Real foundations of financial crisis

Bill Gibson and Mark Setterfield*

Abstract

There exists a well-developed, rigorous and elegant literature that uses the multi-agent systems (MAS) approach to analyze the financial crises. This literature focuses exclusively on the financial sector to the neglect of the real economy. In this paper, we build on the real-sector MAS model of Setterfield and Budd (2008), relaxing some of the restrictive assumptions regarding the goods market in that model, and adding a financial sector. The latter is inspired by the MAS financial model of the work of Johansen et al. (2000), Voit (2005),Vandewalle et al. (1999) and others, but again relaxes some of the assumptions found in the original. The result is a more realistic model of the economy in which the real and financial sectors are integrated and interact with one another. It is shown that the real economy contributes to the evolution of the financial sector in two ways. First, it generates the initial conditions required for the log-periodic fluctuation of asset prices leading up to a crash, as observed by Johansen et al. (2000) and others. Second, the real economy appears to seed financial crises, so that even if financial markets are efficient in the sense of Fama (1970), they need not be crash-proof. Financial crises are exceedingly rare in the model and this raises the question of the proper amount of resources to devote to preventing a financial crisis.

1 Introduction

The motivation for this paper is the "great recession", the single most important macroeconomic event in recent US economic history. The paper employs a multi-agent systems model to study the effects of adding a real sector to a relatively well- developed and now standard analysis of the financial system based on the percolation model. It builds on an existing real-side MAS model, due to Setterfield and Budd (2008), which has roots in the structuralist tradition in growth modeling (see Taylor (1983), Taylor (1991),and Setterfield (2010)). It is shown that the real economy contributes to the evolution of the financial sector in two ways. First, it generates the initial conditions required for the log-periodic fluctuation of asset prices leading up to a crash, as observed by Johansen et al. (2000) and others. Second, the real economy appears to seed financial crises, so that even if financial markets are efficient in the sense of Fama (1970), they need not be crash-proof. Financial crises are exceedingly rare in the model and this raises the question of the proper amount of resources to devote to preventing a financial crisis.

The paper is organized as follows. The next section reviews the existing literature on financial crises. The third section discusses herds and cascades as the theoretical framework on which the empirical model is built. The real and financial sides of the model are elaborated in sections four and five. The fourth section presents some simulations of the model and the fifth section concludes.

2 Existing literature

Immediately prior to the onset of the great recession, macroeconomics was dominated by the view that the US was in the midst of a great moderation–a marked reduction in the volatility of the economy, at least as

*April 2010, version 1.02; John Converse Professor of Economics, University of Vermont, Burlington, VT 05405; e-mail bill.gibson@uvm.edu; http//www.uvm.edu/wgibson and Mark Setterfield, Professor of Economics, Trinity College, Hartford, CT 06106 mark.setterfield@trincoll.edu; agent-based models, JEL codes D58, G01, G12, B16, C00.

compared to the 1970s and 80s. Some economists are now rethinking their earlier assumptions. Acemoglu (2009), for example, observes

> It is an opportunity for us, and here I mean the majority of the economics profession, unfortunately myself included, to be disabused of certain notions that we should not have accepted in the first place. It is also an opportunity for us to step back and consider the most important lessons we have learned from our theoretical and empirical investigations that remain untarnished by recent events and ask if they can provide guidance in further policy debates. The era of aggregate volatility had come to an end. We believed that through astute policies, including better methods of communication and inventory control, business cycles were conquered. Our belief in a more benign economy made us more optimistic about the stock market and housing market. If any contraction came, it must be soft and short-lived and it becomes easier to believe that financial intermediaries, firms and consumers should not worry about large drops in asset values.

The relative attenuation of business cycle volatility over the century is now a fact of economic history, whether because of monetary policy or because of technological advances that permit more rapid adjustment in goods and factor markets to demand shifts. In particular, resources flow more easily from declining firms or sectors without creating as much unemployment, delay and capital loss. Moreover, the financial sector is better able to diversify risk in ways that enable individual firms to explore technological advances, absorb idiosyncratic shocks and discover profitable opportunities that would have been impossible earlier when finance was more rigidly allocated.

But while the financial system has always been the signature institution of dynamic capitalist economies, and the binding constraint on less successful ones, recent events clearly indicate that the financial sector is not always entirely subservient to the real. In the current crisis, the financial sector did not simply *reflect* more fundamental problems of the real side, such as inventory or trade cycles, shortage of critical raw materials or inadequate infrastructure or supplies of skilled labor, to name but a few. Instead, it *amplified* them. Rather than attenuate the effects of an excess supply of housing, the financial sector multiplied their intensity, leading to a credit freeze and the necessity of emergency fiscal intervention of literally historic proportions. The financial system did not simply fail to resolve problems on the real side of the economy, it became their cause.

More than a few observers noticed this latency prior to the crash. The *financialization* literature is an admirable attempt to formally model the financial sector and its effects on growth (See the review by Hein and van Treeck (2010)). This literature was only implicitly critical of previous models that had all but ignored the financial sector by effectively assuming perfect accommodation. In traditional models, the financial sector only served to "oil the wheels" of real sector dynamics in a frictionless manner. Any downturn in the system had its origin in the real rather than financial side. The skeptics were few: Kregel (1985), for one, famously charged that they were guilty of staging "Hamlet without the Prince". Minsky is another important counterexample, but he did not propose a tractable analytical framework. Taylor and O'Connell (1985) was an early attempt to formalize some of his insights, but loses much of the richness of Minsky's work due to necessary simplifications.

The new-wave models of the effects of financialization on growth were certainly a step in the right direction. The approach was to add more structural parameters to the real side to mediate the influence of the financial sector. Taking as the point of departure the canonical one-sector model of growth, the literature typically amends the model by parameterizing interest rates, debt-to-income ratios, corporate retention rates and shareholder value orientation among others. The net effect of financialization can then be studied by calculating a *total* derivative that can be decomposed into a string of partial derivatives. A consistent finding in the literature is that the effect of increasing financialization on growth is *ambiguous*. As in the IS-LM model of Hicks, in which the interest rate signals the need to slow investment, the financial side of the market reacts to the excess demand for loanable funds by reflexively impinging on planned investment one way or another.

As Acemoglu notes, the problem in *all* the pre-crisis literature was that it abstracted away from complications that were imagined to be of secondary importance. The argument is that when the financial system was less developed, less able to parse risk and diffuse real-side shocks, it was also less likely to become a source of instability itself. To understand and appreciate the power and increasingly crucial role the financial system plays in supporting the real side, and indeed in moderating the business cycle itself, it is now necessary to reintegrate it in a more fundamental way into the analysis. Specifically, human agency within the financial sector must be taken into account.

To function effectively in the ways we observe in the modern economy, the financial system must evolve "complex webs of counterparty relations," to use Acemoglu's term. When these new structural networks appear in the economy, it is logical to think they would bring with them a range of problems related to their own evolution. Following Albert and Barabási (2002), networks are characterized by their degree distribution, connectedness and growth. This imparts to networks their stability and resistance to attack, in short their susceptibility to external shock. As is well known, networks that expand according to preferential attachment exhibit very different properties than random networks and so it is important to examine the paths along which these structures develop to understand their properties. In particular, if one abstracts from the formation of the financial network as a complex system, then the idea that there could arise some potentially destructive emergent properties of that system is lost.

In other words, if the strength and resilience of the financial system stem from the web of counterparty relationships established among traders, it is within the grasp of contemporary theory to model the system as a real graph. This approach differs from the "first generation" microfoundations project based on the omniscient representative agent in its emphasis on human agency and its explicit attention to decision making in the financial sector.[1]

Note that the central issue here is *not* the shortcomings of standard neoclassical theory and the implicit critique offered by heterodox economists in the financialization literature. What Acemoglu and others find wrong is far more profound and applies almost as much to the critics as it does to standard theory. Thus, in *Why Markets Fail*, largely a history of 20th century economic thought, Cassidy (2009) identifies a "band of heterodox economists" as recently making the most important contribution to economics since Keynes. But as the argument unfolds, it becomes clear that by "heterodox," Cassidy does not mean structuralists, post-Keynesians, Marxians, etc., but rather *behavioralists* inspired by the work of Kahneman and Tversky (2000),Olson (Olson), Thaler (1993), Schleifer (2000), Ariely (2008) and others. More realistic models based on the principles of agent heterogeneity, bounded rationality and complex interaction seem to be required.

The central complaint about introducing systemic complexity into economic modeling is that the ability to derive general theorems appears to be lost. Computational methods lack the precision and generality of traditional approaches based on closed-form mathematical models, and have been called "works of fiction" by Cartwright and other appropriately skeptical philosophers of science. Much of the heterodox literature had already begun in some ways to abandon the axiomatic method in any case. Models that rely on heavy parameterization of important variables and non-analytical conditions for stability, such as coefficients that are "not too big", or that a sequence of multiplied parameters has a particular sign, are *essentially simulation models*. They are at least hybridized in that their analytical results often depend on combinations of measured quantities (the hallmark of computational models), highly specific conditions that can hardly be stipulated axiomatically or taken as self-evident. This is not to say that these models are unrealistic or that their parameters are not locally self-evident, but only that any claim to generality is compromised by the way they are constructed.

The multi-agent systems approach developed here embraces much of the prior work on the nature of the interaction of financial and real models. These models are explicitly calibrated to parameters that are chosen for their realism. Behavior is modeled as much according to the behavioralists noted above as on strict principles of optimization. The MAS approach to analyzing such systems involves focusing on heterogeneous agency (rather than proceeding directly to analysis at a higher level of aggregation, as in structuralism), and then analyzing the results of the interaction of these agents statistically, looking for emergent aggregate relationships. Epstein (2006) calls this a "generative" approach to social science, noting

[1] While the standard theory, tightly focused on real-side interactions based on a representative agent, might have enjoyed an illustrious past, Kirman and many others have argued that the inherent limitations of this approach are increasingly self-evident (see Kirman (1989),Kirman (1992)).

that to fully understand social systems researchers must be able to *build them* and build them in ways that result in realistic emergent properties.

It follows that there is no fundamental distinction between *structure and agency*, at least in this paper. It is taken as self-evident that for social structure to be *social*, it must have arisen as the product of agency. This removes at a stroke questions about the source of social structure such as, for example, "shareholder value" and social norms that are intrinsic to structuralist growth models. Social structure presupposes the prior existence of human action–intentional or otherwise. Accordingly, that social structure is the product of agency is exactly what makes it social (as opposed to natural) in character. This is a view shared by many social theorists with no obvious connection to the MAS approach. Hence, for example, Lawson (Lawson, p. 167) argues that "if the human race were to disappear tomorrow, social structure would disappear along with it. Social structure depends upon human beings."

Obviously, this does not deny the reverse line of causality, that structure and agency interact as part of a recursive dynamic. Hence even as social structure is the product of agency, so existing social structure both constrains and facilitates human action, and hence its subsequent influence on social structure. Formal rules governing traffic flows were created by the very agents that must then obey them and as such these structures are nothing but ossified agency. Inter-agent communication presupposes the existence of a common language and thus social structure serves to expand the set of activities available to agents. At the same time, social structure also infuses agency, in the sense of determining (with stochastic variation) its internal data structure–the beliefs, values, aspirations, preferences, etc. that individual agents hold and that inform their decision making and hence their subsequent determination of social structure.

3 Herding, cascades and crises

According to the efficient market hypothesis, only the revelation of new and significant information could cause a financial crisis. But as many authors in this field point out, *ex post* analyses of actual historical crashes do not reveal any particular event that sparks a crisis. Instead, the conditions that exist at the time of the crisis are also observed to exist prior to the crisis, but are not acted upon. As was remarked earlier, these same conditions characterized the financial crisis of 2008. Although associated in retrospect with a precipitous rise in household indebtedness, the trend in household balance sheets long pre-dated the crisis itself (Cynamon and Fazzari (2008).

