



Emergent Inequality and Endogenous Dynamics in a Simple Behavioral Macroeconomic Model

Yuki M. Asano, Jakob J. Kolb, Jobst Heitzig and J. Doyne Farmer

5th July 2019

INET Oxford Working Paper No. 2019-11

Complexity Economics Programme



EMERGENT INEQUALITY AND ENDOGENOUS DYNAMICS IN A SIMPLE BEHAVIORAL MACROECONOMIC MODEL

Yuki M. Asano^{a,b,c}, Jakob J. Kolb^{b,d}, Jobst Heitzig^b and J. Doyne Farmer^{e,f,g,1}

^aDept. of Engineering Science, University of Oxford, OX1 3PJ, UK

^bFutureLab on Game Theory and Networks of Interacting Agents, Potsdam Institute for Climate Impact Research (PIK), PO Box 60 12 03, D-14412 Potsdam

^cFernUniversität in Hagen, Universitätsstraße 47, D-58097 Hagen

^dHumboldt-Universität zu Berlin, Unter den Linden 6, D-10099 Berlin

^eInstitute for New Economic Thinking at the Oxford Martin School, University of Oxford, OX1 3UQ, UK

^fMathematical Institute, University of Oxford, OX2 6GG, UK

^gSanta Fe Institute, Santa Fe, 87501 NM, USA

¹Corresponding author: doyne.farmer@inet.ox.ac.uk

July 5, 2019

ABSTRACT

Standard macroeconomic models assume that households are rational in the sense that they are perfect utility maximizers, and explain economic dynamics in terms of shocks that drive the economy away from the steady state. Here we build on a standard macroeconomic model in which a single rational representative household makes a savings decision of how much to consume or invest. In our model households are myopic boundedly rational heterogeneous agents embedded in a social network. From time to time each household updates its savings rate by copying the savings rate of its neighbor with the highest consumption. If the updating time is short, the economy is stuck in a poverty trap, but for longer updating times economic output approaches its optimal value, and we observe a critical transition to an economy with irregular endogenous oscillations in economic output, resembling a business cycle. In this regime households divide into two groups: Poor households with low savings rates and rich households with high savings rates. Thus inequality and economic dynamics both occur spontaneously as a consequence of imperfect household decision making. Our work here supports an alternative program of research that substitutes utility maximization for behaviorally grounded decision making.

economic growth and inequality are important problems in economics [1, 36]. Standard macroeconomic models are based on the assumption of a single representative rational utility maximizing agent and assume that the dynamics of business cycles are driven by exogenous shocks. However, empirical evidence from behavioral economics indicates that real households are heterogeneous and make substantial deviations from rationality. This has led to new directions of research, including the incorporation of heterogeneous or boundedly rational agents into macroeconomic models. This is typically done by allowing the agents to differ in terms of factors such as education while preserving the assumption of rationality [12, 31, 44], or alternatively allowing for bounded rationality but maintaining utility maximization [26]. Realism is injected through imposing frictions, such as sticky wages. These models require shocks to generate economic dynamics.

However, it has long been known that endogenous dynamics are possible in economic models [9, 10, 11, 17, 38], and more recently Beaudry *et al.* have shown how limit cycles can emerge in a standard framework where agents are perfect utility maximizers [5, 6]. An alternative approach bases household decision making on simple heuristics rather than rationality [18, 19]. This leads to waves of optimism and pessimism, generating irregular business cycles and giving fat-tailed distributions for economic outcomes such as GDP. We further develop this line of research by demonstrating how a very simple heterogeneous behavioral macroeconomic model leads to an endogenous business

cycle that is not driven by externally imposed shocks. Our purpose is not to make a fully realistic model, but rather to demonstrate that rich emergent behavior can occur even under very simple assumptions.

Here we extend the Ramsey-Cass-Koopmans (RCK) model, which is one of the foundational models of economic growth theory. In this model a representative agent rationally chooses a savings rate in order to maximize discounted consumption. However, there is ample evidence that households do not act as intertemporal optimizing agents and often respond myopically [7, 15, 33]. Evidence from lab experiments suggests that individuals perform poorly in finding optimal consumption paths. In reference [14] subjects deviated from optimal consumption choices by roughly 30 percent on average, increasing to roughly 50 percent when subjects were shown the average consumption level in the previous period. Learning from past generations' consumption paths is somewhat more successful, but the errors are still substantial [3, 13].

We take the opposite approach and assume a strong form of bounded rationality. In our model households are embedded in a social network and make their savings decisions by simply copying their most successful neighbor. They do this episodically and myopically: From time to time they check all their neighbors and adopt the savings rate of the neighbor with the highest consumption.

Although we do not claim that this behavior is fully realistic, there is empirical justification for considering a simple rule of this type. Our agents can be viewed as short-sighted, profligate “conspicuous consumers”, and the tendency of households to copy one another has been well-documented since the time of Thorstein Veblen [41]. Imitate-the-best is one of the decision-making heuristics often applied in settings of high uncertainty and variability [27] and is observed in economic experiments [40]. Savings behavior is highly dependent on social interaction with peers [25, 30, 34, 43] and comparing consumption levels incorporates the visibility bias and selection neglect observed in savings rate decisions [21]. Our implementation by copying based on consumption alone is partly motivated by the fact that a neighbor's consumption is more visible than its capital. This makes it particularly surprising that sometimes their average behavior can be close to optimal.

We find that a key parameter governing economic behavior is the average time interval τ at which households update their savings rate, which we call the *social interaction time*. When τ is small, meaning the households update frequently, the savings rate is low, and the performance of the economy is suboptimal in terms of aggregate consumption. When τ is sufficiently large, in contrast, the economy-wide aggregate savings rate, which equals the income-weighted average household savings rate, becomes close to the optimal rate. For small τ the population of households remains homogeneous, but as τ increases there is a sharp phase transition at a critical value τ_c where the population becomes strongly bimodal, dividing into rich households with high savings rates and poor households with low savings rates. Correspondingly, for low values of τ the GDP and other economic indicators are constant with only small fluctuations, whereas above the critical transition there is an endogenous aperiodic oscillation, resembling a business cycle, in which the aggregate savings rate fluctuates, the population of households alternately becomes richer or poorer, and economic output varies substantially over time.

Our model shows that the use of heterogeneous agents following explicit behavioral rules can produce aggregate behavior that is qualitatively different from that of rational agents. Our model is only qualitative, but our results suggest that an approach that explicitly incorporates empirical behavioral knowledge into household decision making may naturally lead to an explanation of business cycles in terms of endogenous dynamics.

The standard RCK model

The RCK model considers a closed economy in which a representative household provides both labor and capital for the production of a single good by a representative firm. The household receives wages w for labor and a nominal rate of return r on its investment. It spends a fraction $1-s$ of its income on consumption and invests the remaining fraction s (see Materials and Methods).

Given the current value of per-capita capital k , the household chooses a current value of per-capita consumption $c(s)$ determined by intertemporal optimization, leading to an optimal consumption path that maximizes the household's long-term discounted aggregate utility. This determines the time evolution of k and c towards a steady state at (k^*, c^*) .

