Real and financial crises: A multi-agent approach

Bill Gibson and Mark Setterfield

University of Vermont and New School for Social Research[1]

Abstract

Previous analyses of macroeconomic imbalances have employed models that either focus exclusively on real-side effects or financial-side disturbances. Real-side models usually make the unrealistic assumption that firms that save more than they invest effortlessly and costlessly transfer those surpluses to deficit firms, firms that require additional savings to sustain their plans for capital accumulation. On the other hand, there exists a well-developed, rigorous and elegant literature that uses the multi-agent systems (MAS) approach to analyze the recent financial crisis. These stand-alone models of the financial sector focus on the network structure of financial interplay but typically ignore real side interactions. In this paper paper, we develop a MAS model that integrates real and financial elements. The focus remains on the network structure and it is seen that randomly connected networks are more crash prone than are preferentially attached networks of financial agents. when real-financial interactions are taken into account. The results cast doubt on the connection between systemic risk and financial entities that are "too big or too linked to fail."

Keywords: Systemic risk; Crash; Herding; Bayesian learning; Endogenous money; Preferential attachment; Agent-based models., JEL codes D58, E37, G01, G12, B16, C00

1. Introduction

The motivation for this paper is the "great immoderation", the rise in the number of significant stock-price drawdowns over the last three decades.[2] The paper employs a multi-agent systems (MAS) model to study the effects of adding a real sector to a relatively well-developed and now standard analysis of the financial system that arises out of network analysis (Acemoglu et al., 2012). The model builds on an existing real-side MAS model, due to Setterfield and Budd (2011), which has roots in the structuralist tradition (see Taylor

---

[1]Thanks to Amitava Dutt, Jerry Epstein, Diane Flaherty, Arjun Jayadev, Blake LeBaron, Victor Lesser, Suresh Naidu, Andre Neveu, Rajiv Sethi, Gil Skillman, Peter Skott, Daniel Thiel and Roberto Veneziani for useful comments on earlier versions. Programming assistance by Paul Wright is also gratefully acknowledged. We with to thank Jeannette Wicks-Lim for making available advanced computational facilities of the Political Economy Research Institute for the project.

[2]Financial crises in various guises are usefully surveyed in Reinhart and Rogoff (2009), who identify the build up of debt, either public or private, as the principal cause of financial collapse.

(1983), Taylor (1991), and Setterfield (2010)). A financial sector, inspired by the MAS financial models of Johansen et al. (2000), Sornette (2003), Harras and Sornette (2011), Tedeschi et al. (2012), LeBaron (2012), Thurner et al. (2012), and others, is then added. The result is a model of the economy in which financial intermediation can constrain investment spending by firms, and hence the pace of growth in the real sector. Meanwhile, the profits and savings generated in the real sector affect the ability and willingness of financial intermediaries to lend. The conclusion is that real-financial interactions increase the likelihood of crises, while preferentially attached financial networks decrease the probability of financial collapse. The results suggest that the often supposed connection between systemic risk and financial entities that are "too big or too connected to fail" may be oversimplified."

The paper is organized as follows. The next section describes the simulation model and how it is designed to capture real-financial interactions. The role of endogenous money is identified and the way in which financial forecasts and asset prices are modeled is elucidated. A major focus of this section is on network structure, whether financial agents are linked randomly or by way of preferential attachment. Both weighted and unweighted networks are considered. The third section presents the simulation design and results, with attention to how crashes are identified, simulation settings and descriptive statistics of the large volume of output data. The fourth section concludes. An appendix contains the pseudo code of the model.

2. The simulation model

Real side models of any economy are complicated as are stand-alone financial models. It follows that a model that seeks to combine the two can easily become unwieldy without some well chosen simplifying assumptions. While there are many ways in which an amalgam of the real and financial sectors could be configured, the model of this paper is set out on a grid with agents divided into two subsets, firms and financial agents. Simple behavioral rules are defined for these agents and the macro performance of the model arises as an emergent property of the interaction of the agents. There is no policy authority and all adaptation is through reinforcement rather than directed learning.

Figure 1 illustrates the three principal structural features of the multi-agent model, firms, financial agents, and financial network architecture. Each firm, represented as a colored square, operates only one production process, combining capital and labor to satisfy demand for a homogeneous good.[3] The dilemma for firms is whether to invest. This decision is based on expectations about future market conditions and firms' ability to cover any short-fall in planned investment over accumulated savings through borrowing from the financial system.

Firms generate financial surpluses when they save more than they invest and these financial surpluses are held by their financial agents, represented by numbered circles in figure 1. Financial agents must then decide how to allocate these surpluses across the grid. Their

---

[3]This is the easiest way to think about the adjustment mechanisms in the model. Equivalently, there could be heterogeneous goods with prices adjusting behind the scenes. Price movements would shift investible surplus from one firm to another, but the overall macro properties of the model would remain unchanged.

Figure 1: Principal structural features of the model
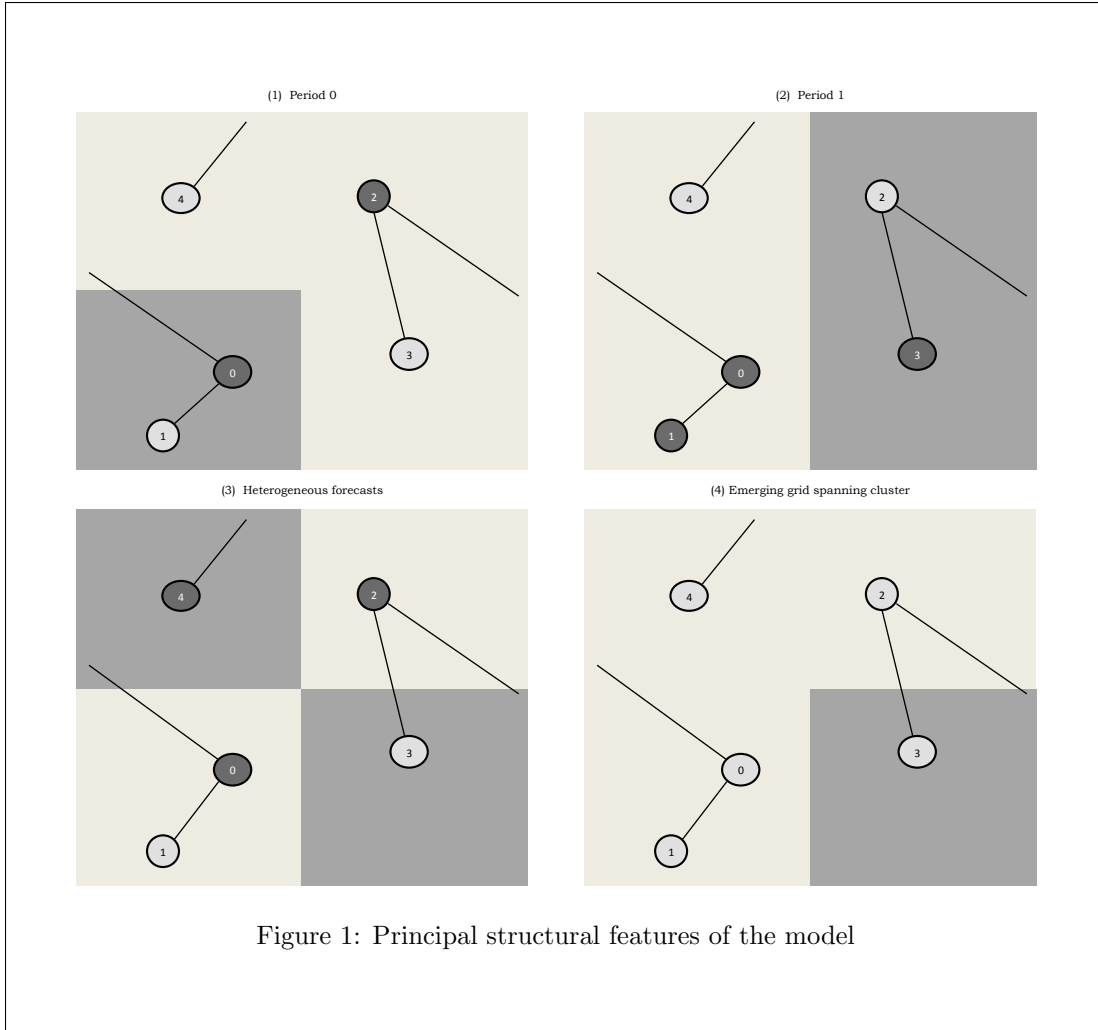
dilemma is also binary: lend/do not lend. This decision depends on financial agents' forecasts of market conditions. Like firms, the financial agents are elementary agents, with only one on-off decision to make. Consequently each financial agent can at most offer financing for one production process. Figure 1 shows, however, that firms can have access to more than one financial agent. Financial agents can be assumed to have first ranked their clients and will only deal with their most preferred. Finance for firms, by contrast, is fungible and given a fixed and constant interest rate, firms need only be concerned with the quantity of finance available.

Network architecture is represented by the edges joining the numbered financial agents in figure 1, and influences how financial surpluses are allocated.[4] Suppose that the surplus firm in the first panel of figure 1, shown in gray, makes a deposit with financial agent 0. For the

---

[4]Network architecture also influences financial agents' expectations, which are based (in part) on the forecasts of their linked network neighbors. See section 2.2.2

moment, take financial agents' expectations as given, denoted as light ("bearish") and dark ("bullish"). Financial agent 0 is connected to financial agent 1 and through torus wrapping to financial agent 2. Financial agent 1 is bearish but in any case has nothing to do, since the firm to which she is willing to lend is in surplus. Financial agent 2, however, serves a deficit client and is connected to financial agents 0 and 3, who is also serving a deficit firm. Financial agent 2 then calls on financial agent 0 to make the surplus available for her client and savings is thereby channeled from a surplus to a deficit firm. Financial agent 3, however, is stymied, since indirect access to financial agent 0's funds is not permitted in the model.[5] In any event, financial agent 3 is bearish and would block any flow of finance that happened to be available. Financial agent 4 is not part of the cluster of financial agents discussed here, but is connected to financial agents on other parts of the grid not shown in figure 1.

