Forum for Macroeconomics and Macroeconomic Policies

IMK Macroeconomic Policy Institute

The IMK is an institute of the Hans-Böckler-Stiftung

FMM WORKING PAPER

No. 72 • October 2021 • Hans-Böckler-Stiftung

THE NETWORK ORIGINS OF AGGREGATE FLUCTUATIONS: A DEMAND-SIDE APPROACH

Emanuele Citera[1], Shyam Gouri Suresh[2], Mark Setterfield[3]

ABSTRACT

We construct a model of cyclical growth with agent-based features designed to study the network origins of aggregate fluctuations from a demand-side perspective. In our model, aggregate fluctuations result from variations in investment behavior at firm level motivated by endogenously-generated changes in `animal spirits' or the state of long run expectations (SOLE). In addition to being influenced by their own economic conditions, firms pay attention to the performance of first-degree network neighbours, weighted (to differing degrees) by the centrality of these neighbours in the network, when revising their SOLE. This allows us to analyze the effects of the centrality of linked network neighbours on the amplitude of aggregate fluctuations. We show that the amplitude of fluctuations is significantly affected by the eigenvector centrality, and the weight attached to the eigenvector centrality, of linked network neighbours. The dispersion of this effect about its mean is shown to be similarly important, resulting in the possibility that network properties can result in `great moderations' giving way to sudden increases in the volatility of aggregate economic performance.

[1] Department of Economics, New School for Social Research, New York, NY 10003 USA.
email address: emanuelecitera@newschool.edu

[2] Department of Economics, Davidson College, Davidson, NC 28035 USA, and Department of Economics, FLAME University, Maharashtra, India. Email address: shgourisuresh@davidson.edu

[3] Department of Economics, New School for Social Research, New York,
email mark.setterfield@newschool.edu; FMM fellow.

The network origins of aggregate fluctuations: a demand-side approach

Emanuele Citera[*] Shyam Gouri Suresh[†] Mark Setterfield[‡]

October 8, 2021

Abstract

We construct a model of cyclical growth with agent-based features designed to study the network origins of aggregate fluctuations from a demand-side perspective. In our model, aggregate fluctuations result from variations in investment behavior at firm level motivated by endogenously-generated changes in 'animal spirits' or the state of long run expectations (SOLE). In addition to being influenced by their own economic conditions, firms pay attention to the performance of first-degree network neighbours, weighted (to differing degrees) by the centrality of these neighbours in the network, when revising their SOLE. This allows us to analyze the effects of the centrality of linked network neighbours on the amplitude of aggregate fluctuations. We show that the amplitude of fluctuations is significantly affected by the eigenvector centrality, and the weight attached to the eigenvector centrality, of linked network neighbours. The dispersion of this effect about its mean is shown to be similarly important, resulting in the possibility that network properties can result in 'great moderations' giving way to sudden increases in the volatility of aggregate economic performance.

JEL codes: C63, E12, E32, E37, O41

Keywords: Aggregate fluctuations, cyclical growth, animal spirits, state of long run expectations, agent-based model, random network, preferential attachment, small world.

[*]Department of Economics, New School for Social Research, New York, NY 10003 USA. Email address: emanuelecitera@newschool.edu

[†]Department of Economics, Davidson College, Davidson, NC 28035 USA, and Department of Economics, FLAME University, Maharashtra, India. Email address: shgourisuresh@davidson.edu

[‡]Department of Economics, New School for Social Research, New York NY 10003 USA. Email address: mark.setterfield@newschool.edu

1. Introduction

Despite the individualistic orientation of much economic theory, it is an undeniable fact that humans are social beings and that human interaction is a central feature of social reality. This interaction can be more or less superficial. On one hand, individuals can (in principle) engage in collective activities such as production, or interact with one another in processes such as exchange, without this affecting their 'internal structure' (their beliefs, objectives, preferences, aspirations, etc.). This is the basis on which mainstream economics conventionally treats social interaction. On the other hand, exchanges of information and opinions may well shape goals, ambitions and the like, and so the internal structure of individuals itself may depend on social interaction. The resulting 'socialization' or 'acculturation' is frequently emphasized in social sciences other than economics. Whether more or less superficial, however, an important feature of social interaction is that it is structured by *social networks* – social structures made up of agents (whether individuals or organizations) who are linked to one another by dyadic ties that represent some form of inter-personal (or inter-organizational) 'connection', such as friendship, a trade relationship, or a shared interest in some particular outcome or objective.

Networks so-defined matter because of the possibility that they affect behaviour and so have a bearing on observed economic and social outcomes (Jackson, 2014; Jackson et al., 2017). Structural features of networks, such as their density (the average number of connections to other agents of each agent in the network, or the *degree* of each agent), homophily (the extent to which agents link to other agents who are similar to themselves), and clustering (the extent to which the network neighbours to which any agent is linked are, themselves, linked to one another) can affect, for example, the diffusion of information and the extent to which a society displays patterns of segregation. The potential significance of network structure extends to its effects on macroeconomic outcomes. For example, Acemoglu et al. (2012)

argue that aggregate fluctuations in output can be traced to the structure of the networks that link individual firms. One important feature of the analysis in Acemoglu et al. (2012), however, is that it is informed by real business cycle (RBC) analysis and is therefore focused exclusively on supply-side linkages between firms that belong to production networks. In this paper, we explore the network origins of aggregate fluctuations from a demand-side perspective, building on the model of Gouri Suresh and Setterfield (2015), in which individual firms invest based on their 'animal spirits' or state of long-run expectations (SOLE) that are affected, in part, by the experience of other firms to whose fortunes they pay attention.

In the model developed in this paper, a network edge or link indicates two firms that pay attention to one other when computing their SOLE.[1] Since the link does not represent a physical flow, two firms could be linked for any of several different reasons: they could be headquartered in the same location; they could be rivals in the same sector; they could share some of the same members on their boards of directors; their CEOs could be connected socially; and/or they could be linked through input-output connections, *inter alia*. The relevant network features in our model are therefore those that determine patterns of information flows and information aggregation. While aggregating information, an individual firm will recognize that not all of its connections are equally important — some connections are likely to be more informative because these connections themselves either have more connections or are connected to other more-informative sources. Thus, the importance of a particular connection depends on the importance of that connection's connections, leading to a recursive definition of importance, much like the original formulation of the 'beauty-contest' setting in Keynes (1936). As will become clear, the upshot of these considerations is that in our model, the *centrality* of a linked network neighbour is the feature of network architecture that is of particular importance for our analysis of the network origins of aggregate fluctuations.

[1] As will become clear, our analysis is based on undirected networks, so any pair of linked firms influence one another.

The remainder of the paper is organized as follows. In section 2, we outline how concern with the network origins of aggregate fluctuations has arisen from concern with the 'granular' origins of aggregate fluctuations as expressed by Gabaix (2011), and how it also reflects the basic schism in macroeconomics between supply-side and demand-side approaches to the study of major macroeconomic pathologies. In section 3, we summarize the structure of the original model developed by Gouri Suresh and Setterfield (2015) on which we build, while in section 4, we extend this model to incorporate network connections between individual firms. Section 5 presents and analyzes the results of simulation experiments based on the model developed in the preceding section. Finally, section 6 concludes.