An agent-based version of RCK

Economic Model

We introduce a heterogeneous agent model in the tradition of agent-based modeling [8, 16, 20, 22, 28, 32, 39], using agents that follow a very simple behavioral rule. Our model contains N households labeled by i with heterogenous

capital K_i . For simplicity, all households supply the same labor $L_i = L/N$. (Introducing heterogeneous labor has little effect on the results). As in the original RCK model, total economic production is given by the Cobb–Douglas production function, in this case applied to the aggregate input factors $K = \sum_{i=1}^N K_i$ and $L = NL_i$. As in the original model, capital returns r and wages w equal marginal returns (Methods, Eq. 7), but incomes I_i now differ between households,

$$I_i = rK_i + wL/N. \quad (1)$$

Our key assumption is that each household individually and dynamically sets its time-dependent savings rate $s_i(t)$ according to a behavioral decision rule introduced below, leading to household capital dynamics

$$\dot{K}_i = s_i I_i - \delta K_i = (rs_i - \delta)K_i + ws_i L/N. \quad (2)$$

At the steady state where $\dot{K}_i = 0$, the steady state value K_i^* for household i 's capital is a function of the aggregate capital K via its dependence on w and r , nonlinearly interconnecting all the agents' savings rates and consumption levels.

Household Decision Making

While the standard RCK model is a one-dimensional dynamical system in which consumption is a deterministic function of the total capital, the agent-based version is $2N$ -dimensional, and aggregate consumption depends on all households. We assume that each household updates its savings rate at random times¹ according to a Poisson process with rate $1/\tau$. We will see that τ plays a crucial role for the model's behavior.

Households are embedded in a social network in which each household i has neighbors $\mathcal{N}(i)$. Whenever household i updates its savings rate, it compares the consumption rates of its neighbors and applies the 'imitate-the-best' heuristic, copying the savings rate of the neighbor with the highest current consumption with a small deviation that can either be interpreted as an error or as an exploration [35]. More precisely, when the consumption of a neighbor is higher, it adopts a new savings rate of

$$s_i^{\text{new}} = s_{\underset{j \in \mathcal{N}(i)}{\text{argmax}}(C_j)} + \epsilon, \quad (3)$$

where ϵ is distributed uniformly in the interval of $\pm 1\%$. (The behavior is insensitive to this as long as there is some diversity).

Results

We simulate the model for a variety of different parameters such as the average social interaction time τ and the network topology. In Fig. 1 we show the distribution of the final savings rates as a function of the social interaction time τ for a complete network with the other parameters fixed. The figure compares this to the optimal, 'golden rule' savings rate s_{gold} , corresponding to the rational expectations equilibrium where the consumption of the representative agent is maximized. There are two distinct regimes, separated by a critical social interaction time $\tau_c \approx 250$. In the *stable regime*, corresponding to $\tau < \tau_c$, the savings rates of the households are unimodally distributed around a low savings rate. For very small values of τ the savings rates are close to zero, and the economy is stuck in a poverty trap in which its output is very low. As τ increases, the savings rate and output increase, but the distribution remains unimodal, with a sub-optimal aggregate saving rate.

For $\tau > \tau_c$ we enter what we call the *oscillatory regime*, where the behavior is dramatically different. In this regime the savings rate distribution is bimodal – some households have high savings rates and are quite wealthy, while others have low savings rates and are very poor. We thus observe the spontaneous emergence of extreme inequality, with a lower class and an upper class².

Strikingly, as long as $\tau > \tau_c$, the ensemble average of the aggregate savings rate \tilde{s} is within 1% of the optimal value $s_{\text{gold}} = 0.5$, even when the individual distributions are bimodal. Furthermore, the time averages of total economic output $Y(t) = 10.15$ and consumption $C = 4.99$ are close to their optimal values $Y^* = s_{\text{gold}}L/\delta = 10$ and $C^* = (1 - s_{\text{gold}})Y^* = 5$ in the standard RCK model. It seems surprising that such a simple, near zero-intelligence learning rule can maintain the system this close to its optimal behavior.

The system dynamics become clearer when we look at the economy as a function of time, as illustrated in Fig. 2. For $\tau > \tau_c$ there is an endogenous oscillation in many of the aggregate properties of the economy, including the aggregate

¹ This leads to smoother transitions than synchronous updates [24, 42].

²Very near τ_c the distribution in Fig. 1 is tri-modal. This is due to intermittent oscillations between the unimodal and bi-modal regimes. Thus the system either exhibits a middle class, or a lower class and an upper class, but never all three at once.




Figure 1: The critical transition from the stable regime to the oscillatory regime. We perform an ensemble of simulations at different values of the social interaction time τ , with other parameters held fixed (see SI). We show a heatmap indicating the probability density of the distribution of individual households savings rates for each value of τ , along with the aggregate savings rate \tilde{s} . We compare it to the golden rule savings rate $s_{\text{gold}} = 0.5$ and the savings rate s^* predicted by Equation (4).

savings rate $\tilde{s}(t)$ and output $Y(t)$. This oscillation is also visible in the behavior of individual households. If we follow any single household it goes through epochs with a high savings rate, near $s_i \approx 90\%$, and a low savings rate, near $s_i \approx 5\%$. At any point in time there is typically an imbalance between rich households and poor households, so that the aggregate savings rate and the aggregate output fluctuate. We loosely refer to this endogenous oscillation as a “business cycle”.

Understanding the stable regime ($\tau < \tau_c$)

Although our behavioral rule requires minimal intelligence, the selection process of copying the household with the highest consumption provides a simple mechanism of collective search that becomes more effective as the social updating time τ increases. This is perhaps counter-intuitive, as it means that inattention results in superior collective outcomes. The underlying explanation is as follows: The savings rate of the household that is copied has on average been fixed for a time interval of order τ . When τ is small, planning is too myopic, “short term thinking” dominates, and the households cannot escape using low savings rates with high consumption. As τ gets bigger, however, the time between updates becomes long enough that there is more time to accrue an advantage by saving, which drives the savings rate up and increases economic output. The competitive selection process guarantees that for a sufficiently large population and large τ the savings rates are close to optimal.

We use this intuition to derive an approximate formula for the aggregate savings rate s^* as a function of τ . We take advantage of the fact that in the stable regime the distribution is unimodal and assume that all households have essentially the same savings rate, and derive the optimal savings rate for time horizon τ . As explained in detail in the Supplementary Information, for capital elasticity $\alpha = 0.5$ the optimal savings rate under these conditions is

$$s^*(\tau) = \frac{1 - e^{-\delta\tau/2}}{2 - e^{-\delta\tau/2}}. \quad (4)$$

This approximation is shown in green in Fig. 1 and provides a good fit throughout the stable regime.

