

for all considered periods t, where $\mathcal{X}^N$ denotes the generic residual for all N entities, $\mathcal{X}^K$ for the largest K ones and $\mathcal{X}^{N-K}$ for the remaining smallest $N-K$ industries or firms. From this definition, $\hat{\beta}_1$ can be expressed as

$$\hat{\beta}_1 = \frac{\sum_t^T (\mathcal{X}_t^K - \bar{\mathcal{X}}^K) \cdot (y_t - \bar{y})}{\sum_t^T (\mathcal{X}_t^K - \bar{\mathcal{X}}^K)^2.}$$

From $\psi \cdot (\mathcal{X}_t^K + \mathcal{X}_t^{N-K}) = y_t$ and therefore $\psi \cdot (\bar{\mathcal{X}}^K + \bar{\mathcal{X}}^{N-K}) = \bar{y}$, we get

$$= \frac{\sum_t^T (\mathcal{X}_t^K - \bar{\mathcal{X}}^K) \cdot \psi \cdot ((\mathcal{X}_t^K + \mathcal{X}_t^{N-K}) - (\bar{\mathcal{X}}^K + \bar{\mathcal{X}}^{N-K}))}{\sum_t^T (\mathcal{X}_t^K - \bar{\mathcal{X}}^K)^2}.$$

Rearranging gives

$$\begin{aligned}
&= \psi \cdot \frac{\sum_t^T (\mathcal{X}_t^K - \bar{\mathcal{X}}^K) \cdot ((\mathcal{X}_t^K - \bar{\mathcal{X}}^K) + (\mathcal{X}_t^{N-K} - \bar{\mathcal{X}}^{N-K}))}{\sum_t^T (\mathcal{X}_t^K - \bar{\mathcal{X}}^K)^2} \\
&= \psi \cdot \frac{var(\mathcal{X}^K) + cov(\mathcal{X}^K, \mathcal{X}^{N-K})}{var(\mathcal{X}^K)} \\
&= \psi \cdot (1 + \frac{cov(\mathcal{X}^K, \mathcal{X}^{N-K})}{var(\mathcal{X}^K)}).
\end{aligned}$$

Since we assume that all sectors draw independently from productivity distributions with equal first and second moments, $cov(\mathcal{X}^K, \mathcal{X}^{N-K})$ and $var(\mathcal{X}^K)$ are independent. Also, by the same independence assumption, it has to hold that $E[cov(\mathcal{X}^K, \mathcal{X}^{N-K})] = 0$. Taking expectations yields

$$\begin{aligned}
E[\hat{\beta}_1] &= E[\psi \cdot (1 + \frac{cov(\mathcal{X}^K, \mathcal{X}^{N-K})}{var(\mathcal{X}^K)}) \\
&= \psi + \psi \cdot E[\frac{cov(\mathcal{X}^K, \mathcal{X}^{N-K})}{var(\mathcal{X}^K)}]
\end{aligned}$$

and by invoking the independence of $cov(\mathcal{X}^K, \mathcal{X}^{N-K})$ and $var(\mathcal{X}^K)$

$$= \psi + \psi \cdot (E[cov(\mathcal{X}^K, \mathcal{X}^{N-K})] \cdot E[var(\mathcal{X}^K)^{-1}])$$

which yields with $E[cov(\mathcal{X}^K, \mathcal{X}^{N-K})] = 0$ again from independence

$$= \psi.$$

Consider now the coefficient of determination for the given simple linear regression with

$$R^2 = \hat{\beta}_1^2 \cdot \frac{var(\mathcal{X}^K)}{var(y)}.$$

Taking expectations yields

$$E[R^2] = E[\hat{\beta}_1^2 \cdot \frac{var(\mathcal{X}^K)}{var(y)}]$$

and since $var(y)$ is observable, $E[var(y)] = var(y)$ and by invoking that $\hat{\beta}_1^2$ and $var(\mathcal{X}^K)$ are uncorrelated,[24]

$$= var(y)^{-1} \cdot E[\hat{\beta}_1^2] \cdot E[var(\mathcal{X}^K)].$$

Substituting gives

$$\begin{aligned}
&= \frac{\psi^2 \cdot \sigma_\pi^2 \cdot H(K)}{\psi^2 \cdot \sigma_\pi^2 \cdot H(N)} \\
&= \frac{H(K)}{H(N)} \\
&= \sum_{i=1}^{K} \frac{(w_i)^2}{\sum_{j=1}^{N} (w_j)^2}.
\end{aligned}$$