An essential feature of the percolation based financial sector models and the model of this paper is that they are consistent with this most basic "signature" of financial crises: a crash can occur without any new information arising (or at least in response to only a minor exogenous shock). This is because of the interplay of the private and social signals in the determination of traders' forecasts and, in particular, the possibility that the social signal will come to dominate the private signal completely, resulting in herding behavior.

An example serves to illustrate this point. A large number of passengers randomly distributed on the deck of ship will not affect its ability to remain on an even keel, even in heavy seas. But if passengers are free to move about, and if they all move in the same direction, the ship may lurch in that direction and even (in the extreme case) capsize. But why would the passengers ever contribute to their own undoing? The answer lies in their response to the private and social signals they receive from their experience aboard the ship. Each of the passengers is aware of the rolling motion of the ship, and may react to this private signal. They also observe other passengers reacting and may follow suit (a social signal). If so, the passengers themselves *will contribute to the ship's instability* but in so doing will reinforce their own beliefs that *some kind* of reaction is necessary.

Shin (2005) recalls the opening of the Millennium Bridge over the Thames in London in June 2000 when, shortly after the bridge's opening, it began shaking so violently that it was subsequently closed for more than 18 months. Engineers discovered that the event was triggered by a lateral disturbance at 1 hertz–which could result from a normal human footfall, or a gust of wind. And as Shin notes, once the initial disturbance occurred, because people react to their environment, they adjusted to the disturbance in a similar fashion at the same time. This synchronous behavior fed back to the bridge, causing it to move in a more exaggerated fashion, which caused further synchronous adjustments by the people on the bridge, and so on–resulting in the violent shaking recorded by BBC cameras covering the opening of the bridge. In short, the initial shock was amplified by adjustments within the system. Shock absorbers to damp the initial vibration were all that was needed.

In essence, the behavior we are now contemplating is similar to that considered by Keynes (1936) in his discussion of decision making under uncertainty, based on the psychological propensity of the individual to assume other agents have superior information. Financial markets are the supreme example of an environment where individuals react to what‘s happening around them, and where individual‘s actions affect the outcomes themselves (Shin, 2005, p. 383). When a trader sells a share because his private signal is weak, other traders learn from this event. The sale indicates to others that the trader in question no longer believes it is worthwhile to hold the asset he is selling, a process of social learning. (Chamley, 2004, p. 66).

But the same process of forecast formation can result in other behavioral outcomes. Hence an individual trader is said to engage in herding if she buys an asset purely in response to the observation that other traders are buying. In other words, herding by an individual agent involves acting in response to a social signal and independently of her private signal. For example, the herding individual‘s private signal may suggest a bullish outlook is warranted, but the trader will nevertheless adopt a bearish position due to the influence of colleagues. The behavior of the herding trader reveals no private information.

Social influence gives rise to the possibility that a herd will form. A herd is defined as a subset of traders who all take the same action after some date (Chamley, 2004). In financial markets herding creates *order* as opposed to the disorder necessary for the proper functioning of the market. Not all of the individuals who make up a herd are necessarily herding: they are simply making the same forecast (and therefore choosing the same action) as other members of the herd, although they could have reached a different forecast and therefore chosen to act differently. There can still be social learning in a herd, then, although the amount of this social learning must be very small in order for the herd to be sustained (Chamley, 2004, p. 61).

Ultimately, if the degree of orderliness in the market reaches a point where all traders engage in herding, the result is an informational cascade (Bikhchandani et al., 1992). In a cascade, the social signal dominates the private signal of any individual trader: no trader‘s actions are influenced by their private signal. As such, no private information is revealed by market behavior. No agent learns anything over time, as a result of which their behavior is repeated period after period, and the cascade continues indefinitely. As noted by (Chamley, 2004, p. 61), in such circumstances “the failure of social learning is spectacular”.

4 Model

The real side of the model is based on Setterfield and Budd (2008) which treats individual agents as essentially closed economies, not unlike the hunter-gather communities recently modeled by Choi and Bowles. (2007).[2] This paper incorporates the effects of economy-wide aggregate demand on individual agents and introduces financial intermediaries to balance the excess savings of some agents with the deficits of others.

The main results of Setterfield and Budd (2008) are:

1. Capacity utilization and growth will fluctuate in an aperiodic manner due to the evolution of a learning variable, characterized there as the state of long-run expectations (SOLE).

2. In the process of this cyclical growth, a Mandelbrotian (scalable) distribution of firm sizes will emerge.[3] Setterfield and Budd (2008) is not a model of crisis, however. The model evolves smoothly, with no herding behavior and there is no explicit description of the financial sector to impede growth. Finance is implicitly assumed to perfectly accommodate the dynamics of real-side accumulation;

[2]These societies generated their own aggregate demand in the goods market and the focus of the paper was on relative rates of accumulation of capital with an evolving state of long-run expectations that depended on both the local as well as economy-wide performance. The model was stylized and did not account explicitly for exports and imports, but nonetheless captured these essential features of dynamic, demand-led economies.

[3]The hypothesis of a Gaussian distribution of firms is rejected frequently in the literature (Axtell, 1999)

3. When investment plans of heterogeneous firms depend on each other, there are multiple equilibria. The model of Setterfield and Budd (2008), like most models with multiple equilibria, is mathematically intractable and thus requires statistical/computational methods to proceed.

4. In a closed economy with no monetary sector, saving is the only source of funding for current investment. But planned or intended investment is determined independently of savings. Whether investment is less than, equal to, or greater than savings is determined endogenously. If required investment is greater than what can be produced with the current period capital stock, then accumulation must proceed at the pace set by the savings out of the capital-limited production. In this case, macroeconomic outcomes can be said to be classical inasmuch as savings determines investment. If required investment is less than full capacity production, accumulation will proceed at the pace set by firms' intended investment. In this case, outcomes can be said to be Keynesian, the realized change in the capital stock determined independently of decisions to save. The macroeconomic structure of the model developed here, then, can be termed a Classical-Keynesian hybrid, in the sense that it admits the possibility of both savings- and investment-driven growth (Gibson, 2007).

While the potential for saving-constrained expansion may not appeal to all structuralists as a description of the accumulation process, the multi-agent model developed here has two main advantages. First, macro properties that depend on the micro structure of the model emerge from the micro structure itself without having to be pre-ordained. There is no need to decide whether an economy is savings or investment driven, since the answer is endogenous. Second, even if the growth regime is Keynesian (constrained in the first instance by intended investment), the possibility exists that actual investment may fall short of what is intended by virtue of the failure of individual firms to find sources of funding for their planned investment projects. Finance thus plays a vital role in mediating growth outcomes regardless of the precise character of the aggregate growth regime.

Agents in the model are categorized as either firms on the real side, or traders in the financial sector. Following Setterfield and Budd (2008), planned investment of the ith firm, I_i, is determined by a combination of an animal spirits term α_i, which differs among firms and reflects the state of long-run expectations, the profit rate, r_i and capacity utilization, u_i. For the ith firm

$$I_i = g_i K_i$$

$$g_i = \alpha_i + g_{ri} r_i + g_{ui} u \tag{1}$$

where g_{ri} is the partial derivative of the growth rate of the capital stock g_i with respect to the profit rate and g_{ui} is the partial with respect to u_i. In an economy with no financial sector, each firm must rely on its idiosyncratic savings rate out of profits, s_i, to finance its accumulation

$$\begin{aligned} s_i r_i K_i &= g_i K_i \\ s_i r_i &= \alpha_i + g_{ri} r_i + g_{ui} u_i \end{aligned}$$

Note that r_i can always be written as

$$r_i = \frac{\pi_i X_i}{K_i} \tag{2}$$

so that

$$\begin{aligned} r_i &= \frac{\pi_i X_i}{Q_i} \frac{Q_i}{K_i} \\ &= \frac{\pi_i u_i}{v_i} \end{aligned}$$

so that

$$u_i = \frac{r_i v_i}{\pi_i}$$

where π_i is profit per unit of output and $v_i = K_i/Q_i$ is the capital-output ratio with output measured at capacity, Q_i. This implies that

$$
\begin{aligned}
s_i r_i &= \alpha_i + (g_{ri} + g_{ui}\frac{v_i}{\pi_i})r_i \\
r_i &= \frac{\alpha_i}{s_i - g_{ri} - g_{ui}\frac{\pi_i}{v_i}} \\
g_i = r_i s_i &= \frac{s_i \alpha_i}{s_i - g_{ri} - g_{ui}\frac{\pi_i}{v_i}}
\end{aligned}
$$

Hence, growth increases with the α, the profit and utilization coefficients and the profit coefficient and falls with the savings rate and the capital-output ratio. All these are standard comparative static results.

From the second of this last set of equations, it is obvious that holding other parameters constant, a rise in s will lower the profit rate. An improvement in profits will occur for any savings rate $s' < s$. Note, however, that there is a maximum profit rate given by

$$
r_{max} = \frac{\pi_i}{v_i}
$$

with $u = 1$. Hence rational agents will improve their profitability by lowering the savings rate until $s = s^*$

$$
\begin{aligned}
s_i^* \frac{\pi_i}{v_i} &= \alpha_i + g_{ri}\frac{\pi_i}{v_i} + g_{ui} \\
s_i^* &= \frac{\alpha_i + g_{ui}}{\pi_i} v_i + g_{ri}
\end{aligned}
$$

Any $s > s^*$ should either be consumed or invested in some other production process.

Following Setterfield and Budd (2008), define a discrete dynamic process by

$$
K_{t+1} = (1 + \delta)K_t + (\alpha + g_r r_t + g_{ui} u_t)K_{t-1}
$$

where δ is the rate of depreciation and investment and the i subscript has been suppressed for clarity. The dynamics of this system are exceedingly simple and its trajectory is shown in figure 1.

If for some reason $s_i = s_i^* \, \forall i$, each process would become a fractal image of the economy as a whole. Since each cell would grow at full capacity utilization the economy as a whole would also grow at full capacity utilization. It is a model of balanced growth in that each individual firm produces, saves and then makes an investment decision based on how much of its output is saved and how much is consumed *locally* by its owners and workers. Those involved in firm i are paid in the product of firm i and may in fact trade with agents involved in firm j, but trade must always be in balance. If too much is saved by the workers and owners of firm j, then capacity utilization, u_j, falls and plans for further accumulation are scaled back. If too little is saved, then u_j rises and accumulation rises with it. In Setterfield and Budd (2008) the share of aggregate demand is variable and the fractal nature of firms essentially divorces production from economy-wide consumption. It is as if each set of agents associated with each firm were paid a product wage or return to capital. In a second stage, agents then trade among themselves according to their own consumption preferences. The model is always in a fully adjusted long-run equilibrium in the sense that, up to a stochastic variation, equation 3 holds. Were there no such variation, the model follows a von Neumann turnpike, with each sector plowing back into its capital stock precisely the amount necessary to meet its share of aggregate demand.

In Setterfield and Budd (2008), each firm's error term is correlated to a different degree with business conditions in the recent past. These conditions include both the individual firm's recent performance, and that of the entire economy. Firms influence each other indirectly via their contribution to aggregate economic performance: an indicator that is posted on a blackboard to which all firms pay varying degrees of attention as in Carver and Lesser (1992). The correlation of the random term in Setterfield and Budd (2008)is based

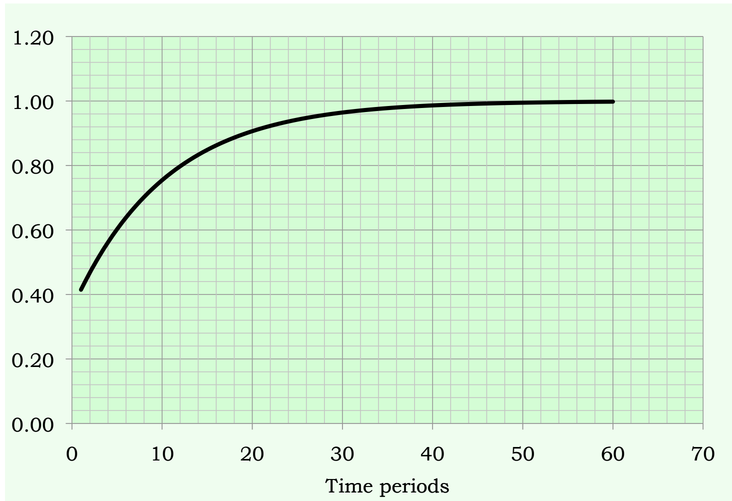

Figure 1: $u(t) \to 1$ with $s = s^*$

exclusively on a private signal, which they derive from from their own conditions of demand. Individual firms do not directly pay attention to one another (there is no social signal), so social learning is present in the model at all times in that other firms can observe the behavior of the ith firm, confident that it bears information about the state of nature rather than an imitation of other firms (Chamley, 2004). Moreover, the process by which the error term is revised when coupled with the process by which expectations are formed in Setterfield and Budd (2008) means that the dynamics of the model are entirely adaptive.