The optimal savings rate in the classical RCK model depends on the discount rate ρ , which is a free parameter. As shown in the SI, substituting s^* from Eq. [4] into the relation for the classical RCK model gives an effective discounting rate for our model in terms of the social interaction time τ and the depreciation rate δ ,

$$\rho(\tau) = \frac{\delta/2}{e^{\delta\tau/2} - 1}. \quad (5)$$

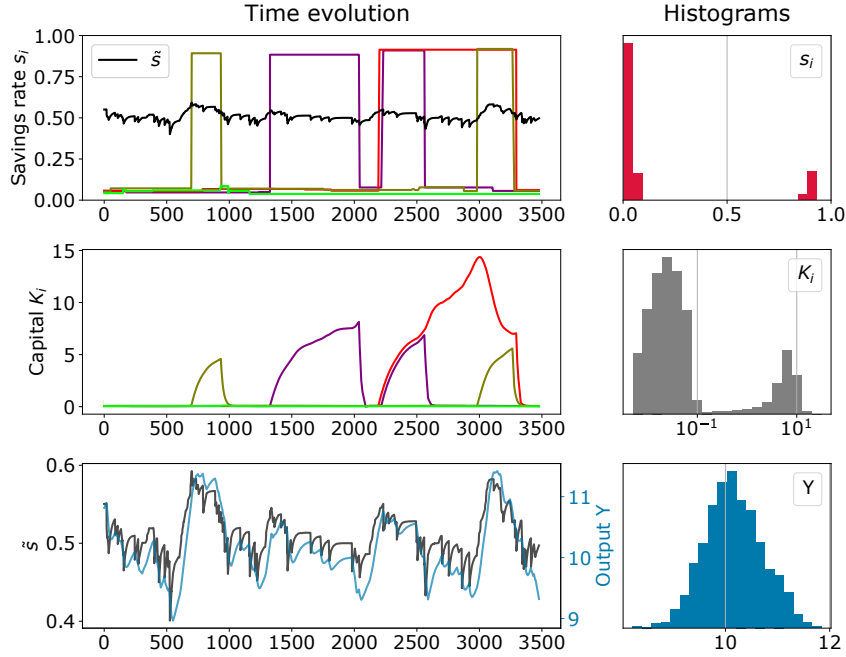


Figure 2: The endogenous business cycle in the oscillatory regime. We show several time series when $\tau > \tau_c$. The top left panel shows the savings rates $s_i(t)$ for four randomly chosen households as a function of time, as well as the aggregate savings rate \bar{s} . The middle left panel shows the capital $K_i(t)$ of the same four households as a function of time. The bottom left panel shows the cyclic behavior of the aggregate output superimposed on the aggregate savings rate. The panels on the right are histograms of the indicated variables, accumulated over a longer interval.

In the limit as $\tau \rightarrow 0$, the discount rate $\rho \rightarrow \infty$, consistent with the observed collectively myopic behavior. But for $\tau \rightarrow \infty$, $\rho \rightarrow 0$. Thus in this case the individually myopic households act collectively “as if” they were farsighted, with an emergent effective discounting rate $\rho(\tau)$ which is not a free parameter but is rather a function of the social interaction time τ .

Understanding the oscillatory regime ($\tau > \tau_c$)

To get a deeper understanding of what is happening in the oscillatory regime, where $\tau > \tau_c$, in Fig. 3 we illustrate the collective and individual dynamics. In Fig. 3A we show the average per capita consumption rate c as a function of the average capital k . This illustrates how the aggregate consumption and capital orbit around the optimal steady state (k^*, c^*) of the standard RCK model, generating a business cycle. In relation to the optimal savings rate $s^* = 0.5$ of the RCK model, the effective aggregate savings rate \bar{s} is typically greater than s^* when the system is below the optimality curve and less than s^* when it is above the optimality curve. This is interesting as the optimality curve is obtained via optimizing household consumption for an infinite horizon, whereas our model has no explicit optimization.

To understand what is going on at the individual level, in Fig. 3B we plot a snapshot of the capital vs. the savings rate for all households at two different times, t_1 and t_2 . At time t_1 the economy is just beginning to recover from a recession. There are two clusters of households, corresponding to rich households in the upper right corner and poor households in the lower left corner. More households are poor, and because the return r is inversely proportional to total capital according to $r \propto K^{-1+\alpha}$, where here $\alpha = 0.5$, this means that returns to investment are high. When the household shown in blue gets its chance to update its savings rate, it copies the higher savings rate of one of the rich households, transitions to the right as indicated by the arrow, and begins accumulating capital by saving more. Other households follow, and eventually the economy reaches the state shown in the lower panel at time t_2 , where many houses have high savings rates and are rich. The resulting excess capital makes the returns on savings low, which when combined with their high savings rates, drives the consumption of these households down. As a result, when one of the rich households gets its turn to update, it copies a household with a low savings rate and goes on a spending spree. At this point its consumption rate becomes very high, and all of its neighbors copy it, creating a boom in consumption while decreasing the aggregate savings rate. A majority of households eventually become impoverished and the cycle repeats itself. These dynamics are also given as an animation in the Supplementary Material.

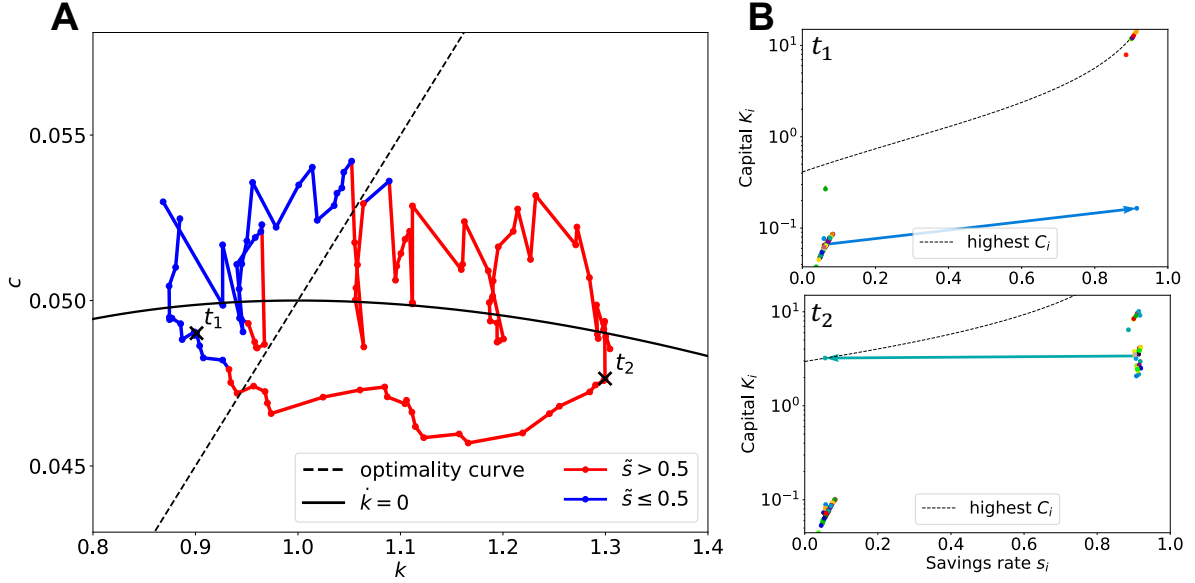


Figure 3: Endogenous dynamics in the oscillatory regime. **A:** We plot the average per-capita consumption c against the average per-capita capital k and show the aggregate saving rate \tilde{s} as red when it is greater than 0.5 and blue when it is less than 0.5. The trajectory orbits around the optimal steady state (k^*, c^*) of the standard RCK model, which is at the intersection of the dashed optimality curve and the solid black $\dot{k} = 0$ line. Each dot corresponds to one timestep; the orbit is counterclockwise. **B:** An illustration of the cause of the oscillatory dynamics. The two panels show snapshots at two different times as indicated in figure A. At time t_1 the aggregate savings rate is low, aggregate capital is low and the economy is in a depression; at time t_2 the opposite is true. The capital and savings rates of individual households are shown as dots with different colors. There are two clusters, corresponding to rich and poor households. The household that is currently switching its savings rate is indicated by an arrow connecting its previous state to its current state. The dashed black curve indicates the iso-consumption curve for the household i with the highest consumption.

