

Real and financial crises in the Keynes–Kalecki structuralist model: An agent-based approach

Bill Gibson¹ | Mark Setterfield²

¹ University of Vermont

² New School for Social Research

Correspondence

Mark Setterfield, Economics Department,
New School for Social Research,
New York, NY 10003.
Email: mark.setterfield@newschool.edu

Abstract

Agent-based models are inherently microstructures—with their attention to agent behavior in a field context—and only aggregate up to systems with recognizable macroeconomic characteristics. One might ask why the traditional Keynes–Kalecki or structuralist (KKS) model would bear any relationship to the multi-agent modeling approach. This paper shows how KKS models might benefit from agent-based microfoundations, without sacrificing traditional macroeconomic themes, such as aggregate demand, animal spirits and endogenous money. Above all, the integration of the two approaches gives rise to the possibility that a KKS system—stable over many consecutive time periods—might lurch into an uncontrollable downturn, from which a recovery would require outside intervention. As a by-product of the integration of these two popular approaches, there emerges a cogent analysis of the network structure necessary to bind real and financial agents into an integrated whole. It is seen, contrary to much of the existing literature, that a highly connected financial system does not necessarily lead to more crashes of the integrated system.

1 | INTRODUCTION

At first glance, Keynes–Kalecki or structuralist (KKS) macroeconomics and agent-based models (ABMs) would seem to be strange bed-fellows. Structuralists since Taylor (1983) have integrated the financial and real sides of the models but mostly through the interest rate (Taylor, 1981) and the profit rate (Taylor & O’Connell, 1985). Since then the literature on ‘financialization’ has emphasized the bloat, waste and inequality ushered in by growth of the financial services sector (Epstein, 2005; Palley, 2014). This literature has been less concerned with the central role played by finance in channeling savings into investment. Early Keynesians were certainly concerned and it was immediately recognized by Robertson (1940) and later by Chick (1983) that at the core of the Keynesian economy was a

well-functioning financial system capable of creating money.¹ The system transfers purchasing power from those who wish to invest less than they save to their counterparts, those who want to invest *more*. Any interruption in this flow of funds spreads through the rest of the economy through the multiplier-accelerator process, creating havoc.

Constructing convincing models of real-financial interaction, however, has been made difficult by the inherent complexity and interlinked nature of the financial system. Recent advances in agent-based modeling with its focus on concentration and contagion across structured networks offer an opportunity for relatively deep integration with KKS models. This paper shows how such integration might be accomplished and how the central features of both approaches can be preserved in the merging of the two schools of thought.

Benefits accrue to both sides: structuralists have traditionally waved off the criticism that the microfoundations of their models are weak and not explicitly tied to standard optimization models (Skott, 2012). Microfoundations in multi-agent models, however, are not always based on strict rationality. Agents are, by definition, boundedly rational; they are heterogenous, sometimes myopic and generally incapable of making rigorous inter-temporal trade-offs. Random elements play an important role, as does asymmetric information, incomplete contracts and other institutional constraints. These are all features that structuralists commonly include in their *macro* models, but by assumption, rather than their having been derived from some microfoundation. On the other hand, the benefit to standard agent-based models is the integration of the real side into an explicitly financial model. This has been ably attended to in the past, of course, but the Keynesian quantity adjustment process, with its emphasis on aggregate demand, is less visible in the existing literature (Ashraf, Gershman, & Howitt, 2011; Gaffeo, Delli Gatti, Desiderio, & Gallegati, 2008). Moreover, tracing the flows from *surplus* firms, those that invests less than they save, to *deficit* firms, those that do the opposite, is not as easily seen in the standard literature as in the model to follow. This perspective shows explicitly the conditions under which money must be endogenously created and moreover, allows for animal spirits to be an explicit driver of the level of economic activity. Prominent and popular financial models, such as Gai, Haldane, and Kapadia (2011), focus on contagion in the interbank market. The framework presented here is more concerned with the Keynesian features of modern capitalism, most specifically that an independent investment function drives the level of savings instead of the reverse as well as the irrepressibility of endogenous money.

Setterfield and Budd (2011) make a first step toward integrating the KKS perspective with agent-based models. In that model, each agent is essentially its own KKS economy, with its own aggregate demand and savings-investment balance. Agent interaction is limited to a 'blackboard' communication system, in which an agent's incentive to invest is affected by the performance of its colleagues. Aggregate demand is not shared among firm agents, however, nor is there any explicit modeling of the financial sector.

Beginning with a multi-agent KKS model, this paper integrates a financial sector, inspired by the agent-based literature (Harras & Sornette, 2011; Johansen, Ledoit, & Sornette, 2000; LeBaron, 2012; Sornette, 2003; Tedeschi, Iori, & Gallegati, 2012; Thurner, Farmer, & Geanakoplos, 2012, and many others). The result is a model of the economy with two agent sets, real sector firms and financial sector intermediaries. It is seen that 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 microfoundations are based on Bayesian learning (Chamley, 2004) in which a prior public signal is updated by way of a private signal based on the performance of the firms to which the financial agent is linked.

¹See also Gibson and Setterfield (2015).

The paper is organized as follows. The next section outlines the design concepts and coordination environment of the two agent sets. Section 3 discusses the main results of the model, endogenous money and the possibility of a real-financial crisis. It is seen that the integration of the two approaches allows an important role for network structure, whether financial agents are linked randomly or by way of preferential attachment. The fourth section concludes. An appendix contains the pseudo code of the model.²

2 | DESIGN CONCEPTS

One way to initiate the transition from the pure KKS to a framework that integrates the multi-agent approach is to individually model workers and capitalists and the full range of economic decisions they make. This approach could be extended to the government, the financial sector and the central bank. A fully granular model would show all the inner-workings of the economy in exquisite detail and one would hope and expect that the agent model would fully align with all the main features of the structuralist macromodel.

In a closed economy with no public sector, however, the structuralist model often invokes the assumption that workers do not save and capitalists save out of profits. This confers a significant degree of simplification since Walras' law, which implies that the sum of agent savings is equal to investment, obviates the need to model workers' behavior directly. In general, there would be no reason to make the assumption that workers do not save in an agent-based model. Here, however, the project is to link the structuralist macro and agent-based methods and so the same assumption is maintained. Workers are still present in the model, but only implicitly.

This simplification also enables the model to account for the interaction of just two kinds of agents, firms and financial agents, without the complexity of a full agent-based system that might obscure the comparison with the conventional one-sector KKS system. Moreover, structuralist models can already adequately handle issues arising from the interaction of workers and capitalists. Attention here is paid to another dimension along which the structuralist model is less competent and well worked out, that of the relationship between firms and their financial agents.

In an economy with a single firm, the savings generated by that firm must be equal to total investment. When there are heterogeneous firms, some generate *financial surpluses* when they save more than they invest and *financial deficits* when the opposite is true. The primary role of the financial system is intermediation, facilitating the flow of funds from surplus to deficit firms.

As noted, in the model of this paper there are two agent sets: firms, which undertake all savings and investment, and financial agents that provide intermediation services. Financial agents neither save nor invest. Simple behavioral rules are defined for these two agent classes and the macro performance of the model arises as an emergent property of their interaction. There is no policy authority or explicit monetary policy and all adaptation is through Bayesian 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 square, operates only one production process, combining capital and labor to satisfy demand for a homogeneous good.³ The key

²The replication code for the model (written in NetLogo) is available at <http://www.uvm.edu/~wgibson/Research/GS.nlogo>. The model can be run in 'headless' configuration on a cluster or standalone using only the NetLogo software.

³This 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.