2. From the granular origins to the network origins of aggregate fluctuations

A common trope in macroeconomics is that in any economy with a sufficiently large number of firms, idiosyncratic firm-level disturbances will 'wash out' having no effect (or at most, an effect so small as to be negligible) on aggregate outcomes (Lucas, 1977). Following Gabaix (2011), consider a population of n independent firms, each of which is subject to an idiosyncratic shock $\varepsilon \sim N(0,1)$. Define the growth rate of each firm as:

$$\frac{\Delta y_{jt}}{y_{jt-1}} = \sigma \varepsilon_{jt}$$

where y_j is the output of the j^{th} firm and σ is the standard deviation of individual firm output (assumed common to all firms). Aggregate output (y) therefore grows at the rate:

$$g_t = \frac{\Delta y_t}{y_{t-1}} = \frac{\sum_{j=1}^{n} \Delta y_{jt}}{y_{t-1}} = \sum_{j=1}^{n} \sigma \rho_{jt-1} \varepsilon_{jt} \tag{1}$$

where $\rho_{jt-1} = \frac{y_{jt-1}}{y_{t-1}}$ is the j^{th} firm's share of aggregate output.

Now consider the standard deviation of aggregate output growth, σ_g. Drawing on the result in (1), this can be written as:

$$\sigma_g = \left[\sigma^2 \sum_{j=1}^{n} (\rho_{jt-1})^2 \mathrm{Var}(\varepsilon_{jt})\right]^{1/2} = \sigma \left[\sum_{j=1}^{n} (\rho_{jt-1})^2\right]^{1/2} = \sigma h \tag{2}$$

since $\mathrm{Var}(\varepsilon_{jt}) = 1$ by hypothesis and where $h = \left[\sum_{j=1}^{n} (\rho_{jt-1})^2\right]^{1/2}$ is the square root of the economy-wide Herfindahl index.

Suppose now that all firms are of identical size, so that $y_{jt-1} = \frac{y_{t-1}}{n}$, $\forall j$. Then:

$$\rho_{jt-1} = \frac{\frac{y_{t-1}}{n}}{y_{t-1}} = \frac{1}{n}, \forall j\,;$$

so that:

$$h = \left[n\left(\frac{1}{n}\right)^2\right]^{1/2} = \frac{1}{\sqrt{n}}\,.$$

Substituting this last expression for h into equation 2 yields:

$$\sigma_g = \frac{\sigma}{\sqrt{n}} \tag{3}$$

In equation (3), the term $1/\sqrt{n}$ acts as a scaling factor, relating the volatility common to all individual firms (σ) to the volatility of aggregate output growth (σ_g). Clearly, the effect of σ on σ_g diminishes with n and will ultimately become negligible for sufficiently large n.

However, this argument overlooks properties of the 'micro-structure' of an economy that can propagate idiosyncratic shocks to individual firms into non-negligible effects on aggregate outcomes (Stiglitz and Gallegati, 2011, p.8). It assumes that firms are both *independent* (firms are treated as islands and are subject to idiosyncratic shocks that are i.i.d.) and *identical*. This second feature is reflected in the assumptions that $\sigma_j = \sigma$ and $\rho_j = \rho\,, \forall j$ – the latter implying that the size distribution of firms is uniform, which distribution is platykurtic (thin tailed). But Gabaix (2011) shows that idiosyncratic shocks do not 'cancel out' if the size distribution of firms is leptokurtic (fat tailed). In particular, if the size distribution of firms conforms to a power law – as has long been suggested by empirical evidence (Steindl, 1965; Ijiri and Simon, 1977; Axtell, 2001) – then shocks to individual firm output will affect aggregate output by an order of magnitude that is much larger than $1/\sqrt{n}$.[2] Intuitively, this is because of the disproportionate contribution to aggregate output of a small number of very large firms. The result is what Gabaix (2011) refers to as the 'granular origins of aggregate fluctuations'.

One problem with the approach taken by Gabaix (2011) is that it focuses exclusively on 'supply granularity' (Dosi et al., 2019, p.69). According to Dosi et al. (2019), the basic granularity result in Gabaix (2011) is beyond dispute, but the *nature* of the granular shock to which he appeals is not. Hence Gabaix (2011) focuses on granular productivity shocks to individual firms in an explicit effort to provide a microfoundation for real business cycle theory, in which aggregate supply shocks resulting from the uneven pace of technological change are understood to explain the business cycle. Dosi et al. (2019) show that Gabaix's results are sensitive to his treatment of the data and amount to assuming that shocks themselves are *not* leptokurtic, contrary to the empirical evidence. When this assumption is dropped, granular productivity shocks lose importance – but *demand* granularity (proxied by investment growth) is important as a source of aggregate fluctuations. Building on this insight and using the 'Schumpeter meeting Keynes' agent-based model pioneered by Dosi et al. (2010), the authors perfom regressions on simulated data that demonstrate the statistical significance of demand granularity (local investment growth shocks) and the statistical insignificance of supply granularity (local productivity shocks) as sources of aggregate fluctuations. They conclude that the granular origins of aggregate fluctuations are fundamentally Keynesian in nature.

[2]The scaling factor in the case of a power law distribution of firms is $1/\ln n$.

Meanwhile – and equally congruent with the basic insight of Stiglitz and Gallegati (2011) regarding the importance of the economy's micro-structure – Acemoglu et al. (2012, p.1978) argue that the *interconnectedness* of firms (rather than features of their size distribution) can ensure that idiosyncratic shocks affect aggregate outcomes. Allowing for such interconnectedness involves departing from the assumption, used in the derivation of equation (3), that firms are independent. Hence Acemoglu et al. (2012) look beyond the granular origins of aggregate fluctuations into the *network* origins of aggregate fluctuations. Appeal to network theory involves explicitly modeling the connections between firms and thus moving away from the assumption that they can satisfactorily be treated as 'islands'. In particular, Acemoglu et al. (2012) show that the size of aggregate fluctuations is increasing in the asymmetry of network connections (differences in the degree of individual firms) rather than the sparseness of a network (low average degree), and in the extent of second-order connectivity (that is, the degree of the neighbouring firms to which each individual firms is linked). This second result they associate with the capacity of a network to propagate *cascades*.[3]

However, and like Gabaix (2011), Acemoglu et al. (2012) adopt a supply-side perspective rooted in RBC theory. They study networks rooted in input-output connections between firms, and once again neglect the independent role of demand in contributing to aggregate fluctuations. In this paper, inspired by Dosi et al. (2019) and their reformulation of the 'granular origins' story from a demand-side perspective, we explore the network origins of aggregate fluctuations in a demand-led growth model. The approach taken involves building on the model of Gouri Suresh and Setterfield (2015), in which the growth of investment at firm-level depends (in part) on the revision of 'animal spirits' or the 'state of long-run expectations' (SOLE), which is (in turn) influenced by firms' observations of both their own recent performance, and that of others. Gouri Suresh and Setterfield (2015) furnish what is

[3] Cascades arise in networks when behaviour is sequential, and the observed behaviour of others holds sway over the behaviour of any individual decision maker regardless of any private information that the latter possesses.

essentially a 'complete graph' model of firm interaction.[4] In what follows, we depart from this set-up so as to consider a range of incomplete graphs arising from different network structures. Simulations of our model are then used to generate data, the analysis of which allows us to draw conclusions about the influence of network structure on aggregate fluctuations.