[24]For $\hat{\beta}_1^2$ and $var(\mathcal{X}^K)$ to be uncorrelated, we need to show that the correlation of $var(\mathcal{X}^K)$ and $cov(\mathcal{X}^K, \mathcal{X}^{N-K})/var(\mathcal{X}^K)$ is zero. Since $var(\mathcal{X}^K)$ and $cov(\mathcal{X}^K, \mathcal{X}^{N-K})$ are independent, $cov(cov(\mathcal{X}^K, \mathcal{X}^{N-K})/var(\mathcal{X}^K), var(\mathcal{X}^K)) = E[cov(\mathcal{X}^K, \mathcal{X}^{N-K})] \cdot (1 - E[var(\mathcal{X}^K)] \cdot E[1/var(\mathcal{X}^K)])$ and by $E[cov(\mathcal{X}^K, \mathcal{X}^{N-K})] = 0$, we thus get $cov(cov(\mathcal{X}^K, \mathcal{X}^{N-K})/var(\mathcal{X}^K), var(\mathcal{X}^K)) = 0$, i.e., β_1^2 and $var(\mathcal{X}^K)$ being uncorrelated.

Thus, under the generalized version of Gibrat's law and assuming that we correctly identified the idiosyncratic productivity shocks to each firm or industry, the expected R^2 for the granular residual for K industries should be given by the ratio of the weight Herfindahl for the K largest firms or industries to the weight Herfindahl for the total economy. To illustrate this result, we perform a simple Monte Carlo simulation. The weight distribution is given by a power-law with tail exponent $\alpha = 1.9$. There are $N = 100$ industries in total and the idiosyncratic productivity shocks are drawn from a Laplacian distribution with location parameter $\mu = 0$ and scale parameter $\sigma = 0.05$. We consider 100 time periods each and repeat the simulation 50 times with different realizations for the idiosyncratic shocks. Figure 5 shows the mean R^2 for these 50 simulation runs and a simple linear regression model as presented above as well as the expected R^2 by the results below. This is given by $\frac{H(K)}{H(N)} = \frac{\zeta_K(\frac{2}{\alpha})}{\zeta_N(\frac{2}{\alpha})}$ where $\zeta_N(\alpha)$ is a truncated Zeta function at N with $\zeta_N(\alpha) = \sum_{k=1}^{N} k^{-\alpha}$ and with $\alpha = 1.9$.

As can easily be seen, the expected R^2 from the ratio of Herfindahls is remarkably close to the mean R^2 from the simulation exercise. Notice that our results are not materially sensitive to the specific choice of the distribution for idiosyncratic productivity shocks. The results are robust for several different parameter values of μ, σ and κ for a (symmetrical) Subbotin distribution that includes the outlined Laplacian as a special case for $\kappa = 1$. The approximation only breaks down for very small shape parameters $\kappa < 0.3$, for which the variance and all higher moments explode such that 100 different realizations of idiosyncratic productivity shocks no longer warrant convergence of the mean to the expected value by the law of large numbers. The result is also robust to various specifications of the weight distribution (measured by the tail exponent α).[25] Of course, this is only the expected or mean functional form of the R^2. Indeed, depending on the specific realization of draws from the distribution for idiosyncratic shocks, the realized R^2 might differ, as Figure 6 shows. Given a sufficiently heterogeneous weight distribution, however, the R^2 plot cannot deviate too much from the expected R^2, as shown in Figure 5, since it should increase almost monotonically, thereby excluding large downturns in the empirical R^2 plot.

Finally, consider the boundary condition Gabaix (2011) gives for granular firms to

[25] Material available upon request.