The fractal nature of the Setterfield and Budd (2008) model is inadequate in ways that are important to understand. First, firms must finance their own investment since there is no borrowing. Second, output is determined only by the consumption of the owner of the firm. If the savings rate out of profit, s_i, for firm i falls, then u rises since the owner's demand has increased. If u rises, then there is more investment in the next period. Thus, the same individual, the owner of the firm, is making a decision about how much of his own product to consume.[4] Output and capacity utilization are entirely subjectively driven by the preferences of the owner of the firm. The owner may observe the blackboard and see that other owners are also consuming a lot of their output. But this has no logical impact on his own demand, since their own consumption is unrelated. It is only a psychological boost that the owner gets. In particular, if other sectors reach full capacity utilization, then there is no spillover to the ith firm. The Setterfield and Budd (2008) model admits multiple equilibria, of course, since investment depends on the interaction of preferences of non-workers. This is similar to Romer's endogenous growth model with coordination failures: firms expand only when they believe other firms will also expand.

In this paper the real side of the model differs from Setterfield and Budd (2008) in two important ways. First, the assumption of sectorally balanced growth is dropped, so that demand shocks affect the level of capacity utilization in each firm. Supply is still perfectly elastic up to full capacity utilization based on the evolution of the capital stock in each sector and fixed capital-output ratios. The fraction of output set aside

[4]There is no need to forecast demand since the owner of the firm controls it. There is nothing in Setterfield and Budd (2008) that says that the owner cannot exchange his own output for that of another owner. If the owner's demand rises for some other good, then output could also rise and the savings rate shrink.

for accumulation depends positively on the utilization ratio, but this is determined by aggregate demand in the economy as a whole, not just the sector itself. When demand exceeds capacity for any sector j, excess demand spills over to other sectors raising their capacity utilization. The share of aggregate demand, θ_i, absorbed by ith sector would in principle depend on price, but in this version prices are fixed and so θ_i, is

$$\theta_i \sim K_i/v_i \tag{3}$$

The share θ_i is revised, however, when firms encounter the full capacity constraint. If no firm reaches full capacity utilization in one iteration, shares remain fixed for the next. But if some subset $f \subset N_f$ where N_f is the number of firms, operates at full capacity, demand for the jth firms rises to absorb the overflow. Output equations can thus be expressed as

$$\begin{bmatrix} \mathbf{u} \\ \mathbf{1} \end{bmatrix}_t = \begin{bmatrix} \theta_t & 0 \\ 0 & (1-\theta_t) \end{bmatrix} \left[\begin{bmatrix} \mathbf{D_i} & \mathbf{D_j} \\ \mathbf{D_i} & \mathbf{D_j} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{1} \end{bmatrix}_{t-1} + \begin{bmatrix} \alpha_\mathbf{i} + \alpha_j \\ \alpha_\mathbf{i} + \alpha_\mathbf{j} \end{bmatrix} \right] \tag{4}$$

where $D_i = 1 - s_i \pi_i + g_u + g_r u\pi/v_i$ is vector of demands from sector i and $\mathbf{u}$ is vector of capacity utilization for sectors with $u < 1$. The model solves this vector equation by way of the Gauss-Seidel method. At each iteration, the vector of θs 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 sector that had. The equation can be written as

$$\theta_t = \frac{u_i K_i/v}{\sum_{i=1}^{n} u_i K_i/v_i}$$

The order in which the sectors are allocated the spillover from any one of the sectors that reaches full capacity utilization is random so that no particular sector 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.

Firm savings rates are assigned according to the identity $s_i = g_i/r_i$ where g is defined by equation 1 above and r by equation 2. Note that while g_u and g_r are defined on a uniform probability distribution, the central limit theorem implies that the distribution of the s_i is approximately normal.

The parameter values shown in figure 2 define a *financial surplus* for each firm

$$S_{fi} = (s_i \pi_i u_i/v_i - g_i)K_i \tag{5}$$

as the difference between savings out of profits and planned investment. Macroeconomic equilibrium requires that

$$\sum_{i}^{n} S_{fi} = 0 \tag{6}$$

In the fractal approach to investment in Setterfield and Budd (2008), $S_{fi} = 0$, $\forall i$. Here it is possible that in any period, an individual firm generates sufficient savings of its own to fund current planned investment. But if not, then the firm must borrow from other firms. Under the assumption that equation 6 holds for all but a pair i, j, it follows that they must have equal and opposite financial surpluses: such that $S_{fi} = -S_{fj}$. If these firms are able to locate each other on a finite grid it is in the interest of both agents to make the loan. The agent for which the financial surplus is positive cannot invest the surplus in her own firm since that would generate excess capacity and cause the rate of profit to fall. With a deficit, the second firm cannot invest and thus the capacity utilization will fall for *all* firms, including its own. Let z denote a jump process that is one if counterparties (i, j) can agree on a transaction and zero if not. It follows that total investment will be reduced by $\Delta I \geq -S_{fi}$. The reason for the inequality is that in the model each agent has a determinant amount of investment required to reach full capacity utilization according to the fractal investment function as shown in figure 1. It is in the best interest of each firm to stay on the balanced

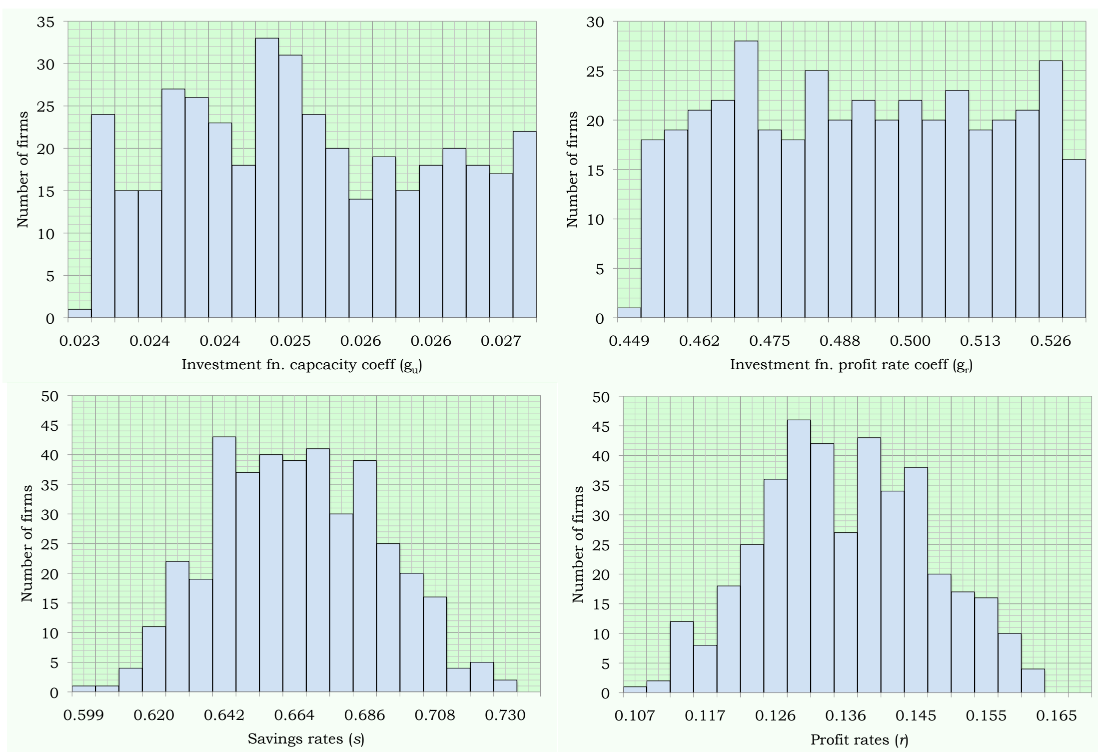

Figure 2: Parameter values

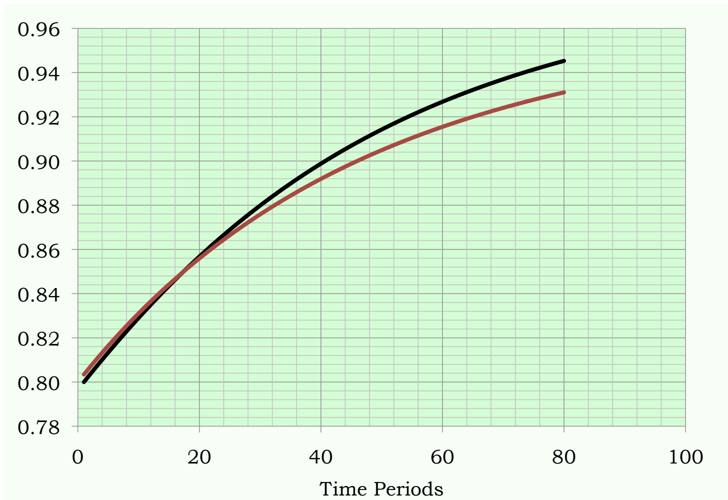

Figure 3: Capacity utilization in a binary firm structure with lossless finance

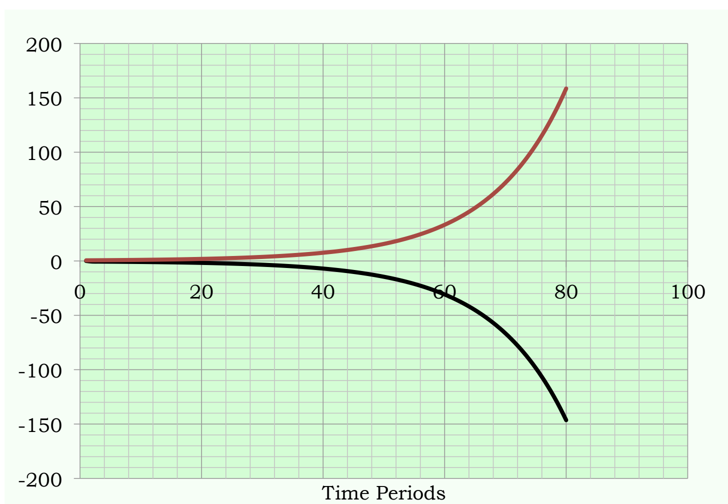

Figure 4: Financial surplus in a binary firm structure with lossless finance

Table 1: Capacity utilization u after trading day

	Average of 100 runs					
	1	10	25	50	75	100
Firm 1	0.806	0.793	0.835	0.869	0.905	0.947
Firm 2	0.806	0.793	0.835	0.869	0.905	0.947
Firm 3	0.805	0.771	0.880	0.864	0.936	0.916
Firm 4	0.821	0.827	0.858	0.921	0.907	0.925
Firm 5	0.806	0.802	0.858	0.885	0.948	0.961
Firm 6	0.806	0.799	0.838	0.918	0.931	0.939
Firm 7	0.808	0.795	0.846	0.872	0.918	0.962
Firm 8	0.813	0.806	0.889	0.879	0.933	0.955
Firm 9	0.801	0.751	0.861	0.884	0.932	0.956
Firm 10	0.802	0.825	0.869	0.891	0.911	0.948

Source: Author's calculations.

growth path, conditional on the assumption that other firms will do the same. Figures 3 and 4 shows the evolution of a pair of firms such as counterparties (i, j). Both firms are on their designated growth paths to full capacity utilization, but the financial surplus of firm 1 is always negative while that of its counterparty is positive. The figure illustrates a perfect financial market, frictionless, with no ts associated with lending or borrowing surplus funds. Firm 1 saves too little for the chosen growth path, but firm 2 is happy to supply her excess savings to finance the other firm's expansion. Both firms are better off under this cooperative relationship since the path of the firm with $S_{fi} > 0$ depends on the success of the second firm in filling its financial deficit. Both know that if the other fails to invest the proper amount, profits in both industries will suffer. The smooth functioning of the financial system in the background assures that the *coordination* problem will be resolved.