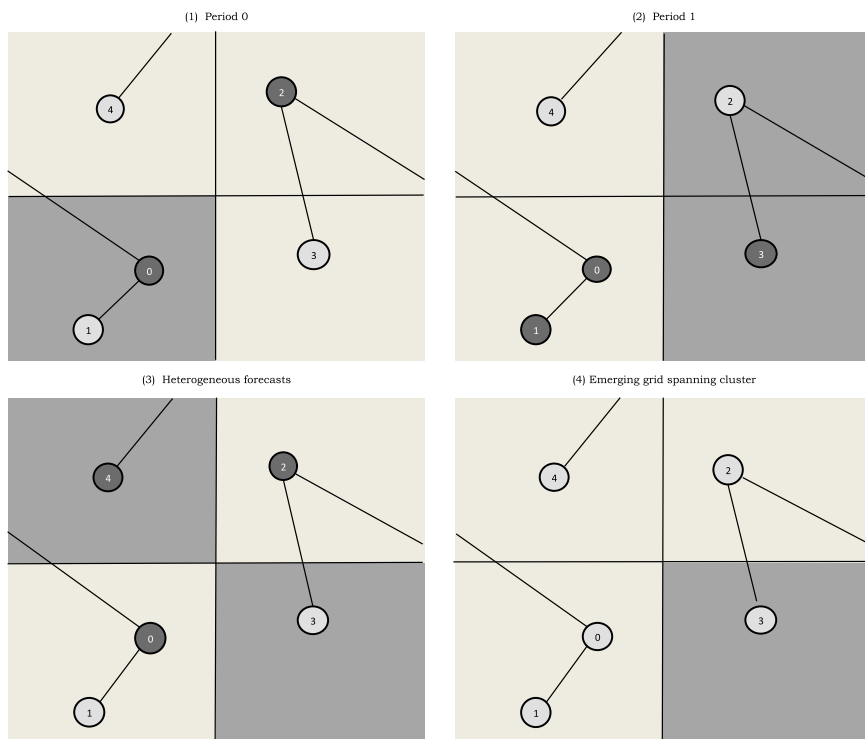


FIGURE 1 Structural features of the multi-agent model

decision for firms is whether and how much to invest. Investment is based on expectations about future market conditions and firms' ability to cover any short-fall in savings, relative to planned investment, through borrowing from other firms with the help of one or more financial agents. At each sweep of the model, surplus firms, shown as dark squares, make deposits with their resident financial agents, represented by numbered circles. Financial agents must then decide whether to grant loan requests to deficit firms, shown as light squares. The decision of financial agents is also binary: lend/do not lend. This decision depends on their forecasts of the probability that the loan will be paid back, dark for 'bullish' and light for 'bearish'.

Each financial agent can at most offer *direct* financing for one production process, the one with which it is permanently associated. That same agent can, however, be called upon to provide *indirect* financing to another firm by way of another financial agent. Thus a bullish financial agent that lacks sufficient liquidity to meet the demand for loans from its associated firm, may arrange for a loan from one of its neighbors.

To capture the locality of intermediation, information constraints and boundedness on the rationality of the financial agents, however, indirect borrowing must be limited to some degree. This is achieved by imposing a restriction on the connectedness of the system through the number of plys to which each financial agent has indirect access. A simple but useful assumption is to limit the number of plys to one, so that financial firms can only borrow from their *linked* neighbors. This 'one-ply' assumption is the point at which network analysis—agents dispersed on grid—adds to the standard conceptual framework of the KKS model. Without the one-ply assumption, firms could chain across the entire grid, borrowing directly or indirectly from all available sources of savings. No financial system dislocations could arise to interrupt the growth of the system and the model would effectively

revert to the standard one-sector KKS model, in which there is always finance available for any desired investment.⁴

The linked neighbor may have refused a loan to its own (deficit) firm on the grounds of a bearish forecast. It is nonetheless willing to accept the indirect request of its neighbor. This feature of the model is designed to capture the ‘originate and distribute’ characteristic of the financial system. Indirect loans are viewed as less risky since deficit firms are vetted by their own financial agent and indirect financing is, therefore, never refused. Figure 1 shows, however, that firms can have direct access to more than one financial agent. It may seem from Figure 1 that there are more financial agents than firms, inverting reality. Note, however, that financial *agents* are not conterminous with financial *firms* (which are associations of financial agents), so that the number of financial firms is not defined within the model. In any case, the latter would already be included in the agent set ‘firms’, which earn profits and pay wages. Consequently, there is no explicit interest rate since prices are not explicit due to the assumption of a single homogeneous good.

If the number of financial agents is *less* than the number of firms, large deficits of aggregate demand arise, since many deficit firms are unable to find financing for their projects. A far less restrictive assumption is that the number of financial agents, m , is greater than the number of firms, n . This assumption produces reasonably robust growth and prevents subsets of firms from experiencing very low levels of effective demand and capacity utilization when the rest of the grid is booming. The two critical assumptions of the model are, then, the one-ply assumption and $m > n$, and it should now be clear that the model would behave in a highly unrealistic way if either of these assumptions were relaxed.

Network architecture is represented by the edges joining the numbered financial agents in Figure 1, and influences how financial surpluses are allocated.⁵ Suppose that the surplus firm, shown in dark gray in the first panel of Figure 1, makes a deposit with financial agent 0 which 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 it would lend is already in surplus. Financial agent 2, however, serves a deficit client and is connected to financial agent 0, associated with a surplus firm. Agent 2 is also connected to 3, who is associated with a deficit firm and has no liquidity. Financial agent 2 then calls on financial agent 0 to make the surplus available for its client. Savings is thereby channeled from a surplus to a deficit firm.⁶

Agent 3 is bearish, but even had 3’s forecast been bullish, it might have been stymied in its effort to provide indirect financing for its deficit firm. Whether 3 has *indirect* access to financial agent 0’s funds is a setting of the model and determines the degree of connectedness of the financial system. KKS models that have no constraint on interbank borrowing essentially allow each financial agent complete access to *all* other financial agents on the grid. This is the assumption that implicitly underpins KKS models *without* a financial sector, models in which money is endogenous.

Even under the one-ply restriction, money is still endogenous. To see this, first note that behavioral decision rules in multi-agent systems are executed asynchronously. This asynchronicity allows the possibility of *collisions*, that is, conflicting claims on financial resources. This works as follows: at the

⁴The reader could reasonably wonder whether a two- or three-ply assumption might be more realistic, coming closer to the way in which financial firms are linked in the real world. On reflection, however, it is apparent that the crucial distinction is between limited ply borrowing and full grid access. Multi-ply systems can always be converted to one-ply systems, by ‘eliminating the middleman’, which has been done here for simplicity.

⁵Network architecture also influences financial agents’ expectations, which are based (in part) on the forecasts of their linked network neighbors. See Section 2.5.

⁶Note that financial agent 4 is not linked to agents 0–3, but is connected to financial agents on other parts of the grid not shown in Figure 1.

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 its funds, a collision is created. The conflict can only be resolved by allowing financial agents to create money.⁷ 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. This change occurs because aggregate demand alters the ability of firms to save and their incentive to invest. 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, or contagion, sets the stage for a financial, and possibly, a real crash. The probability of the formation of a grid-spanning cluster measures the systemic risk of the system. The following sections describe more precisely how firms and financial intermediaries decide whether to invest and lend, respectively.

The key to understanding the integration of the structuralist and agent-based perspective lies in the relationship between savings and investment. While it might seem natural to model agents as making a decision between consumption and investment and then either accumulating physical or financial assets with their savings, this would lead to a savings driven economy at the macro level. The alternative, pursued here, is to explicitly model the investment process, with an exogenous term for animal spirits and then allow savings to adjust to investment in equilibrium. The agent making this central decision of the model is the firm and one of the firm's drivers is animal spirits, a term not linked to savings.⁸

2.1 | Model synergy: What the ABM perspective adds to KKS

With a network-based financial structure added to the standard KKS macro adjustment mechanism, the model can be used to study contagion among financial agents that can severely interrupt the flows of savings into investment. In this respect the unification of the ABM and KKS approaches confers two benefits. First, it shows how financial contagion can be reinforced by the interaction of the financial sector with the real economy. Contagion in the model is based on Bayesian learning in which a prior *public signal* is updated by way of a *private signal* based on the performance of the real-sector firms to which financial agents are linked (Chamley, 2004). Second, the model shows how contagion resulting from real-financial interactions depends on network structure, the precise architecture of the links that are formed between financial agents.

The implications of combining the KKS and ABM approaches run deep. In a standard structuralist model with no financial sector, a reduction of one unit of aggregate demand reduces income and employment by way of the multiplier analysis. When attention is given to the links of this chain, the traditional focus is on unintended inventory accumulation. Since aggregate demand is the only driver, a unit decrease in aggregate demand always affects the sum of output and employment in the same

⁷The model thus reflects the hidden presence of a monetary authority that does not allow credit creation, except when these collisions occur. See Section 3.1 below for a fuller discussion of this assumption.