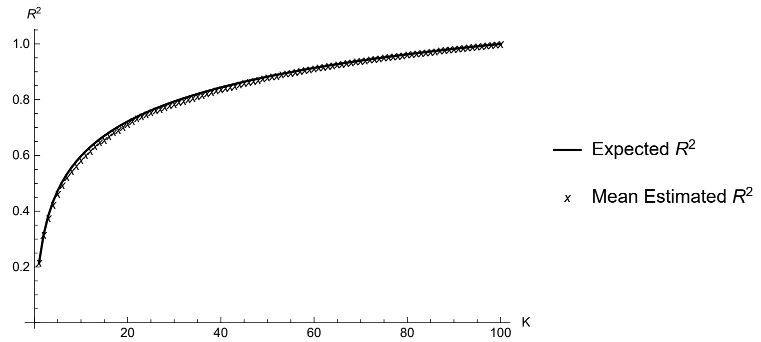

Figure 5: Mean R^2 for Monte Carlo simulation excercise and expected R^2 derived from the weight distribution.

exist - they have to be distributed by a power law with a tail exponent α at most equal to 2.[26] This implies that if all K weights are granular, $H(K) \geq \zeta_K(1)$. This, however, is just the harmonic series and thus, for any given $H(N) = H$ for the total economy, the ratio $\frac{H(K)}{H} \approx \frac{1}{H(N)}(ln(K) + \gamma)$, where γ is the Euler-Mascheroni constant. Thus the R^2 plot can be expected to increase at least in proportion to the natural logarithm of K.

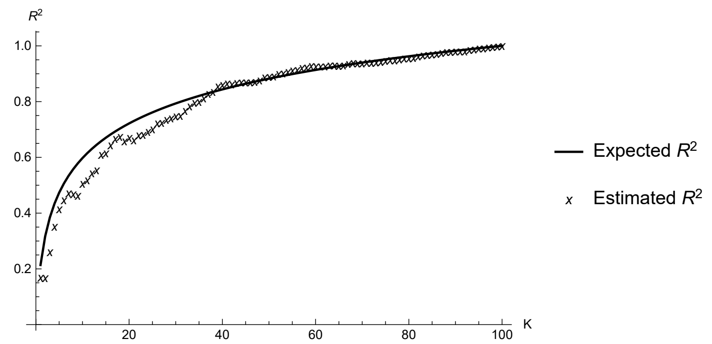

Figure 6: R^2 for a specific realization of idiosyncratic productivity shocks and expected R^2 derived from the weight distribution.

B Winsorizing, β estimates and R^2

Consider the same framework as in Appendix A. However, assume now that the productivity growth rate distribution of the K largest industries was winsorized at the pth and $1 - pth$ quantile. We define $\nu(p) = \tilde{\sigma}^2_\pi(K, p)/\sigma^2_\pi(K)$, that is, ν gives the ratio of variances for the p-winsorized and initial distribution of productivity growth rates for the K largest industries with $\nu(p) \in (0, 1)$ for $p \in (0, 0.5)$. It has to hold that $\lim_{p \to 0} \nu(p) = 1$ and

[26]See Online Appendix E for further details.

$\lim_{p \to 0.5} \nu(p) = 0$. If we assume the generalized version of Gibrat's law to hold, it follows that $E[\sigma_\pi^2(K)] = E[\sigma_\pi^2(N)] = \sigma_\pi^2$ and therefore $E[\tilde{\sigma}_\pi^2(K,p)] = \nu(p) \cdot \sigma_\pi^2$.

The expected, winsorized variance of $\tilde{\mathcal{X}}^K(p)$ of the generic residual defined in the above section is then

$$E[var(\tilde{\mathcal{X}}^K(p))] = \nu(p) \cdot \sigma_\pi^2 \cdot H(K).$$

For analytical tractability, we assume that the winsorized generic residual is affected by some time-invariant scalar $\delta^K(p) \in \mathbb{R}$ such that $\mathcal{X}_t^K = \tilde{\mathcal{X}}_t^K(p) + \delta^K(p) \forall t$, which might vary with the winsorizing probability p and with the number of considered entities K. Following the discussion in Appendix A, we thus get

$$\begin{aligned} y_t &= \psi \cdot \mathcal{X}_t^N \\ &= \psi \cdot (\mathcal{X}_t^K + \mathcal{X}_t^{N-K}) \\ &= \psi \cdot (\tilde{\mathcal{X}}_t^K(p) + \delta^K(p) + \mathcal{X}_t^{N-K}) \end{aligned}$$

and analogously for the time-series averages, where $\bar{\mathcal{X}}^K(p)$ is the time-series average of the generic residual winsorized at p

$$\bar{y} = \psi \cdot (\bar{\mathcal{X}}^K(p) + \delta^K(p) + \bar{\mathcal{X}}^{N-K}).$$

