



City Research Online

City St George's, University of London

Citation: Iori, G. & Porter, J. (2012). Agent-Based Modelling for Financial Markets (12/08). London, UK: Department of Economics, City University London.

This is the submitted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/1744/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).



School of Social Sciences
CITY UNIVERSITY LONDON

Department of Economics

Agent-Based Modelling for Financial Markets


Giulia Iori

City University London

James Porter^{*}

City University London

**Department of Economics
Discussion Paper Series
No. 12/08**



^{*} Corresponding author: *Department of Economics, School of Social Science, City University London, Northampton Square; James.Porter.1@city.ac.uk*

AGENT-BASED MODELLING FOR FINANCIAL MARKETS

GIULIA IORI AND JAMES PORTER

*Department of Economics, School of Social Sciences, City University London,
Northampton Square, London EC1V 0HB.*

Chapter prepared for the Handbook on Computational Economics and Finance. Eds. Shu-Heng Chen and Mak Kaboudan, OUP Oxford University Press, Forthcoming

1. INTRODUCTION

An agent-based model (ABM) is a computational model which can simulate the actions and interactions of individuals and organisations, in complex and realistic ways. Even a simple agent-based models can exhibit complex behavior patterns and provide valuable information about the dynamics of the real-world system which they emulate. Agent Based Models transcend the numerous restrictive assumptions underlying most main-stream models and can create the emergent properties arising from complex spatial interaction and subtle interdependencies between prices and actions, driven by learning and feedback mechanisms. The theoretical assumption of mathematical optimization by agents in equilibrium is replaced by the less restrictive postulate of agents with bounded rationality adapting to market forces. Many approaches have been adopted in modelling agent behaviour for financial agent-based models. Agents can range from passive automata with no cognitive function, to active data-gathering decision makers with sophisticated learning capabilities. Indeed, agents are not only heterogeneous and interacting but also adaptive; they have different circumstances, different histories and adapt continuously to the overall situation they create. Agents can engage in comprehensive forms of learning that include inductive reasoning (experimentation with new ideas) as well as aspects of reinforcement learning, social mimicry, and forecasting of future events. When agent interaction is contingent on past experience, and especially when there is a continuous adaptation to that experience, mathematical analysis is very limited in its ability to derive the dynamic consequences.

Traditional economic models can get reasonably good insights into the economy by assuming that human behaviour leads to stable, self-regulating markets, with prices never departing too far from equilibrium. But the theory of complex systems shows that although a system may have an equilibrium state, its basin of attraction may be very narrow and the system rarely settles there. The equilibrium may also be very sensitive to small perturbations, and therefore becomes less relevant for the understanding of the system. The highly stylised, analytically tractable

E-mail address: Giulia.Iori@city.ac.uk, James.Porter.1@city.ac.uk.

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement CRISIS-ICT-2011-288501.

traditional models in economics and finance are not well-suited to study crisis situations (Bouchaud [2008], Farmer and Foley [2009], Kirman [2010]); in fact there is no framework in classical economics for the understanding of crises. ABMs on the contrary can represent unstable systems with crashes and booms that develop out of non-linear responses to proportionally small changes.

Economists have developed powerful tools to understand the role of strategic interaction among a limited number of agents, but the embeddedness of economic activity in social settings has been largely ignored by the economic profession until the early 1990s. Since then the study of socio-economic networks has exploded, with main focus the development of models of strategic network formation. The underlying assumption in these models is that the payoffs to each individual provide the incentives to form or sever links, and the basis for a welfare evaluation. By focusing on the optimal behavior of agents in forming links, these theoretical models provide useful insights to understand why certain network structures emerge. Nonetheless the networks that emerge as stable or efficient are too simple (such as star networks) and rarely observed in reality. Thus, these models are not well suited in terms of matching the properties of observed large social networks, characterised by considerable amount of heterogeneity in the network structures. As suggested by Jackson [2007], agent-based simulation could be a valuable tool to study more realistic network formation models that could capture, both, more node heterogeneity and randomness in behaviour.