3. The basic model

The behaviour of the j^{th} firm is represented as follows:

$$g^i_{jt} = \alpha_{jt} + g_r r^e_{jt} + g_u u^e_{jt} \tag{4}$$

$$g^s_{jt} = s_\pi r_{jt} \tag{5}$$

$$r_{jt} = \frac{\pi u_{jt}}{v} \tag{6}$$

$$g^s_{jt} = g^i_{jt} \tag{7}$$

$$r^e_{jt} = r_{jt-1} \tag{8}$$

$$u^e_{jt} = u_{jt-1} \tag{9}$$

$$\Delta\alpha_{jt} = \alpha([u_{jt-1} - u^e_{jt-1}], \Delta\bar{u}_{t-1}) \tag{10}$$

where g^i and g^s are the actual rate of accumulation and the rate of accumulation consistent with investment being equal to savings, respectively, α denotes the SOLE, r is the rate of profits, u is the rate of capacity utilization, π is the profit share of income, v is the (fixed) full capacity capital-output ratio, and e superscripts denote the expected values

[4]In graph theory, a complete graph is one in which every pair of distinct vertices (the 'nodes' or agents in a network) is connected by a unique edge, and these connections or edges are undirected. This means that the connections 'run both ways'. For example, an undirected connection between two parties might involve each party sharing information with the other. A directed connection, on the other hand, would be asymmetric – only one of the two parties would receive information form the other.

of variables. This system of equations essentially describes an agent-based version of an otherwise-standard Kaleckian growth model. The only significant modification is that the SOLE of each firm is endogenous: both to its own past performance (following Kregel (1976), the j^{th} firm's SOLE responds positively or negatively to disappointment of its own short-term expectations); and to that of its 'linked network neighbours' ($\Delta\bar{u}$ representing the change in the average rate of capacity utilization of the other firms to which firm j pays attention). Note that apart from their interaction via revision of the SOLE, the model does little to relax the treatment of firms as 'islands' that was intrinsic to the derivation of equation (3). Although agent interaction is a signal feature of agent-based modelling, this parsimony is deliberate: following Dibble (2006, p.1538), it is designed to facilitate understanding and interpretation of the model's workings and its simulated output. For this reason, the model is best thought of as an 'aggregate structural model with agent-based features' rather than a traditional agent-based model (ABM).

Upon substitution, the model can be expressed in terms of the recursive interaction of the three-equation system:

$$g^i_{jt} = \alpha_{jt} + \left(g_u + \frac{g_r \pi}{v}\right) u_{jt-1} \tag{11}$$

$$u_{jt} = \frac{v}{s_\pi \pi} g^i_{jt} \tag{12}$$

$$\Delta\alpha_{jt} = \alpha(\Delta u_{jt-1}, \Delta\bar{u}_{t-1}) \tag{13}$$

where:

$$\Delta u_{jt-1} = u_{jt-1} - u_{jt-2} = u_{jt-1} - u^e_{jt-1}$$

Given the rate of change of firm j's own rate of capacity utilization in the previous period, together with the rate of change of the average utilization rate among its linked network neighbours in the previous period, the change in the SOLE (from the previous period to the current period) is determined in equation (13). The current SOLE and the rate of capacity utilization from the previous period are then used to determine the current rate of growth (equation (11)), which then determines the current rate of capacity utilization (in equation (12)). The model then moves forward one period, and the process just described begins again. Note that because u is a bounded variable, it follows from equation (12) that there exists a maximum rate of growth $g^i_{max} = \frac{s_\pi \pi}{v}$ for $u_j = 1$.[5] We therefore add to the system (11) – (13) the equation:

$$g^a_{jt} = \min\left(g^i_{jt}, \frac{s_\pi \pi}{v}\right)$$

and replace g^i_{jt} with g^a_{jt} in equation (12) for purposes of simulation, in order to ensure that the model always remains consistent with its logical bounds.

Table 1 details the setting and revision of the SOLE summarized by the implicit function in (13). The first row rules out negative values of the SOLE;[6] the second row sets α_j to the value that satisfies $g^i_{jt} = g^i_{max}$ whenever $g^a_{jt} = g^i_{max}$. The remaining rows describe how the SOLE is otherwise revised from period to period. κ denotes the 'degree of isolation' (the extent to which firm j pays less or more attention to the outcomes experienced by its linked network neighbours when revsiing its SOLE), c is a constant conventional threshold value for determining when recent economic performance should be considered particularly encouraging or discouraging, ε is a constant representing the normal degree to which the SOLE is revised between any two periods, and ϕ denotes the number of periods over which firms conventionally assess their ability to expand in a manner that involves raising capacity

[5]Consistent with the conditions for the existence and stability of a positive steady-state rate of growth in the Kaleckian growth model, we set $\alpha_{jt} \geq 0\,, \forall t$. Note, then, that even with $u_{jt-1} = 0$, this ensures that $g^i_{jt} \geq 0$ in equation (11). This is consistent with the logical minimum value of the growth rate that can be derived from equation (12) when $u_j = 0$.

[6]Recall that, as previously noted, a non-negative value of α is one of the conditions for the existence and stability of a positive steady-state rate of growth in the Kaleckian growth model.

utilization towards its maximum possible value.

Table 1: Setting and revising the state of long run expectations (SOLE).

Criteria	Effect on SOLE
$\alpha_{jt} < 0$	$\alpha_{jt} = 0$
$g^a_{jt} = \frac{s_\pi \pi}{v}$	$\alpha_{jt} = \alpha'_{jt} = \frac{s_\pi \pi}{v} - \left(g_u + \frac{g_r \pi}{v}\right)$
$1 - u_{jt-1} < \phi \Delta u_{jt-1}$	$\Delta \alpha_{jt} = -\varepsilon$
$\kappa \Delta u_{jt-1} + (1-\kappa)\Delta \bar{u}_{t-1} \geq c$ *and* $1 - u_{jt-1} \geq \phi \Delta u_{jt-1}$	$\Delta \alpha_{jt} = \varepsilon$
$\kappa \Delta u_{jt-1} + (1-\kappa)\Delta \bar{u}_{t-1} \leq -c$	$\Delta \alpha_{jt} = -\varepsilon$
$\kappa \Delta u_{jt-1} + (1-\kappa)\Delta \bar{u}_{t-1} > -c$ *and* $\kappa \Delta u_{jt-2} + (1-\kappa)\Delta \bar{u}_{t-2} \leq -c$ *and* $1 - u_{jt-1} \geq \phi \Delta u_{jt-1}$	$\Delta \alpha_{jt} = \varepsilon$
$\kappa \Delta u_{jt-1} + (1-\kappa)\Delta \bar{u}_{t-1} < c$ *and* $\kappa \Delta u_{jt-2} + (1-\kappa)\Delta \bar{u}_{t-2} \geq c$	$\Delta \alpha_{jt} = -\varepsilon$

Table 1 essentially comprises a menu of 'if ... then' statements that describes the conventional bases for revising the SOLE. Its contents can be summarized as follows: the j^{th} firm revises its SOLE by comparing actual changes in capacity utilization to a conventional 'normal' rate of change, c, subject to a sustainability criterion.[7] This sustainability criterion takes the form of a heuristic that requires firms to check the feasibility of maintaining the recent rate of increase in their capacity utilization rate, Δu_{jt-1}, over the coming ϕ periods against

[7]Note that in any period t, *none* of the conditions in the first column of table 1 need be satisfied for any given firm j. In such cases, the SOLE remains constant and the firm will begin to converge to a steady state rate of capacity utilization consistent with its now-constant SOLE. This is consistent with the 'ordinary' workings of the principle of effective demand in the presence of a constant SOLE.

currently-available spare capacity, $1 - u_{jt-1}$. If this condition is not satisfied, firms are discouraged and the SOLE is automatically revised downward (as in the third row of table 1).

If the sustainability criterion *is* satisfied, however, then firms next check $\Delta\tilde{u}_{jt-1} = \kappa\Delta u_{jt-1} + (1-\kappa)\Delta\bar{u}_{t-1}$ against the benchmark value c. If $\Delta\tilde{u}_{jt-1} \geq c$ ($\Delta\tilde{u}_{jt-1} \leq -c$), the SOLE is revised upward (downward) – as in rows four and five of table 1. Otherwise, firms evaluate events in period $t-1$ with reference to events one period earlier, establishing whether or not events in period $t-1$ were as good (bad) as they previously appeared to be (in period $t-2$). Specifically, if $\Delta\tilde{u}_{jt-2} \geq c$ whereas $\Delta\tilde{u}_{jt-1} < c$ ($\Delta\tilde{u}_{jt-2} \leq -c$ whereas $\Delta\tilde{u}_{jt-1} > -c$), firms conclude that 'the best (worst) is now over', and in response to this determination, they revise their SOLE downward (upward). Behaviour of this type is the substance of rows six and seven of table 1.