For the regression coefficient, it follows that

$$\begin{aligned} \hat{\beta}_1(p) &= \frac{\sum_t^T (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p)) \cdot (y_t - \bar{y})}{\sum_t^T (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p))^2} \\ &= \Bigg[\sum_t^T (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p)) \cdot \psi \cdot ((\tilde{\mathcal{X}}_t^K(p) + \delta^K(p) + \mathcal{X}_t^{N-K}) - (\bar{\mathcal{X}}^K(p) + \delta^K(p) + \\ &\bar{\mathcal{X}}^{N-K}))\Bigg] / \Bigg[\sum_t^T (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p))^2 \Bigg] \\ &= \psi \cdot \frac{\sum_t^T (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p)) \cdot (\tilde{\mathcal{X}}_t^K(p) + \mathcal{X}_t^{N-K} - \bar{\mathcal{X}}^K(p) + \bar{\mathcal{X}}^{N-K})}{\sum_t^T (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p))^2} \end{aligned}$$

and by rearranging

$$
\begin{aligned}
&= \psi \cdot \frac{\sum_t^T (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p)) \cdot (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p)) + (\mathcal{X}_t^{N-K} - \bar{\mathcal{X}}^{N-K}))}{\sum_t^T (\tilde{\mathcal{X}}_t^K(p) - \bar{\mathcal{X}}^K(p))^2} \\
&= \psi \cdot \frac{var(\tilde{\mathcal{X}}^K(p)) + cov(\tilde{\mathcal{X}}^K(p), \mathcal{X}^{N-K})}{var(\tilde{\mathcal{X}}^K(p))} \\
&= \psi \cdot (1 + \frac{cov(\tilde{\mathcal{X}}^K(p), \mathcal{X}^{N-K})}{var(\tilde{\mathcal{X}}^K(p))}).
\end{aligned}
$$

If the winsorizing procedure did not generate spurious dependence between industries, the terms $cov(\tilde{\mathcal{X}}^K(p), \mathcal{X}^{N-K})$ and $var(\tilde{\mathcal{X}}^K(p))$ remain independent and $E[cov(\tilde{\mathcal{X}}^K(p), \mathcal{X}^{N-K})] = 0$. Thus, as shown in Appendix A, the estimate for the β_1 coefficient in a simple linear regression model remains unbiased, and $E[\hat{\beta}_1(K)|p = 0] = E[\hat{\beta}_1(K)|p] = \psi$. However, the expected R^2 as the explanatory power for K industries should for any degree of winsorizing p be biased downwards. This follows purely from the fact that the variance of the winsorized generic residual is downwards biased in expectations. To see this, consider $E[R^2(K)|p]$ as

$$
\begin{aligned}
E[R^2(K)|p] &= \psi^2 \frac{E[var(\tilde{\mathcal{X}}^K(p)]}{var(y)} \\
&= \frac{\psi^2 \cdot \nu(p) \cdot \sigma_\pi^2 \cdot H(K)}{\psi^2 \cdot \sigma_\pi^2 \cdot H(N)} \\
&= \nu(p) \cdot \frac{H(K)}{H(N)} \\
&= \nu(p) \cdot \sum_{i=1}^{K} \frac{(w_i)^2}{\sum_{j=1}^N (w_j)^2}
\end{aligned}
$$

and that $\nu(p) < 1$ for all $p \in (0, 0.5)$. It follows that, compared to the benchmark expected R^2 without winsorizing, the estimate is downwards biased by $\Lambda(p)$ percent given by

$$
\begin{aligned}
\Lambda(p) &= \frac{E[R^2(K)|p = 0] - E[R^2(K)|p]}{E[R^2(K)|p = 0]} \\
&= 1 - \nu(p).
\end{aligned}
$$

The degree of bias crucially depends on the kurtosis of the underlying growth rate distribution, as we illustrate in Online Appendix C.