The last few years have seen significant popular coverage of the potential of agent-based models for preventing financial crises and better understanding the economy. A article in the Economist (Economist [2010]) suggests that Agent Based Models might do better than conventional economic models, such as dynamic stochastic general equilibrium models, in foreseeing financial crises. In a recent interview on Institutional Investors Farmer [2012] advocates the use of ABM simulations to understanding the economy and financial markets as complex evolving systems. Buchanan [2009] asks whether, in analogy to traffic forecasting models, it may be possible to build a control centre (or war room) for financial markets, where policy makers could be alerted to potential crises and run appropriate simulations in order to understand how best to respond to ongoing events. Regulators and policy makers (Trichet [2010], Haldane [2009]) have also been calling for novel approaches and tools to monitor the state of the economy, that recognise its complex and interconnected nature.

Despite the widespread interest in ABM approaches, agent-based models remain at the fringe of mainstream economics. Some critics argue that ABMs are too narrow in focus. Agent-based modelling in financial markets has devoted a lot of attention to providing a behavioural explanation of a number of universally observed facts (or stylized facts) of financial time series which were inconsistent with standard asset pricing models; while it has achieved considerable success in this area, there has been less engagement with the more general topics of interest to the more traditional financial market research communities.

Another criticisms which is made of agent-based modelling is the lack of clarity about how one can do policy with them. Farmer and Foley [2009] describe the substantial progress that has been made using ABMs to model large parts of an economy; however, they acknowledge the need to go further and to apply the ABM

methodology to the creation of larger models which can incorporate multiple markets. In this direction promising results have been achieved by the Eurace project (Deissenberg et al. [2008]) which represents the first attempt to create a large scale model of the European economy. More recently, the CRISIS project¹ has undertaken the challenging task to build an integrated finance/macroeconomic ABM to produce a quantitative understanding of financial crisis. A visionary project, the FuturICT Knowledge Accelerator² is another major effort towards large scale ABM and foresees, among its goals, the development of a sophisticated simulation platform, with models driven and calibrated by data aggregated in real-time. This might be used to address issues such as risk, trust, resilience and sustainability and support the policy making, along with business and individual decisions.

The most *fundamental* critique from economists is that ABMs lack microfoundations for agents' economic activities unlike traditional intertemporal optimization models³. The aim of this review is to show how agent based models in financial markets, have evolved from simple zero intelligence agents, that follow rather arbitrary rules of thumb, into more sophisticated agents described by better microfounded rules of behaviour. We then look at the key issue of model calibration. Finally we look at some cases where ABMs have been successful at providing insight for policy making.

2. EARLIER ABM REVIEWS

A number of ABM reviews have been published over the last 10 years testifying the growing academic interest for these methodology, both in the economics and physics communities. In this review we focus on showing the development of models of financial markets and in the increasing structural and behavioral sophistication, thinking about empirical issues for such models, thinking about policy issues and future possibilities. The below reviews offer complementary foci.

A relatively early and comprehensive survey of agent-based modelling for finance is LeBaron [2006]. LeBaron concentrates on questions of design before surveying the types of existing models and some empirical issues. The design section is of particular interest to those pursuing agent-based modelling of financial markets from an economics perspective. It covers issues such as preferences (and time), price formation, evolution, learning, how to represent information and social learning. The importance of having 'benchmarks', or parameters for which the model is well understood, is highlighted. LeBaron survey covers a range of models running from 'few type' models to very dynamic, heterogeneous models. The 'few types' models analyze a small number of strategies, typically technical or fundamental, that are used by agents to trade a risky asset. The proportion of agents adopting different strategies is determined by the strategies past performance. These models tend to be more analytic than computational. In 'many type' models the small sets of tractable trading rules are replaced with larger sets of strategies. Model remain close to a well-defined theoretical framework but extend the framework by including learning agents. The next set of artificial market models moves away from testing specific theoretical models. These models are characterised by a dynamic ecology of trading strategies and simulations are used to determine which strategies will

¹<http://www.crisis-economics.eu/home>

²<http://www.futurict.eu/>

³Of course many traditional 'microfounded' models have their own significant limitations.

emerge and survive, and which will fail. The Santa Fe Artificial Stock Market, which we return to below, is one of the earliest examples of such models.