Note that recursive revision of the SOLE in the manner described above imbues our model with the potential to create self-reinforcing booms and slumps in the SOLE of, and hence the rate of accumulation undertaken by, any individual firm, punctuated by turning points that transform booms into slumps (and *vice versa*). In other words, unlike the RBC model of Acemoglu et al. (2012) (which is an impulse-response model, in which the significance of network properties for simulated aggregate fluctuations is reflected in the rate of decay of these fluctuations), ours is a model of (potential) cyclical growth. As will become clear below, the model's parameters can be (and are) set up to generate ongoing fluctuations in the simulated pace of growth. In this context, the significance of network properties for aggregate fluctuations will show up in the *amplitude* of simulated growth cycles.

Finally, because our interest is in aggregate fluctuations, and because our firms are 'islands' except for the interaction among them in the process of revising the SOLE when $\kappa \neq 1$ in table 1, the basic model is completed by the addition of following simple aggregation procedure:

$$K_t = \sum_{j=1}^{n} (1 + g^a_{jt}) K_{jt-1}$$

We can then calculate the economy-wide rate of growth as:

$$g^a_t = \frac{K_t - K_{t-1}}{K_{t-1}}$$

Finally, since the values of v, s_π, and π are common to all firms, the aggregate rate of capacity utilization follows as:

$$u_t = \frac{v}{s_\pi \pi} g^a_t$$

4. The extended model

Our extension of the original model introduces three *undirected* and *unweighted* network structures through which firms are connected to one another: a random graph; a graph based on preferential attachment; and a small world. Even though the original model can be thought of as an undirected random graph, interaction among firms in that model occurs only *indirectly* through a 'blackboard system', in which individual firms' behavior is influenced by shared information about macroeconomic outcomes derived from the 'blackboard' (see, e.g., Wooldridge (2002, pp. 301-309)). The ambition in this paper is to model direct interactions among firms in what is akin to a 'beauty contest' setting (Keynes, 1936, Chapter 12), in which every agent is influenced by the average behaviour of a well-defined peer group in which some peers are more salient than others.[8] The type of system we have in mind is a 'complex adaptive system' (Arthur, 2014, Chapter 1) wherein agents update their internal model parameters in light of realized results. Specifically, agents form expectations

[8] See Davis (2017) for an analysis of the 'beauty contest' in terms of complex and reflexive dynamics.

not only in regard to economic 'fundamentals', that is their own previous rates of capacity utilization u_{jt-1}, but also by paying attention to other agents' investment decisions, as captured by the average rate of capacity utilization among linked neighbouring firms, $\bar{u}_{\Theta jt-1}$. From this process, through individual and aggregate capital formation, an aggregate rate of capacity utilization, u_t, emerges, variations in which give rise to aggregate fluctuations in the economy. Our interest is in the sensitivity of these fluctuations to properties of the network structures that link individual firms and, in particular, the centrality of a firm's linked network neighbours, and the weight that a firm attaches to the centrality of its linked network neighbours.

As previously noted, our concern with centrality arises from its association with a recursive definition of the importance of a firm's linked network neighbours that is akin to Keynes's famous 'beauty contest'. The Bonacich family of centrality measures (Bonacich, 1987) is based on precisely this type of recursive definition, where the extent to which a node's centrality depends on the centrality of its neighbors can be varied based on the parameter β. When $\beta = 0$, a node's relative centrality depends solely on the number of its neighbors and not on the neighbors' centralities. When β reaches its upper limit, the reciprocal of the largest eigenvalue of the adjacency matrix of the network,[9] a node's relative centrality is equal to its eigenvector centrality. The eigenvector centrality of a node is thus indicative of how important (or central) each node is in a manner that takes the entire network structure into consideration to the extent possible.

Eigenvector centrality is therefore a key property of each node in network analysis, and has been found to matter immensely in several different studies relevant to our own. In the field of financial networks and systemic financial failures, for example, Yun et al. (2019) find that PageRank centrality (a variant of eigenvector centrality applicable to directed net-

[9] If β is permitted to be larger than the reciprocal of the largest eigenvalue, the recursive definition of centrality will result in an infinite sum.

works) captures the prevalence of systemic risk better than other well-known systemic risk measures. Markose et al. (2012) study the US credit default swaps market and find that a 'superspreader' tax based on eigenvector centrality can mitigate losses that might arise from systemic fragility. In the field of information diffusion within a network, eigenvector centralization can function as a measure of structural bias in information aggregation (Bienenstock and Bonacich, 2021). Using a dataset of social networks within villages in Indonesia, Alatas et al. (2016) find that households with higher values of eigenvector centrality are significantly better informed about the economic circumstances of other households in their village. In the field of opinion dynamics, Buechel et al. (2015) find that agents with higher eigenvector centrality have greater influence on long-run group opinions. Finally, Jia et al. (2015) study a model where nodes share their opinions over a sequence of issues, and find that the extent of social influence enjoyed by a node asymptotically approaches its value for eigenvector centrality if each node's self-appraisal is based on its relative control over prior issues.

4.1. Measuring and weighting centrality, and its role in firm decision making

In order to take account of centrality, we need to model the influence on firms' investment decisions of the firms to which they are linked. In so doing, we limit our attention to first-order network links, and model firms as paying attention to the weighted average rate of capacity utilization of their linked network neighbours, $\bar{u}_{\Theta_j}$, where Θ_j is the set of all first-order degree linked network neighbours of firm j. Hence with $\bar{u} \equiv \bar{u}_{\Theta j}$, the SOLE reaction function of the j^{th} firm, previously summarized in equation (10), now becomes:

$$\Delta\alpha_j = \alpha(\Delta u_{jt-1}\,, \Delta\bar{u}_{\Theta_{jt-1}}) \tag{14}$$

with $\Delta\bar{u}_{\Theta jt-1}$ replacing $\Delta\bar{u}_{jt-1}$ in the the criteria stated in table 1. When calculating $\bar{u}_{\Theta_{jt-1}}$, we weight the capacity utilization of linked network neighbors by the weighted centrality of these neighbours. In other words, in addition to using the centrality of a linked network neighbour to weight the attention that firm j pays to this neighbour's outcomes, we also allow firm j to attach different weights to centrality when calculating $\bar{u}_{\Theta_j}$. The precise measure of centrality used in this calculation is the eigenvector centrality, x_j, defined as:

$$x_j = \frac{1}{\lambda_1} \sum_{l=1}^{N} a_{jl} x_l \,, \tag{15}$$

where, using j and l to denote two different nodes (firms) in the network, and with a_{jl} each element of the adjacency matrix $\mathbf{A}$, λ_1 denotes the largest eigenvalue of the adjacency matrix $\mathbf{A}$. As previously noted, we use this measure of centrality, as opposed to, for example, the degree centrality,[10] because even if a firm is connected to only a few other firms (thus having a low degree centrality), its neighbors may be important, and therefore the node is important too, giving it a high eigenvector centrality. Accordingly, a firm's eigenvector centrality can be large either because it has many linked neighbours or because it has important linked neighbours, or both. In addition, eigenvector centrality is particularly suited to measuring centrality in undirected graphs, on which our analysis is based (see Newman (2010, Chapter 7)).