⁸In many multi-agent models, such as Gaffeo et al. (2008), exogenous spending is linked to prior savings of the unemployed or consumers frustrated in their search for goods and services. Savings, in a sense, still drives investment even though it is not linked to the current level of output.

way. The method by which the investment is financed does not always come into the picture. When it does, with, say, an upward sloping *LM* schedule, the effect is *always* to soften the blow of a decline in demand, in response to which the interest rate falls.

In the conjoined ABM–KKS structure, the scenario could unfold very differently. Depending on the economic structure that exists when the aggregate demand shock materializes, the impact could be attenuated or amplified by the financial system. Figure 1 above shows how either could happen. If the demand shock hits *surplus* firm *j*, the deposited surplus falls; less finance is available for deficit firm *k* and now the investment plans for the latter are blocked by the financial structure. Aggregate demand is now lower in some parts of the grid and a negative public signal begins to emerge that contradicts the private signal that other financial agents receive from their own firms. A refusal to propagate loans on the part of these financial agents may spread to the rest of the grid as the outlook of financial agents turns bearish. Whether this happens depends on the structure of the economy, as embodied in network linkages among financial agents and the assumptions about how agents learn. There are now two channels through which the negative demand shock might propagate, the real side (as in the standard KKS model) and the financial sector (as modeled in the stand-alone financial sector agent-based finance literature).

Conversely, the financial system may counteract the negative demand shock, attenuating its impact on the real economy. A surplus firm that fails to invest as a result of the demand shock now deposits what would have been spent on investment goods, raising the available finance for deficit firms that might otherwise have been frustrated. A stronger public signal could emerge that frees finance for other deficit firms who then further amplify the signal, with the result that the initial contraction in aggregate demand is offset or even reversed.

2.2 | Macrofoundations: What KKS adds to ABMs

The traditional KKS perspective invokes microfoundations that are not linked to strict optimizing behavior. They are rules of thumb, behavioral rules, that apply to a class of individuals rather than an individual. These rules of thumb can change as circumstances warrant. The ABM approach is wholly consistent with this view in that few employ formal dynamic optimization models to guide individual behavior. More commonly applied are principles of behavioral economics, or appeals to limited or costly computational capacity or bandwidth on the part of agents. Rules of thumb can be more numerous, varied and complex in ABMs, but conceptually the approach adds nothing new to the KKS system.

What the KKS system imposes upon the microeconomy, populated by autonomous agents, is perhaps far more valuable. Structure in an economic system is, at the most general level, ossified agency, the product of past decisions made by preexisting agents (some of whom survive to the present). These evolving structures may be viewed as *macrofoundations*, constraints imposed on agents by the structure itself, that impinge on the further evolution of the system. A simple example of a macro constraint is the notion of a ‘public signal’ itself. The overall macroeconomy provides a data stream of information that not only enters into individual decision making but also, by virtue of its very boundedness, constrains the behavior of individual agents, both firms and financial agents. Workers, meanwhile, are constrained by their class identity in the KKS, and cannot choose to become firms or financial agents. Indeed, they are so ‘macroconstrained’ that they have few choices to make and can therefore be left out of the modeling effort.

2.3 | Firms as decision-making agents

The firm agent’s data structure is described in Tables 1 and 2. Firms are instantiated and initialized as shown in Table 1, where it is seen, for example, that capital-output ratios vary according to a uniform

TABLE 1 Firm parameters

Name	Symbol	Distribution	Range
<i>Observer parameters</i>			
Number of firms	n	—	676
<i>Fixed parameters</i>			
Initial capacity utilization	\bar{u}	—	1
Depreciation	δ	—	0.05
<i>Distributed parameters</i>			
Capital-output	v_i	U	2–4
Labor coefficient	l_i	U	0.1–0.6
Firm savings ^a	s_i	N	0.265–0.356
Initial growth rate of capital stock	\bar{g}_i	U	0.05–0.09
Initial capital stock ^b	\bar{K}_i	E	1–10
<i>Computed parameters</i>			
Initial capacity	\bar{Q}_i	\bar{K}_i/v_i	0.25–5
Initial output	\bar{X}_i	\bar{Q}_i	0.25–5
Initial profit rate	\bar{r}_i	$(1-l_i)/v_i$	0.1–0.45
Animal spirits coef.	α_{0i}	$\bar{g}_i/2$	0.025–0.045
Accelerator coef.	α_{1i}	$\bar{g}_i/4$	0.0125–0.0225
Profitability coef.	α_{2i}	$(g_i - \alpha_{0i} - \alpha_{1i}\bar{u})/\bar{r}_i$	0.006–0.633

U = Uniform distribution; E = exponential; N = normal. *Source:* Authors.

^aRange is mean \pm 2 *SD*.

^bTruncated.

probability distribution over the range 2–4. Saving rates, however, are initially normally distributed, within the 2 *SD* range as shown.⁹

As noted, firms produce a single homogeneous good in the model by way of a Leontief production function. The technology, which depends on a capital-output ratio and labor coefficient, is fixed once firms are instantiated. Aggregate demand depends only on spending on consumption and investment. The usual assumption in KKS models is that all 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. Let X_{it} be the output of firm i in period t . Aggregate demand in period t is then

$$Y_t = \sum_{i=1}^n [(1 - s_i \pi_i) X_{it} + I_{it-1}]. \quad (1)$$

The output of each firm is determined by a share of aggregate demand, λ_{it} , such that

$$X_{it} = \lambda_{it} Y_t \quad (2)$$

where $\sum_{i=1}^n \lambda_{it} = 1$. Investment is given by an independent function for each firm

⁹They are then scaled to give total savings equal to total investment so that the initial solution for the model as a whole is in macroeconomic equilibrium.

TABLE 2 Firm updating methods

Name	Symbol	Equation	Equation No.
<i>Observer variables</i>			
Aggregate demand	\mathbf{Y}_t	$\sum_{i=1}^n [(1-s_i\pi_i)X_{it} + I_{it-1}]$	1
<i>Endogenous variables</i>			
Share of \mathbf{Y}_t	λ_{it}	X_{it}/\mathbf{Y}_t	2
Output	X_i	$\begin{cases} \lambda_{it}\mathbf{Y}_t & X_{jt} \leq Q_{jt-1} \\ Q_{jt-1} & X_{jt} > Q_{jt-1} \end{cases}$	2, 8
Investment	I_{it}	$g_{it}K_{it}$	3
Growth rate of capital stock	g_{it}	$\alpha_{0i} + \alpha_{1i}u_{it} + \alpha_{2i}r_{it}$	4
Profit rate	r_{it}	$\pi_i X_{it}/K_{it-1}$	5
Capacity utilization	u_{it}	X_{it}/Q_{it-1}	6
Capacity	Q_{it}	K_{it}/v_i	7
Financial surplus	F_{it}	$\sum_{i=1}^n s_i\pi_i X_{it} - I_{it-1}$	9
<i>State variables</i>			
Capital stock	K_{it}	$K_{it-1}(1-\delta + g_{it-1})$	

Source: Authors.

$$I_{it} = g_{it}K_{it} \quad (3)$$

where K_{it} denotes the capital stock and g_{it} is an accumulation function that depends on animal spirits, capacity utilization and the rate of profit, or

$$g_{it} = \alpha_{0i} + \alpha_{1i}u_{it} + \alpha_{2i}r_{it} \quad (4)$$

where the α 's are constants and the rate of profit is

$$r_{it} = \pi_i X_{it}/K_{it-1} \quad (5)$$

The values assigned to the parameters of this equation are shown in the lower part of Table 1. The animal spirits term, α_{0i} , is fixed for each firm, but given the randomness of other settings in the model animal spirits is uncorrelated with any of the endogenous variables of the model.

It might seem more natural to model the firm in its traditional role of profit maximizer, combining the factors of production in variable proportions to achieve least cost per unit of output. In fact, technical efficiency in the multi-agent perspective is second-order and it is rather the *interaction* of many firms and financial agents that attracts the attention of the agent-based method.