Figure 5 shows the GDP for 100 trading days over 100 runs of the base model with no financial constraints. The average growth rate is approximately 1.81 percent on an annual basis (0.0072 per trading day with approximately 253 trading days per year). It is evident from table 1 that the model is not ergodic. Initial conditions make an important difference in the time path of the model, especially the initial capacity utilization. One of the major problems of the underlying real-side model is that its dynamics are extremely simple in that it reaches a steady state quickly and is absurdly robust to shocks (Gibson, 2010). In particular, as firms reach full capacity utilization their financing needs per unit of output gv/u fall. The financial surplus or deficit reaches a limit

$$S_{fi} = (s_i \pi_i / v_i - g_i) K_i \tag{7}$$

which evidently grows at the same rate as the capital stock. Prior to $u = 1$, the rate of growth of the financial surplus increases as is evident from equation 7. In figure 3 both firms start from capacity utilization of 0.8 and the simulations underlying table 1 begin from roughly the same level. Since firms depend upon all other firms for their aggregate demand, having them all start from this low level of utilization creates a slightly distorted image of how the economy actually operates. Since in reality there is no common start time, ergodicity would seem to be a desirable property of the model.

On the other hand, it is equally problematic to assume that all firms are fully adjusted and operating at $u = 1$. To address this dilemma a firm mortality rate is introduced. For every tick of the model, some firms

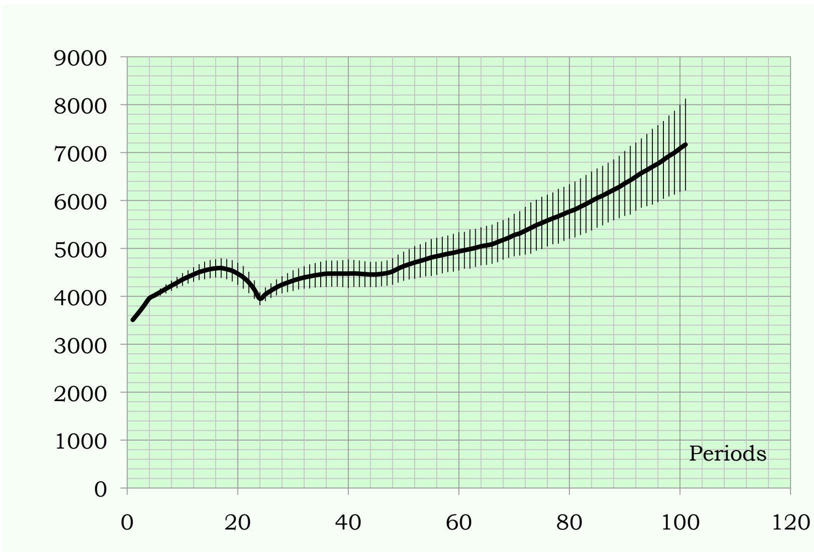

Figure 5: GDP with no financial constraint-2σ error bars

go out of business and reorganize with half the capital stock they previously owned. This might capture, for example, a product life cycle or other demand shock that reduces the value of the firm. These cheap productive assets are then acquired by another producer, who begins production of a new good with a finite life cycle and the accumulation process continues. To ensure that these firms are forced to start from a low level of capacity utilization, their share of aggregate demand is also reduced to 25 percent of its previous value. This is necessary since if their θ was not appropriately reduced along with their capital stock, their capacity utilization would rise in an unrealistically rapid fashion. This modification creates a mix of firms that are well established, with nearly full capacity utilization, and those that are restarting the accumulation process from a low and uncertain level of capacity utilization.[5]

5 The financial sector

The financial sector of the model closely follows an emerging literature on percolation models of financial crisis but crucially relaxes some of the assumptions (Sornette, 2003). The agent set is augmented to include "traders", responsible for bringing together the borrowers and lenders of Setterfield and Budd (2008). Traders accept deposits from firms. They do not simply accommodate all petitions for credit from firms, however. Their propensity to loan to other firms is based on the earnings performance of their own portfolio relative to that of the aggregate economy, a *private* signal. But they also have an incentive to imitate other traders, a *social* signal, inspired by Black's classic treatment of noise traders (Black, 1986). Traders exert control over the realization of investment plans and so, ultimately, on the aggregate performance of the real side of the economy. It is always possible for the financial sector to constrain investment at the microeconomic level even if, in the aggregate, the growth regime is Keynesian in the sense that there exists excess capacity.

While firms are subject to animal spirits, traders are more rational. Firms are driven by local technological developments, market conditions and "hunches" about the applications of their products and whether they will gain widespread acceptance in the market.[6] Traders take a more dispassionate view and manage the exuberance of firms by applying restraint on the flow of financing. They form beliefs on the basis of subjectively determined probability distributions about the state of nature (Chamley, 2004). They then receive signals about firm performance–measured by the relative performance of the firms to which they are lending. Each individual trader then uses this private signal to update beliefs about the likely future performance of firms in a Bayesian fashion (with an exogenously specified lag). At the same time, traders do trust other traders to some degree, believing them to have equally objective views of firms. Hence each trader learns from the behavior of other traders in the financial system, interpreting this behavior as providing information about the likely future performance of firms, their social signal.

For the population of traders as a whole, the martingale property is assumed to hold: agents may update their beliefs, but the expectation of the change in beliefs as a result of a new observation is zero. As noted above, no single event can trigger a financial crisis.[7] This means that the *ex ante* probability of a crash must be positively related to the rate of return that traders earn: traders must be compensated with a higher rate of return if they are to hold an asset that will lose value if the financial system crashes. It also means that, from the traders' perspective, a crash is an emergent property and must be taken as an exogenously determined event.

All models of this type deduce the presence of a critical parameter from the "coupling" or imitation of one trader by others. As cohorts of like-minded agents form, a crisis precipitates only when all elect to make the same trade at the same time. The dynamics of this coupled influence (the disappearance of social learning

[5] The reason that the initially low level of capacity utilization may remain low is related to the nature of the computation of the basic feasible solution to model discussed above. There is no reason that a young firm would necessarily be chosen to benefit from spillover aggregate demand caused by other firms having reached full capacity utilization. This may of course occur, but nothing the model makes it a certainty.

[6] The *failure* to perform this social function of allocating credit to only worthy borrowers is largely identified as one of the root causes of the housing market bubble and subsequent collapse.

[7] Essentially, then, crises are analogous to Southern Californian wildfires: they are known to be possible, but individual agents cannot pin down their causes in advance. Of course, if agents did know what events were responsible for crises, they would be able to profit from their knowledge, and would act accordingly.

according to Chamley (2004)) are crucial to the realism of the model. Here we follow Gonçalves (2003) in allowing the coupling parameter to differ from agent to agent and adopt Sutton and Bartow's account of reinforcement learning Sutton and Barto (1998) to describe their behavior. In short, financial agents couple more closely when the prediction of their colleagues is reinforced by the movement of share prices.

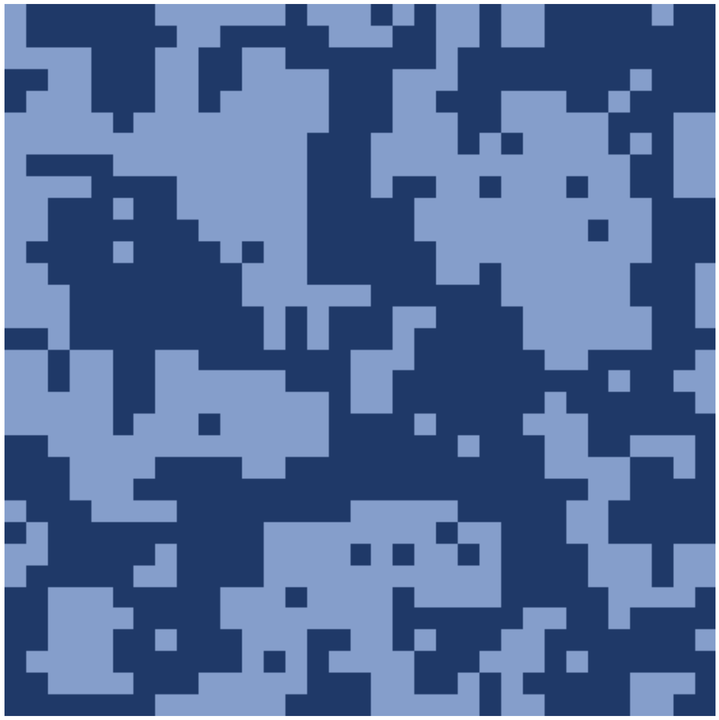

Figure 6: Percolation in a finite lattice

5.1 Mean field theory and percolation

Consider a finite lattice structure as shown in figure 6 with a total 1,089 cells. The light blue cells block flow so percolation requires a connected path of dark blue cells spanning the dimension of the grid (either vertically or horizontally depending on the problem). If the probability that the cell is occupied (dark blue) rises, percolation is more likely and vice versa. Let z_L be the probability that a path exists across the $L \times L$ grid when the proportion of dark blue cells is p. Note that an increase in L raises the probability of percolation so that when L increases without bound, the probability rises. On the other hand, when the proportion of dark blue cells is low, a smaller grid admits a greater spanning probability. When a spanning cluster exists, it is said that a *phase transition* has occurred. Knowledge of where the phase transition takes place measured on some metric relevant to the system is obviously of great value. In the case of H_2O, the important metrics are temperature and pressure.

The most interesting feature of the phase transition is that the state of a given cell as the p rises is *not* constant. It is not as if a cell joins one side (dark blue) or the other (light blue) and stays that way until a spanning cluster appears. On finite grids, the transition is not instantaneous and there is no sharp step in the phase transition. Boundary conditions matter and since there are obvious practical computational problems with infinite grids, periodic boundary conditions can be designed to approximate large systems.

Near the critical p, the boundaries between clusters of different colors are indistinct and some cells switch rapidly back and forth until the transition occurs. In the example of water at the critical point (zero celsius), solid and liquid states coexist and alternate even as the coexistence regions shrink. Other examples include

Table 2: Cluster size in figure 6

Cluster size	Number	Scaled number
1	13	52
2	3	12
3	1	4
:	:	:
8	1	4
:	:	:
26	1	4

Source: Author's calculations

the phase transition, Curie magnetization at very lower temperatures, superconductivity, virus propagation and the like. In particular, structural failure seems to follow the percolation pattern (Ormerod, 2007). Just as a pressurized vessel explodes it can be seen to pass through a brief period of time in which it expands *and* contracts in a periodic fashion. As the critical point nears, the frequency (1 / period) of the cycles increases until structural failure results.[8]

Now consider the probability density function that describes the number of clusters of size s at any given dark-blue proportion p. Is it possible that this PDF is Gaussian? In general the answer is no since just looking at the cluster distribution in figure 6 it is evident that there are many clusters of size 1 when, if the distribution were Gaussian, the number of clusters of any extreme size would be small. The distribution of cluster sizes is approximately given in table 2.