Drawing on these considerations, our measure of the weighted average capacity utilization of the first-degree linked network neighbors of firm j, $\bar{u}_{\Theta_j}$, can be stated as:

$$\bar{u}_{\Theta_j} = \frac{\sum\limits_{l \in \Theta_j} u_l x_l^{\omega}}{\sum\limits_{l \in \Theta_j} x_l^{\omega}} \,, \tag{16}$$

where the exponent ω denotes the weight attached to the centrality of linked network neighbours by the j^{th} firm. When $\omega = 0$, $\bar{u}_{\Theta_j}$ is the simple (unweighted) average of linked network

[10] The degree k_j of firm j is calculated as $k_j = \sum_{l \in \Theta_j} a_{jl}$.

neighbours' utilization rates. But as ω increases, the value of $\bar{u}_{\Theta_j}$ increasingly reflects linked neighbours' eigencentralities. Hence when $\omega = 1$, the weights x_l^{ω} are exactly equal to the eigenvector centralities of linked network neighbours, while for values of $\omega > 1$, the weights ascribe disproportionately more (less) importance to linked neighbors with higher (lower) eigencentralities.

4.2. Network structures

Our interest in network structure and, in particular, the importance of the centrality of linked network neighbours in generating aggregate fluctuations, demands that we specify the type of network that is assumed to link individual firms. We consider three basic network architectures: a random network; a preferentially attached network; and a small world. The characteristics of these network structures are outlined in appendix A.

5. Simulation structure and analysis of results

5.1. Simulation details

For each type of network structure, we simulate the impact of x_j and ω on aggregate fluctuations ($\hat{\sigma}_u$) by setting $\kappa = 0.5$, so that firms attach equal weight to changes in their own circumstances and changes in the weighted average rate of capacity utilization of linked neighbours when revising their SOLE. We then vary $\omega \in [0, 4]$ in discrete increments of 0.1. We perform 50 repetitions, thus generating 2,050 (50 repetitions $\times$ 41 values of ω) observations. A new network configuration is generated for each repetition of the model, and this configuration then remains unchanged over the entire range of variation in ω. Finally, every simulation involves 50 firms, and the initial shock to the individual rate of capacity, Δu_{j0}, is set as a draw from $\Delta u_{j0} \sim N(0, \sigma_u), \forall j$. Following Gouri Suresh and Setterfield (2015), we treat the convention c as an exogenous, fixed parameter derived from observed properties of the aggregate rate of capacity utilization u:

$$c = \beta \sigma_u \,, \tag{17}$$

where σ_u is computed as the annual average standard deviation of the seasonally adjusted monthly value of total industry capacity utilization in the United States, 1967–2019 and $0 < \beta < 1$. We then set the remaining parameter values of the model in accordance with those found in Gouri Suresh and Setterfield (2015). The complete set of parameter values is reported in table 2.

Table 2: Parameter values.

Parameter	Value
g_r	0.49
g_u	0.025
π	0.33
v	3
s_π	0.8
u^*	0.8242
g^i_{max}	0.088
$\bar{\alpha}$	0.0075
ϵ	0.00075
ϕ	2
σ_u	0.04075
β	0.09
κ	0.5

5.2. The impact of centrality on aggregate fluctuations

We now turn to analyze the impact of the centrality of a firm's linked network neighbours, and the weight attached to this centrality, on aggregate fluctuations. Table 3 reports the mean, standard deviation (S.D.), and coefficient of variation (C.V.) of the values of the degree and eigenvector centrality for each of the three network structures we use, as computed over all of the actual network configurations generated in the course of our simulations. As is conventional, eigenvector centralities for each network are normalized such that the highest value for eigenvector centrality in the network equals 1. As previously noted, eigenvector centrality is our preferred measure of centrality: we report descriptive statistics for the degree in table 3 to demonstrate the importance of this distinction, as reflected in the observable differences between the means, S.D.s and C.V.s of these two measures of centrality.

Table 3: Average network metrics (standard deviations across simulations in parentheses).

	$\bar{k}_j$			$\bar{x}_j$		
	Mean	S.D.	C.V.	Mean	S.D.	C.V.
Random network	24.501	3.448	0.141	0.766	0.106	0.139
	(0.702)	(0.342)	(0.015)	(0.034)	(0.009)	(0.014)
Preferential attachment	3.880	3.252	0.838	0.252	0.189	0.755
	(0.000)	(0.358)	(0.092)	(0.029)	(0.017)	(0.034)
Small world	4.000	0.607	0.152	0.500	0.200	0.418
	(0.000)	(0.105)	(0.026)	(0.082)	(0.029)	(0.119)

Looking at the descriptive statistics for $\bar{x}_j$ in table 3, the important point to note is that the value of the C.V. reported for the preferentially attached and small world networks is several orders of magnitude larger than that for the random network. This is important because it suggests that the two network structures that are more congruent with features of real-world social networks display much greater variation in centrality among nodes, as compared to the 'baseline' random network – an observation that is consistent with there being fewer more central nodes, that can act as aggregators and disseminators of information, in these networks.[11] As will become clear, this heterogeneity in the centrality of firms has important implications for aggregate volatility in the economy.

Figure 1 then illustrates the effect on the mean amplitude of aggregate fluctuations in capacity utilization ($\hat{\sigma}_u$) of varying the weight (ω) that individual firms attach to the centrality of linked neighbours – and hence the weight they attach to the utilization outcomes of more central linked neighbours – when calculating $\bar{u}_{\Theta j}$ and revising their SOLE. As can be seen from figure 1, in all three network structures, the greater the value of ω, the lower is the amplitude of the cycle. In other words, paying more attention to more central linked network neighbours at the micro level has a stabilizing effect in the aggregate. The extent of this stablizing effect does appear to differ between network structures, being less marked in the random network than in either the preferentially attached or small world networks, where we see much larger declines in the amplitude of the cycle as ω increases. Recall, however, that preferentially attached and small world networks possess features that more closely resemble real-world social networks.

[FIGURE 1 HERE]

The stabilizing effect of centrality and the weight attached to it across all three of the network structures featured in figure 1 is borne out by table 4, which reports the results of linear regressions of the form:

$$\hat{\sigma}_u = \gamma_0 + \gamma_1\omega + \gamma_2\omega^2 + \gamma_3 CV(\bar{x}_j) + \xi \tag{18}$$

where $CV(\bar{x}_j)$ denotes the coefficient of variation of the eigenvector centrality of firm j's

[11] The reader is again referred to the appendix A to this paper for discussion of the characteristics of the three different types of networks on which our analysis rests.

linked network neighbours and where, to capture the seemingly nonlinear relationship between $\hat{\sigma}_u$ and ω in figure 1, we include both ω and ω^2 as regressors. Our data set comprises the 2,050 observations generated by our simulations, with averages of the independent variables in equation (18) calculated over each run. The results reported in table 4 show that network structure has a statistically significant bearing on aggregate fluctuations consistent with the stabilizing effect of centrality that is apparent in figure 1 and discussed above. In all three networks, based on the coefficients for both ω and ω^2, the weight attached to eigenvector centrality exerts a negative and statistically significant impact on $\hat{\sigma}_u$,[12] as does the coefficient of variation of eigenvector centrality $CV(\bar{x}_j)$. This last result suggests that *ceteris paribus*, greater heterogeneity in the centrality of firms within a network – i.e., a greater mixture of firms that are either more or less central to opinion formation – produces a less volatile macroeconomic environment. Together, these results are consistent with the notion that a network featuring relatively few highly central firms, to whose outcomes other firms attach considerable weight when revising their SOLE, will have a stabilizing effect on the aggregate economy.