The methods by which firms update their data structures are shown in Table 2. The basic rationality behind the firm's decision-making is to increase (rather than maximize) profits by investing, creating productive capacity through capital accumulation. This avoids loss of market share because of capacity constraints. When firms are capacity constrained, aggregate demand they cannot satisfy is allocated to other firms in an iterative processes within each sweep of the model, described in more detail below.

Capacity utilization, u_{it} , is the firm's variable that signals the need for more or less investment and is defined as

$$u_{it} = X_{it}/Q_{it-1} \quad (6)$$

where capacity output, Q_{it} , is

$$Q_{it} = K_{it} / v_i \quad (7)$$

where K_{it} is a state variable determined by the usual stock-flow relationship, as shown in Table 2. Here Q_{it} is the maximum quantity of production given the capital stock, K_{it} . Thus

$$X_{it} \leq Q_{it-1} \quad (8)$$

or equivalently, $u_{it} \leq 1$. Firms adapt to changes in the utilization of capacity by adjusting their rate of investment.¹⁰

Note that there are, implicitly, multiple layers of decision-making for the firm agent. Entailed in the behavior of the firm are incomes, both of workers and owners. Consumer behavior depends on these incomes and so indirectly depends on firm agents. Aggregate demand depends on the sum of these indirect effects together with the exogenous component of animal spirits in the investment function.

Typically, an agent-based model would take advantage of this multi-layered structure to endow the consumers with their own choice points. Consumers may respond to spatial availability and price as they do in Gaffeo et al. (2008). They may decide to quit or change jobs, to accumulate human capital, retire, emigrate or join the informal sector (Gibson, 2012). In the hybrid model of this paper, these decisions are all directly linked, as they are in the KKS framework, to the decision to invest.

Here the system-wide distribution of aggregate demand endows each firm with a quantity of savings that may or may not be sufficient to finance its planned investment. Inter-agent communication for firms depends on the financial surplus, F_{it} , defined as

$$F_{it} = s_i \pi_i X_{it} - I_{it-1} \begin{cases} < 0 & \text{deficit} \\ = 0 & \text{balanced} \\ > 0 & \text{surplus} \end{cases} \quad (9)$$

for $i=1, 2, \dots, n$. The designation of surplus/deficit is temporary and depends on t , since in order to execute planned investment, a deficit firm must first locate a surplus firm from which to borrow. If this proves impossible, the deficit firm flips to surplus, as described in more detail below, and aggregate demand falls. In a model with perfect information, zero transactions costs, and endogenous money, nothing inhibits the flow of funds from surplus to deficit firms. Inter-agent communication is fast, effective, costless and complete. In reallocation economies, however, inter-agent communications is less perfect and may breakdown, potentially reducing aggregate demand bringing about real-financial crisis.

2.4 | The coordination environment of firm agents

Productive capacity is a state variable in the model, evolving along with the capital stock. In the current model, capacity, Q_i , may limit the amount of savings a firm generates, causing a surplus firm to flip and become a deficit firm.

So far, there is nothing in the model preventing the distribution of aggregate demand from causing output, X_i , to exceed capacity, and therefore u_i to rise above 1. To prevent this unrealistic outcome, a coordination mechanism is built into the model such that aggregate demand is shared and no firm exceeds its own capacity to produce.

¹⁰Since the rate of profit can also be written as $r_{it} = \pi_i u_{it} / v_{it}$, Equation 4 is a function of the single signaling variable, u_{it} .




FIGURE 2 The demand redistribution coordination method

To see how this coordination mechanism functions in the model consider a simple two-agent system in which the share of aggregate demand for firm one is λ and for firm two is $1-\lambda$. There are not yet any financial agents in the model. The equilibrium conditions set consumer and investment demand equal to the supply of each good

$$X_1 = \lambda[(1-s_1\pi_1)X_1 + (1-s_2\pi_2)X_2 + I]$$

$$X_2 = (1-\lambda)[(1-s_1\pi_1)X_1 + (1-s_2\pi_2)X_2 + I]$$

where $I = I_1 + I_2$ is taken as a given parameter determined in the previous period. The system can then be written explicitly as a function of aggregate demand as

$$u_1 = \lambda Y(u_1, u_2)$$

$$u_2 = (1-\lambda)Y(u_1, u_2)$$

where $Y = \sum_{i=1}^2 [(1-s_i\pi_i)X_i + I_i]$ and $Q_1 = Q_2 = 1$ to simplify. Figure 2 shows a graphical solution for this set of equations in the two unknowns, u_1 and u_2 . The solid lines in the diagram show an initial solution in which the capacity constraint is exceeded by the second firm. The demand sharing coordinating mechanism operates as described in the figure and accompanying pseudo code in Figure 3. The share of aggregate demand allocated to firm 1, λ , rises until the level of capacity utilization in firm 2 falls to 1. At the same time, the share of aggregate demand allocated to firm 1 rises, as shown in Figure 2. The intersection of the dotted lines indicates a feasible solution with both $u_1 \leq 1$, and $u_2 \leq 1$.¹¹

In the process of demand reallocation, Y may increase, decrease or stay constant. The intuition is derived from standard macroeconomic reasoning: the sum of savings is equal to investment before and after the reallocation of demand. Investment has not changed since the adjustments are taking place within one sweep of the model. Therefore, if the savings per unit of output by firm 2, $s_2\pi_2$, is greater

¹¹These adjustments are taking place within a single sweep of the model and therefore the t subscript is suppressed.

Since $u_2 > 1$

1. **set** $u_2 = 1$
2. **set** $Y(u_1, u_2) = Y(u_1, 1)$
3. **set** $u_1 = \lambda Y(u_1, 1)$
4. **reset** $\lambda = u_1 / \sum_{j=1}^2 u_j$
5. **If** $u_1 \leq 1$ **and** $u_2 \leq 1$ **halt**
 else repeat for $u_1 > 1$

FIGURE 3 Pseudo code for coordination method in Figure 2 ($Q_1 = Q_2 = 1$)

than of that of firm 1, a unit of aggregate demand shifted from firm 2 to firm 1 will result in an *increase* in aggregate demand.¹²

Nothing prevents the redistribution from causing *both* firms to exceed the capacity constraint. If this were to occur, the closure of the model would endogenously change from Keynesian to supply constrained model. One might think that price is the logical adjusting value here when demand exceeds supply, but there is no price in the model as currently configured. Other options have been employed in multi-agent models, such as an adjustment in inventories, imports, or the saving rate as noted above.

The introduction of the demand reallocation process obviates the need for such a variable, since full capacity utilization for the economy as a whole is never reached. There is one exception. As a thought experiment, consider the special case in which the savings per unit of output is the *same* in both firms. Now shifting a unit of aggregate demand from firm 2 to firm 1 does not change the total. If capacity is now exceeded in firm 1, the algorithm will try to shift the unit of aggregate demand *back* to firm 2! This chattering solution cannot be ruled out and in practice prevents the sweep from converging on a feasible distribution of the u_i .¹³

There are a number of common elements of the agent-based approach that are intentionally ignored in this treatment of the real side of the model. Firm agents could, for example, *learn* to combine factors of production more efficiently, or invest in research and development to adjust their Leontief technologies, or use sophisticated models to forecast demand and therefore alter their investment decision away from Equation 4. None of these features is implemented in the current version of the model since this would diminish the connection to the original KKS model, in particular the role of aggregate demand. Many are, however, implemented on the financial side. There, financial agents learn, communicate with one another and forecast behavior of other agents using sophisticated models, all of which can significantly affect firm investment.

2.5 | Financial decision-making agents

The discussion of firm agents above has set the stage for the introduction of financial agents, who operate between surplus and deficit firms. The parameters defining this agent set are shown in Table 3. While there are $n = 676$ firms, the table indicates that there is $m = 1,000$ financial agents, so that on

¹²This is the so-called paradox of costs that is emblematic of KKS type models. It is, for example, the key to ‘wage-led’ growth, the well-know stagnationist argument that holds that redistribution of income from capitalists to workers will cause an increase in capacity utilization, *ceteris paribus*. Here the redistribution takes place by way of a change in the *composition* of firms rather than within any one firm, but the effect is, of course, the same.