Changing the scale on which the number of clusters is measured here is arbitrary. It should not matter if the number of clusters is doubled, renormalizing the size so that each one is divided into four. The second column shows the rescaled number. The distribution for this exceedingly simple example is shown in figure 7. Assume for the moment that there is a function $g(s)$ that describes the distribution. The question is what does the function look like?

$$g(xs) = k(x)g(s)$$

where $k(x)$ is the scaling parameter, with $x = 4$ in the example above. What matters here is that the scalar parameter does *not* depend on the distribution of s; $k(x)$ shifts the entire distribution up or down. Note that when the average cluster size is 1,

$$g(x) = k(x)g(1)$$

and $g(1)$ is just a number, it follows that the functional form of $k(x)$ must be the same as $g(x)$. Thus

$$g(xs) = g(x)g(s).$$

There are few functions for which this property holds. One of them is the *power law*

$$g(s) = s^{-\tau}.$$

with τ as the *critical exponent*.[9].The distribution of cluster sizes follows a power law at the critical point and this signals a phase transition. power-law distributions are know as "scale free" in that they do not

[8]Note that any simulation that uses a finite sized lattice cannot exhibit a true phase transition because eventually the clusters span the finite lattice.

[9]Note that $g(x)g(s) = x^{-\tau}.s^{-\tau} = (xs)^{-\tau} = g(xs)$

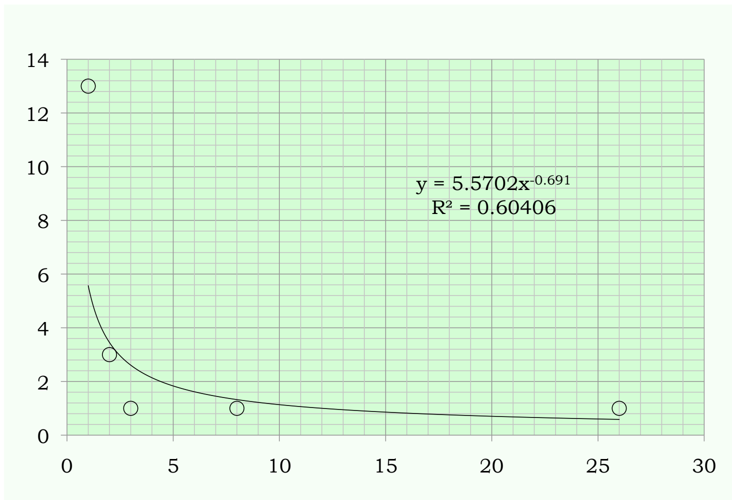

Figure 7: Mean cluster size-original

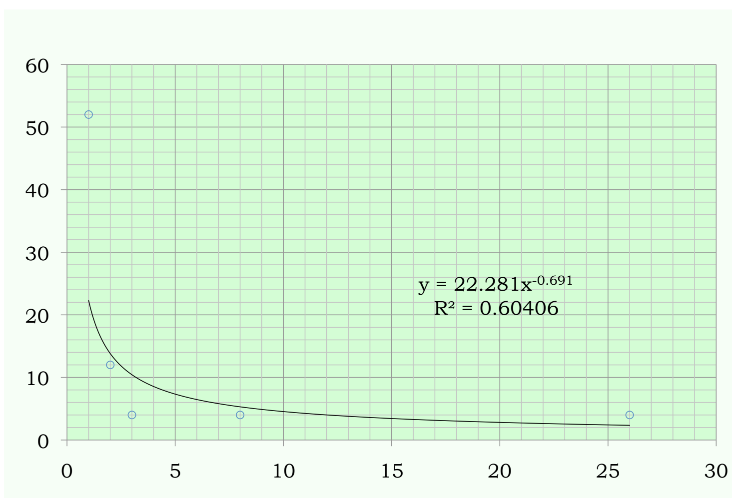

Figure 8: Mean cluster size-renormalized

depend upon a "typical" or most likely value, such as the mean of the normal distribution.[10] Imagine, as the story goes, that there are 4 patrons of bar, all with roughly the same net worth. The distribution can be characterized as Gaussian and the mean is meaningful. But now allow some very wealth individual to join the four, think Gates, Soros, etc. The mean of the resulting distribution loses its importance, wealth per capita rises to some number but it conveys no information about the distribution of net worth among the patrons. Instead the power-law distribution is described by way of "excedences". If in a sample there are only 96 clusters with a correlation length exceeding 50 cells and the critical exponent for the system had been estimated at 1.5, the power-law distribution would imply that the number of cells with correlation length greater than 100 would be $(200/100)^(- 1.5) = 36$. How those 36 are distributed would *again* follow the power law.[11] The simple structure of a power law gives rise to a straight line on a log-log plot since $logg = \tau logs$, but not every log-linear plot is a power law and maximum likelihood tests are available to check for the existence of a power law (Clauset et al., 2009).

The point of this digression into percolation theory is that the properties of the geometrical phase transition in percolation have been widely recognized as qualitatively similar to financial collapses. The long-range correlations observed in percolations correspond to markets that are characterized by a high degree of order rather than disorder. In other words, traders all abandon their private signal, relying instead on the social one.

5.2 Crashes as critical points

Since Sornette et al. (1996) first proposed the connection between crashes and critical points there has been an enormous amount of work on the topic. The general approach is summarized in this section, but the reader is referred to Sornette (2003) for a comprehensive treatment of the topic. Figure 9 shows the fit of the model for the S&P 500 from the middle of 1985 (measured in trading-day fractions of a trading year, 253 calendar days). Note the characteristic oscillations in the smooth trend line leading up to the crash. The frequency of the oscillations increases as the date of the crash approaches.

Were the underlying processes in the crash of 1987 Gaussian, we could expect a decline of more than 22 standard deviations only once in several billion times of the age of the universe. Since observed crashes seemed to be more frequent than that, it follows that some other distribution might better characterize financial markets. The first approximation might be a power law and indeed Sornette et al. (1996) fit a (linearized) power law to this data to find

$$F(t) = 327 - 79(t_c - t)^{0.7} \tag{8}$$

as is shown in figure 10.[12]

The log-periodic function of figure 9 is given by

$$F(t) = 327 - 165(87.74 - t)^{0.33} \left\{ 1 + 12[\cos(7.4 \log \frac{87.74 - t}{2})] \right\} \tag{9}$$

and is a much better fit. Replication yields a value of 44.3 for sum of squared errors. This equation is a power law modified by an oscillatory term that increases in frequency as t approaches the critical value.

To get this result requires a model of the rational trader who operates according to the principle enunciated above. Coupling alone will not generate the predictive signature of log-periodic, or fractal variation, in the behavior of prices before the critical point. To get this, most of the literature follows Derrida (1981) who solves the mean-field model for a particular grid geometry, the diamond lattice. In the random network of

[10] In the Gaussian or normal distribution, large deviations are rare and exceptionally so with 99.7 percent of observations lying in the range of three standard deviations from the mean.

[11] The spanning cluster is a fractal and approaches zero density as the size of the system becomes larger (Gould et al., 2007, p. 472).

[12] The sum of squared residuals for the power law is reported by Sornette (2003) as 107 but replicated following the instructions in the text is actually 123. Thus an exponential of the form $F(t) = 3.39E{-}8e^{0.2621t}$ seems to do a better job at 107 (oddly, the exact same value Sornett obtained for the power law!).

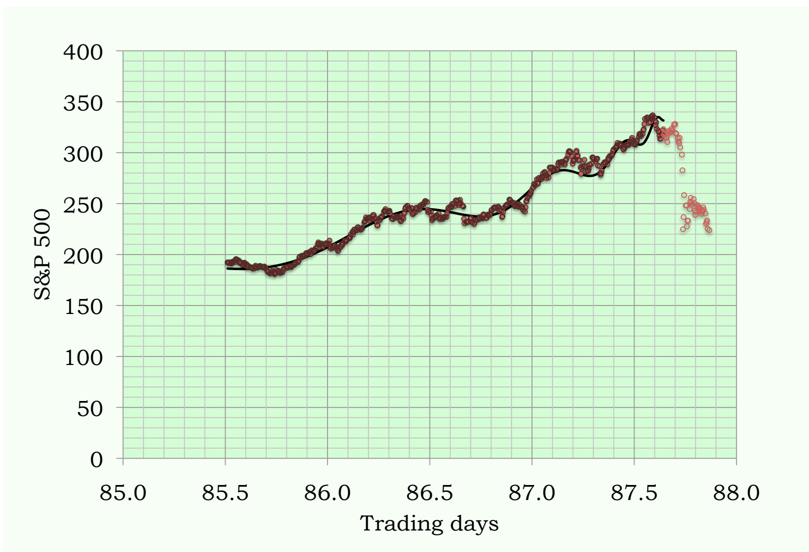

Figure 9: Log-periodic fit for the crash of 1987.

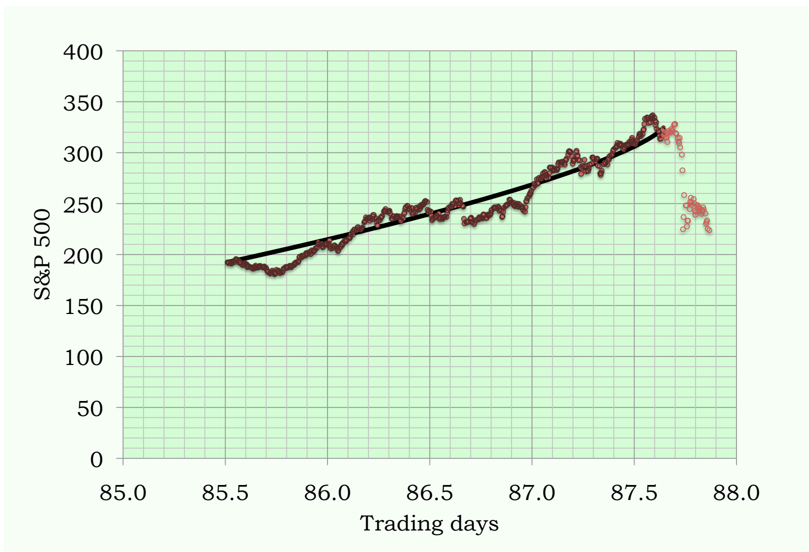

Figure 10: Power law fit for the crash of 1987.

Erdos-Renyi, the distance between any pair of nodes along connecting edges is proportional to the logarithm of the number of vertices.[13] A cluster in network theory is defined as a group of nodes in which every member knows every other member. If that were the case, then the total number of edges would be $k(k-1)/2$ for a node with k edges attached. For less connected nodes a clustering coefficient

$$C = \frac{2E}{k(k-1)}$$

can be defined where E is the *actual* number of edges. It can be shown that the clustering coefficient in a random graph is p. In real-world networks the clustering coefficient is much higher than p (Albert and Barabási, 2002) and the implication is that the *degree distribution*, that is the distribution of nodes with E edges is not Gaussian, or even Poisson as classical random network theory would suggest. Albert and Barabási (2002) have shown that the degree distribution is a power law.

In the literature on financial crises there are two kinds of lattices commonly employed. The first corresponds to the Erdos-Renyi's random graph in which connections between traders are formed randomly. Start with 2 traders and then let each trade with a new partner as the grid grows. After the first iteration, the number of traders is 4 and after the second, the number is 8 and 16 and so on so that after p iterations the number of traders is 2^p. But note that after the 3rd iteration, the first two traders are now connected *to all four neighbors*. As the graph grows, more traders are connected, but no individual trader has more connections than 4 (a von Neumann neighborhood) and the degree distribution is Poisson.

This is the graph that underlies the Ising model of magnetism as shown in figure 6. It is non-hierarchical and produces a power-law distribution of cluster sizes as shown in figure 11. What it does *not* do is produce a log-periodic decoration of the power law (in equation 9) that seems to be characteristic of actual crashes.

To obtain this, some alteration of the trader geometry must be employed. It is has been shown that a "diamond lattice" will produce the desired effect. A diamond lattice amounts to a minimal amount of doping of the regular landscape of pure percolation models and is formed by connecting traders in a slightly more complex fashion. Imagine two traders connected by an edge as shown in the first quadrant of figure 12. In the next iteration add two more nodes to form a diamond as shown in the second quadrant. Note that the degree distribution of the nodes is still uniform, with all nodes connected by two edges. Repeat this procedure in the third quadrant and now the degree distribution begins to change. Figure 12 shows the degree distribution for this structure. Note from figure 13 that after the last iteration, the distribution is beginning to take the shape of a power law.