Critical social interaction time

What determines the critical social interaction time τ_c ? The approximation that the imitate-the-best heuristic results in behavior that is optimal over a time horizon τ helps understand the instability driving the transition. Suppose that an external shock of size Δ perturbs the aggregate savings rate \tilde{s} away from s^* , and suppose that household i is allowed to optimize its savings rate s_i while the others hold theirs constant. A numerical investigation shows that when $\tau \ll \tau_c$ the optimal savings rate s_i computed remains close to \tilde{s} . In contrast, when $\tau \gg \tau_c$, if $\Delta > 0$ then the optimal savings rate is very small, with s_i approaching 0, and if $\Delta < 0$ the optimal savings rate is large, with s_i approaching 1 (see Fig. S3). This happens because when $\Delta > 0$ the aggregate savings rate is high, so the returns on investment are low, which discourages saving. Similarly, when $\Delta < 0$ the aggregate savings rate is low, so returns on investment are high, which encourages saving. This destabilizes the unimodal solution around s^* . The transition occurs sharply at a parameter value near τ_c , though the precise value depends on Δ .

Network size and structure

We have so far used complete networks in the simulations, but in general the behavior depends on the network size and structure. For example, we investigate Erdős–Rényi networks with average degree $\langle k \rangle = Np$, where p is the probability that any two nodes are connected. The critical social interaction time τ_c depends on both $\langle k \rangle$ and N (see Figs. S4 and S5). Starting at any given node and moving one link at a time, the number of neighbors that are reached grows exponentially with time at rate $\langle k \rangle$. The typical distance required for a disturbance to propagate across the network is the average shortest path length χ , defined as the average number of nodes that must be traversed in order to go from any given node to any other node. Motivated by this logic, we investigate the empirical relationship between τ_c , χ , and $\langle k \rangle$, finding the proportionality

$$\tau_c \sim e^{-\chi} / \langle k \rangle. \quad (6)$$

Fig. 4 shows that this makes a good prediction of τ_c . Because χ increases with N , in the large N limit the system is always in the oscillatory regime. Varying the network size and structure parameters also results in qualitative changes in the nature of the oscillation, affecting its frequency, amplitude and variability. (See a few examples in the SI).




Figure 4: The critical social interaction time depends on network properties. The logarithm of the mean number of neighbors $\langle k \rangle$ times the critical interaction time τ_c is plotted vs. the average shortest path length χ for various values of N and p , confirming Eq. [6]. The stable regime ($\tau < \tau_c$) is shaded in blue.

Discussion

Our primary purpose here is to make a conceptual point by demonstrating how emergent inequality and endogenous dynamics can naturally emerge from a heterogeneous behavioral model. Nonetheless, our model makes the prediction that during recessions savings rates increase before output rises (see Fig. 2). This has been observed for private savings in 19 OECD countries [2].

Although our model has two random inputs, they are small and very different in character from the shocks that drive the dynamics of standard models. The first random input determines the time at which individual households update their savings rates under the Poisson process. This must be random to ensure that the order in which households update their savings rates varies. (A fixed order leads to a static economy). The second random input is the copying error for the savings rate. This is small (1%) and its value makes little difference to the behavior. In contrast to standard shocks, which affect the economy as a whole, both of these inputs are at the level of individual households, and affect each household differently. For a large number of households the copying errors cancel out but the endogenous dynamics nonetheless persist. Thus while random inputs are necessary in our model, they do not directly drive booms and recessions as the shocks of standard models do. This is why we say that the economic dynamics in our model are endogenous.

To illustrate the conceptual difference between our model and standard macroeconomic models it is useful to draw an analogy to a simple physical system. Consider the problem of pole balancing, in which a man attempts to move his hand to maintain a pole in a vertical position, as shown in Fig. 5. Short poles tip over more quickly than long poles,



Figure 5: The problem of pole balancing is analogous to the problem of optimizing savings in an otherwise unstable economy. A man attempts to maintain a pole in a vertical position. This is possible if the pole is long enough, but small errors in the control process drive endogenous oscillations in the angle of the pole.

making it impossible to maintain a vertical position because the pole will tip over before the man can react. If the pole is long enough, however, the man can move his hand to compensate, and maintain the pole in a roughly vertical position [29]. There is a sharp critical transition between stability and instability that occurs when the pole is about a meter in length³. Nonetheless, even when the pole is very long, it is not possible to maintain a perfectly vertical position, and the pole oscillates substantially.

An argument in the style of a standard macroeconomic model would posit that the man is a perfect pole balancer, and any deviations in the angle must be driven by external shocks, such as sharp gusts of wind, that suddenly cause the pole to deviate from vertical. Under this view, after each shock the man moves his hand perfectly to make the pole vertical again as fast as possible, but before he can achieve this, another shock strikes it, making the pole oscillate around its vertical position. For pole balancing it is clear that this explanation is wrong. Instead, theories that assume that oscillations are endogenously caused by imperfect control provide a better explanation [29].

³This is trivial to confirm empirically – simply attempt to balance a pole of 60 cm. vs. 130 cm.

Our suggestion here is that we should revisit the conceptual explanation for business cycles as well. In the analogy above, the position of the man’s hand is like the collection of household savings decisions, and the pole/gravity system is like the economy. Our model adds weight to the idea that at least part of the variation in savings and investment that occurs during business cycles emerges endogenously due to the imperfect reasoning of households and firms. Our model also suggests that models with heterogeneity might help illuminate the interaction between business cycles and inequality. The fact that such rich behavior emerges from such a simple model supports a research agenda for macroeconomics based on empirically derived behavioral rules.

Materials and Methods

Code

The source code (Python) will be made available here upon publication.

The standard RCK model

In our formulation of the standard RCK model, we follow [1, p.287–317], [4, p.85–135] and use continuous time, as in the original [37]. We ignore labor growth for brevity. Using a fixed amount of labor L and the varying amount of capital K , the economy produces a single numeraire good Y , assuming a Cobb–Douglas production function $Y = K^\alpha L^{1-\alpha}$ with capital and labor elasticities $\alpha, 1 - \alpha \in (0, 1)$. Per-capita production, $y = Y/L$, is thus a function of per-capita capital, $k = K/L$, only, $y = k^\alpha$. The model assumes fully competitive factor markets and thus the two factors are compensated according to their marginal products, giving wages and capital rents

$$w = \partial_L Y = (1 - \alpha)y, \quad r = \partial_K Y = \alpha y/k, \quad (7)$$

thereby fully redistributing the numeraire good to households and leaving the representative firm with no profits. The main model parameters of interest are the savings rate $s \leq 1$ and capital depreciation rate $\delta > 0$ that govern aggregate and per-capita capital growth,