¹³This would cause the simulation to be removed from the data.

TABLE 3 Financial agent parameters

Name	Symbol	Distribution	Range
<i>Observer Parameters</i>			
Number of FA	m	—	1,000
Initial share price responsiveness to traders forecast	$\bar{\psi}$	—	0.1
Change in ψ with respect to $\sum_{j=1}^m \phi_j$	ψ_f	—	10^{-5}
Change in ψ with respect to \sqrt{K}	ψ_K	—	5×10^{-6}
<i>Fixed parameters</i>			
FA j linked to firm i	m_{ji}	—	1–10
Mean number of FA per firm	\bar{m}_{ji}	—	1.49
Initial confidence parameter in trader's forecast	$\bar{\sigma}$	—	0.5
Initial liquidity	\bar{L}_j	—	1
Initial share weights	$\bar{\omega}_j$	—	1
<i>Distributed parameters</i>			
Network adjacency ^a	a_{ij}	B	0,1
Logistic smoothing parameter in FA forecast	γ_{0j}	U	5–6
Logistic smoothing parameter in FA forecast	γ_{1j}	U	8–10
Random error in FA forecast ^b	ϵ_j	N	–2–2
Lag length in FA forecast ^c	τ_j	U	3–13
<i>Computed parameters</i>			
Initial weight on FA			
Public signal ^d	$\bar{\kappa}_j$	$\{1 + e^{[\gamma_{0j} - \gamma_{1j}(1 - \bar{\sigma}_j)]}\}^{-1}$	0.2–0.5

B = binary; U = uniform distribution; N = normal; FA = financial agent. Source: Authors.

^aBidirectional link of FA i to j ; distribution either random or by preferential attachment (PA) at initialization. Random: number of links: 693 ± 15.2 or PA $1,000 \pm 0$. If linked $a_{ij} = 1$.

^bMean ± 2 SD.

^cLength in periods.

^dLogistic function.

average there are approximately 1.5 financial agents per firm. Financial agents are randomly allocated to firms such that each firm is linked to between one and ten financial agents.

Each financial agent begins life with one unit of liquidity, their initial capitalization. To simplify matters, the capital stock of the firms to which each financial agent is randomly associated is the sum of the initial liquidity of the associated financial agents.¹⁴ Thereafter strict balance sheet consistency is imposed on financial agents: since there is no currency or reserves, loans must be equal to deposits plus initial capitalization. Financial agents cannot fail or die; they can only run out of liquidity and must then refuse a loan request.

As noted above, financial agents only borrow to meet the needs of a deficit firm to which they could not otherwise provide financing. This approach differs from the traditional treatment as in Gai et al. (2011), for example, in that this paper focuses primarily on bank intermediation, rather than whether banks can acquire needed liquidity through short-term interbank borrowing. Here there is no

¹⁴Thus, a firm connected to two financial agents is instantiated with two units of capital and a firm with seven agents has seven units of capital. An initial correlation of firm size, measured by their capital stock, with the number of financial agents is then established.

lender of last resort that might respond to a short-term liquidity crisis. Neither can financial agents freely create money.¹⁵

In the hybrid KKS model studied here, the decision financial agents make is whether to extend loans to deficit firms. This decision depends on two factors: availability of liquidity and willingness to lend. Conditional on having sufficient liquidity (direct finance) or access to liquidity through the inter-bank market (indirect finance), 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. Financial agents make the decision to allow or disallow loans to deficit firms based on a relatively sophisticated forecast composed of two parts: a *private signal* based on private information and a *public signal*, based on the forecasts of their linked neighbors, which is public information. The relative weights the financial agents place on these two sources of information change depending upon the *confidence* they have in their own forecasting ability relative to that of linked neighbors. Table 3 shows that the initial confidence is set at 0.5, which means that financial agents initially weigh private and public information equally.

Financial agent interaction is based on the Bayesian rational network learning model. In this approach financial agents update priors determined by their linked neighbors by reference to the real-side performance of their own client. 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, in other words, the state of the world. Assume that θ_t is an observation from the stationary distribution of share price changes, distributed as $N(\bar{\theta}, \sigma^2)$. The j th financial agent has a prior subjective distribution $N(\phi_{pj}, 1/\rho_{pj})$, where ρ_{pj} is the precision of the prior signal, the inverse of the variance.

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

$$\mu_{jt} = \theta_{jt} + \varepsilon_{jt}$$

with $\varepsilon_{jt} \sim N(0, 1/\rho_{ej})$ where, ρ_{ej} , is the precision of the private signal. Each financial agent uses this signal to update and improve their prior precision

$$\rho_{jt} = \rho_{pj} + \rho_{ej}.$$

To implement this Bayesian updating scheme in the hybrid KKS model, let the private signal be 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 list $d_j(u_{jt})$ and then forms a forecast by regressing a subset, τ_j , of these utilization rates on time according to

$$u_{jt} = \hat{\beta}_{0j} + \hat{\beta}_{1j}t + \epsilon_{jt} \quad (10)$$

with $t = 1, 2, \dots, \tau_j$ and with ϵ_{jt} as a random error term. Since different financial agents weigh history differently, the size of the subset, τ_j , varies randomly between 3 and 13 time periods as seen in Table 3. The private signal, μ_{jt} , is thus

$$\mu_{jt} = \hat{\beta}_{1j}t$$

the trend regression slope coefficient. A positive trend is associated with a bullish private signal. The mean of the prior distribution, ϕ_{pj} , is determined by the weighted average of the forecasts of the financial agent's J' linked neighbors

$$\phi_{pj} = \sum_{i=1}^{J'} \omega_i \phi_i \quad (11)$$

¹⁵As explained in more detail below, money can be created but only when financial agents are forced to do so by the asynchronicity of the system.

where the subscript i , previously used for firms, is employed here to identify a linked neighbor of the j th financial agent and ω_i is the weight that the j th financial agent applies to the network link it shares with financial agent i .¹⁶

In period t , the financial agent's *forecast*, ϕ_{jt} , of the state of the economy is then the updated mean of the Bayesian prior distribution, ϕ_{pj} ,

$$\phi_{jt} = \kappa_{jt} \phi_{pj} + (1 - \kappa_{jt}) \hat{\beta}_{1jt} \quad (12)$$

where κ_{jt} is the current weight on the Bayesian prior. The relative weights of the public and private signals are 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 weights of the private signal and prior, it is necessary to map the forecast error into an open interval (0, 1).

For the moment, take the change in share price in any given sweep of the model as given.¹⁷ Financial agents' weight, κ_{jt} , depends upon confidence and if the agent has predicted the change in share price accurately, its confidence increases. Suppose the slope coefficient is $\hat{\beta}_{1jt} > 0$ and so the private signal is bullish. If the share price increases, $\Delta p_t > 0$, then the financial agent has made a correct forecast based on the private signal alone. 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.

According to Table 4 if the financial agent beats the market in any given period, the financial agent records in the memory list $f_j(\eta_{jt})$ 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, 0.5. The financial agent then computes an average of these signals over the idiosyncratic τ_j -length history to arrive at a raw forecast error. To avoid precipitous swings, the financial agent then attenuates the raw forecast error by way of a logistic smoothing function, as is common in machine learning models of neural networks (Russell & Norvig, 2010).¹⁸ The smoothing logistic function reduces the weight on the private signal to something slightly less than one.

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_{jt} will rise, and its complement will fall at a decreasing rate toward zero. The informational content offered by the agent's firm is increasingly ignored. Since its linked neighbors are taking into account this behavior in their own forecasts, it could be said that *social learning* is breaking down. Herding is beginning and in the limit, when all financial agents are herding, cascades occur (Chamley, 2004). Financial agents become progressively detached from the source of

¹⁶Detailed description of the financial network linking financial agents, including the determination of ω , follows below.

¹⁷See Section 3.2 for details.