These behavioral decision rules are executed asynchronously thereby allowing the possibility of collisions, conflicting claims on financial resources, on the grid. At the end of each period, firms deposit their savings out of profits determined by a given savings propensity. At the beginning of the following period, financial agents respond to demand for this liquidity from both surplus *and* deficit firms. If a deficit firm applies for a loan before the surplus firm has had a chance to invest the funds, a collision is created. The conflict is resolved claims that can only be resolved by allowing financial agents to create money.[6] The endogeneity of the money supply is critical to the macro performance of the model, since if investment were always constrained by savings, there would be no room for expectations-driven growth.

The second panel of figure 1 shows the same neighborhood one period later. Note that the network architecture, the number of firms, and the number and location of financial agents remains fixed. Observe that the surplus firm of the first period is now in deficit and the formerly deficit firms in the north-east and south-east positions are now in surplus. The forecasts of financial agents have also changed. Generally, forecasts remain heterogenous as shown in the third panel of figure 1. It is possible for a grid-spanning cluster of opinion to arise, however, as illustrated in the fourth panel of figure 1 where all financial agents are bearish. As explained in detail below, this spanning cluster sets the stage for a financial crash. The probability of the formation of a grid-spanning cluster measures the systemic risk of the system.

Even under this rash of simplifying assumptions there is necessarily complex machinery supporting how the central decisions of the two agent sets are made and joined by the network structure. The following sections describe more precisely how firms decide to accumulate capital and how financial agents decide whether to assist in the process or throw sand in the gears.

---

[5]One alternative is to allow each financial agent complete access to *all* other financial agents on the grid, an obviously unrealistic assumption. To capture the locality of intermediation, information constraints and boundedness on the rationality of the financial agents, some line must be drawn. For simplicity, the line is drawn here at one ply.

[6]This simple model thus reflects the hidden presence of a monetary authority that does not allow credit creation, except when these collisions occur. See section 2.2.1 below for a fuller discussion of this assumption.

2.1. The real side

The real side of the model is derivative of Setterfield and Budd (2011), in which a standard structuralist model is recast as a multi-agent system. They treat each firm as a separate economy with its own growth dynamics, essentially as in a trade model. Here a second approach is adopted, in which each firm operates a productive process within a macroeconomy, thereby sharing in available aggregate demand at a given instant in time. The re-allocative mechanism is simply to adjust demand shares so that any demand that would cause the capacity of the ith firm to exceed one, spills over to the jth.[7] If the overflow causes the capacity utilization of the selected firm to rise above one, then its overflow is allocated to another randomly selected firm and so on. The mechanism halts when all firms are either at or below full capacity utilization.

Following Setterfield and Budd (2011) assume that labor income per unit of output, l_i, is spent on consumption while a fraction, s_i of capitalists' income, $\pi_i = 1 - l_i$, is saved.[8]

Investment by the ith firm is $I_i = g_i(u_i)K_i$ where K denotes the capital stock and $g(u)$ is an accumulation function that depends on capacity utilization, u_i, and animal spirits, α_0 of the form

$$g = \alpha_0 + \alpha_1 u + \alpha_2 rate \tag{1}$$

where the α's are calibrated constants and the rate of profit is $r = \pi X/K$.[9]

Using carets to denote diagonal matrices, the real side can be expressed compactly as

$$\mathbf{x} = \hat{\mathbf{Q}}\mathbf{u}$$

where $\mathbf{x}$ is an n-element column vector of firm outputs, $\hat{\mathbf{Q}}$ is a diagonal matrix of full capacity output and $\mathbf{u}$ is a column vector of capacity utilization rates. The system can be written

---

[7]Since the solution algorithm allocates demand randomly among firms on each sweep, the ith and jth firms are well defined at runtime and, moreover, no bias is introduced. Of course if no firm reaches full capacity utilization in a time period, demand shares need not be adjusted.

[8]Each firm hires workers according to a uniform probability distribution $U(0.3, 1/30)$. How the savings rate, s_i, is calibrated is discussed in the next footnote.

[9]To calibrate equation 1, the animal spirits intercept, α_0 is set at half the randomized value of g and the coefficient on capacity utilization. The coefficient α_1 is uniformly distributed between 0 and $g/4$. The calibration of the last coefficient is a residual and is undertaken as follows. The initial growth rate of each firm's capital stock is randomly distributed according to $g \sim U(0.07, 1/300)$. Each financial agent has a portfolio, the nominal value of which is set to unity. The initial capital stock, K, of firms then depends on the number of associated financial agents. Output is given by applying a capital-output ratio, ζ randomly distributed according to $\zeta \sim U(3, 1/6)$. Output in this case is capacity, Q; full utilization of capacity, $u = 1$ is initially assumed for simplicity but can vary away from 1 as the model runs. With output, X, and K known, r is determined and so α_2 must be calibrated as a residual in equation 1. To see how the savings rate s is calibrated, note that the rate of investment for the ith firm, is $I_i = g_i K_i + \delta K_i$ so that with a common rate of depreciation, $\delta = 0.05$, the rate of investment for each firm is known. Summing the investment functions over the 676 firms gives a total value for aggregate investment and thus total savings and the aggregate savings rate, $\bar{s}$. Each individual firm's savings rate is then initially normally distributed according to $s_i \sim N(\bar{s}, 0.0025)$, with the mean depending on each initialization. The savings rates are then scaled to give total savings equal to total investment so that the initial solution is in macroeconomic equilibrium.

with output as a function of consumption demand $\mathbf{D}$ and investment, $\mathbf{gK_{t-1}}$ as

$$\hat{\mathbf{Q}}\mathbf{u} = \hat{\lambda}\left[\mathbf{D}\hat{\mathbf{Q}}\mathbf{u} + \mathbf{gK_{t-1}}\right]$$

where λ_i is the share of aggregate demand absorbed by the ith firm, and where

$$\mathbf{u} = \begin{bmatrix} \mathbf{u} \\ \mathbf{1} \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} 1 - s_1\pi_1 & 1 - s_2\pi_2... & 1 - s_n\pi_n \\ 1 - s_1\pi_1 & 1 - s_2\pi_2... & 1 - s_n\pi_n \end{bmatrix}, \quad \mathbf{gK} = \begin{bmatrix} g_1K_1 & g_2K_2 & ... & g_nK_n \\ g_1K_1 & g_2K_2 & ... & g_nK_n \end{bmatrix}$$

Since both $\hat{\lambda}$ and $\hat{\mathbf{Q}}$ are diagonal and therefore symmetric the system can be solved for $\mathbf{u}$

$$\mathbf{u} = \hat{\lambda}\left[\hat{\mathbf{Q}}^{-1}\mathbf{D}\hat{\mathbf{Q}}\mathbf{u} + \hat{\mathbf{Q}}^{-1}\mathbf{gK_{t-1}}\right] \tag{2}$$

The computational model solves the vector equation by way of the Gauss-Seidel method. At each iteration, the vector of λ's is updated to reflect the firm's new share of total aggregate demand that results from either its having encountered the full capacity constraint *or* having to satisfy spillover demand from a firm that has reached full capacity. The order in which firms are allocated spillover demand is random, so that no particular firm benefits from the procedure. Note that there is no implicit optimization of output here; the program simply looks for a basic feasible solution to a simultaneous set of demand equations under the constraint that no level of capacity utilization can exceed one. Once it finds a basic feasible solution, the Gauss-Seidel halts. Without this mechanism, firms could produce an unbounded quantity of output without regard to the necessary factors of production.[10].

2.2. The financial sector

Financial agents are Bayesian rational network learners, updating priors by reference to the real-side performance of their own clients. The key parameter of the financial network is the the degree distribution of edges. Here, both random and scale-free networks are considered. In random networks the degree distribution is Gaussian whereas scale-free network architecture depends on a scaling parameter γ.[11] The probability that a given financial agent is connected to d other financial agents is then given by

$$zp(d) = d^{-\gamma}$$

where z is a scale-invariant calibration constant.

The degree distribution $p(d)$ depends on the specific architecture of the network. Following Erdös and Rényi (1950) a random network is constructed in the multi-agent system of this paper by instantiating financial agents and then linking them to exiting financial agents

---

[10]The initial shares of aggregate demand are set in proportion to Q_i

[11]Scale-free networks have a power-law degree distribution of links, which means that a few nodes have a very large number of edges, while most nodes have a small number. A network in which 20 percent of the nodes are connected by 80 percent of the edges is a Pareto-Zipf distribution or power-law degree distribution (Barabási and Albert, 1999; Bagrow et al., 2008).

randomly. A second mechanism, due to Barabási and Albert (1999), employs *preferential attachment*. In the setup routine, the likelihood that an additional financial agent created is linked to an existing financial agent now varies directly with the degree of the potential partner.[12] Once initialized, the number of financial agents is held constant for simplicity. In a random network, the spread of shocks requires that a certain minimum number of vertices are impacted, whereas in a scale-free network, this threshold is essentially zero (Iori et al., 2008). By generating networks both randomly and through preferential attachment, whether the latter makes the financial system more crisis prone can be tested.[13]

In addition to the degree distribution of links between financial agents, the model of this paper also applies weights to distinguish larger and more influential financial agents from the rest. Each financial agent is associated with a real-side production process but the latter have a distribution of capital stocks, which in turn implies that their associated financial agent may wield more power and influence than the degree distribution alone would capture. The weight, ω_{ij}, attached by the jth financial agent to the link to the ith agent is defined as

$$\omega_{ij} = \frac{K_i}{\sum_{f \in F} K_f} \tag{3}$$

where F is the set of firms and the t subscript has been suppressed. The weight attached to a link by any financial agent is then the share of the total capital stock served by the agent's linked neighbor. Note that in models without associated real sectors, no such weighting scheme naturally suggests itself.

The *weighted* degree distribution integrates the interconnectedness of the financial sector with the real side. The impact on the stability of the system as a whole remains, however, an open question to be decided by numerical simulations. Weights could, for example, dampen random fluctuations in the outlook of smaller financial agents, thereby reducing the propensity to crash in the aggregate. If a few decide to block loans to their own customers, there is little effect on aggregate capacity utilization. On the other hand, a bearish outlook on the part of an influential financial agent could block investment and reduce demand, quickly infecting other agents with its pessimism. These two effects may even cancel each other out.

2.2.1. Endogenous money

As noted a deficit firm may well have already contracted with a given financial agent to borrow the funds that a surplus firm has deposited and plans to invest but, due to the asynchronous nature of the computational model, has not yet done so. If the funds are so

---

[12]For example, suppose that financial agent A is already linked to both financial agents B and C, each of whom are connected only to financial agent A. There are thus four network connections in total. The probability that a new financial agent, D, will link to financial agent A is $2/4 = 1/2$, whereas the probability that D will link to B (or alternatively, to C) is $1/4$.