$$\dot{K} = s(rK + wL) - \delta K, \quad \dot{k} = \bar{r}k + w - c, \quad (8)$$

where $\bar{r} = r - \delta$ is the real return rate and $c = (1 - s)(rk + w)$ is per-capita consumption. The household aims at maximizing its discounted aggregate utility $\int_0^\infty dt e^{-\rho t} u(c(t))$, by choosing an optimal path $s(t)$ for the savings rate, where $\rho > 0$ is its discount rate. For the instantaneous utility, one assumes a constant relative risk aversion (CRRA) function parameterized by $\theta \geq 0, \theta \neq 1, u(c) = (c^{1-\theta} - 1)/(1 - \theta)$. The solution to this problem fulfills the Ramsey–Keynes Equation that gives the relative consumption growth rate as $\dot{c}/c = (\bar{r} - \rho)/\theta$. In particular, this system has two steady states with $\dot{c} = 0$, a trivial one in which $c = k = 0$ and another in which the real return rate equals the discount rate, $\bar{r} = \rho$, corresponding to a modified ‘golden rule’ [1, p.300], with capital, consumption and savings rate given by

$$k^* = \left(\frac{\alpha}{\rho + \delta} \right)^{\frac{1}{1-\alpha}}, \quad c^* = k^{*\alpha} - \delta k^*, \quad s_{\text{RCK}}^* = \frac{\alpha \delta}{\rho + \delta}. \quad (9)$$

For the limit case $\rho \rightarrow 0$, this reproduces the Solow model’s golden rule [4, p.35], $s_{\text{RCK}}^* = s_{\text{gold}} = \alpha$, leading to the largest possible sustainable consumption, $c^* = (1 - \alpha)(\alpha/\delta)^{\alpha/(1-\alpha)}$. For $\rho > 0$, the discount rate pushes the households to save less and shift consumption towards the present.

Acknowledgments

We thank Paul Beaudry, Jean-Philippe Bouchaud, Roger Farmer, Cars Hommes, Ulrike Kornek, Bastian Ott, Martin Braml, Marco Pangallo and Frederik Schaff for valuable comments and suggestions. YMA is funded by the EPSRC CDT in AIMS (EP/L015897/1) and JDF by Baillie Gifford and the Institute for New Economic Thinking.

References

- [1] D. Acemoglu. *Introduction To Modern Economic Growth*. Princeton University Press, 2009.
- [2] Y. Adema and L. Pozzi. Business cycle fluctuations and household saving in oecd countries: A panel data analysis. *European Economic Review*, 79:214–233, 2015.
- [3] T. P. Ballinger, M. G. Palumbo, and N. T. Wilcox. Precautionary saving and social learning across generations: an experiment. *The Economic Journal*, 113(490):920–947, 2003.
- [4] R. J. Barro and X. Sala-i Martin. *Economic growth*. MIT Press, 2004.

- [5] P. Beaudry, D. Galizia, and F. Portier. Reviving the limit cycle view of macroeconomic fluctuations. Technical report, National Bureau of Economic Research, 2015.
- [6] P. Beaudry, D. Galizia, and F. Portier. Putting the cycle back into business cycle analysis. *CEPR Discussion Paper No. DP11647*, 2016.
- [7] S. Benartzi and R. H. Thaler. Myopic Loss Aversion and the Equity Premium Puzzle. *Quarterly Journal of Economics*, 110(1):73–92, feb 1995.
- [8] B. J. L. Berry, L. D. Kiel, and E. Elliott. Adaptive agents, intelligence, and emergent human organization: Capturing complexity through agent-based modeling. *Proceedings of the National Academy of Sciences*, 99(Supplement 3):7187–7188, may 2002.
- [9] L. Blume and D. Easley. Evolution and market behavior. *J. Econ. Theory*, 58(1):9–40, 1992.
- [10] M. Boldrin and L. Montrucchio. On the indeterminacy of capital accumulation paths. *Journal of Economic Theory*, 40(1):26 – 39, 1986.
- [11] M. Boldrin and M. Woodford. *Equilibrium Models Displaying Endogenous Fluctuations and Chaos: A Survey*. Princeton University Press, Princeton, 1992.
- [12] W. A. Branch and B. McGough. A new keynesian model with heterogeneous expectations. *Journal of Economic Dynamics and Control*, 33(5):1036–1051, 2009.
- [13] A. L. Brown, Z. E. Chua, and C. F. Camerer. Learning and visceral temptation in dynamic saving experiments. *The Quarterly Journal of Economics*, 124(1):197–231, 2009.
- [14] E. Carbone and J. Duffy. Lifecycle consumption plans, social learning and external habits: Experimental evidence. *Journal of Economic Behavior & Organization*, 106:413–427, 2014.
- [15] J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick. Reinforcement learning and savings behavior. *The Journal of finance*, 64(6):2515–2534, 2009.
- [16] H. Dawid, P. Harting, and M. Neugart. Economic convergence: Policy implications from a heterogeneous agent model. *Journal of Economic Dynamics and Control*, 44:54 – 80, 2014.
- [17] R. H. Day. The emergence of chaos from classical economic growth. *Q. J. Econ.*, 98(2):201–213, 1983.
- [18] P. De Grauwe. The scientific foundation of dynamic stochastic general equilibrium (dsge) models. *Public Choice*, 144(3):413–443, Sep 2010.
- [19] P. De Grauwe. Animal spirits and monetary policy. *Economic Theory*, 47(2):423–457, Jun 2011.
- [20] G. Dosi, G. Fagiolo, and A. Roventini. Schumpeter meeting keynes: A policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control*, 34(9):1748 – 1767, 2010. Computational perspectives in economics and finance: Methods, dynamic analysis and policy modeling.
- [21] B. Enke. What you see is all there is. 2015.
- [22] J. M. Epstein. *Generative social science : studies in agent-based computational modeling*. Princeton University Press, 2006.
- [23] P. Erdős and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.
- [24] N. Fatès and V. Chevrier. How important are updating schemes in multi-agent systems? An illustration on a multi-turmite model. *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 1:533–540, may 2010.
- [25] R. H. Frank, A. S. Levine, and O. Dijk. Expenditure cascades. *Review of Behavioral Economics*, 1(1-2):55–73, 2014.
- [26] X. Gabaix. A behavioral new keynesian model. Technical report, National Bureau of Economic Research, 2016.
- [27] G. Gigerenzer and W. Gaissmaier. Heuristic Decision Making. *Annual Review of Psychology*, 62(1):451–482, 2011.
- [28] C. Hommes. Behavioral & experimental macroeconomics and policy analysis: a complex systems approach. Technical report, European Central Bank, 2018. Working Paper Series:2201.
- [29] T. Insperger and J. Milton. Stick balancing with feedback delay, sensory dead zone, acceleration and jerk limitation. *Procedia IUTAM*, 22:59–66, 2017.
- [30] M. Kaustia and S. Knüpfer. Peer performance and stock market entry. *Journal of Financial Economics*, 104(2):321–338, 2012.