¹⁸Consider for example the j th financial agent who initially places equal weights on both the private and prior public signals. The horizon for this particular financial agent is $\tau_j = 3$. In the last three periods, the financial agent beat the market, recording $\eta_j = 1$ each time. Does the financial agent stop updating, concluding that the weight on the private signal is then one? To avoid this unrealistic behavior on the part of the financial agent, the model assumes that successful forecasts are smoothed by the logistic function

$$\kappa_{jt} = \frac{1}{1 + e^{[\gamma_{0j} - \gamma_{1j}(1 - \frac{1}{\tau_j} \sum_{t=1}^{\tau_j} \eta_{jt})]}} \quad (13)$$

TABLE 4 Financial agent updating methods

Name	Symbol	Equation	Equation No.
<i>Observer variables</i>			
Total capital stock	\mathbf{K}_t	$\sum_{i=1}^n K_{it}$	
Change in share price ^a	Δp_t	$\psi_t \sum_{j=1}^m \omega_{jt} \phi_{jt} + \psi_K \sqrt{\mathbf{K}_{t-1}}$	15
<i>Endogenous variables</i>			
Private signal ^b	$\hat{\beta}_{1jt}$		10
Public signal ^c	ϕ_{pjt}	$\sum_{i=1}^m a_{ji} \omega_{it} \phi_{it}$	11
Forecast	ϕ_{jt}	$\kappa_{jt} \phi_{pjt} + (1 - \kappa_{jt}) \hat{\beta}_{1jt}$	12
Confidence in public signal	σ_{jt}	$(1/\tau_j) \sum_{i=1}^{\tau_j} \eta_{jti}$	
FA beats market ^d	η_{jt}	1	
Market beats FA ^e	η_{jt}	0	
FA follows market ^f	η_{jt}	0.5	
Weight on			
Public signal ^g	κ_{jt}	$\{1 + e^{[\gamma_{0j} - \gamma_{1j}(1 - \sigma_{jt})]}\}^{-1}$	13
Liquidity ^h	L_{jt}	$\bar{L}_{jt} + F_{it}/m_{jt}$	9
<i>List variablesⁱ</i>			
Data for private signal	$d_j(u_{jt})$		
History of forecast success/failure	$f_j(\eta_{jt})$		
<i>State variable</i>			
Capital share weights	ω_{jt}	$K_{jt-1}/\mathbf{K}_{t-1}$	14

FA = Financial agent. *Source*: Authors.

^aNatural log. Here ψ_t calibrated to S&P 500; see text for explanation.

^bDetermined by OLS regression of last τ_j values of u_j on time in each sweep. See text for explanation.

^cSummed over linked-neighbors of j .

^dPrivate signal predicts share-price movement, while public signal does not.

^ePublic signal predicts share-price movement, while private does not.

^fNeither public signal nor private signal predicts share price or both do.

^gLogistic function.

^hFinancial surplus equally shared among FA linked to firm i .

ⁱLength τ_j .

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 begin to rise (Hansen, 2012).

As noted above, financial agents communicate across adjacency lines, to their linked neighbors, to the benefit of firms. Specifically, if a financial agent is unable to satisfy the demand for borrowing by a deficit client firm, it may call on linked neighbors to ask for a loan. Linked financial agents who have sufficient funds agree to loan the originating financial agent funds to be channeled to the deficit firm.

Figure 4 shows how financial agents communicate. If the forecast of the j th financial agent is positive *and* it has sufficient liquidity or *access to* liquidity by way of its adjacent neighbors, it will make the loan to the deficit firm. If however, either of these conditions fail to hold, the loan is withheld and the planned investment is unrealized.

```

Since  $F_i < 0$                                 deficit firm  $i$  linked to  $FA_j$ 

if  $\phi_j > 0$                                  $FA_j$  willing to lend
  and  $L_j + \sum_{k=1}^{J'} L_k \geq |F_i|$      $FA_j$  able to lend

  set  $I_i = g_i K_i$                         firm  $i$  invests

else

  set  $I_i = 0$                                  $FA_j$  blocks loan
                                              firm  $i$  does not invest
                                              flips from deficit
                                              to surplus

```

FIGURE 4 Pseudo code for financial agent communication

The reduction in planned investment does not have the same impact as it would in the standard KKS model with accommodating money. There, one would expect a multiplier effect on aggregate demand that would exceed the reduction in planned investment denied in the second branch of the if-else statement in Figure 4. In the multi-agent world however, when the investment is blocked, the deficit firm becomes a surplus firm and its savings is available to other deficit firms that may themselves have been denied access to financing. Total investment need not fall, but total investment and number of blocked loans are inversely correlated in numerical simulations.

2.6 | The coordination environment of financial agents

Financial agents coordinate through a financial network that is imposed rather than emergent. Many stand-alone models of the interbank market associate links with borrowing and so the links are directed. The hybrid KKS model is not, strictly speaking, a model of the fragility of interbank lending.¹⁹ Since a given firm can oscillate between surplus and deficit due to its share in the overall level of aggregate demand, as well as capacity constraints, the firm's relationship to financial agency is necessarily bidirectional. The firm is sometimes a lender and sometimes a borrower. The focus here, however, is on the financial services that financial agents offer in the intermediation between surplus and deficit firms. It follows that for two such firm agents to interact, they must have at least one linked pair of financial agents in common. As the number of financial agents in a network increases, so too does the level of connectivity between firm agents. Thus, higher connectivity of financial agents reduces the degree to which any individual market forecast may affect financial flows. Moreover, large firms (measured by their capital stock) tend to have access to more sources of finance, as noted above, and thus are less susceptible to the blocking effect of bearish forecasts.

To gauge how network structure affects the hybrid KKS model's performance, several coordination mechanisms are considered: random and preferential attachment, with and without weights. Following Erdős and Rényi (1950), a random network is constructed by instantiating financial agents and then randomly linking them to existing financial agents. A second, due to Barabási and Albert (1999),

¹⁹Most models of interbank lending essentially depend on the mechanisms of virus propagation across networks, epidemiological in nature, that trace patterns of susceptibility, infection and recovery. See Toivanen (2013) and Weng, Menczer, and Ahn (2013). The analogy to the epidemiological models is obvious and correct: as an infected agent circulates on a grid, it transmits the disease to a variety of partners depending on their susceptibility. These models are realistic and much valuable work has been done using this framework to understand the recent financial crisis.

employs *preferential attachment* in which the likelihood that each additional financial agent links to an existing financial agent depends on the degree of the potential partner.²⁰ Once initialized as either a random or preferentially attached, the financial network remains fixed. Which coordination structure is more susceptible to financial crises can then be tested.²¹

Weights may also be applied to financial network links to distinguish larger and more influential financial agents from the rest. Each financial agent is associated with financing 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_{ji} , attached by the j th financial agent to the link to the i th agent is defined as

$$\omega_{ji} = \frac{K_{it-1}}{\sum_{i=1}^n K_{it-1}}. \quad (14)$$

The weight attached to a link by any financial agent is then the share of the total capital stock served by that 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 hybrid KKS model has come a long way from its progenitor with smooth, unerring intermediation of financial surpluses. Certainly, this is not the only way to integrate the financial sector but it shows how agent-based methods enhance the richness of the KKS approach.

3 | OBSERVATIONS ON THE HYBRID KKS MODEL

Has the introduction of the agent-based framework obliterated traditionally recognizable macroeconomic features of the KKS model? Happily, it has not. Aggregate demand and animal spirits continue to play a central role in the analysis, as do well-established features of KKS, such as the possibility of wage-led growth. Also present is one of the canons of post-Keynesian thought, the endogeneity of money.

3.1 | 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 preempted, the financial agent is left with no legal choice but to create liquidity when surplus firms demand their own deposited 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.

²⁰For 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 $1/4$.

²¹In 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, Masi, Precup, Gabbi, & Caldarelli, 2008). 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.

Observe that rather than *assuming* a KKS framework in which endogenous money allows animal spirits to play a defining role in the determination of investment and growth, the agent-based perspective allows a *derivation* of the KKS 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 results in the endogenous money creation that imparts to the model a KKS flavor. The reason is now clear: only 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* would a pure savings driven model assert itself.

No real agent would behave this way in a system that protects private property in deposits. As a result, the institutional properties of the financial system allow aggregate investment in period t to exceed savings in period $t - 1$. In other words, the agent-based model provides a microfoundation for endogenous money creation and in turn the KKS nature of the system itself.²² The emergent point is that any institutional framework that binds autonomous agents with financial surpluses and deficits in *asynchronous* exchanges will produce a quantity of endogenous money greater than or equal to zero and allow for an expanding economy driven by animal spirits (Gibson & Setterfield, 2015).