[13]Note that preferential attachment is a simplification that comes at the cost of precluding *clustering*. This is ruled out since no existing financial agent can add a link to some other existing financial agent that has already attached itself to the network.

preempted, the financial agent is left with no choice but to create liquidity when surplus firms demand their own funds. Deficit firms can and do crowd out other deficit firms due to the limitation given by the total surplus. They cannot, however, crowd out surplus firms.

Observe that rather than *assume* a Keynesian framework in which animal spirits have room to play a defining role in the determination of investment and growth, the agent-based perspective allows a derivation of the Keynesian nature of the economy. Specifically, the regime of property rights in banking–that agents own and control their deposits–combined with an asynchronous stream of borrowing requests by deficit firms imparts to a model a Keynesian flavor. The reason is now clear: it would only be if a surplus firm were to realize and accept that its investment plans were blocked because a deficit firm *had beaten it to its own money* that the neoclassical savings in advance model could assert itself. No conceivable real agent would behave this way in a system that protects private property in deposits. Animal spirits ultimately allow aggregate investment in period t to exceed savings in period $t-1$ and, at least in part, these animal spirits are authorized by the institutional properties of the financial system. In other words, the agent-based model provides a micro-foundation for *partially* endogenous money and in turn the Keynesian system itself.[14]

Of course there would be no reason to distinguish surplus and deficit firms if credit or money were *fully* endogenous. Firms that lacked sufficient savings from the previous period would simply borrow for investment from financial agents who are, in turn, able to create money. Creating money is nothing more than making a loan unsupported by previous deposits and the point of this section is that any institutional framework that binds autonomous agents with financial surpluses and deficits in asynchronous exchanges will necessarily produce some residual quantity of endogenous money.

Figure 2 provides a sensitivity analysis of the basic assumptions of the model. The upper trajectory corresponds to full capacity utilization, with fully endogenous money, and ebullient animal spirits. The lower trajectory is a bound along which no deficit firm preempts a surplus firm and so no money is created.[15] While highly unlikely, the savings constrained trajectory is included for completeness.[16] With asynchronicity giving rise to partially endogenous money, a typical simulated economy will find itself somewhere between the extremes of figure 2. The mean trajectory shown in the figure is derived from a large number of paths of the simulation

---

[14]Note that as modeled, the financial sector reflects the spirit of Kalecki's principle of increasing risk, by making the execution of planned investment easier for surplus firms, which are investing their own capital, than for deficit firms that need to borrow in order to invest (Kalecki, 1937).

[15]The economy described by the lower path will eventually turn down. If any deficit firm is rejected by a financial agent investment is blocked and savings re-equilibrates at lower level. The decline in aggregate demand can then cause firms that would otherwise have been in surplus to fall into deficit, increasing the probability that investment plans will be blocked by a shortage of surplus, bearish sentiment, or local availability on the grid. An isolated event then propagates and in so doing imparts a negative trend to GDP as seen in figure 2.

[16]The upper bound is, in reality, just as unlikely since it presupposes that *no* financial agent will ever block the investment plans in the real sector. The potential instability provoked by such a system of unrestricted credit was made abundantly clear in the recent financial crisis.

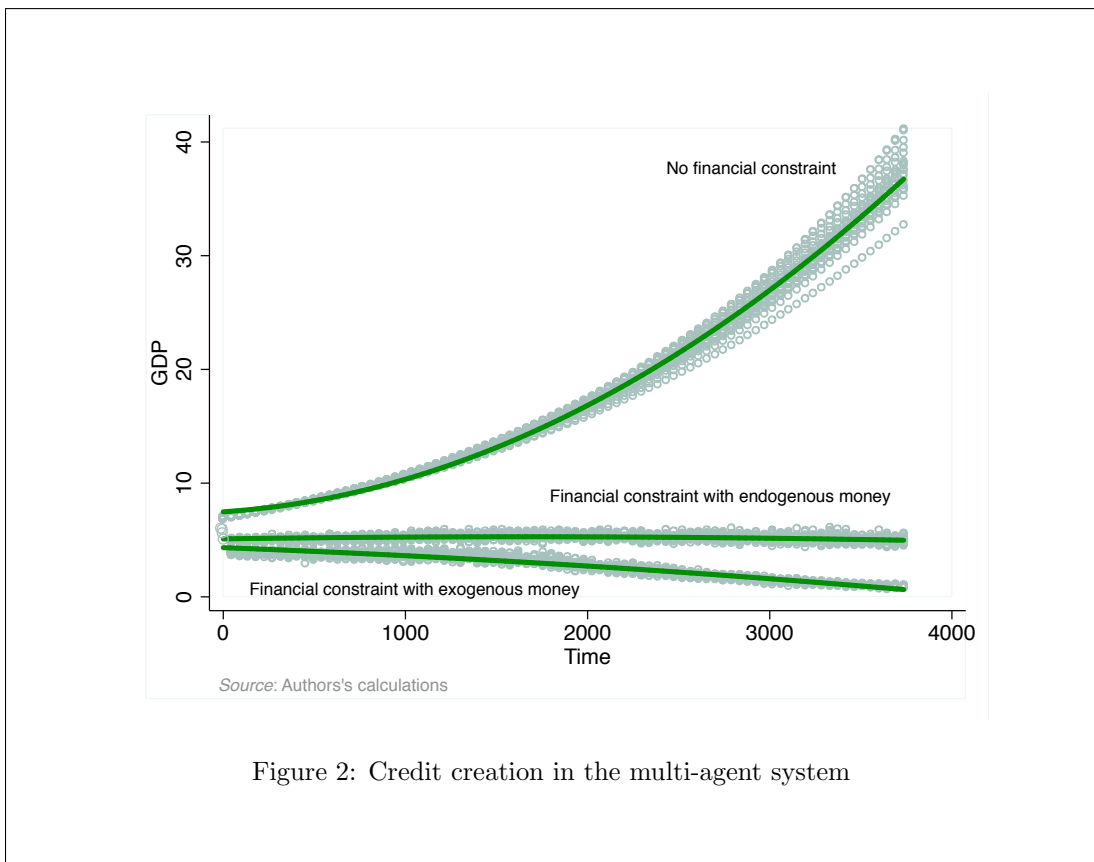model described in more detail below.[17]

Figure 2: Credit creation in the multi-agent system

2.2.2. The financial agent's forecast

Conditional on having sufficient liquidity, a financial agent will provide finance to a client firm if the financial agent is optimistic that the value of these shares will rise over time. Lacking perfect foresight, financial agents form an opinion about the distribution of the change in share prices in the form of a forecast. In each period, t, financial agents update their forecasts consistent with the Bayesian model. In the Gaussian version of the Bayesian model, this means that each financial agent has a subjective probability distribution of the change in share prices.

The updating proceeds as follows: θ_t is an observation from the stationary distribution of share price changes, $N(\bar{\theta}, \sigma^2)$. Each financial agent (identifying subscript suppressed) has a prior subjective distribution $N(\phi_p, 1/\rho_p)$, where ρ_p is the precision of the prior signal, the inverse of the variance.

---

[17]The simulations do not claim to be representative of any actual path in a real economy. The trajectories in figure 2 were engineered to have a nearly zero slope in order to reduce the probability of a purely real side crises that would then propagate to the financial sector, as explained in more detail below.

As in standard Gaussian Bayesian updating, each financial agent then receives an informative private signal, not observable to other financial agents

$$\mu_t = \theta_t + \varepsilon_t$$

where $\varepsilon_t \sim N(0, 1/\rho_\varepsilon)$ where, ρ_ε, is the precision of the private signal.

Each financial agent uses this signal to update their prior precision

$$\rho_t = \rho_p + \rho_\varepsilon$$

Note the precision of the prior signal is improved by an informative private signal. In period t, the financial agent's *forecast*, ϕ_t, of the state of the economy, θ_t, is then just the updated mean of the Bayesian prior distribution, ϕ_p,

$$\phi_{jt} = \kappa\mu_t + (1 - \kappa)\phi_p$$

where where $(1 - \kappa)$ is the weight on the Bayesian prior and

$$\kappa = \frac{1}{\frac{\rho_p}{\rho_\varepsilon} + 1}$$

There are many ways an agent's private signal, μ_t, could be modeled. In modern financialized economies agents frequently employ sophisticated data analysis. A simple approach would then be to model the private signal as a function of the current capacity utilization of the financial agent's client. Here the financial agent records the level of capacity utilization in a private data base and then forms a forecast by regressing a subset, τ, of these utilization rates on time according to

$$u_t = \hat{\beta}_0 + \hat{\beta}_1 t + \epsilon_t$$

with $t = 1, 2, ..., \tau$, with ϵ_t iid. Since different financial agents weight history differently, the size of the subset, τ, varies randomly between 3 and the 13 time periods. The private signal, μ_t, is thus

$$\mu_t = \hat{\beta}_1 + \nu_t$$

the trend regression slope coefficient plus random error, ν_t. A positive trend is associated with a bullish private signal. The error term allows financial agent intuition to dampen or amplify the signal and thus allow for contrarian behavior in the model.

The mean of the prior distribution, ϕ_p, is determined by the share weighted average of the forecasts of the financial agent's linked neighbors, denoted as set J'

$$\phi_p = \sum_{i \in J'} \omega_{ij}\phi_i$$

where ω_{ij} are the share weights in equation 3. A positive forecast is, again, a bullish prior of intensity depending on the same Bayesian updating undertaken by each of the agent's linked neighbors.

The relative precision of the private signal is determined by forecasting error, which in turn depends on how the market actually behaves. Given the unboundedness of a forecast error for any agent and the requirement that the forecast error determine the relative precisions of the signal and prior, it is necessary to map the forecast error into an open interval $(0, 1)$.