The substance of our regression results is illustrated in figure 2 which, based on the estimated coefficients reported in table 4, illustrates the dampening effect on the amplitude of aggregate fluctuations of increases in ω for three different values of $CV(\bar{x}_j)$ – the maximum, mean and minimum values derived from our simulations. As can be seen in figure 2, across the range of variation in ω as a whole, increasing the weight attached to centrality reduces the amplitude of fluctuations for any value of $CV(\bar{x}_j)$ regardless of network structure and (consistent with figure 1) to a greater extent in preferentially attached and small world networks than in random networks. Figure 2 also captures the inverse relationship between the size of $CV(\bar{x}_j)$ and the amplitude of aggregate fluctuations. It illustrates how raising (lowering) the value of $CV(\bar{x}_j)$ above (below) its mean value reduces (increases) the amplitude

[12] This claim is borne out by figure 2 which is introduced and discussed below.

of the cycle for any given value of ω.

[FIGURE 2 HERE]

Table 4: Regression results

	RN	*PA*	*SW*
ω	0.001^{*}	-0.122^{***}	-0.050^{***}
	(0.0005)	(0.002)	(0.002)
ω^2	-0.001^{***}	0.020^{***}	0.002^{***}
	(0.0001)	(0.001)	(0.001)
$CV(\bar{x}_j)$	-0.053^{***}	-0.045^{**}	-0.081^{***}
	(0.010)	(0.021)	(0.006)
Constant	0.344^{***}	0.371^{***}	0.377^{***}
	(0.001)	(0.016)	(0.003)
Observations	2,050	2,050	2,050
Adjusted R^2	0.210	0.750	0.708

Notes: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

RN, *PA* and *SW* respectively refers to the random network, the preferentially attached and the small world networks.

Returning now to figure 1, we observe a remarkable difference in the range of the amplitude of aggregate fluctuations (the shaded gray area in the left-hand panels of figure 1), associated with an increase in the standard deviation of the amplitude of aggregate fluctuations, as the weight attached to centrality increases. This is particularly evident in the preferentially attached and small world networks, where increasing ω reduces the mean amplitude of the cycle but also markedly increases the standard deviation around this mean. The values of the skewness of $\hat{\sigma}_u$ (reported in the right-hand panels of figure 1) suggest that the distributions of $\hat{\sigma}_u$ about its mean values are approximately symmetric in both preferentially attached and small world networks and for all values of ω. This means that random draws from these distributions are about as likely as not to produce substantial increases in the amplitude of aggregate fluctuations – which observation is suggestive of environments in which, for suitably high values of ω, 'great moderations' (associated with low mean values of $\hat{\sigma}_u$) can suddenly give way to periods of markedly increased macroeconomic volatility. Indeed judging by the results illustrated in the left-hand panels of figure 1, this transition can occur either by virtue of a decline in the weight attached to centrality, or simply because of the materialization of the 'downside risk' (associated with the high variance of $\hat{\sigma}_u$) inherent in higher values of ω. The second of these two channels implies that we do not need to observe a change in firms' information aggregation processes (associated with a reduction in ω) for volatility to arise from a period of tranquility: for sufficiently high values of ω, the economy is innately vulnerable to the onset of such increased volatility.

6. Conclusions

Motivated by previous studies of the granular and network origins of aggregate fluctuations, and the comparative neglect of the demand side in this literature, this paper explores the network origins of aggregate fluctuations in a demand-led model of cyclical growth. In our model, aggregate fluctuations are caused by variations in investment spending that are, in turn, driven by changes in firms' 'animal spirits' or state of long run expectations (SOLE). The latter are revised with reference to both firms' own past performance and that of linked network neighbours. When reckoning the performance of their neighbours, firms attach more weight to outcomes associated with more central linked network neighbours, and also vary the weight they attach to centrality itself.

Using eigenvector centrality as our preferred measure of the centrality of any given firm, our results suggest that in general, centrality has a stabilizing effect on the aggregate economy. More specifically, we find that when, in the process of revising their SOLE, firms attach more weight to centrality – and hence more weight to the utilization outcomes of more central linked network neighbours – macroeconomic volatility declines. This is especially so when centrality is heterogeneous: the amplitude of aggregate fluctuations falls as the coefficient of variation of eigenvector centrality rises, *ceteris paribus*. In other words, networks in which there are fewer more central firms acting as aggregators and disseminators of information have a stabilizing effect on the aggregate economy, regardless of the weight attached by firms to centrality itself in the process of revising their SOLE. These effects of network structure on volatility are observed consistently across qualitatively different network structures, and are particularly evident in preferentially attached and small world networks that better approximate features of real-world social networks.

Our results are consistent with a growing literature that emphasizes the importance of centrality in the determination of outcomes influenced by social interaction that is structured by network connections (see, for example, Markose et al., 2012; Buechel et al., 2015; Jia et al., 2015; Alatas et al., 2016; Yun et al., 2019). Furthermore, and in keeping with earlier work (see, for example, Dosi et al., 2010; Setterfield and Budd, 2011; Gouri Suresh and Setterfield, 2015; Gibson and Setterfield, 2018; Dosi et al., 2019), our analysis is also suggestive of ways in which typically aggregative demand-led models of growth and fluctuations can be imbued with micro-foundational features that advance our understanding of the dynamics of a demand-led economy.

Appendix A Network structures

A.1 Random network

The random network we adopt is the $G(n, p)$ variant of the Erdős–Rényi model (Erdős and Rényi, 1959). $G(n, p)$ is the ensemble of networks with n vertices and m edges in which each graph G appears with probability:

$$P(G) = p^m (1-p)^{\binom{n}{2}-m} \, .$$

We set $p = 0.5$ – the case where all $2^{\binom{n}{2}}$ graphs on n vertices are chosen with equal probability – and $n = 50$. For small values of n, the degree distribution of a random network is binomial:[13]

$$P(k) = \binom{n-1}{k} p^k (1-p)^{n-1-k} \, ,$$

where p denotes the independent probability that a given k vertex in the graph is connected to $n - 1$ other vertices.

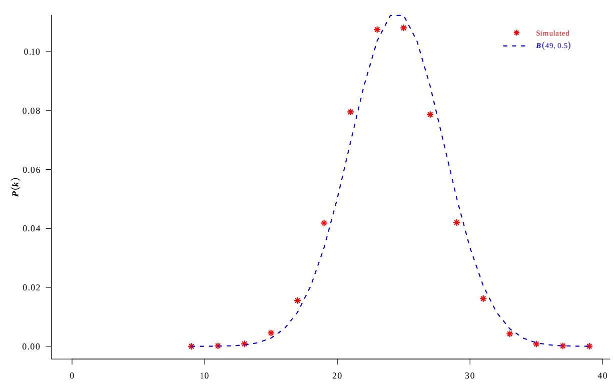

Figure 3: Degree distribution of the random network.

[13]Note that the behavior of random graphs is usually studied for $n \to \infty$. In this case, the degree distribution follows a Poisson distribution. However, for computational reasons, we choose $n = 50$.

As we can see in figure 3, the degree distribution of the simulated random network, represented by the red asterisks, is properly approximated by a binomial distribution $B(49, 0.5)$, denoted by the dashed blue line.

Unfortunately, random networks have severe shortcomings with respect to social networks observed in reality. Firstly, they exhibit almost null clustering coefficients, as opposed to real-world networks where clustering coefficients are typically quite high. Secondly, random networks show no correlation between the degrees of adjacent vertices (given that the edges are randomly generated), whereas in reality such correlation can be observed. Finally, real-world networks have right-skewed degree distributions, meaning that most vertices have low degree while a minority are high-degree 'hubs'. As it is clear from figure 3, random graphs (independently of the precise form of their degree distributions) do not exhibit this property.

In light of the above observations, we use the random network as a benchmark against which to judge the importance of two additional network structures that exhibit more realistic features.

A.2 Preferential attachment

The preferential attachment network we generate is a scale-free network based on the Barabási–Albert (BA) model (Barabási and Albert, 1999). The BA model works as follows.