3.2 | Financial collapse in the hybrid KKS model

In the standard KKS model, there is little possibility of an endogenously generated crash. Certainly, there can be a deficiency in aggregate demand, but with the seamless transfer of financial surpluses from one firm to another, by way of accommodating money supply growth, any crisis would have to originate in and remain confined to the real side. In the hybrid KKS model, a drawdown in the share price, by contrast, can cause a decline in the rates of utilization on the real side. Nothing, however, has yet been said about how share prices are determined.

The simplest possible way to define the share price is to define the change in the natural log of the share price as the sum of the forecasts of financial agents

$$\Delta p_{t+1} = f\left(\sum_{j=1}^m \phi_{jt}\right) \quad (15)$$

where f is a functional form calibrated to the scale and variability of the S&P 500. The functional form used is (a) quadratic in the sum of forecasts; (b) augmented by a trend that depends on the aggregate capital stock; and (c) designed to allow capital share weighting of the financial agents' forecasts for selected simulations. While this is clearly an ad hoc approach to the pricing of shares, it is nonetheless consistent with the paper's focus on intermediation.²³ Table 4 shows the specific functional forms and Table 3 gives the values of the associated parameters. The share price is then a random walk during 'normal' times, but breaks out during organized bull or bear markets to produce a bubble and then possibly a crash. Financially induced collapses do occur in the hybrid KKS model, but they are relatively rare.²⁴

²²Note 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).

²³For a more sophisticated model, see LeBaron (2012), where agents explicitly buy, sell or hold equities.

²⁴Using 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 . The criterion for a build is designed to rule out a series that declines for a long period and then accelerates its decline. Thereafter, a crash is a decline of 50% in the share price within the final 25 weeks. This roughly corresponds to the single worst 6-month performance in the history of the S&P 500 index. The decline in question took place during the 2008–2009 financial crisis.

TABLE 5 Passive financial system

	Preferential attachment?			
	<i>No</i>		<i>Yes</i>	
	Weighted network?		Weighted network?	
	<i>No</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
Crash frequency ^a	1.1	3.3	0.00	0.00
GDP growth ^b	4.62E-04	4.62E-04	4.63E-04	4.62E-04
Loans denied ^c	0	0	0	0
Capacity utilization ^d	100	100	100	100
Prior weight ^e	0.31	0.31	0.31	0.31
Total runs	900	899	900	900

Source: Authors' computations.

^aCrashes per 1,000 runs.

^bAverage rate of weekly growth from logarithmic regression.

^cPercent of total firms.

^dMoving average over last 100 runs averaged over all runs.

^eAverage over all runs.

This observation is from the computational model that simulates real and financial performance over 1,500 weeks (30 year) periods.²⁵ 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 is then a total of 14,178 runs recorded in the database, including 3,599 control runs with no agent-based financial system present and 10,579 treatment runs with a financial system with the one-ply constraint discussed above.²⁶

Within this set of 14,178 runs, there were a total of 172 crashes with an active financial system, but only 4 without. Table 5 shows the crash frequency per 1,000 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 of a lack of intermediation. The table also shows capacity utilization and the average Bayesian prior weight. Table 6 shows the same data when the agent-based financial system is active.

Comparing the tables confirms that with a passive, fully accommodating, financial system there are fewer crashes in the financial sector ($t = -11.2$), faster growth ($t = 35.7$) and higher capacity utilization in the real sector ($t = 72,000$), results that agree with the theory developed above.²⁷ When the financial

²⁵This simulation length corresponds to discrete episodes of growth, such as the Golden Age (1948–1973) or Neoliberal growth regime (1980–2007) (Maddison, 2007).

²⁶For computational reasons, the total number of runs per batch was 900 and there were 16 batches run for a total of 14,400 simulations. If a crash were observed in the first 2,250 weeks, the run was deleted. The larger number of runs with an active financial system is needed to study the various configurations of network architecture, as seen in the following section.

²⁷The 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 6, only their relative levels. Moreover, the framework allows a consideration of the effect of financial structure alone, without mixing in the effect of an exogenously expanding (or contracting) economy on crash propensity.

TABLE 6 One-ply active financial system

	Preferential attachment?			
	No		Yes	
	Weighted network?		Weighted network?	
	No	Yes	No	Yes
Crash frequency ^a	26.7	24.1	6.7	8.2
GDP growth ^b	2.40E-04	2.44E-04	2.54E-04	2.53E-04
Loans denied ^c	19.8	20.0	20.5	20.5
Capacity utilization ^d	70	70	69	69
Prior weight ^e	0.29	0.30	0.30	0.30
Total runs	2,581	2,612	2,693	2,693

Source: Authors' computations.

^aCrashes per 1,000 runs.

^bAverage rate of weekly growth from logarithmic regression.

^cPercent of total firms.

^dMoving average over last 100 runs averaged over all runs.

^eAverage over all runs.

system is passive, 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 initiate a crisis.

These results demonstrate the importance of real-financial interactions for financial instability in the KKS model. When firms are financially constrained by lack of intermediation, 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 5 confirms that Brownian crises arise, but they are rare as would be expected.

The weight on the Bayesian prior is low, around 30% 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, Hirshleifer, & Welch, 1998).

Observe from the tables that the prior weight is slightly higher in the model with a one-ply active financial system. In Table 5 capacity utilization never deviates from 100% 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 sector is passive, since capacity utilization is always full and so there is nothing to be learned, except from random parameter shocks.

With the one-ply financial system present, however, deviations from full capacity utilization do indeed convey useful information. In the data set as a whole, the variance of the asset price with the

TABLE 7 Regression results

	(1) Crashes ^a	(2) Crashes ^a	(3) GDP Growth ^b	(4) GDP Growth ^b
Weighted network	−3.421E-07 (1.667E-06)		1.653E-06 (1.199E-05)	
Preferential attachment		−1.221E-05*** (1.689E-06)		1.114E-05 (1.198E-05)
Constant	1.110E-05*** (1.190E-06)	1.718E-05*** (1.495E-06)	2.470E-04*** (8.438E-06)	2.422E-04*** (8.320E-06)
R^2	0.00000000	0.00000341	0.00000000	0.00000005
Observations	1.57E+07	1.57E+07	1.57E+07	1.57E+07
F -stat	0.042	52	0.019	0.865

Source: Authors' computations.

^aCrashes per 1,000 runs of 1,500 weeks (30 years).

^bReal GDP growth per week.

Standard errors in parentheses. *** $p < .01$. ** $p < .05$. * $p < .1$.

active financial system is 1.02 but only 0.66 otherwise. A comparison of means shows a significant effect of the financial system 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 one-ply financial system present than not. This is verified by the extremely low correlation (-0.009) between the private and prior signals with the passive financial system. 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. A financial agent's informative signal is more likely to be correlated with the informative signals of its linked neighbors, unless all firms are at full capacity.

3.3 | The impact of network structure

Augmenting the traditional KKS model with an agent-based financial sector opens up a range of subtle but important issues. Here the question of the effect of network structure is addressed to illustrate the potential of the hybrid framework in advancing the KKS program. 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 (-0.004) so that the t -statistics are not affected by omitted variable bias. The effects of the two structural elements are shown in Tables 7 and 8.

The first two columns of Table 7 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 7 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 8 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 8 Regression results

	(1)	(2)	(3)	(4)
	Loans denied ^a	Loans denied ^a	Bayesian prior ^b	Bayesian prior ^b
Weighted network	0.8*** (0.009)		0.0029*** (4.989E-05)	
Preferential attachment		4.132*** (0.009)		0.009*** (4.988E-05)
Constant	136.1*** (0.007)	134.4*** (0.007)	0.292*** (3.550E-05)	0.289*** (3.591E-05)
R ²	0.00047796	0.01273840	0.00021442	0.00204497
Observations	1.57E+07	1.57E+07	1.57E+07	1.57E+07
F-stat	7526	2.03E + 05	3375	32233

Source: Authors' computations.

^aPercent of total number of firms.

^bAverage weight on Bayesian prior.