Johansen et al. (2000) observe that after p iterations, there are $2/3(1+4^p)$ traders and 4^p edges joining them. Most traders in this lattice are of degree 2, as shown above. The original traders have degree 2^P but the degree drops quickly.

Although this lattice is clearly more realistic than the randomly seeded grid in figure 6, the degree distribution is *not* the result of preferential attachment as in Albert and Barabási (2002). It characterizes a crystalline structure doped by impurities in a systematic way, rather than the result of a scale-free distribution that arises out of social structures. Hierarchical networks such as the diamond lattice have "discrete scale invariance," in effect, a fractal architecture (Sornette, 2003, 186-7). Preferential attachment is the process by which something like the diamond lattice *might actually come about*. So therefore, if log-periodic behavior cannot be generated by a more realistic model of the social space than the diamond lattice, then perhaps the financial crisis literature is seeing something that is not really there, at least as far the actual economy is concerned.

Remarkably it is possible to solve analytically the equations describing the diamond lattice, or least approximately so. This was done by Derrida (1981), as noted by Sornette (2003). As a percolation bed, the diamond lattice gives rise to the same behavior with a critical point at which the connection length becomes unbounded. The difference is that the critical exponent is a complex number and this is responsible for the periodic term in equation 9.[14]

[13]This model starts with N nodes connected with probability p. The resulting graph has approximately $pN(N-1)/2$ edges that are distributed randomly.

[14]A lucid and thorough explanation of what is involved here is given in Sornette (2003). The details are beyond the scope of the present paper.

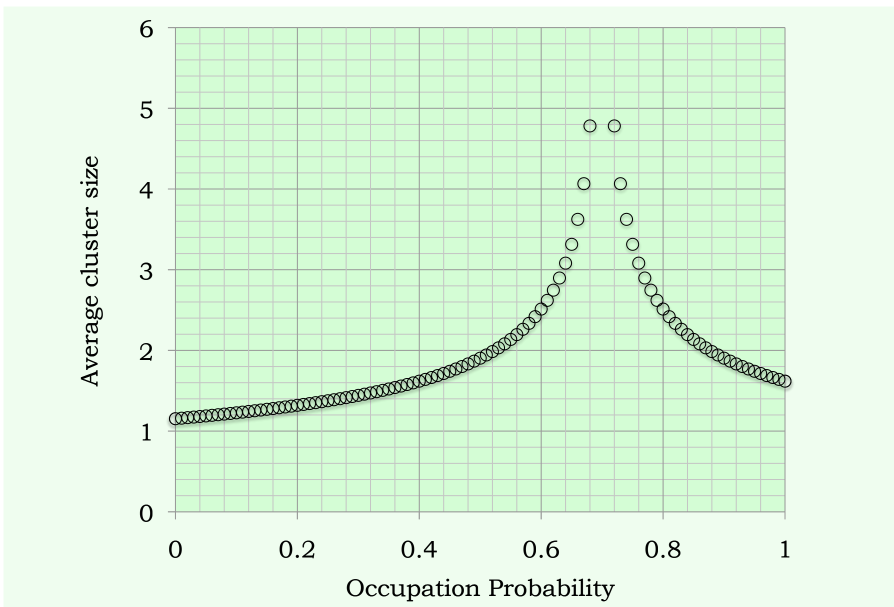

Figure 11: power-law distribution

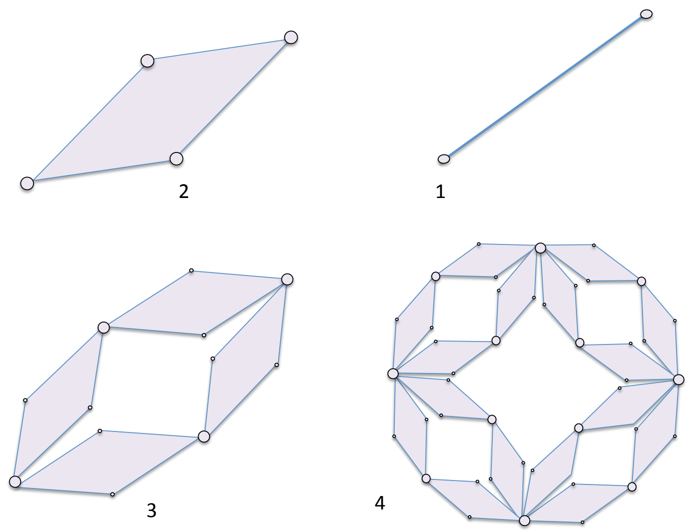

Figure 12: Diamond lattice

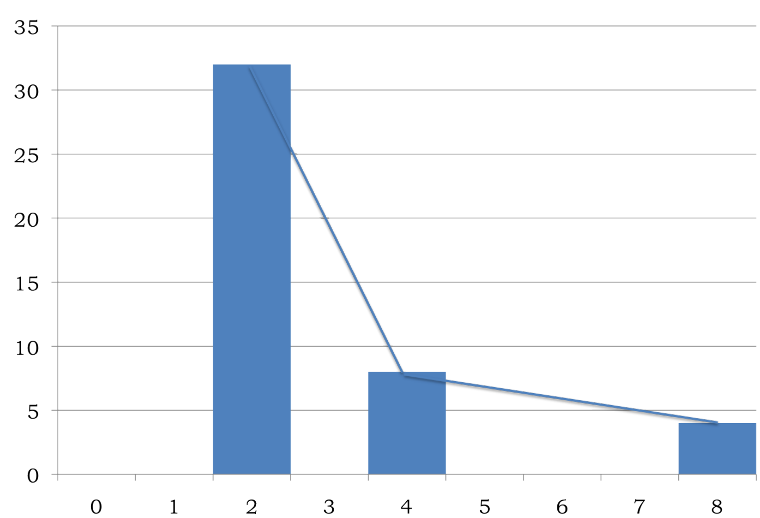

Figure 13: Degree distribution of an emerging diamond lattice

The procedure for fitting equation, 9 written generally as

$$F(t) = A + B(t_c - t)^{\tau} \left(1 + C\cos\{\omega \log[(t_c - t)/T]\}\right) \tag{10}$$

is to guess values for t_c, τ, ω and T. With these in hand, the equation is then linear in the coefficients $A, BandC$ and so a regression can be run. The sum of squared residuals is then computed as a function of the guessed parameters. A search can then be performed over the space of initial guesses to find the minimum of the sum of squared residuals. This is how the parameters of equation 9 were estimated. The procedure is fraught with problems, however, mainly because there is no way to ensure that a given solution is a global rather than local minimum. Also small changes in even large data sets, such as was used to estimate the equation, can have a significant impact on the estimated parameters. While the power law signature of a critical point is quite well established, the idea that log periodic oscillations will precede a crash is a bit more tenuous.

6 Simulation results

The purpose of the simulations is to investigate the possibility of a log-periodic signal that a critical point is approaching in a multi-agent system. First, the diamond lattice structure is replaced by something somewhat more realistic, a structural model of the real side of the economy.[15] Even if a diamond lattice can in theory produce log-periodic oscillations, it is of greater interest to know if real economies actually self-organize in a way that is consistent with the predictions of log periodicity. In particular, it is insufficient to say that if a functional form such as 10 can be fit to the data then the underlying theory is validated. Since in general a Fourier transformation can replicate *any* wave form, it is not at all surprising that the literature is able to estimate the parameters of equation 10 (or indeed even a second order approximation, as is done in Johansen et al. (2000)).

Even if every crisis were preceded by log-periodicity, it may be that there are many episodes of log periodic fluctuations that do not lead to a crash. This would imply that insofar as the model is used to predict actual crises, it is an instrument that is rather blunt. While the literature is clear on the possibility of a safe landing after a decorated power law bubble, it should lead to a crash in the *majority* of cases.

The agent-based model as described above does consistently produce a power-law degree distribution of traders. Simulations show that the interaction of the real and financial sectors is sufficient to generate preferential attachment and a consequent hierarchy of trader liquidity. The result arises as an emergent property of the design of the interplay of the real and financial sectors.

6.1 The emergence of preferential attachment

As noted above, there is nothing in the basic structuralist model of growth that ensures any firm will self-finance at $s = s_i^*$, the rate that leads to full capacity utilization. If this were not so, then investment would be fractal in the sense that each firm's accumulation would be a scaled down but self-similar representation of the investment process for the economy as a whole. Since a wide variety of savings rates might apply to any given firm, macroeconomies would follow a more chaotic path than they seem to do.

Traders regulate the flow of finance to firms with negative financial surpluses according to the following rule: do not allow any one firm more capital than would be required to enable the firm to converge to full capacity utilization, assuming that the shares of the capital stock of all firms σ_i remain approximately constant. Note that traders do not have full information about the future path of the economy. They cannot, therefore, predict the steady-state value of σ_i. But they can, given the current distribution of capital stocks for all firms, calculate the s_i^* consistent with any given firm moving toward full capacity utilization and the maximum rate of profit.

[15] The use of the diamond lattice can be seen as a particularly vivid instance of Friedman's "if it works use it" instrumentalist methodology!

At the beginning of simulated time no trader has any assets. At the first sweep, firms with positive financial surpluses make deposits with the nearest trader. Liquidity ℓ of the jth trader is defined as the cumulative sum of deposits less loans

$$\ell_j = \int_{i \subset J} S_{fi} \tag{11}$$

where J is the set of firms who are the clients of trader j. Firms with negative financial surpluses must then locate a trader with adequate liquidity to finance the next phase of capital accumulation for the deficit firm. If a firm with $S_f < 0$ cannot locate a trader with sufficient liquidity, then the firm hoards its savings and waits for the next sweep. In the current sweep no investment is made by the frustrated firm and thus aggregate demand falls and with it average capacity utilization for all firms.

Investment is *quantized* in the model, since in principle, a deficit firm could invest at something less than the rate required by the expansion plan. This is not, however, permitted in the model and is justified by appeal to the lumpiness that most real investment plans exhibit and the problems that accrue to undercapitalized firms. This is admittedly a somewhat arbitrary assumption that could easily be relaxed in future research.

6.2 The financing decision

Traders face a known demand for liquidity at each sweep of the model based on the number of firms and $\sum_i S_{fi}$ But the question still remains as to whether finance will actually be forthcoming. The financing decisions depend on a "forecast", essentially whether trader sentiment is either bullish or bearish about the future. Following Gonçalves (2003), the forecast of each trader depends on a combination of two signals. The first is an idiosyncratic private signal, defined as the difference between the profit rate of the clients of that trader and the aggregate rate of return for the economy as a whole, $\bar{r}$.

$$r_{di} = r_i - \bar{r} \tag{12}$$

Note that the inclusion of $\bar{r}$ implies the presence of a *blackboard* element in the private signal through which macroeconomic events influence private behavior(Carver and Lesser, 1992). Following (Wooldridge, 2002) a multi-agent system is essentially a collection of pattern or action rules together with a *working memory* of facts and a sweep of the model is a forward chaining of those rules. The key problem with MAS prior to the work of Lesser and others was that learning was unstructured in that the grid was a random collection of amorphous agents. Some independent agent or agents must act as a "knowledge source" and have the ability to communicate that knowledge in encoded rules. Knowledge sources write data on a shared data structure called a blackboard and agents in the model are able to read everything on it, but react in an idiosyncratic way. In a more sophisticated model knowledge sources would include rating agencies such as Standard and Poor and Moody's, but the common rate of profit (supplied by the observer) is an approximation.

The second component of each trader's forecast is the subjectively perceived social signal, discussed in detail above, that arises from the traders' perceptions of the forecasts of other traders to whom the individual trader is coupled. This social signal is included to capture Keynes' well known "beauty contest" or noise trading. In an efficient market, asset prices should convey all available information, but with noise trading, the price can rise to reflect what traders believe other traders believe. Following Schleifer (2000) it is not irrational for informed traders to follow noise so long as they can exit the market before a crash occurs.