- [31] J. Leahy, Y. Achdou, M. Aguiar, F. Alvarez, A. Auclert, J. Benhabib, L. Dedola, E. Farhi, M. Gertler, N. Kocherlakota, K. Kuester, D. Lagakos, E. Nakamura, L. Schmidt, J. Steinsson, M. Wiederholt, G. Kaplan, B. Moll, and G. L. Violante. Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743, 2018.
- [32] B. LeBaron, W. Arthur, and R. Palmer. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23(9):1487 – 1516, 1999.
- [33] G. Loewenstein. Emotions in Economic Theory and Economic Behavior. *American Economic Review: Papers and Proceedings*, 90(2):426–432, 2000.
- [34] T. J. Lu and N. Tang. Social interaction effects and individual portfolio choice: Evidence from 401 (k) pension plan investors. *Ning, Social Interaction Effects and Individual Portfolio Choice: Evidence from*, 401, 2015.
- [35] K. Mehlhorn, B. R. Newell, P. M. Todd, M. D. Lee, K. Morgan, V. A. Braithwaite, D. Hausmann, K. Fiedler, and C. Gonzalez. Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3):191, 2015.
- [36] T. Piketty. About capital in the twenty-first century. *American Economic Review*, 105(5):48–53, 2015.
- [37] F. P. Ramsey. A Mathematical Theory of Saving. *The Economic Journal*, 38(152):543, dec 1928.
- [38] J. A. Scheinkman. Nonlinearities in economic dynamics. *The Economic Journal*, 100(400):33–48, 1990.
- [39] M. Simon, C. Schwartz, D. Hudson, and S. D. Johnson. A data-driven computational model on the effects of immigration policies. *Proceedings of the National Academy of Sciences*, 115(34):E7914–E7923, 2018.
- [40] A. Traulsen, D. Semmann, R. D. Sommerfeld, H.-J. Krambeck, and M. Milinski. Human strategy updating in evolutionary games. *Proceedings of the National Academy of Sciences*, 107(7):2962–2966, 2010.
- [41] T. Veblen. *The Theory of the Leisure Class*. McMaster University Archive for the History of Economic Thought, 1899.
- [42] G. Vizzari and S. Bandini. Coordinating change of agents’ states in situated agents models. *Proceedings of the International Conference on Autonomous Agents*, pages 1395–1396, 2005.
- [43] A. C. Zhang, J. Fang, B. Jacobsen, and B. R. Marshall. Peer effects, personal characteristics and asset allocation. *Journal of Banking and Finance*, 90, 2018.
- [44] W. Zhao, Y. Lu, and G. Feng. How many agents are rational in china’s economy? evidence from a heterogeneous agent-based new keynesian model. *Computational Economics*, pages 1–37, 2018.

Supplementary Material:
Emergent inequality and endogenous dynamics in a simple behavioral macroeconomic model

Dated: July 5, 2019

Analytical approximation of the stable regime's steady state

Here we derive a simple approximate formula for the mean overall savings rate that emerges as a stochastic steady state in the stable regime where $\tau < \tau_c$.

The main idea is to study which member of an ensemble of households starting at similar but slightly different savings rates and capital stocks will have the largest consumption after the short time interval τ , and then assume all households will copy the savings rate of this best household with some error. This is only an approximation since in the actual model, households do not simultaneously imitate and not after exactly time τ , and the approximation will only be good when households have already converged to similar savings rates and capital stocks. However, it turns out that it describes rather well the joint motion towards a steady state once households have converged towards each other. In particular, the steady state savings rate predicted by the approximation can be seen to match the one observed in the numerical observations quite well.

To see how individual households' consumption at time τ depends on their individual savings rate, we need to approximate the evolution of r and thus of total capital K first. Assume all households' savings rates s_i are close to the overall savings rate s and stay constant between time zero and time τ . Then K evolves as

$$\begin{aligned}\dot{K} &= (rs - \delta)K + wsL \\ &= \left(\alpha \left(\frac{L}{K} \right)^{1-\alpha} s - \delta \right) K + \alpha \left(\frac{K}{L} \right)^{1-\alpha} sL.\end{aligned}$$

For $\alpha = 1/2$ this simplifies to

$$\dot{K} = s\sqrt{LK} - \delta K. \quad (S1)$$

Assuming that s does not change before time τ , this has two solutions given by

$$\begin{aligned}K(t) &= \left(\frac{B - Ee^{-\delta t/2}}{\delta} \right)^2, \\ r(t) &= \sqrt{L/K}/2 = \frac{A}{B - Ee^{-\delta t/2}}, \\ w(t) &= \sqrt{K/L}/2 = \frac{B - Ee^{-\delta t/2}}{4A}\end{aligned} \quad (S2)$$

for all $t < \tau$, where $A = \delta\sqrt{L}/2$, $B = s\sqrt{L}$, and E has the two possible values $s\sqrt{L} \pm \delta\sqrt{K_0}$. Since we are interested in the case where r is positive, we have $E = s\sqrt{L} - \delta\sqrt{K_0} < B$.

Knowing $r(t)$ and $w(t)$, we can now determine which household consumes most after time τ . Household i 's capital $K_i(t)$ evolves as

$$\begin{aligned}\dot{K}_i &= (s_i r(t) - \delta)K_i + ws_i L_i \\ &= \left(\frac{s_i A}{B - Ee^{-\delta t/2}} - \delta \right) K_i + \frac{B - Ee^{-\delta t/2}}{4A} s_i L_i.\end{aligned}$$

This has an analytical solution involving complicated hypergeometric functions. For small values of τ , we can simplify the problem by approximating $r(t)$ and $w(t)$ for $t \in [0, \tau]$ by their mid-term values $r(\tau/2)$ and $w(\tau/2)$, giving

$$\dot{K}_i \approx G_i K_i + F_i$$

with $G_i = \frac{s_i A}{B - Ee^{-\delta\tau/4}} - \delta$ and $F_i = \frac{B - Ee^{-\delta\tau/4}}{4A} s_i L_i$, which solves as

$$K_i(t) \approx (K_i(0) + F_i/G_i)e^{G_i t} - F_i/G_i.$$

The corresponding consumption of household i at time τ is then

$$\begin{aligned} C_i(\tau) &= (1 - s_i)(r(\tau)K_i(\tau) + w(\tau)L_i) \\ &\approx (1 - s_i) \left(H((K_i(0) + F_i/G_i)e^{G_i\tau} - F_i/G_i) + L_i/4H \right) \end{aligned}$$

with $H = \frac{A}{B - Ee^{-\delta\tau/2}} = \frac{A}{s\sqrt{L(1 - e^{-\delta\tau/2})} + \delta\sqrt{K_0}e^{-\delta\tau/2}}$. Since we assume all households imitate at time τ that s_i which has led to the largest $C_i(\tau)$, we can determine whether s will increase or decrease by identifying whether the s_i that gets copied is larger or smaller than s . Since we also assume households' savings rates s_i are distributed closely around s and all $K_i(0), L_i$ are similar, this question can be answered by seeing whether $C_i(\tau)$ increases or decreases when s_i is increased from below s to above s , i.e., by studying the derivative $\partial C_i(\tau)/\partial s_i$ at the point $s_i = s$. Up to a factor of N , this derivative is

$$(1 - s)H \left[\frac{e^{G\tau} - 1}{G} (L/4H - HF/G) + e^{G\tau}\tau H(K_0 + F/G) \right] - H[(e^{G\tau} - 1)F/G + e^{G\tau}K_0] - L/4H \quad (S3)$$

where $F = \frac{B - Ee^{-\delta\tau/4}}{4A} sL$ and $G = \frac{sA}{B - Ee^{-\delta\tau/4}} - \delta$. As long as the above expression is positive or negative, s will increase or decrease over time, respectively.