This is done in a two-step process. Suppose that the slope coefficient on a financial agent's regression is positive. If the price does increase, $\Delta p > 0$, then the financial agent has made a correct forecast. If the financial agent's linked neighbors forecast a price *decrease*, the financial agent has "beaten the market", demonstrating expertise relative to the linked neighbors.[18] If the financial agent beats the market in any given period, the financial agent records binary success of 1 and 0 if the market beats the financial agent. In the case in which the financial agent and linked neighbors forecast the same price movement, the financial agent records the average of the binary signals. The financial agent then computes an average of these signals over the idiosyncratic τ history to arrive at a raw forecast error. In the second step, the financial agent attenuates the raw forecast error by way of a logistic smoothing function, as is common in machine learning models of neural networks (Russell and Norvig, 2010).[19]

Specifically, the evolution of κ_t is modeled as follows. First, the financial agent processes the raw forecast error as follows

$$\mu_t \times \Delta p > 0 \quad \begin{cases} \phi_p \times \Delta p < 0 & \text{set } \eta_t = 1 \\ \phi_p \times \Delta p > 0 & \text{set } \eta_t = 0.5 \end{cases}$$

$$\mu_t \times \Delta p < 0 \quad \begin{cases} \phi_p \times \Delta p < 0 & \text{set } \eta_t = 0.5 \\ \phi_p \times \Delta p > 0 & \text{set } \eta_t = 0 \end{cases}$$

Next, the average of the η_t over the last τ periods determines the position on the logistic smoothing function

$$\kappa_t = \frac{1}{1 + e^{[\gamma_1 - \gamma_2(1 - \frac{1}{\tau} \sum_{t=1}^{\tau} \eta_t)]}}$$

where γ_1 and γ_2 are calibrated to center the logistic function on 0.5 with some random variation for each agent.[20] The weight on the private signal varies logistically between 0 and

---

[18]There is a long tradition in economic theory that holds that market prices reflect all (publicly) available information, and apart from insider trading, "expertise" does not really exist. Financial agent are successful only because they are lucky. Agents in models with endogenously determined variable signal precision cannot be said to be actually learning anything about their reward field, but they may behave as if they think they can. Linked neighbors may be fooled by their apparent success without actually learning anything.

[19]Consider for example a financial agent who initially places equal weights on both the private and prior signals. The horizon for this particular financial agent is $\tau = 3$. In the last three periods, the financial agent beat the market, recording $\eta = 1$ each time. Does the financial agent stop updating, concluding that the precision of the private signal is then one? To avoid this unrealistically volatile behavior on the part of the financial agent, the model assumes that successful forecasts are "squashed" by the logistic function of equation 2.2.2.

[20]The constants are uniformly distributed $\gamma_1 \sim U(5.5, 1/12)$ and $\gamma_1 \sim U(9, 1/3)$.

1, and is inversely related the forecast error. The logistic function prevents abrupt swings in the relative weights of the private and public signals in any given period.

The evolution of κ described above is consistent with standard "spin-glass" models. Note that if the financial agent's network neighbors consistently forecast asset price movements correctly while the financial agent's private signal is always at odds with asset price movements, relative signal precision and hence κ will fall, and its complement will increase at a decreasing rate toward one. The informational content of the capacity utilization of firms is ignored as the spanning cluster begins to form and social learning breaks down, giving way to herding behavior and in the limit, when all financial agents are herding, cascades (Chamley, 2004). Financial agents become progressively detached from the source of informative signals, their own firms and the firms of their linked neighbors. The more the informational bond between the real and the financial sector breaks down, the more likely order, the precursor of financial collapse, becomes. Systemic risks begins to rise (Hansen, 2012).

2.2.3. Preferential attachment and lending

If lending agents do not have the liquidity necessary to satisfy the demand for borrowing by deficit firms, they may call on their linked neighbors to ask for a loan. Linked financial agents who have sufficient funds can agree to loan the originating financial agent the balance to meet the demand for liquidity of the deficit firm. The depth of this lending relationship is limited to one ply as noted above.

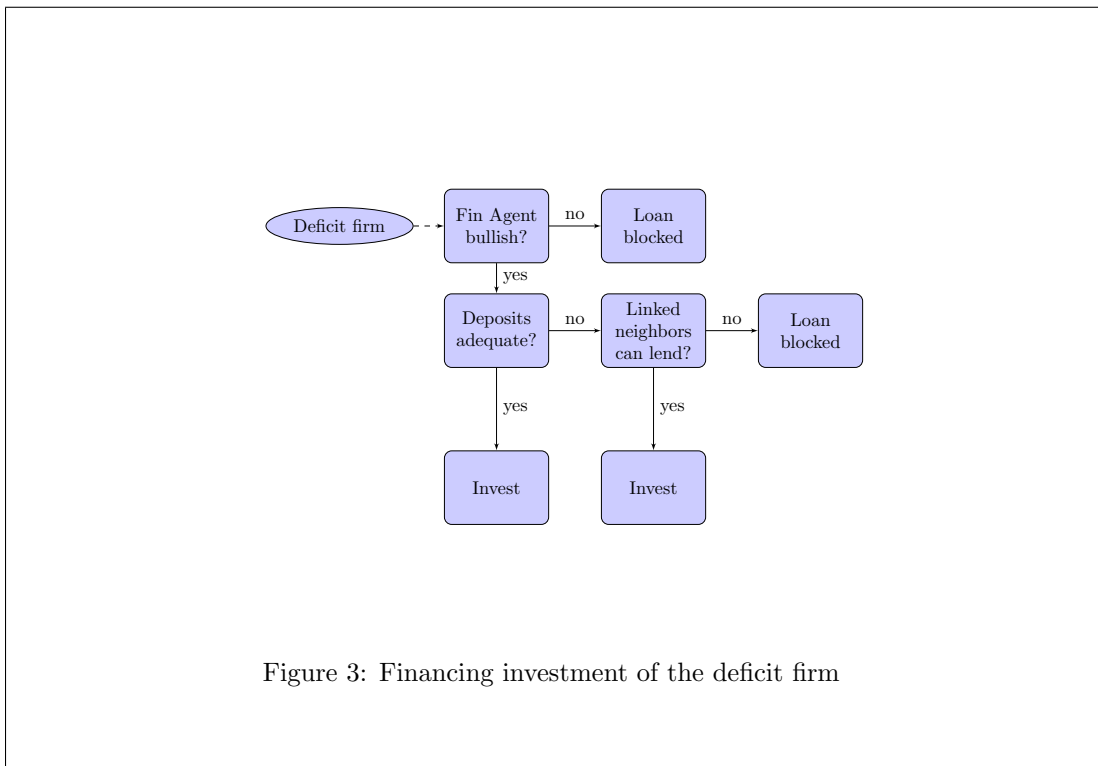Figure 3: Financing investment of the deficit firm

Figure 3 describes the decision tree with a deficit firm petitioning a local financial agent for a loan. If the financial agent is bearish, with a negative forecast, then the loan is denied and the investment is blocked. If the financial agent is bullish but short of liquidity, the financial agent asks his linked neighbors if they can make the funds available. Crucially, the loan does not depend on the forecast of the linked neighbors. The loan originator in this way bears all the risk, while counter-party surveillance of the originator by the linked neighbor is effectively nil.[21]

Again the structure of the network is hypothesized to be crucial to systemic risk. With preferential attachment, bullish financial agents with many linked neighbors will be more able to make loans, and so their ability to finance deficit firms will rise. A firm associated with one of these financial agents is much more likely over time to find finance for any deficit that might arise. As a result, firms with highly linked financial agents will tend to accumulate capital stock more easily and grow larger over time.

As a high-degree hub turns bearish, however, it blocks loans to its own client if at that instant, the client happens to be in deficit. This, of course, reduces the level of effective demand in the system and to a greater degree since the firm's capital stock will likely be large. As large firms reduce their investment demand, other surplus firms may well go into deficit. Preferential attachment concentrates this kind of demand shock to the system in the hands of a few financial agents. On the other hand, highly linked, bearish financial agents having refused loans to their own clients then have more liquidity to pass along to their bullish linked neighbors. Sorting out the net effect of these currents and counter currents on aggregate demand, share prices and systemic risk is the job of the simulation model below.

2.2.4. Asset prices

The share price is defined as an index of financial agents' forecasts plus a trend based on the size of the aggregate capital stock.[22]

If the sum of forecasts is bullish, the share price rises and vice-versa, relative to the trend that depends on the accumulation of capital stock. The equation for the share price is

$$\Delta p_{t+1} = \psi \sum_{i \in J} \omega_{ijt} \phi_{it} + \psi_K \sum_i K_{it}^\alpha$$

where ψ is a parameter determining the scale of the marginal sensitivity of the share price to variations in financial agents's forecasts, ψ_K a constant is set at 5×10^{-5} and the elasticity

---

[21]This absence of of counter-party surveillance was one of the hallmarks of the recent financial crisis. Its deprecation could have been extended to a more lengthy and complex set of financial associations, but is limited to one ply here for model simplicity.

[22]For a more sophisticated stand-alone financial model, which enables financial agents to explicitly buy, sell or hold equities, see LeBaron (2012). In LeBaron, asset prices are set to clear the market for a fixed supply of equities. The simplified approach adopted here also captures most common stylized properties of asset prices, as seen below.

$\alpha = 0.5.$[23] The share price is a random walk during "normal" times, but breaks out during organized bull or bear markets to produce a bubble and then possibly a crash.

3. Simulations

3.1. Simulation design

To summarize the model, agents are divided into two disjoint subsets, firms and financial agents. A firm is instantiated with randomly assigned capital stock, savings rate, direct labor coefficient, and parameters of an investment function that depends on both animal spirits and capacity utilization in the previous period. Financial agents are instantiated with a forecast, an initial relative weight of private and prior signals, and linked neighbors within a network structure that depends on whether the model is run with or without preferential attachment and share weights. Financial agent must also keep track of liquidity and shares in the firms to which they make loans. Trades and production are tracked on a weekly basis.

The model generates real and financial performance over 1,500 week (30 year) periods.[24] The simulations are approximately ergodic since they are run for 2,250 trading weeks (45 years) prior to the 1,500 trading-week (30-year) period for which the data is recorded in order to mitigate the influence of initial conditions on the results. No crashes are counted during the first 2,250 weeks. There are then a total of 14,178 runs recorded in the data base, 3,599 runs with the financial constraint off and 10,579 runs with the financial constraint on.[25]

3.1.1. Identifying crashes

Using historical series for the S&P 500, it was determined that a typical build and crash involved some 225 weeks in total. A build is an increase in the share price from a period 225 to 25 weeks before a crash in period t.[26] Thereafter, a crash is a decline of 50 percent in the share price within the final 25 weeks. This roughly corresponds to the single worst six-month performance in the history of the S&P 500 index.[27]

---

[23]Initially, $\psi = 0.1$; thereafter, it evolves endogenously according to

$$\psi_{t+1} = \psi_t + \bar{\phi}_t$$

where $\bar{\phi}_t$ is the average forecast. This last term is included so that the stronger the financial agents' forecasts, in either direction, the greater the change in the share price.