The model is then a combination of the standard "mean field" theory of collective systems (Goldenfeld, 1992) and Reddy and Lesser's parallel blackboard system in which firms contribute to the blackboard data base through the formation of the average rate of profit. The simplest way to describe an imitation process is to assume that the hazard rate, the probability of a crash conditional on a crash not having occurred, follows the log of the price of a uniform asset traded by financial agents. A higher price then confers information: the hazard rate has increased and so has the probability that a crash will occur in the next period. For informed traders, this implies that the price must accelerate to compensate for the risk of holding the asset for an additional period.

Traders are indexed $i = 1, 2, ..., I$ and $N(i)$ denotes the neighborhood of agents connected to agent i. When agents are coupled, they exchange information. There are only two possible forecasts, f, for the ith trader,

$$\begin{aligned} f &= -1 \text{ bear or pessimistic} \\ f &= 1 \text{ bull or optimistic} \end{aligned}$$

Clients and their profit rates are initially randomly distributed: more profitable firms do not preferentially attach to larger traders. However, the quantized nature of investment effectively requires that firms search out liquid traders. In the financial crash literature, this idiosyncratic signal is described as " intuition" but here there is some real information conveyed by the firms to their traders.

Note that an informed trader maximizes her return by having taken the "right" signal from her colleagues. This is $\sum_{i \subset N_{ij}} f_i$ the sum of the j neighbors to whom i is connected. The signal is the "right" signal if in fact her neighbors are representative of the trader agent set as a whole. In that case, the price will have moved in the same direction as the neighbors' forecast. Were the trader to act on fundamentals alone, only the idiosyncratic signal, which combines local and blackboard information would constitute the signal.

The traders' forecast combines the idiosyncratic signal, the profit rate of the clients of that trader relative to the profit rate of the economy as whole, and the subjectively perceived signal that arises from perceptions of other traders behavior. The forecast is

$$f_j = K(w) \sum_{i \subset N_{ij}} f_i + \sum_{i \subset C_{ij}} (\phi_i - \bar{r}) + \epsilon$$

where N_{ij} is the von Neumann (8) neighbors of agent j and C_{ij} is the set of clients of agent j. Here ϕ_i is the weighted average of individual client profit rates in proportion to the total capital stock of all clients and $\bar{r}$ is the average rate of profit in the economy as a whole. A random error term allows traders to override their private signal and/or social signals and set a forecast based on intuition.[16] Note that the implicit weight of the private signal is one so that $K(w)$, the coupling strength, is normalized by the value of the strength of the private signal. The K is a measure of market "depth", the degree to which traders are connected to each other.

The degree of influence exerted by other traders is controlled by the "coupling" parameter, $K(w)$. The weight, w, implies that all traders are not the same size, and more weight is attached by any given trader to the forecast of large traders in the derivation of the social signal. If profitability of the economy as a whole is rising (capacity utilization rising toward one) then more traders will be receiving a positive rather than negative idiosyncratic signal. This will raise the probability that the price increases in the next round.

The model thus conforms to the classical Ising or "spin" model, as discussed above. It is not yet a "spin-glass" however, since there is nothing that prevents the Ising model from reaching a minimum energy configuration, although it can take a significant amount of time. The spin would then correspond to a pure Walrasian system that has a unique equilibrium with its well established welfare theorems. Spin glass is second or third best, one of a number of multiple equilibria to which the model can converge, but without any clear welfare implications.

What imparts the "glass" to the "spin" model is precisely the coupling parameter K, which is obviously missing in the Walrasian system. The coupling parameter is taken as a constant in Johansen et al. (2000), but it can vary subject to confirmation bias: if the share price, p, moves in the direction consistent with the direction of influence, $K(t)$ rises and vice-versa. Price predictions are thus self-reinforcing because an uptick in the price will increase the coupling parameter. What this means intuitively is that each trader attaches more weight to what she perceives as superior information possessed by other traders.

The share price is defined as an index of traders' forecasts, $f(t)$. If forecasts are all bullish, the share

[16] Without this random error term, the modeled economy would never recover from a crash: cascades go on forever.

price rises and vice-versa.

$$\begin{aligned} \sum_{i=1} f_{i,t-1} &> 0 \text{ price rises} \\ \sum_{i=1} f_{i,t-1} &< 0 \text{ price falls} \end{aligned}$$

The equation for the share price is

$$p(t+1) = p(t) + \frac{1}{n} \sum_{n} f(t)$$

If there is a balance between bulls and bears, the share price remains constant. If bulls outnumber bears $p(t+1)$ rises and vice versa. The share price is a random walk during "normal" times and breaks out during organized bull or bear markets to produce a bubble and then either a crash or soft landing.[17]

The path of the coupling parameter is then

$$K(t+1) = K_0 + \frac{1}{n} \sum_{n} f(t)$$

Finally, if traders were bullish last period, the probability is slightly greater that they will be bullish this period, everything else equal.

It goes without saying that if no firm were ever denied funding for their planned investment, there would be no financial constraint: the economy would evolve in a fashion similar to the frictionless neoclassical and structuralist models described earlier, where finance is perfectly accommodative. Instead, the number of firms denied funding reduces the level of investment and thus drives down aggregate demand as noted. Let the number of firms be N_d. It is easy to show that N_d is strongly (negatively) correlated with GDP growth,defined as the $\sum_i v_i K_i$. Figure 14 averages 100 runs of the model over 500 trading days, with and without the financial constraint that $N_D = 0$.[18] Alternative versions of the model show less influence of traders and the reasons are still under investigation.

One final feature of the financial sector is the possibility that traders will "originate and distribute". If for example $f_i = -1$ for trader i, her forecast is bearish. If petitioned by a firm for finance, she would ordinarily deny. If, however, one of her neighbors j has both sufficient liquidity *and* a bullish forecast $f_j = 1$, the firm's investment demand will be satisfied by trader i borrowing from trader j. A measure of system "leverage" is then increased by the value of the loan and this could be taken as an indicator of overall financial fragility. In this situation, traders are unwilling to risk their own capital, but have no trouble "spending somebody else's money".[19] The model is therefore consistent with the idea that financial firms abandon their positive role as overseers of financial allocation in exchange for a cut from a "greater fool".

Table 3 summarizes the model. It shows both breeds, firms and traders, and provides a comprehensive list of the variables in each of their data structures.

6.3 Model structure

Each sweep of the model represents one *week*, despite the fact that the literature on stand-alone models takes a trading day as a unit of time as done above in 9. The integration of the real and financial sectors requires the meshing of two separate time-scales separated by an order of magnitude. Any given firm on the

[17] Of course as the uber-critic Taleb (2007) notes, knowing whether we are in normal times or not *should* be the whole point of the analysis.

[18] In the latter, firms are able to invest according to their idiosyncratic savings rates. In effect, there is a perfectly functioning financial system that accommodates the investment plans of agents. Chick (1983) notes that some "bridge financing" must be available until the investment creates its own savings.

[19] At the moment there are no specific implications of the rise in leverage, but it could shorten the lifespan of firms in the next generation. This is a topic for future elaboration.

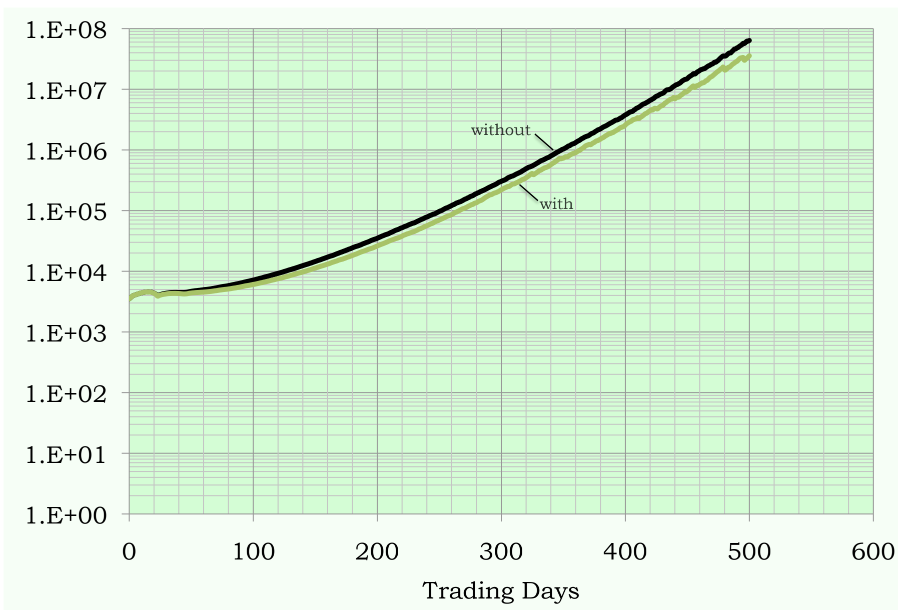

Figure 14: Base run with and without financial constraint

Table 3: Agent data structures

Firms		Traders		
v	capital-output ratio	f	forecast	*(+1)* bullish *(-1)* bearish
s	savings rate out of profits	ℓ	liquidity	sum of deposits minus loans
u	capacity utilization	ϵ	forecast volatility	variance of forecast
r	profit rate	κ	coupling	sensitivity to public signal
g_u	change in g wrt u	κ_0	base coupling	initial κ
g_r	change in g wrt r	ν	private signal	forecast sensitivity to r_d
a	animal spirits	r_d	profit differential	
g	growth of capital stock	η	clients	
K	capital stock	σ_ℓ	liquidity share	
s_1	savings rate need for $u = 1$			
$\mathcal{F}_l$	loanable funds $(+/-)$			
l	labor coefficient			
θ	share of demand			
h	accumulated uninvested savings			
T	life span			

Note: Agent specific subscripts suppressed.

real side of the model will invest only sporadically and is certainly not in the market on the time scale that traders observe. Weeks as the unit of time is thus a compromise that requires that firms visit the market in *cohorts*.

The agent-based model of this paper has 6k firms on the real and 49 traders on the financial side. There are 50 trading weeks per year and firms are assumed to invest once per year in order to maintain some connection to the common assumptions of the theoretical literature.[20] At each sweep 120 firms are in the market. Those with positive financial surpluses make deposits with traders. Firms with negative financial surpluses attempt to borrow from traders and are either granted credit or they are denied. Those with sufficient financing then invest and those without hoard, waiting for the next opportunity to invest. Currently denied firms must wait 50 weeks.[21]

After all successful firms have made their investments, a new level of capacity utilization in the economy is calculated for *all* by way of equation 4 above. On the basis of the new u, updated values for the profit rate and planned investment are calculated. Small changes result in all the endogenous variables from week to week.

At the same time traders parse their private and public signals on the basis of which they form their weekly forecasts. If a given firm approaches a given trader in some week and the trader's forecast is negative, the loan application is then denied. The firm may still find financing from a trader's neighbor, but if not the firm is frustrated and hoards its savings. At the present these savings are *not* part of loanable funds.[22] It is important to see that the asset price in the model is only weakly linked to real side fundamentals. If the profit rate on the real side falls, traders may still produce a bearish forecast and be willing to provide

[20]This is not entirely realistic and could be modified in later versions of the model.

[21]While this may seem unrealistic, the lag is taken as representative of all kinds of frictions in the financial system that are not explicitly modeled.

[22]Again this may not seem realistic but is taken to represent the less than full access to all the liquidity in the market (model liquidity *plus* hoarding) firms enjoy.

financing.

6.4 Findings

The first observation is that the presence of the real side of the model does indeed seem to give rise to a power-law distribution of liquidity without the need for a diamond lattice hierarchical degree distribution. Figure 15 shows a typical distribution of liquidity for a run of the model. The liquidity distribution does start off in a fairly uniform way, but as preferential attachment begins to build so does the structure of liquidity begin to resemble a scale-free distribution. Thus a few traders control most of the liquidity and the model seems to be consistent with the view that some traders become too big to fail.