The algorithm begins with an initial connected network of m_0 nodes and new nodes are added one at a time. Each new node j is connected to $m \leq m_0$ existing nodes with the following probability p_j:

$$p_j = \frac{k_j}{\sum_{l=1}^{n} k_l} .$$

As a result, new nodes 'prefer' to attach to nodes that are already highly linked. The accumulation of links by the latter is thus self-reinforcing, which results in them becoming high degree hubs. The degree distribution of the BA model follows a power law:

$$P(k) \sim k^{-3},\, k \to \infty\,.$$

Figure 4 shows the degree distribution of the preferentially attached network we generated (red asterisks) along with the degree distribution of the BA model (dashed blue line) on a log-log scale. The difference between the two distributions is due to the magnitude of k, which in our model has an average value of $\langle k \rangle = 3.88$. If we consider the degree distribution $P(k) \sim k^{-2}$, however, which is represented by the dashed black line in figure 4, we can see that it better approximates the simulated network structure. In fact, the exponents of real-world networks usually lie within the interval $(2, 3)$.[14]

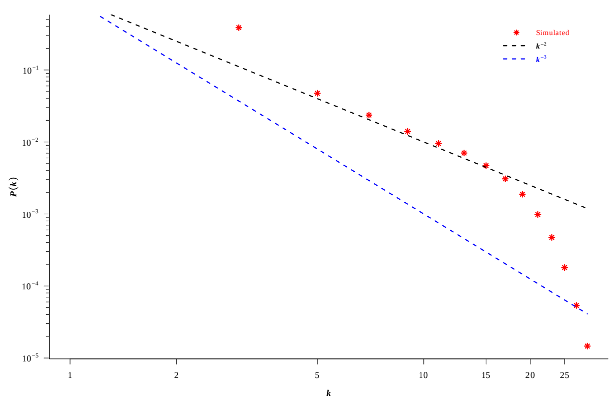

Figure 4: Degree distribution of the preferential attachment (log-log scale).

A.3 Small world

The final network structure we consider is known as a *small world* (Watts and Strogatz, 1998). Small worlds combine two properties of real-world networks – high clustering and short path length – by interpolating between a regular ring lattice and a random graph. The

[14]Note that, to detect a scale-free network, its degree distribution needs to span through at least 2-3 orders of magnitude, meaning that $k_{max} \sim 10^3$. Preferential attachment alone does not guarantee that the degree distribution will be scale free.

algorithm of the model works as follows.

The network is constructed with a ring lattice with n nodes in which every node is connected to an average of $\langle k\rangle$ neighbours, $\langle k\rangle/2$ on each side. Each edge of the lattice is randomly rewired with probability p, such that self-connections and duplicate edges are excluded. By varying p, the network can be transformed from a completely ordered ($p = 0$) to a completely random ($p = 1$) structure.[15] Note that N and $\langle k\rangle$ are chosen so that $N \gg \langle k\rangle \gg \ln N \gg 1$ and thus the graph remains connected ($\langle k\rangle \gg ln(N)$).

The degree distribution of the small world model, in the case of the ring lattice, is represented by the following probability distribution:

$$P(k) = \sum_{n=0}^{f(k,\langle k\rangle)} \binom{\langle k\rangle/2}{n} (1-p)^n p^{\langle k\rangle/2-n} \frac{(p\langle k\rangle/2)^{k-\langle k\rangle/2-n}}{(k-\langle k\rangle/2-n)!} e^{-p\langle k\rangle/2} \,,$$

where $k \geq \langle k\rangle/2$ and $f(k, \langle k\rangle) = \min(k - \langle k\rangle/2, \langle k\rangle/2)$. As figure 5 shows, the shape of the degree distribution is similar to that of a random graph and has a pronounced peak at $k = \langle k\rangle = 4$, which decays exponentially for large $|k - \langle k\rangle|$. Hence a shortcoming of this model is that it generates a degree distribution that is not typical of real-world networks. It does, however, give rise to a high clustering coefficient – unlike the BA model which, on the contrary, fails to produce realistically high clustering.

[15] In our simulation, we set $n = 50$, $p = 0.1$ (which allows for the small world property).

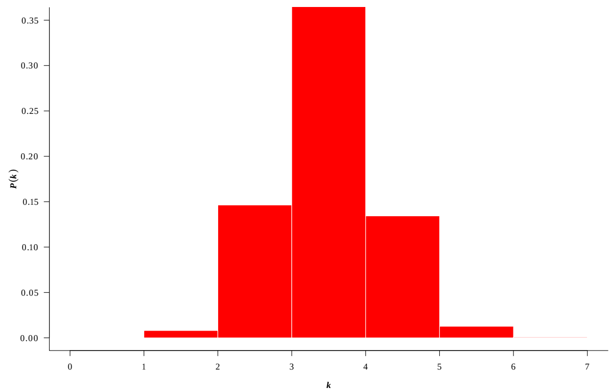

Figure 5: Degree distribution of the small world.

References

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

Alatas, V., A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken (2016). Network structure and the aggregation of information: Theory and evidence from indonesia. *American Economic Review 106*(7), 1663–1704.

Arthur, W. B. (2014). *Complexity and the Economy*. Oxford University Press.

Axtell, R. L. (2001). Zipf distribution of U.S. firm sizes. *Science 293*(5536), 1818–1820.

Barabási, A.-L. and R. Albert (1999). Emergence of scaling in random networks. *Science 286*(5439), 509–512.

Bienenstock, E. J. and P. Bonacich (2021). Eigenvector centralization as a measure of structural bias in information aggregation. *The Journal of Mathematical Sociology 0*(0), 1–19.

Bonacich, P. (1987). Power and centrality: A family of measures. *American Journal of Sociology 92*(5), 1170–1182.

Buechel, B., T. Hellmann, and S. Klössner (2015). Opinion dynamics and wisdom under conformity. *Journal of Economic Dynamics and Control 52*(C), 240–257.

Davis, J. B. (2017). The Continuing Relevance of Keynes's Philosophical Thinking: Reflexivity, Complexity and Uncertainty. *Annals of the Fondazione Luigi Einaudi. An Interdisciplinary Journal of Economics, History and Political Science 51*(1), 55–76.

Dibble, C. (2006). Computational laboratories for spatial agent-based models. In L. Tesfatsion and K. Judd (Eds.), *Handbook of Computational Economics, Vol. 2*, pp. 1511 – 1548. Philadelphia, PA: Elsevier.

Dosi, G., G. Fagiolo, and A. Roventini (2010). Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control 34*(9), 1748–1767.

Dosi, G., M. Napoletano, A. Roventini, and T. Treibich (2019). Debunking the granular origins of aggregate fluctuations: from real business cycles back to Keynes. *Journal of Evolutionary Economics 29*(1), 67–90.

Erdős, P. and A. Rényi (1959). On random graphs. *Publicationes Mathematicae 6*, 290–297.

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

Gibson, B. and M. Setterfield (2018). Real and financial crises in the Keynes-Kalecki structuralist model: an egent-based approach. *Metroeconomica* (forthcoming).

Gouri Suresh, S. and M. Setterfield (2015). Firm performance, macroeconomic conditions, and ‘animal spirits’ in a Post Keynesian model of aggregate fluctuations. *Journal of Post Keynesian Economics 38*(1), 38–63.

Ijiri, Y. and H. Simon (1977). *Skew Distributions and the Sizes of Business Firms*. Amsterdam: North-Holland.

Jackson, M. O. (2014, November). Networks in the understanding of economic behaviors. *Journal of Economic Perspectives 28*(4), 3–22.

Jackson, M. O., B. W. Rogers, and Y. Zenou (2017, March). The economic consequences of social-network structure. *Journal of Economic Literature 55*(1), 49–95.

Jia, P., A. MirTabatabaei, N. E. Friedkin, and F. Bullo (2015). Opinion dynamics and the evolution of social power in influence networks. *SIAM Rev. 57*(3), 367–397.

Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan.

Kregel, J. A. (1976). Economic Methodology in the Face of Uncertainty: The Modelling Methods of Keynes and the Post-Keynesians. *Economic Journal 86*(342), 209–25.

Lucas, R. E. (1977). Understanding business cycles. *Carnegie-Rochester Conference Series on Public Policy 5*(1), 7–29.

Markose, S., S. Giansante, and A. R. Shaghaghi (2012). ‘Too interconnected to fail’ financial network of US CDS market: Topological fragility and systemic risk. *Journal of Economic Behavior & Organization 83*(3), 627–646.

Newman, M. E. J. (2010). *Networks: an introduction*. Oxford: Oxford University Press.

Setterfield, M. and A. Budd (2011). A Keynes-Kalecki model of cyclical growth with agent-based features. In P. Arestis (Ed.), *Microeconomics, Macroeconomics and Economic Policy: Essays in Honour of Malcolm Sawyer*. Palgrave Macmillan.

Steindl, J. (1965). *Random Processes and the Growth of Firms*. New York: Hafner.

Stiglitz, J. E. and M. Gallegati (2011). Heterogeneous Interacting Agent Models for Understanding Monetary Economies. *Eastern Economic Journal 37*(1), 6–12.

Watts, D. J. and S. H. Strogatz (1998). Collective dynamics of ”small-world” networks. *Nature 393*(6684), 440–442.

Wooldridge, M. (Ed.) (2002). *MultiAgent Systems*, West Sussex. John Wiley and Sons, LTD.

Yun, T.-S., D. Jeong, and S. Park (2019). “Too central to fail” systemic risk measure using PageRank algorithm. *Journal of Economic Behavior & Organization 162*(C), 251–272.

Figure 1: Aggregate fluctuations and the impact of weight.

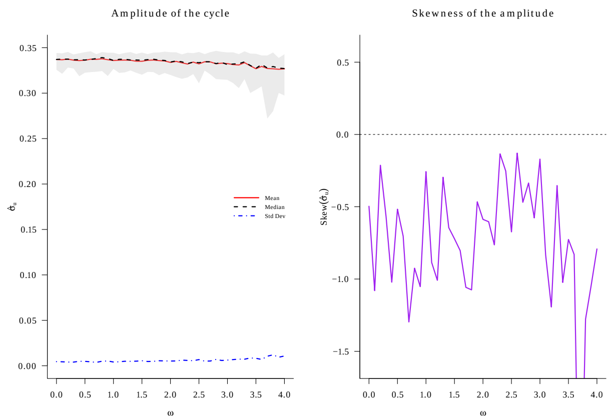

(a) Random network

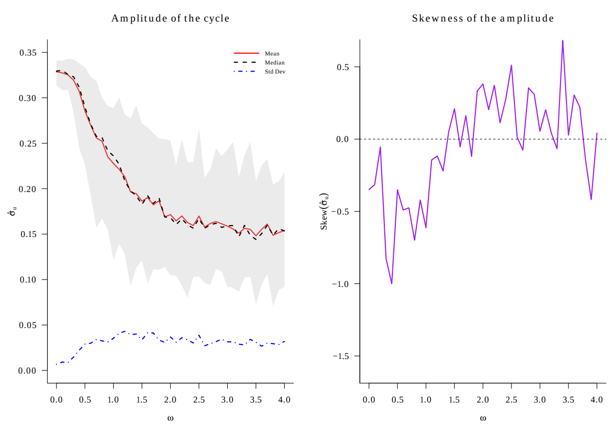

(b) Preferential attachment

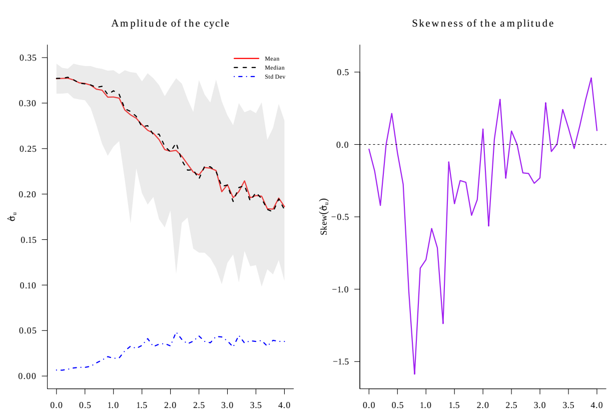

(c) Small world

In the left panel, the shaded gray area represents the range of $\hat{\sigma}_u$.

Figure 2: The joint effects of ω and $CV(\bar{x}_j)$ on aggregate volatility.

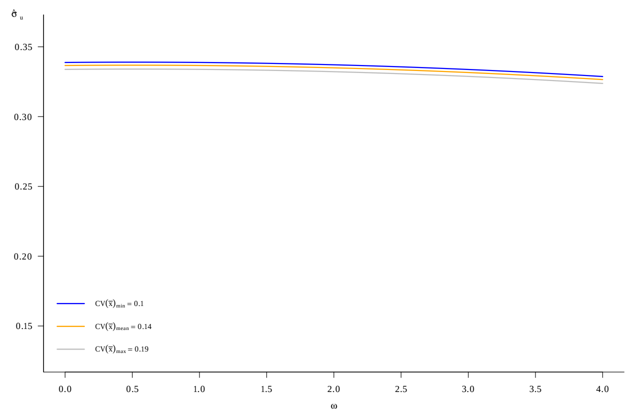

(a) Random network

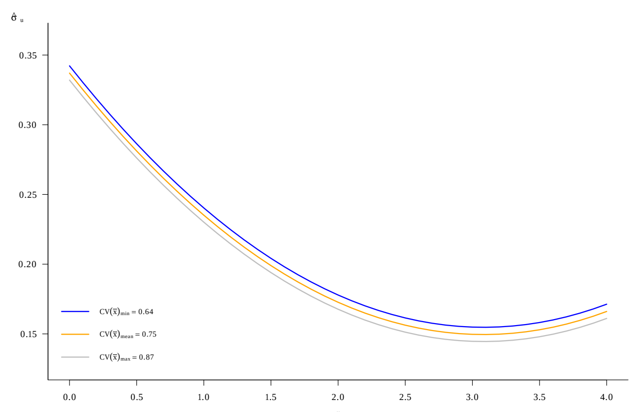

(b) Preferential attachment

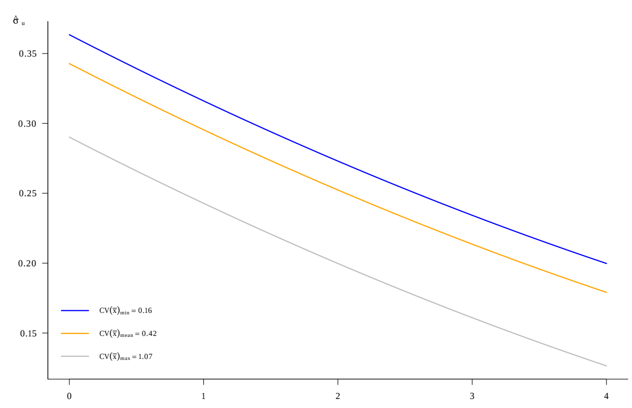

(c) Small world

Imprint

Publisher
Macroeconomic Policy Institute (IMK) of Hans-Böckler-Foundation, Georg-Glock-Str. 18,
40474 Düsseldorf, Contact: fmm@boeckler.de, https://www.fmm-macro.net

FMM Working Paper is an irregular online publication series available at:
https://www.boeckler.de/de/fmm-working-paper-22457.htm

The views expressed in this paper do not necessarily reflect those of the IMK or the Hans-Böckler-Foundation.

ISSN 2512-8655