[24]This chronology is advised by the observation of discrete episodes of growth, such as the Golden Age (1948-1973) or Neoliberal growth regime (1980-2007), lasting for 25-30 years and ultimately resolving in real and/or financial crises (Maddison, 2007).

[25]For computational reasons, the total number of runs per batch was 900 and there were 16 batches run for a total of 14,400 simulations. From these results, there were 221 runs deleted with the financial constraint on and one deleted with the financial constraint off. The larger number of runs with a binding financial constraint is needed to study the various configurations of network architecture used in modeling the financial constraint.

[26]The criterion for a build is designed to rule out a series that declines for a long period and then accelerates its decline.

[27]The decline in question took place during the 2008-09 financial crisis.

This (admittedly arbitrary) definition applied to every 225-week period in each run. If the pattern of price movements fits the crash specification, the program records a crash and halts.[28]
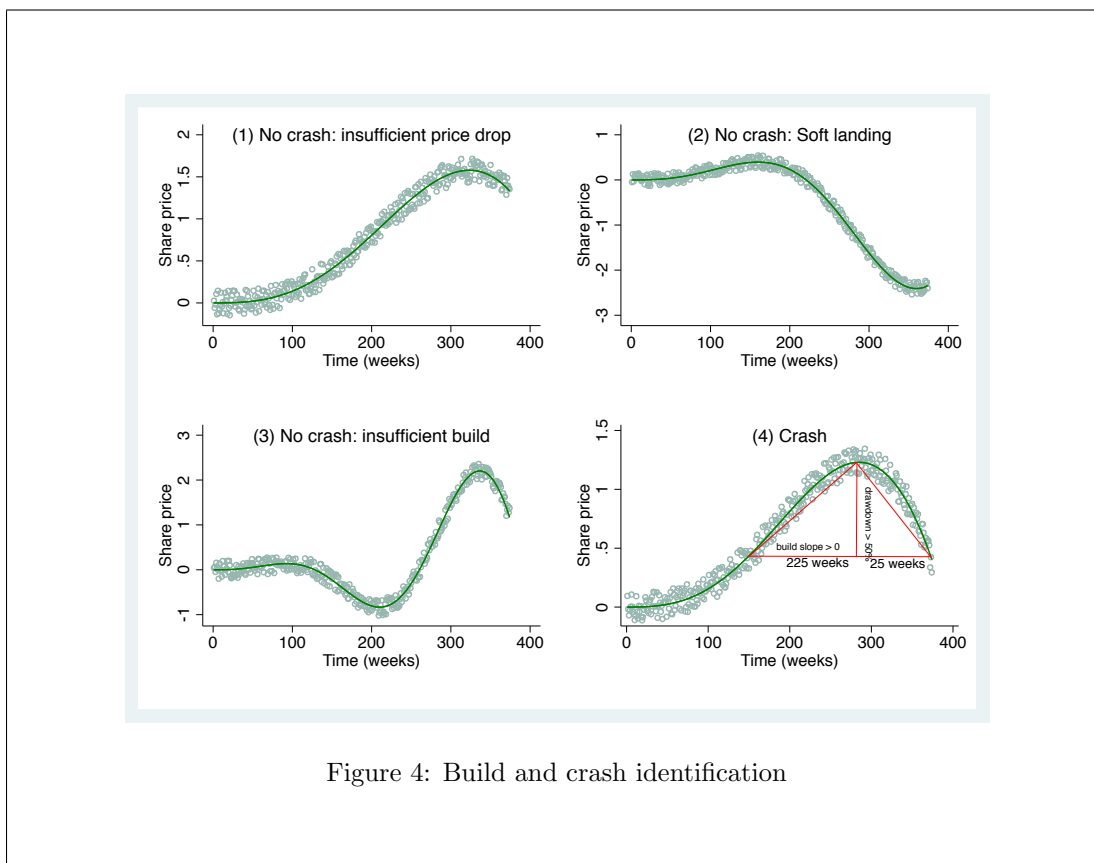

Figure 4: Build and crash identification

Note that the limited horizon for the drawdown, distinguishes a crash from a "soft-landing" orderly correction. The method does not necessarily find peaks or troughs and therefore cannot be used to quantify the size of a given drawdown.[29]

Figure 4 provides some examples of how the crash identification method works in the model. The build in the first panel of figure 4 is sufficient for a crash but the percentage drop of less than 50 percent disqualifies the drawdown. There is a much larger drawdown in the second panel of the figure, but note that the period over which the price declines is more than 150 weeks. In any given 25-week period the drawdown does not exceed 50 percent and moreover, even if it did, the build is insufficient to identify a crash. The third panel presents a more complex situation. There the build to week 100 is inadequate for a crash between

---

[28]Neftci (2008) provides a more sophisticated and comprehensive method to date crashes that could also be implemented. The method requires calibration however and thus could be adjusted to get approximately the same number of crashes as the vastly simpler procedure employed here.

[29]While no crashes are counted in the first 2,250 weeks, crashes that are defined by builds that begin in these weeks are included.

week 100 and 200, even though the drawdown qualifies. Note, however, that in later periods, a strong build is observed but the higher absolute value of the index limits the percentage change, thereby preventing the identification of a crash after week 350. The fourth panel finally identifies a crash. Note the superimposed triangle with the build slope greater than zero for 200 periods and drawdown greater than 50 percent.

3.1.2. A look at the model's dashboard

The primary distinction in the model is whether the financial constraint is binding. When money is exogenous, the real sector contracts, with capacity utilization rates falling below 60 percent as discussed above. This usually produces a financial crisis. The cases are excluded in the results discussed below. The crash identification methodology described above should then be thought of as conditional on average capacity utilization exceeding 0.6.

Financial networks are either random or subject to preferential attachment, and either weighted or unweighted by shares of capital stock. The resulting degree distribution that shows the frequency of high- and low-degree financial agents is illustrated for some 250 initializations in panels 1 and 2 of figure 5. In log-log space the distribution in panel 2 is not linear but concave. A power-law, in contrast, will produce a characteristically linear degree distribution, as shown in panel 1 of figure 5. The critical exponent is approximately 2, a number typical of networks constructed in this way. Panel 3 of figure 5 shows that the initial capital stock approximately follows a power-law distribution as in Axtell (1999). The initial Bayesian prior weight is approximately normal, with mean value slightly below 0.3 and standard deviation of 0.08.

Figure 6 shows the behavior of the natural log of the modeled price in four panels for all simulations in which a crash is recorded. Each displays only the last 250 weeks before the crash and the mean of the trajectory of the log price. The crashes are identified by the triangle method of panel four figure 4 as discussed above.

The are a number of important lessons to be gleaned from a comparison of these crash patterns. The first two panels show crashes when the underlying financial network conforms to panel 1 figure 5, that is when financial agents are preferentially attached. Crashes in both weighted and unweighted networks are shown. Panel 3 shows the pattern for the random network and it is seen that the variance of the log price is higher than with preferential attachment. Note also the density of the third panel. This suggests that preferential attachment does not make the economy more crash prone. Finally, the control case is shown in panel 4 for which there is no financial constraint and so crashes are due to Brownian motion only. There are few crashes and a more quiescent movement of the log price, confirming that there really is nothing of note happening in real models without some kind of financial constraint.

Real-side drivers are shown in figure 7. The first panel shows clearly that the when there is no financial constraint, the model always operates at full capacity utilization. There is no variation in capacity utilization from one run to another. Panel 2, however, tells a very different story. For all simulations with a financial constraint, the effect is a pronounced decline in the average utilization rate, shown by the heavy line. The variance in the utilization rate is substantial from run to run, but in all cases is distinctly different from when there is
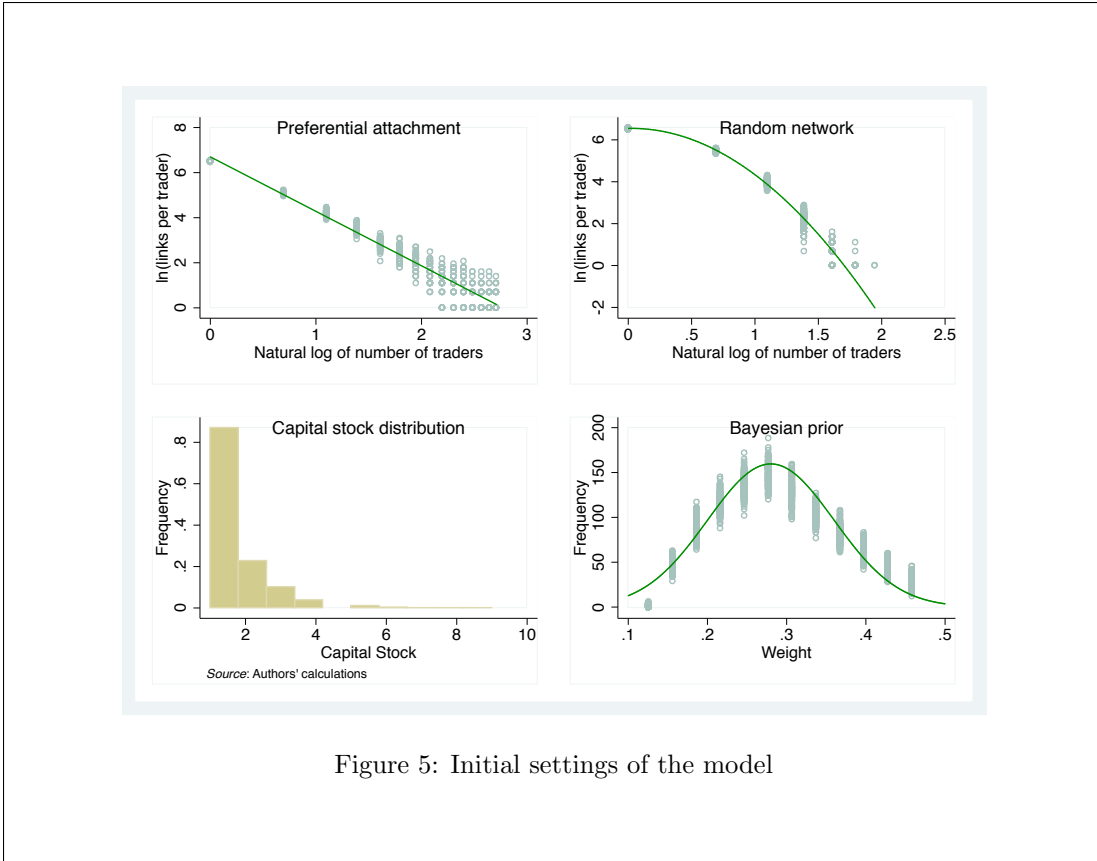
16

Figure 5: Initial settings of the model

no financial constraint.