Standard errors in parentheses. *** $p < .01$. ** $p < .05$. * $p < .1$.

Table 8 indicates that there are more loans denied in both weighted and preferentially attached networks. Yet Table 7 shows no effect of network structure on GDP growth. This is easily understood in the context of the hybrid KKS model, since a higher number of loans denied does not necessarily reduce the total investment and the pace of growth. The large number of loans denied shows that the distribution of growth is affected by network structure, namely, large firms grow faster 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 8, 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 7 indicates that preferential attachment *reduces* crash frequency.

To understand why, consider a large firm associated with a highly linked financial agent. 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 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.

In this case, 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

The central message of the paper is that the traditional KKS model and the agent-based methodology are not necessarily 'strange bed-fellows'. The multi-agent perspective confers a number of analytical advantages as developed in this paper. First, it allows a non-standard microfoundation for the KKS model that is *not* based on the representative agent maximizing an inter-temporal utility function. Heterogenous agents on both the real and financial sides operate with bounded rationality in an informationally constrained environment. Second, the inclusion of an agent-based financial system allows a deep integration of the treatment of intermediation with the fundamental problem of aggregate demand, endogenous money and the balance between savings and investment, all central themes of the KKS program. In this sense the traditional KKS model with a passive financial sector simply ignores much of the rich texture that has emerged as financial markets have become more sophisticated. Kregel (1985) argued that Keynesian real-side models with no monetary and financial sectors were akin to 'Hamlet without the prince' and this paper shows one way in which this shortcoming can be addressed.

Setterfield and Budd (2011) was a first attempt to develop a hybrid KKS model, but there were problems in the demand sharing algorithm that essentially converted the framework into a trade model, with each firm behaving as if it were a country. Other papers, have dealt with the problem of full capacity utilization in various ways. This paper proposes a novel approach to aggregate demand sharing that eliminates the problem that a firm might exceed full capacity utilization. The method is simple and broadly corresponds to the idea that when consumers (or investors) cannot find what they need in one location, the demand spills over to another. As seen, this algorithm can instantaneously change a surplus to a deficit firm and vice-versa. It also preserves basic KKS concepts such as animal spirits and aggregate demand, features that are easily obscured in the agent-based approach.

Two important insights arise from the integration of KKS and the agent-based perspective. First, models that focus on financial crises from an epidemiological perspective have been successful in showing how a highly connected system of interbank lending can facilitate the propagation of financial disturbances. As important as these models have been, however, they appear to be fundamentally incomplete. Banks can cause crisis, but unless that crisis is validated by a real-side contraction in aggregate demand, that then reverberates back into the financial system, the effect will not be as profound. The 1987 crash, for example, differed fundamentally from the recent 2008 financial crisis, which was an outgrowth of deep-seated problems in the housing market.

Second, while connectedness in previous models amplifies financial disturbance, the effect of preferential attachment here seems to be the reverse. A more connected financial system as it interacts with the real side is more resistant to financial disturbance and actually shows fewer crashes than its less connected counterpart. Since the model of this paper differs from those that dominate the agent-based literature on financial crises, it is difficult to conclude that the latter are wrong. Instead, our results perhaps suggest that the complete story of real-financial interaction has yet to be told.

ACKNOWLEDGMENTS

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. Jonas Oppenheim provided an invaluable critical reading of an earlier draft. Programming assistance by Paul Wright is also gratefully acknowledged. We wish to thank Jeannette Wicks-Lim for making available advanced computational facilities of the Political Economy Research Institute at UMass, Amherst. Finally, three anonymous reviewers of the journal made valuable comments that contributed significantly to the paper.

REFERENCES

- Ashraf, Q., Gershman, B., & Howitt, P. (2011). *Banks, market organization, and macroeconomic performance: An agent-based computational analysis* (NBER Working Paper No. 17102). Cambridge, MA: National Bureau of Economic Research.
- Axtell, R. (1999). The emergence of firms in a population of agents. Retrieved from <http://www.brookings.edu/es/dynamics/papers/firms/firmspage.htm>
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (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, UK: Cambridge University Press.
- Chick, V. (1983). *Macroeconomics after Keynes: A reconsideration of the General Theory*. Cambridge, MA: MIT Press.
- Epstein, G. (2005). *Financialization and the world economy*. Cheltenham, England: Edward Elgar.
- Erdős, P., & Rényi, A. (1950). On the evolution of random graphs. *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, 5, 17–61.
- Gaffeo, E., Delli Gatti, D., Desiderio, S., & Gallegati, M. (2008). Adaptive microfoundations for emergent macroeconomics. *Eastern Economic Journal*, 34(4), 441–463.
- Gai, P., Haldane, A., & Kapadia, S. (2011). Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5), 453–470.
- Gibson, B. (2012). Trade, employment and the informal sector: An agent-based analysis. *Margin: The Journal of Applied Economic Research*, 6(2), 277–310.
- Gibson, B., & Setterfield, M. (2015). *Intermediation, money creation, and Keynesian macrodynamics in multi-agent systems* (Working Paper No. 1511). Department of Economics, New School for Social Research, New York.
- Hansen, L. P. (2012). *Challenges in identifying and measuring systemic risk* (Working Paper No. 18505). Cambridge, MA: National Bureau of Economic Research.
- Harras, G., & Sornette, D. (2011). How to grow a bubble: A model of myopic adapting agents. *Journal of Economic Behavior & Organization*, 80(1), 137–152.
- Iori, G., Masi, G. D., Precup, O., Gabbi, G., & Caldarelli, G. (2008). A network analysis of the italian overnight money market. *Journal of Economic Dynamics and Control*, 32(1), 259–278.
- Johansen, A., Ledoit, O., & Sornette, D. (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance*, 3(1), 219–255.
- Kalecki, M. (1937). The principle of increasing risk. *Economica*, 4(16), 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*. Paris, France: OECD Publishing.
- Palley, T. (2014). *Financialization: The economics of finance capital domination*. London: Palgrave Macmillan.
- Robertson, D. H. (1940). Effective demand and the multiplier. In *Essays in monetary theory* (pp. 114–121). London: P. S. King & Son.
- Russell, S. J., & Norvig, P. (2010). *Artificial intelligence - A modern approach*. New York: Pearson Education.
- Setterfield, M., & Budd, A. (2011). A Keynes-Kalecki model of cyclical growth with agent-based features. In P. Arestis (Ed.), *Microeconomics, macroeconomics and economic policy: Essays in honour of Malcolm Sawyer* (pp. 228–250). London: Palgrave Macmillan, UK.
- Skott, P. (2012). Theoretical and empirical shortcomings of the Kaleckian investment function. *Metroeconomica*, 63(1), 109–138.

- Sornette, D. (2003). *Why stock markets crash*. Princeton, NJ: Princeton University Press.
- Taylor, L. (1981). IS/LM in the tropics: Diagrammatics of the new structuralist macro critique. In *Economic stabilization in developing countries* (pp. 465–503). Washington, DC: Brookings Institute.
- Taylor, L. (1983). *Structuralist macroeconomics*. Boston, MA: Basic Books.
- Taylor, L., & O'Connell, S. A. (1985). A minsky crisis. *Quarterly Journal of Economics*, 100, 871–875.
- Tedeschi, G., Iori, G., & Gallegati, M. (2012). Herding effects in order driven markets: The rise and fall of gurus. *Journal of Economic Behavior & Organization*, 81(1), 82–96.
- Turner, S., Farmer, J. D., & Geanakoplos, J. (2012). Leverage causes fat tails and clustered volatility. *Quantitative Finance*, 12(5), 695–707.
- Toivanen, M. (2013). *Contagion in the interbank network: An epidemiological approach*. (Technical report Discussion Paper No. 19/2013). Cambridge, MA: Bank of Finland Research.
- Weng, L., Menczer, F., & Ahn, Y.-Y. (2013). Virality prediction and community structure in social networks. *Scientific Reports*, 3(1), 1–6.

How to cite this article: Real and financial crises in the Keynes–Kalecki structuralist model: An agent-based approach. *Metroeconomica*. 2018;69:566–592. <https://doi.org/10.1111/meca.12201>

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 demand sharing algorithm [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, 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