A steady state will then be reached when both s and K change no longer, i.e., when both \dot{K} as given by Eq. [S1] (with $K = K_0$) as well as $\partial C_i(\tau)/\partial s_i$ as given by Eq. [S3] are zero. The solution of $\dot{K} = 0$ is $K_0 = Ls^2/\delta^2$, at which point we have $E = 0, H = \delta/2s, G = -\delta/2, F/G = -Ls^2/\delta^2 = -K_0, HK_0 = Ls/2\delta$, and $HF/G = -Ls/2\delta = -HK_0$. Substituting all this into Eq. [S3] and setting it zero gives the following surprisingly simple approximate equation for the steady state s :

$$s^* \approx \frac{1 - e^{-\delta\tau/2}}{2 - e^{-\delta\tau/2}} \quad (S4)$$

as stated in the main text.

Simulation details

For all simulations in the main text, we have used the following parameters unless otherwise stated: $\mu_L = 1/N, K(t = 0)_i = 1 \forall i$ and a fully connected network with $N = 100, \delta = 0.05$. We have found that adding small heterogeneity in each household's labor does not change the dynamics significantly and the equilibrium dynamics also remain the same for different initial capital distributions with different $\sum_i K_i = K$. For Fig. 1, savings rate distribution plots are shown at the final state of the simulation at time $5\tau \cdot 10^3$, far beyond the point where the model has reached its asymptotic dynamics. For each value of τ , 200 independent simulations are run and all values of s_i are recorded for each τ to construct a histogram, which is normalized so the values add to one for each set of simulations with the same τ . In Fig. 4, Erdős-Rényi random graphs are constructed such that every node is connected to every other node to avoid households that are static and depreciation rate is $\delta = 0.2$ to speed-up computation.

Savings rate dependency and topology

In Fig. S1, we show the mean values for aggregate savings rates for the two modes of the oscillatory regimes. We can observe values ranging from 2%–23% for the lower savings rates and from 65% to 90% for the higher ones. Furthermore, we can see that the values for the high savings rates are strongly anticorrelated with the values of the lower savings rates, as they still average to $s^* = 0.5$.

Effective time-preference rate

In the RCK model, the discount rate is a free parameter, and the representative household solves the intertemporal optimization problem by choosing a savings rate of

$$s_{\text{RCK}}^* = \frac{\alpha\delta}{\rho + \delta}$$

(see Methods section from the main paper). In contrast, in the stable regime of our agent-based model, the myopic behaviour of the many individual households that learn their individual savings rates by social learning with a mean social interaction time τ leads to an aggregate savings rate s that converges to a steady state s^* which is approximately

$$s_{\text{ABM}}^* \approx \frac{1 - e^{-\delta\tau/2}}{2 - e^{-\delta\tau/2}},$$

as derived above. So, the main household characteristic that determines s_{RCK}^* is the discount rate ρ , while the main household characteristic that determines s_{ABM}^* is the social interaction time τ . Comparing the above two equations, we see that a given social interaction time τ in the ABM leads to approximately the same aggregate savings rate as the discount rate

$$\rho(\tau) = \frac{\delta/2}{e^{\delta\tau/2} - 1}$$

in the RCK model. Conversely, a given discount rate ρ in the RCK model leads to approximately the same aggregate savings rate as the social interaction time

$$\tau(\rho) = \frac{2}{\delta} \ln \left(1 + \frac{\delta}{2\rho} \right)$$

in the ABM. This relationship between ρ and τ is shown in Fig. S2 for various depreciation rates δ . Higher depreciation rates lead to lower discount rates, which is due to the dependence of s_{ABM}^* on δ and the inverse relationship between the classical economically optimal savings rate s_{RCK}^* with ρ . Thus, a higher depreciation leads to an increasing preference for the future. While maybe counterintuitive at first, it can be explained by the fact that s_{RCK}^* is obtained from the optimization of a representative agent which requires a higher savings rate for higher rates of depreciation for optimal consumption.

In the case where either δ or τ are small or ρ is large, we even have $\rho \approx 1/\tau$, which means that the discount rate of the RCK model almost exactly corresponds to the social interaction *rate* of our model.

Single household free-riding

We have seen above that, given an ensemble of households with an aggregate savings rate s , the individual consumptions C_i will approximately evolve as in Eq. S3. If the ensemble of households is large enough, it will contain some household j that uses a savings rate close to the value of s_i which maximizes Eq. S3. Let us call the latter the *best response* to s and denote it by $s_i^*(s)$. Since household j is then the most-consuming member of the ensemble, the social learning implies that other households will start copying $s_j \approx s_i^*(s)$, after which s will move slightly towards $s_i^*(s)$ as well. We have seen that when s equals the value s^* , also $s_i^* = s^*$, so that s^* is an equilibrium. It turns out that for values of s slightly above s^* , we get $s_i^* < s^*$, and for values of s slightly below s^* , we get $s_i^* > s^*$. Since updates occur asynchronously, after some stochastic perturbation of s away from s^* , this implies that s moves back towards s^* with some inertia, so that those households which imitate s_i^* first in such a situation have a temporarily larger consumption than the others.

Interestingly, the difference between the best response $s_i^*(s)$ and the equilibrium s^* depends strongly and highly nonlinearly on τ , as can be seen in Fig. S3. For τ below some critical value, $s_i^*(s)$ is very close to s^* , which explains why not only the aggregate savings rate s but also all individual savings rates s_i remain close to s^* in the stable regime of our model. However, for τ above the critical value, $s_i^*(s)$ is close to zero if $s > s^*$ and very large if $s < s^*$, converging to unity as τ increases. So as long as $s > s^*$, more and more households will switch to $s_i \approx 0$, making s decrease until $s < s^*$, after which more and more households will switch to very large s_i , making s increase again, leaving fewer and fewer households at intermediate values of s_i . This explains the emergence of two separate classes and extreme inequality in the oscillatory regime of our model. In both cases, those households with a countercyclical savings rate can be seen as “free-riding” on the behaviour of the others. When $s < s^*$, those households with large s_i accumulate capital and get high rents due to an overall scarcity of capital. When $s > s^*$, wages and thus incomes increase due to the overall abundance of capital, so those household with small s_i profit from consuming a larger share of their temporarily large income.

Dependence of the critical interaction time

In Fig. S4 and Fig. S5, we show an alternate views of Fig. 4 in the main text. Here we show the dependency of τ_c on the mean number of neighbors $\langle k \rangle$ and the number of households N and the link density p . We see no clear alignment without using the average shortest path length as done in the main text.

Examples of qualitatively different oscillations

Varying the network architecture and the number of agents simulated not only changes the value of τ_c but also the frequency of the oscillations. We show 300 time step intervals in the oscillatory regime for various simulations using Erdős–Rényi random graphs [S1] in Fig. S6 and Watts-Strogatz “small world” [S2] networks in Fig S7. We can find

that decreasing the number of agents and the social interaction time τ increases the roughness of the aggregate capital curves indicating multiple oscillation frequencies with different amplitudes.