Panels 3 and 4 show the weekly growth rate of GDP. The model was set to produce no net growth over time to prevent any bias in the tendency to crash that might arise. The result is a framework the departs from the recent macroeconomic data of the U.S. or for that matter any other country, but was deemed necessary to isolate the effect of the presence of the real side on the probability of a financial crisis. Clearly, a rapidly expanding or contracting real side can authorize financial panics or manias, but the question of interest in this paper is whether the mere presence of real-financial interactions is material to the explantation of financial collapse.

Despite the model's rather remote connection to any given interval of real side growth as captured by NIPA data, the financial side of the model is true to several underlying features of the actual economy, as represented by the S&P 500 index. Reality figured prominently into the definition of a crash above and a glance at the simulated log price shows that the "fat tail" property of the actual S&P is reproduced in the runs of the model. Table 1 shows the kurtosis of the log price for the runs with and without the financial constraint under various network settings.

The overall kurtosis of the distribution of the log price in the model is 4.28 with the financial
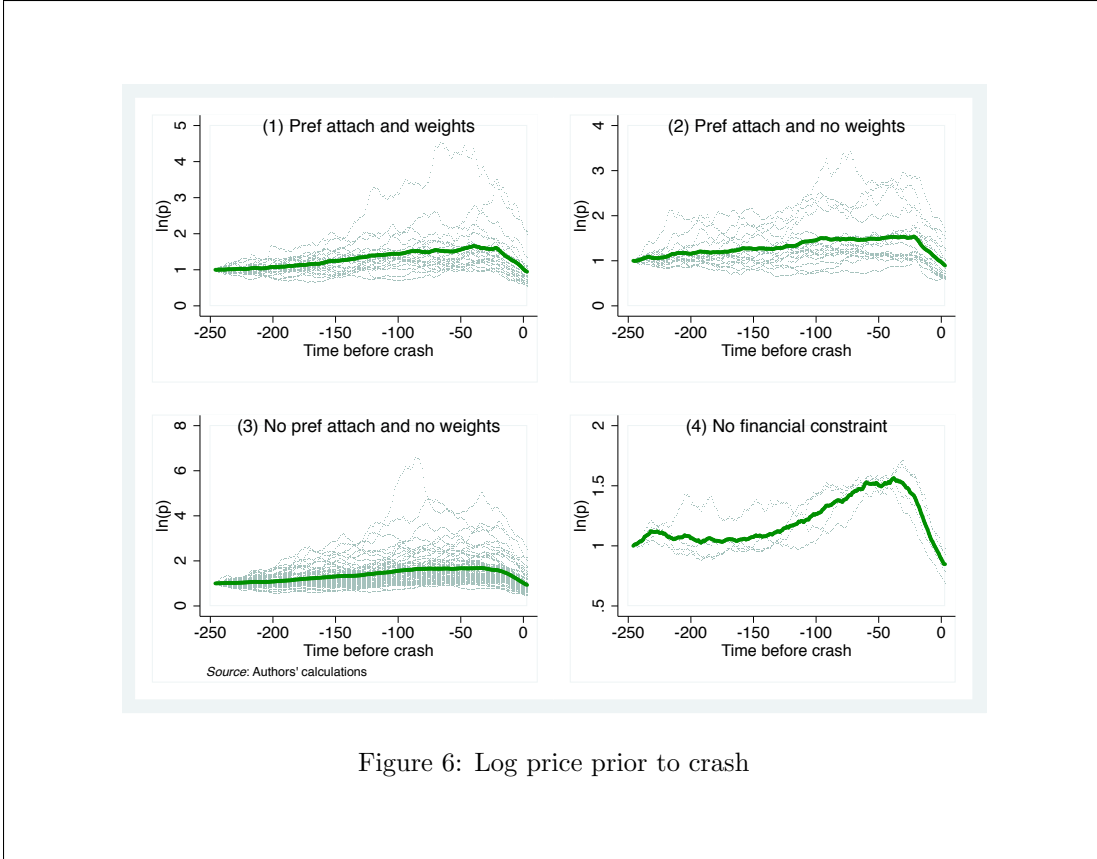
Figure 6: Log price prior to crash

constraint on and only 3.22 without it. Observe that with the financial constraint on, the kurtosis is closer to that of the S&P. With no financial constraint the kurtosis is only slightly above that of the normal distribution, 3.

Two other important properties of the S&P 500 are also captured in the model. The published series displays what is widely known as clustered volatility, the fact that bursts in prices, whether up or down, are clustered together. Interspersed are fluctuations of much lower amplitude, which again seem to be serially correlated. The overall pattern appears in figures 8 and 9.

That these series share a similar pattern of volatility is further supported by the results shown in table 2 and the regressions in columns 1 and 3 of table 3. The average volatility of the S&P over the last 30 years is 0.001414 as reported in table 2. Table 2 also shows that with preferential attachment on and share weights on, volatility of the model's price series is closest to that of the S&P 500 data. The $AR(1)$ table 3 regressions both have significant coefficients on their lagged terms. This contradicts the null that large fluctuations in the log price in the previous period have no effect on the probability of fluctuations this period.

The second and fourth columns test table 3 for the presence of a random walk in both the simulated and S&P 500 time series. Here the augmented Dickey-Fuller test shows that

18

Table 1: Fat tails (kurtosis) in the model and the S&P 500

	Preferential attachment?			
	No		Yes	
	Weighted Network?		Weighted Network?	
	No	Yes	No	Yes
Financial constraint?				
Yes	5.43	4.63	3.76	3.77
No	3.20	3.17	3.10	3.10
Ref: S&P 500	6.92	6.92	6.92	6.92

Source: Authors' computations based on Shiller (2013)

Table 2: Volatility (monthly data with t-tests of difference from S&P 500)

	Preferential Attachment?			
	No		Yes	
	Weighted Network?		Weighted Network?	
	No	Yes	No	Yes
Financial constraint on[1]	4.50E-04	6.06E-04	7.87E-04	8.65E-04
t-test	-5.02	-4.21	-3.27	-2.86
Financial constraint off[2]	3.82E-04	4.83E-04	4.73E-04	5.78E-04
t-test	-5.38	-4.85	-4.90	-4.36
S&P 500[3]	1.41E-03	1.41E-03	1.41E-03	1.41E-03

Source: Authors's computations.
1. Average for financial constraint on: $6.81E-04$ $t = -3.82$, $(n = 3,924,434)$. 2. Average for financial constraint off: $4.79E-04$, $t = -4.87$ $(n = 1,345,275)$. 3. $n = 375$
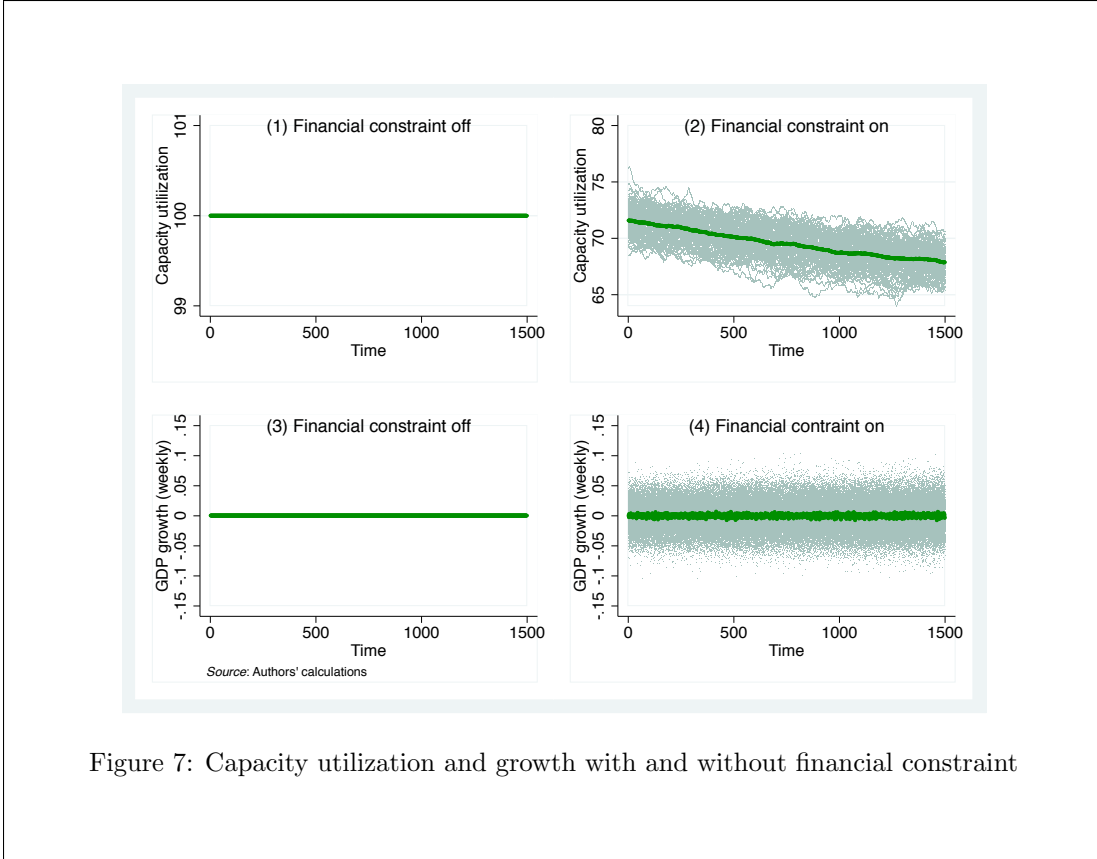
.

Figure 7: Capacity utilization and growth with and without financial constraint

a unit root cannot rejected in either series. Neither are stationary and neither are mean reverting.

3.2. Simulation results

In all runs the first 2,250 weeks were dropped for ergodicity. A total of 10,579 run were retained after dropping those runs which crashed in these initial weeks. Within this set of runs of 1,500 weeks (30 years), there were a total of 172 crashes with the financial constraint on but only 4 with it off.