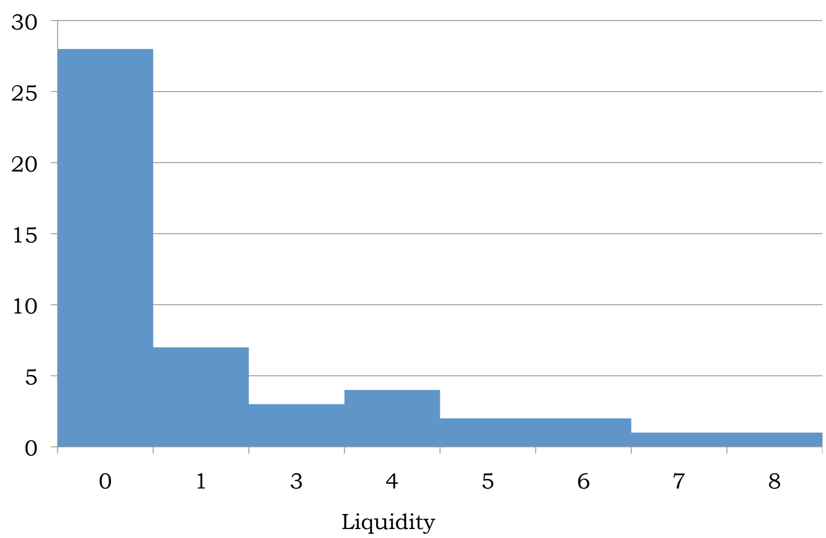

Figure 15: Liquidity distribution

Figure 16 shows a typical "business as usual" trajectory for the logarithm of the asset price simulated over 500 trading weeks, or 10 years. Note the absence of any sustained bubble. Efficient market theory suggests that the volatility of the price should be approximately constant. From figure 16, among many others, it is evident that this not exactly the case. Note periods in which the market is organized and the price rises or falls with no volatility, between weeks 120 and 230 in the figure.

Although the results are preliminary, figure 17 seems to show a bubble building from week 120. Note that in about a year and half, the market roughly doubles in value. This rapid run up is approximated by a weak power law as shown in the diagram. The only missing piece is the log periodic modification of the power law. There are small fluctuations toward the peak but as Lo and Mueller (2010) noted, basing real trading decisions on this forecast might be hazardous to one's wealth. The crash is an amalgam of a true crash, as shown in 9 above, and a soft landing. The market loses as much as 5.8 percent in a single trading week, as shown in table 4 but takes two months to lose approximately 18 percent of its value.

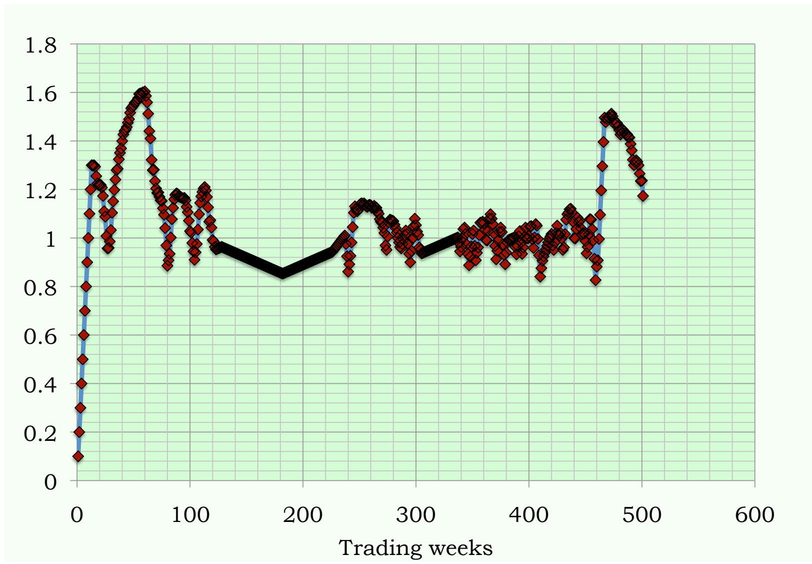

Figure 16: Ln(price)

Table 4: Price decline in simulated crash

1	5.57
2	1.12
3	1.46
4	2.16
5	4.31
6	1.58
7	5.81

Source: Author's calculations

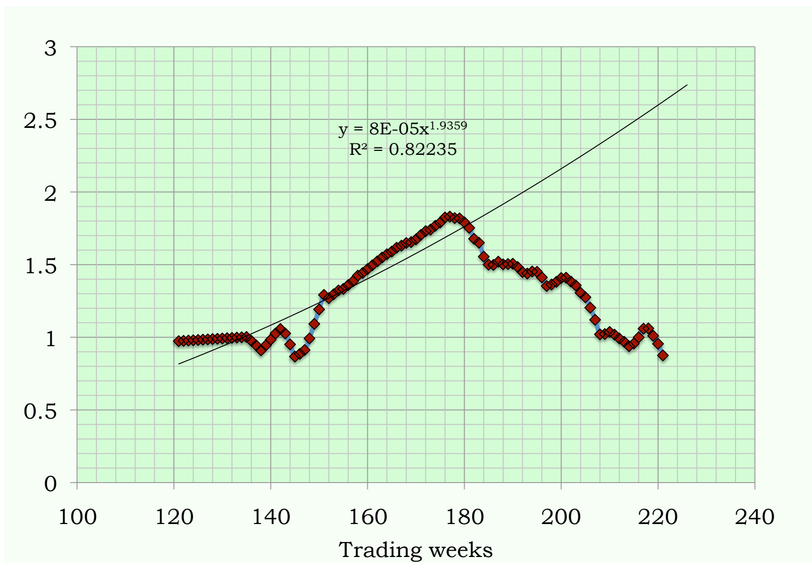

Figure 17: A bubble followed by a crash?

7 Conclusions and conjectures

The first and main conclusion is that the agent-based structure does not impose crisis on the data: as in reality (and in Sornett's GARCH generated data) crises are extremely rare, not to say endangered, species. Bubbles *are* evident in the data generated by the agent-based model, but they are not always followed by a crash and *often* by soft landings. The bubbles seem to show the power law acceleration predicted by the percolation model and to some extent these accelerations are decorated by log-periodic behavior.

The results of this paper show that bubbles-cum-soft landing are actually much more frequent than bubbles-cum-crash. The Sornette (2003) model implies that rational agents cannot predict crashes. They must be random and the critical point is only the peak probability of a crash. A power law acceleration should, argues Sornette, be accompanied by a crash some 70 percent of the time. But this conclusion is not confirmed (so far) by the simulations of this paper. Instead of a 70-30 crash/soft landing ratio, it seems that the proportions are reversed. This is the most damaging criticism of the financial crash literature proposed by the data generating process studied in this paper.

Why might this be the case? The answer is in the title of the paper: the real side does in fact matter. When a power law acceleration in the price of shares occurs, and the economy is below full capacity utilization, the bubble generates *real* output and improves the profitability of the economy. The closer firms are to full capacity utilization the less need for external finance. The financial sector then can frequently have its own crisis *without* seriously damaging the real side of the model. Since profits can remain strong, some traders observe this and abandon their bearish forecast and the landing softens. The 1987 crash, and others, have demonstrated the distance between the real economy and financial asset panics.

References

Acemoglu, D. (2009). The crisis of 2008: Structural lessons for and from economics.

Albert, R. and A.-L. Barabási (2002). Statistical mechanics of complex networks. *Reviews of Modern Physics 74*, 47–97.

Ariely, D. (2008). *Predictably Irrational.* HarperCollins.

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

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *The Journal of Political Economy 100*(5), 992–1026.

Black, F. (1986). Noise. *The Journal of Finance 41*(3), 529–543.

Carver, N. and V. Lesser (1992). Blackboard Architectures for Knowledge-Based Signal Understanding. *Symbolic and Knowledge-Based Signal Processing*, 205–250.

Cassidy, J. (2009). *How Markets Fail: The Logic of Economic Calamities.* New York, NY: Farrar and Straus and Giroux.

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

Chick, V. (1983). Macroeconomics after keynes: a reconsideration of the general theory. Open Access publications from University College London http://eprints.ucl.ac.uk/, University College London.

Choi, J.-K. and S. Bowles. (2007). The co-evolution of parochial altruism and war. *Science 318*(5850), 636 – 640.

Clauset, A., C. Shalizi, and M. Newman (2009). Power-law distributions in empirical data. *SIAM Review 4*(661-703).

Cynamon, B. and S. Fazzari (2008). Household debt in the consumer age: Source of growth - risk of collapse. *Capitalism and Society 3*(2).

Derrida, B. (1981). Random-energy model: An exactly solvable model of disordered systems. *Physical Review B 24*(5), 2613.

Epstein, J. (2006). *Generative Social Science: Studies in Agent-Based Computational Model.* Princeton, NJ: Princeton University Press.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance 25*(2), 383–417.

Gibson, B. (2007). A multi-agent systems approach to microeconomic foundations of macro. University of Massachusetts, Department of Economics Working Paper Series.

Gibson, B. (2010). The structuralist growth model. In M. Setterfield (Ed.), *Handbook of Alternative Theories of Economic Growth*, Chapter 1. Cheltenham, UK: Edward Elgar.

Goldenfeld, N. (1992). *Lectures on Phase Transitions and the Renormalization Group.* Addison Wesley.

Gonçalves, C. P. (2003, NetLogo User Community Models). Artificial financial market.

Gould, H., J. Tobochnik, and W. Christian (2007). *An Introduction to Computer Simulation Methods: Applications to Physical Systems.* Pearson Addison-Wesley.

Hein, E. and T. van Treeck (2010). "financialisation" in post-keynesian models of distribution and growth: a systematic review. In M. Setterfield (Ed.), *Handbook of Alternative Theories of Economic Growth.* Edward Elgar.

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

Kahneman, D. and A. Tversky (2000). *Choices, Values and Frames.* Cambridge Univerisity Press.

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

Kirman, A. (1989). The intrinsic limits of modern economic theory: The emperor has no clothes. *Economic Journal 99 (Supplement)*, 126–39.

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

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

Lawson, T. *Economics and Reality (Economics As Social Theory).* Routledge.

Lo, A. and M. T. Mueller (2010). Warning: Physics envy may be hazardous to your wealth. arXiv:1003.2688v3 [q-fin.RM].

Olson, M. *The Logic of Collective Action: Public Goods and the Theory of Groups.* Harvard University Press.

Ormerod, P. (2007). *Why Most Things Fail: Evolution, Extinction and Economics.* Wiley.

Schleifer, A. (2000). *Inefficient Markets: An Introduction to Behavioral Finance.* Oxford University Press.

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

Setterfield, M. and A. Budd (2008). A Keynes-Kalecki model of cyclical growth with agent-based features. Department of Economics, Trinity College.

Shin, H. S. (2005). Commentary: Has financial development made the world riskier? In *Proceedings of the Federal Reserve Bank of Kansas City Symposium at Jackson Hole.*

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

Sornette, D., A. Johansen, and J. Bouchaud (1996). Stock market crashes, precursors and replicas. *Journal de Physique I, France 6*, 167–175.

Sutton, R. S. and A. G. Barto (1998). *Reinforcement Learning.* Cambridge, MA and London: MIT Press.

Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable.* Random House.

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

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

Taylor, L. and S. A. O'Connell (1985). A minsky crisis. *Quarterly Journal of Economics 100*, 871–875.

Thaler, R. (Ed.) (1993). *Advances in Behavioral Finance.* Russell Sage Foundation.

Vandewalle, N., M. Ausloos, P. Boveroux, and A. Minguet (1999). Visualizing the log-periodic pattern before crashes. *The European Physical Journal B - Condensed Matter and Complex Systems 9*(2), 355–359.

Voit, J. (2005). *The Statistical Mechanics of Financial Markets (Theoretical and Mathematical Physics)* (3rd ed.). Springer.

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