Using the depreciation rate per time step δ to roughly correspond to the real-world depreciation per year, we find that the large oscillation frequencies are generally in the range of 50-100 years and smaller oscillations (e.g. lower panel in Fig. S7) in the range of 15-20 years. These results show that, while the main oscillation frequency is too low for a typical business cycle, this model can generate various oscillations with different frequencies that are superimposed on each other, similar to economic waves.

References

- [S1] Erdős P, Rényi A (1959) On random graphs. *Publicationes Mathematicae* 6:290–297.
- [S2] Watts DJ, Strogatz SH (1998) Collective dynamics of ‘small-world’ networks. *nature* 393(6684):440.

Supplementary Figures




Figure S1: Dependence of the high and low savings rates on the network topology. We show the mean values of the savings rates in the lower mode with $s < 0.5$, (left) and the higher mode with $s > 0.5$ as a heatmap. We can see a range of different values that depend on the topology of the network.




Figure S2: Effective discount rate as a function of the social interaction time. The corresponding discount rate $\rho(\tau) = \frac{\delta}{2}(e^{\delta\tau/2} - 1)^{-1}$ of our model is shown for various depreciation rates δ and low social interaction times $\tau < \tau_c$.




Figure S3: Best response dynamics. Optimal individual savings rate s_i^* (circles and triangles) given a very small perturbation of either $+\Delta$ (blue circles) or $-\Delta$ (orange triangles) of the aggregate savings rate s away from its equilibrium value s^* (green line) as a function of the social interaction time τ , for $\Delta = 10^{-7}$. Below a critical value of τ , best responses are very close to the equilibrium value, while above the critical value, they diverge quickly towards very small or very large values in the opposite direction of the perturbation, thereby stabilising s but leading to temporary profits due to the countercyclical exploitation of large rents (if $s < s^*$) or large wages (if $s > s^*$).




Figure S4: Critical interaction time τ_c as a function of the network mean degree. Here, τ_c is shown as a function of network size N for different mean degrees, given by $\langle k \rangle = Np$ for Erdős-Rényi graphs [S1] and $\alpha = 0.5$, $\delta = 0.2$ and the critical rate increases with the mean degree $\langle k \rangle$.




Figure S5: Critical interaction time τ_c against number of households. Here, τ_c is shown as a function of network size N for different link densities ρ for Erdős-Rényi graphs [S1] and $\alpha = 0.5$, $\delta = 0.2$. The critical interaction time scales proportionally with the number of households and inversely with the link density.

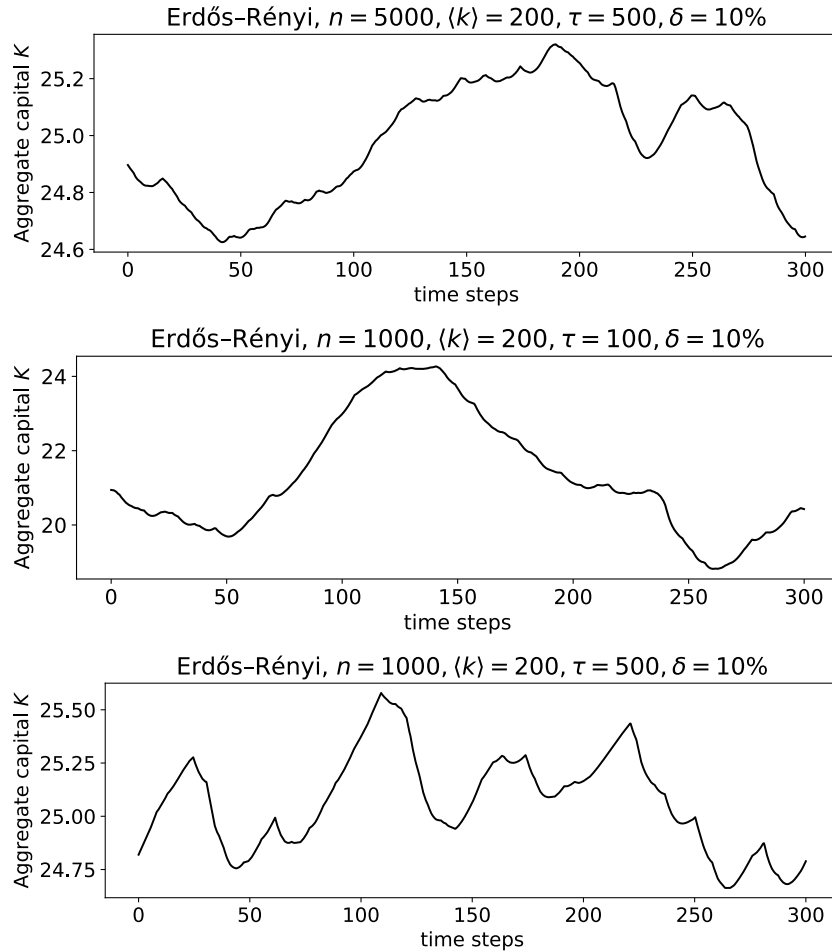


Figure S6: Various oscillation patterns for different Erdős-Rényi graphs. The parameters are given in the titles of the figures. We can see that decreasing the number of agents and decreasing the social interaction time τ increases the roughness of the aggregate capital curves. Note that labor is initialized as $1/N$ so that the equilibrium capital is constant for all simulations.

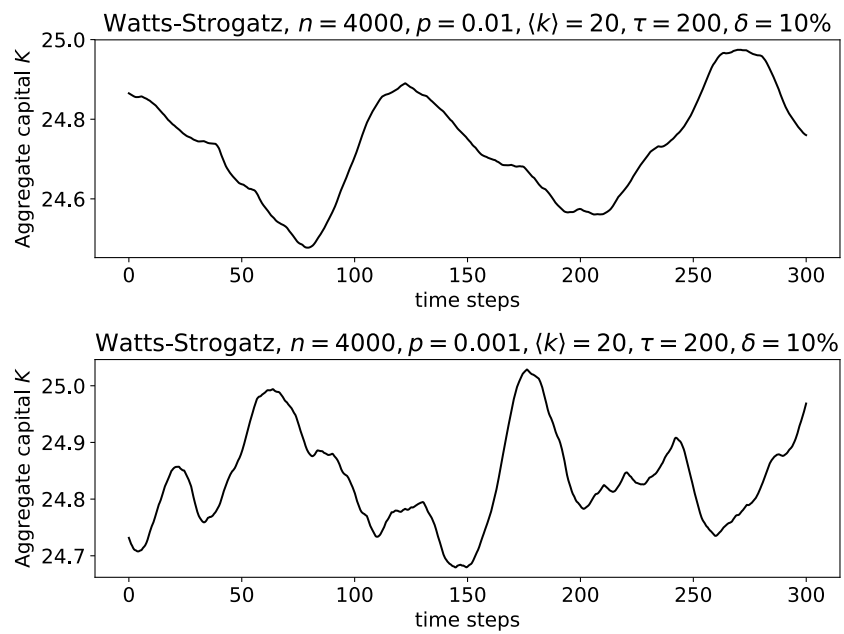


Figure S7: Various oscillation patterns for different Watts-Strogatz graphs [S2]. The parameters are given in the titles of the figures. We can see that decreasing the shortest paths length (by decreasing the rewiring parameter p , leads to a more rough aggregate capital curves with multiple frequencies. Note that labor is initialized as $1/N$ so that the equilibrium capital is constant for all simulations.