Table 4 show the crash frequency per 1000 runs, the average rate of growth of GDP and the percentage of "loans denied", which measures the number of deficit firms that are unable to execute their investment plans because they are financially constrained. The table also shows capacity utilization and the average Bayesian prior weight. Table 5 shows the same data when the financial constraint is binding.

Comparing the tables confirms that with no financial constraint there are fewer crashes in the financial sector ($t = -11.2$), faster growth ($t = 35.7$) and higher capacity utilization

Table 3: Clustered volatility and random walk in the model and S&P 500[1]

	(1) Model Volatility[2]	(2) Model $\Delta\, ln(p)$	(3) S&P Volatility[2]	(4) S&P $\Delta\, ln(p)$
Lagged volatility	0.375***		0.166**	
	(0.002)		(0.050)	
First lagged price		0.533**		0.265***
		(4.216e-04)		(0.046)
Second lagged price		-0.489***		-0.317***
		(8.006e-04)		(0.697)
Third lagged price		-0.002***		0.0434
	⋮	⋮		⋮
Tenth lagged price		-0.004***		
		(3.417e-04)		
Time		1.768e-07***		5.029e-05
		(5.799e-09)		(2.825e-05)
Constant	4.56e-05***	0.002**	0.001***	0.035
	(8.784e-09)	(8.503e-06)	(1.104e-04)	(0.018)
$R^2 - adjusted$	0.140	0.325	0.026	0.068
R^2	0.130	0.325	0.027	0.074
Observations	1.57e+07	1.56e+07	657	656
F-stat	43,200	182,376	11	10

Robust standard errors in parentheses. *** $p < 0.01$, ** $p< 0.05$, * $p< 0.1$
1. Lag structure for model is $AR(10)$ and $AR(3)$ for the S&P data
since model data is weekly and S&P data is monthly. The total lag is then
three months for both data series.
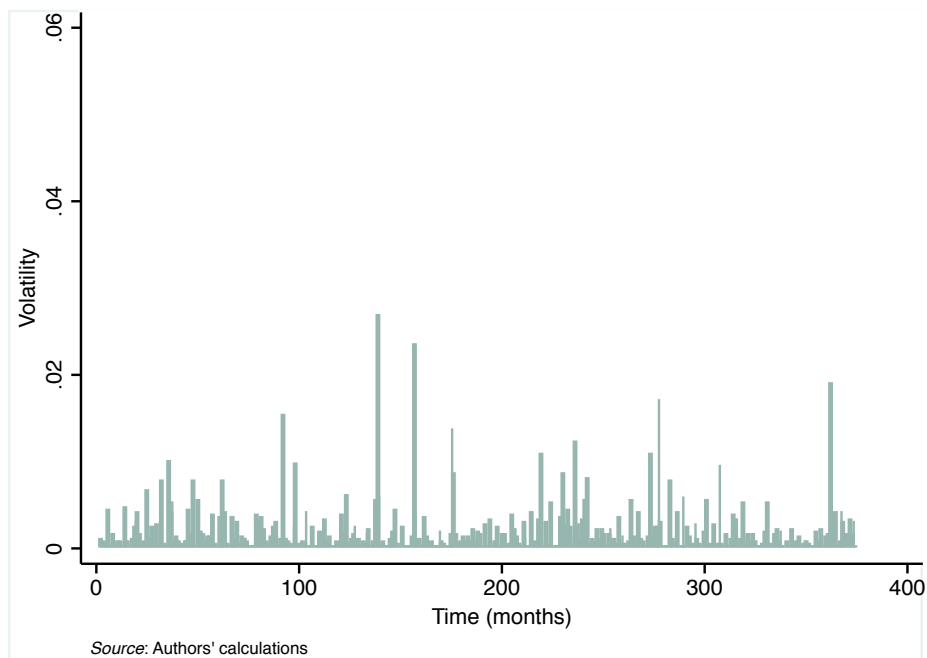2. Demeaned volatility or $\{\Delta ln(p) - E\left[\Delta ln(p)\right]\}^2$

Figure 8: Clustered volatility in a typical run of the model

$(t = 72,000)$ in the real sector, results that agree with the theory developed above.[30] When there is no financial constraint there are no loans denied and so there is more investment, higher capacity utilization and faster growth. There are fewer crashes since the real sector does not transmit "bad news" to the financial sector that can cause crises.

These results demonstrate the importance of real-financial interactions for financial instability. When firms are financially constrained real economic performance deteriorates (growth slows and capacity utilization falls), producing a response in the financial sector where the frequency of crashes increases. When the real and financial sides of the model are dissociated, crises arise from Brownian motion alone. Table 4 confirms that Brownian crises arise, but they are rare as would be expected.

---

[30]As noted, the absolute growth rates are an artifact of the model's parameterization and so there is no particular significance to the size of the growth rate in table 5.

Table 4: Financial constraint *off*

	Preferential Attachment?			
	No		Yes	
	Weighted Network?		Weighted Network?	
	No	Yes	No	Yes
Crash frequency[1]	1.11	3.34	0.00	0.00
GDP growth[2]	4.62E-04	4.62E-04	4.63E-04	4.62E-04
Loans denied[3]	0	0	0	0
Capacity utilization[4]	100	100	100	100
Prior weight[5]	0.31	0.31	0.31	0.31
Total runs	900	899	900	900

Source: Authors's computations.
Notes: 1. Crashes per 1000 runs. 2. Average rate of weekly growth from logarithmic regression. 3. Percent of total firms. 4. Moving average over last 100 runs averaged over all runs. 5. Average over all runs.

Table 5: Financial constraint *on*

	Preferential Attachment?			
	No		Yes	
	Weighted Network?		Weighted Network?	
	No	Yes	No	Yes
Crash frequency[1]	26.73	24.12	6.68	8.17
GDP growth[2]	2.40E-04	2.44E-04	2.54E-04	2.53E-04
Loans denied[3]	19.8	20.0	20.5	20.5
Capacity utilization[4]	70	70	69	69
Prior weight[5]	0.29	0.30	0.30	0.30
Total runs	2581	2612	2693	2693

Source: Authors's computations.
Notes: 1. Crashes per 1000 runs. 2. Average rate of weekly growth from logarithmic regression. 3. Percent of total firms. 4. Moving average over last 100 runs averaged over all runs. 5. Average over all runs.
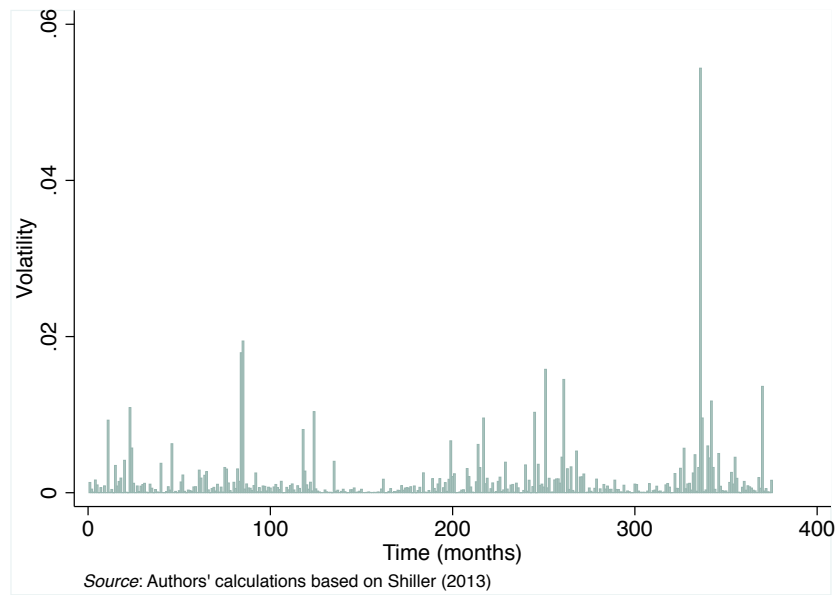
Figure 9: Clustered volatility in the S&P 500

The weight on the Bayesian prior is low, around 30 percent in both tables, which implies that the relative precision of the private signal is, on average, high. This implies that social learning is taking place in all runs, even those that lead to a crash. In the latter financial agents progressively attach less weight to their informative private signals. Social learning begins to break down and herding behavior takes over. In the limit, this gives rise to asset price bubbles fueled by rapid growth that are then followed by crashes Bikhchandani et al. (1998).[31]

---

[31]Observe from the tables that the prior weight is slightly higher with the financial constraint off. In table 4 capacity utilization never deviates from 100 percent and so the private signal never deviates from "bullish" by assumption. The variance of the private signal is zero. Any time the asset price falls, however, the private signal will be incorrect momentarily. The forecast error will then impart some loss of precision of the private signal, increasing the weight on the prior. A noisy signal increases the prior weight. This noise contributes nothing to social learning when the financial constraint is off since capacity utilization is always full and so there is nothing to be learned. With the financial constraint on, however, deviations from full

Table 4 can be seen as the counterpart to stand-alone financial models that heavily populate the literature on the financial system. Table 5, on the other hand, shows the effects of real and financial interactions, arguably the proper environment in which to study the impact of financial network structure. With the financial constraint on, network architecture makes subtle but important differences that demand closer investigation.

Both structural elements of the network, preferential attachment and weighted links, can be seen as treatments relative to a control in which the network is randomly attached and unweighted. The treatments are uncorrelated (-0.0042) so that the t-statistics are not affected by omitted variable bias. The effects of the two structural elements are shown in tables 6 and 7.

The first two columns of table 6 show the impact of financial network architecture on crash frequency. The effect of weighted network links is statistically insignificant, while that of preferential attachment is negative and highly significant. The second two columns of table 6 reveal that neither network weights nor preferential attachment has a significant effect on GDP growth, helping to isolate the effect of network structure on crash frequency. The first two columns of table 7 indicate that both weighted networks and preferential attachment significantly increase loans denied. The second two columns show the same effect on the weight on the Bayesian prior.

Table 7 indicates that there are more loans denied in both weighted and preferentially attached networks. Yet table 6 shows no effect of network structure on GDP growth. This is because a higher level of loans denied does not necessarily reduce the quantity of investment and the pace of growth. The large number of loans denied shows that the distribution of growth is affected by network structure, *viz.*, most growth occurs in large firms while small firms have more limited access to the financial system. Concentration in the financial sector goes hand-in-hand with the concentration in the real as described in Axtell (1999).

Next consider the last two columns of table 7, which show that both weighted and preferentially attached networks increase the weight on the Bayesian prior. Generally, a large weight on the Bayesian prior is a harbinger of crisis, inasmuch as social learning begins to deteriorate as the financial system decouples from the real side. Yet, column two of table 6 indicates that preferential attachment *reduces* crash frequency.

To understand why, consider a large firm associated with a high-degree hub. Flush with finance, this firm makes a significant investment and in so doing, increases capacity utilization throughout the system. The counterpart to this large firm is a smaller producer

---

capacity utilization do indeed convey useful information. In the data set as a whole, the variance of the asset price with the financial constraint on is 1.024 but only 0.663 when it is off. A comparison of means shows a significant effect of the financial constraint in the weighting of private versus prior signals ($t = 520.2$). It is important to see that there is more social learning taking place with the financial constraint on than off. This is verified by the extremely low correlation, (-0.009), between the private and prior signals when there is no financial constraint. When the private signal is variable, the correlation between the private and prior signals rises dramatically to 0.273. This is simply an expression of the interconnectedness of the rates of capacity utilization in the economy. If financial agents are not always predicting full capacity utilization, then their informative signal is more likely to be correlated with the informative signals of their linked neighbors.

Table 6: Regression results

	(1) Crashes[1]	(2) Crashes[1]	(3) GDP Growth[2]	(4) GDP Growth[2]
Weighted network	-3.421e-07		1.653e-06	
	(1.667e-06)		(1.199e-05)	
Preferential attachment		-1.221e-05***		1.114e-05
		(1.689e-06)		(1.198e-05)
Constant	1.110e-05***	1.718e-05***	2.470e-04***	2.422e-04***
	(1.190e-06)	(1.495e-06)	(8.438e-06)	(8.320e-06)
$R^2 - adjusted$	-0.000	0.000	-0.000	-0.000
R^2	0.000	0.000	0.000	0.000
Observations	1.57e+07	1.57e+07	1.57e+07	1.57e+07
F-stat	0.042	52	0.019	0.865

Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$
Notes: 1. Crashes per 1000 runs of 1,500 weeks (30 years). 2. Real GDP growth per week.
Source: Author's computations.

Table 7: Regression results

	(1) Loans Denied[1]	(2) Loans Denied[1]	(3) Bayesian Prior[2]	(4) Bayesian Prior[2]
Weighted network	0.8***		0.0029***	
	(0.009)		(4.989e-05)	
Preferential attachment		4.132***		0.009***
		(0.009)		(4.988e-05)
Constant	136.1***	134.4***	0.292***	0.289***
	(0.007)	(0.007)	(3.550e-05)	(3.591e-05)
$R^2 - adjusted$	0.000	0.013	0.000	0.002
R^2	0.000	0.013	0.000	0.002
Observations	1.57e+07	1.57e+07	1.57e+07	1.57e+07
F-stat	7526	2.03e+05	3375	32233

Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$
Notes: 1. Percent of total number of firms. 2. Average weight on Bayesian prior.
Source: Author's computations.

whose capacity utilization is buffeted by random effects. A negative forecast on the part of the financial agent associated with a small firm might well lead to a denied loan and an inability to invest. The impact on the rest of the economy, however, is minimal and, moreover, the small firm's financial agent will often realize *ex post* that its forecast had been defective. This will lead the small firm's financial agent to put somewhat more weight on the forecast of the larger financial agent, trusting less its own abilities. In the aggregate the Bayesian prior weight will show a slight tendency to rise above its value for a corresponding random network. The rise in the Bayesian prior weight does *not* imply a decoupling of the financial from the real sector, however, and social learning is in fact enhanced. Financial agents are simply learning to pay more attention to those firms who have greater impact on the macroeconomy as a whole. Crash frequency, for these reasons, declines with preferential attachment. Preferential attachment, at least under the assumptions of the present model, does not appear to increase systemic risk.

4. Conclusions

This paper constructs a simple model of real-financial interactions to study the effect of real performance on the frequency of financial crises. When the performance of the real side contributes to the frequency with which financial crises are observed, the model suggests that preferentially attached financial market structures are less fragile and crisis prone. This result confirms the importance of network architecture for the fragility of financial systems. The intuition is that disagreement between high-degree hubs preserves disorder in financial markets. Linked neighbors of high-degree hubs reduce the incidence of grid-spanning clusters of like-minded financial agents that give rise to bubbles and crashes. High-degree hubs are evidently still subject to systemic risk of an ordered sequence of asset prices breaking out of the usual disorder of the market. This is clearly seen in the elevated number of crashes in preferentially attached networks relative to when the financial constraint is off.

The broader theoretical conclusion of the paper is that the recent collapse of the financial system does not seem to be due to its preferentially attached network structure *per se*. The model is highly simplified, however, and further research may uncover a destructive role for preferentially attached financial agents in a more complete description of the actual financial system. The conclusion suggested by the model of this paper, however, is that the financial sector did not become too connected or too *linked* to fail, so reducing the interconnectedness of financial agents should not be the objective of policy designed to bring greater stability to financial markets. To put this into context, even if it were sensible to allow Lehman Brothers to fail in September 2008 as punishment for misconduct, the results here suggest that this may have had unintended negative consequences for financial network architecture by removing an otherwise stabilizing high-degree hub.

The central message of the paper is the critical importance of real foundations of financial crisis. Just as Kregel (1985) argued that Keynesian real-side models with no monetary and financial sectors were akin to "Hamlet without the prince", so, too, it seems that stand-alone financial models that neglect the real side are incomplete. Ultimately, neither stand-alone real sector nor stand-alone financial sector models are suitable instruments.

5. Appendix: Pseudo code

The program can be expressed as:

1. Initialize data structures and runtime options
2. Set key parameters
   (a) Set share weights–boolean
   (b) Set preferential attachment–boolean
   (c) Set financial constraint–boolean
   (d) Set run years–30 × 50 weeks
3. Set up and initialize network
4. Reassign financial agents such that each firm has at least one financial agent
5. Set shareholders as count financial agents for each firm
6. Initialize surplus of each firm based on randomly assigned parameters
7. Run main
8. If financial constraint = FALSE: set invest = TRUE for all firms
9. If financial constraint = TRUE:
   (a) Ask financial agents: make forecast based on last period's private and public signals
   (b) Ask firms: if surplus > 0 set invest = TRUE
   (c) Ask firms: if surplus < 0 ask one of financial agents if loanable funds $|surplus|$
        i. If yes: set invest = TRUE
        ii. If no: ask linked neighbor: if loanable funds > $|surplus|$
            A. if yes: set invest = TRUE
            B. if no: set invest = FALSE
            C. update denied-loan counter
10. Run Gauss Seidel [sum of investment of firms with invest = TRUE]
    (a) Set demand shares of firms
    (b) Set capacity utilization of firms
    (c) Set savings of firms
    (d) Set planned investment
    (e) Set surpluses of firms
    (f) Set loanable funds = surpluses of surplus firms
11. If share-weight = TRUE
    (a) Re-weight links by accumulated capital stocks
12. Stop for crash
13. Stop for year limit
14. Stop for capacity utilization lower limit (0.6)
15. Return to main
16. Process output

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica 80*(5), 1977–2016. MIT: http://econ-www.mit.edu/files/7124.

Axtell, R. (1999). The emergence of firms in a population of agents. http://www.brookings.edu/es/dynamics/papers/firms/firmspage.htm.

Bagrow, J. P., J. Sun, and D. ben Avraham (2008). Phase transition in the rich-get-richer mechanism due to finite-size effects. *Journal of Physica A: Mathematical and Theoretical 41*.

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

Bikhchandani, S., D. Hirshleifer, and I. Welch (1998). Learning from the behavior of others: Conformity, fads, and informational cascades. *The Journal of Economic Perspectives 12*(3), 151–170.

Chamley, C. P. (2004). *Rational Herds: Economic Models of Social Learning.* Cambridge University Press.

Erdös, P. and A. Rényi (1950). On the evolution of random graphs. *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, 17–61.

Hansen, L. P. (2012). Challenges in identifying and measuring systemic risk. Working Paper 18505, National Bureau of Economic Research.

Harras, G. and D. Sornette (2011). How to grow a bubble: A model of myopic adapting agents. *Journal of Economic Behavior & Organization 80*(1), 137–152.

Iori, G., G. d. Masi, O. Precup, G. Gabbi, and G. Caldarelli (2008, January). A network analysis of the italian overnight money market. *Journal of Economic Dynamics and Control 32*(1), 259–278.

Johansen, A., O. Ledoit, and D. Sornette (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance 3*(1).

Kalecki, M. (1937). The principle of increasing risk. *Economica 4*(4), 440–446.

Kregel, J. A. (1985). Hamlet without the prince: Cambridge macroeconomics without money. *American Economic Review 75*, 133–139.

LeBaron, B. (2012). Heterogeneous gain learning and the dynamics of asset prices. *Journal of Economic Behavior & Organization 83*(3), 424–445.

Maddison, A. (2007). *The World Economy: A Millennial Perspective.* Development Centre Studies. Paris: OECD Publishing.

Neftci, S. N. (2008). *Principles of Financial Engineering.* Amsterdam: Elsevier.

Reinhart, C. M. and K. Rogoff (2009). *This Time Is Different: Eight Centuries of Financial Folly.* Princeton University Press.

Russell, S. J. and P. Norvig (2010). *Artificial Intelligence - A Modern Approach.* New York: Pearson Education.

Setterfield, M. (Ed.) (2010). *Handbook of Alternative Theories of Economic Growth.* Edward Elgar.

Setterfield, M. and A. Budd (2011). *Microeconomics, Macroeconomics and Economic Policy: Essays in Honour of Malcolm Sawyer*, Chapter A Keynes-Kalecki model of cyclical growth with agent-based features, pp. 228–250. London and New York: Palgrave Macmillan.

Shiller, R. (2013). Online data.

Sornette, D. (2003). *Why Stock Markets Crash.* Princeton, NJ: Princeton University Press.

Taylor, L. (1983). *Structuralist Macroeconomics.* Boston, MA: Basic Books.

Taylor, L. (1991). *Income Distribution, Inflation, and Growth.* Cambridge, MA: MIT Press.

Tedeschi, G., G. Iori, and M. Gallegati (2012). Herding effects in order driven markets: The rise and fall of gurus. *Journal of Economic Behavior & Organization 81*(1), 82–96.

Thurner, S., J. D. Farmer, and J. Geanakoplos (2012, February). Leverage causes fat tails and clustered volatility. *Quantitative Finance 12*(5), 695–707.