Hommes [2006] surveys heterogeneous agent models (HAM) with an emphasis on models which are at least somewhat tractable by analytical methods. HAM are simple, stylized versions of the more complicated and computationally oriented ABM but share with them the paradigm shift, from the representative agent approach, towards a behavioral approach in which heterogeneous, boundedly rational agents follow rule of thumb strategies. Such strategies, while simple, perform well and lead to sophisticated macro level structure. Attention is initially focused on early models which include ‘fundamentalist’ and ‘chartist’ agents; the former forming their expectations on market fundamentals and the later on trends in historical price patterns. Other topics covered include examples of disequilibrium HAMs which present complex market dynamics such as cycles or chaotic fluctuations, systems of agents with stochastic or social interactions and financial market models with herding behavior. Hommes and Wagener [2009a] survey simple HAM models in which financial markets are viewed as complex evolutionary systems. They introduce the main features of adaptive belief systems and discusses a number of examples, discuss their empirical implications, and confront the models with data from laboratory experiments with human subjects. A number of chapters in the same edited collection (Hens and Schenk-Hopp [2009]) overview cutting hedge research on financial markets that model the dynamics of asset prices as driven by the heterogeneity of investors.

In Kirman [2002], and in greater length in Kirman [2011], consideration is given to the way agent-based modellers build models of economic systems (and more generally how economic modelling should be done). An argument is made for models which take into account direct interactions of agents and in particular for approaches which utilize a network to model these interactions.

Samanidou et al. [2007] outlines the main ingredients of some influential early models in financial markets to move to a number of more recent contributions, appearing in both the physics and economics literature. In particular they focus on models that formalise the description of financial markets as multi agent systems, and can reproduce empirically observed universal scaling laws. The authors nonetheless point out how the ability of most of these models to explain interesting empirical facts vanishes for realistically large populations of agents.

A more recent survey is Cristelli et al. [2011] which discuss, in a unified framework, a number of influential agent based models for finance with the objective of identifying possible lines of convergence. Models are compared both in term of their realism and their tractability. The question which model is better has no clear answer because the stylized facts are relatively limited and not too difficult to reproduce in an ABM framework. As already observed by Samanidou et al. [2007], Cristelli et al. [2011] confirms that in most models considered the stylized facts do not correspond to a genuine asymptotic behavior but can only be obtained for a specific number of agents and in a limited region of parameters. For this reason the authors argue that self-organization should be a crucial ingredient in ABM to drive the dynamics of the system spontaneously to the realistic region.

An extensive review of Econophysicists’ work in agent-based models is provided in Chakraborti et al. [2011]. Three key areas are examined: models of order-driven markets, kinetic theory models for wealth distribution and game theoretic models

(particularly for the minority game). The authors conclude that existing models either are simple toy models that cannot be calibrated with real data, or more realistic models, suitable for calibration, but poorly tractable and whose sensitivity to the various parameters is particularly difficult to understand. Finally, they observe that the cancellation of orders is the least realistic mechanism implemented in existing models and that no agent-based model of order books deals with the multidimensional case. Thus fully reproducing empirical observations on correlation and dependence is still an open challenge for ABM.

A broader perspective can be found in Chen [2012] which gives a historical overview of how agent-based computational economics has developed looking at four origins: the market, cellular automata, tournaments (or game theoretic) and experiments. In thinking about financial markets the first is of most obvious relevance but work stemming from all four approaches have played a role in the agent-based modelling of financial markets. The market, understood as a decentralized process, has been a key motivation for agent-based work; Chen argues that the rise of agent-based computational economics can be understood as an attempt to bring the ideas of *many* and *complex heterogeneous* agents back into economic consideration. Zero intelligence (ZI) agents, or randomly behaving agents, have been a key part of finance research. The intuition behind this assumption, is that, given the law of large numbers, no matter what the individual motivations behind agents behaviour are, their aggregate behavior appears equivalent to that of randomly-behaving agents. Other simple programmed agents have included features such as swarming, social intelligence and regime switching. Computer tournaments have been used to solicit human programmed behaviors for complicated dynamic games and to test computer generated solutions. Experiments, or less formal observation of human behavior, have been important for agent-based modelling and calibration. We return to the issue of calibration in section 5.

A recent survey focused on zero-intelligence approaches for finance is Ladley [2012]. ZI models have allowed researchers to gain insight into market dynamics without having to make diverse behavioural assumptions regarding the strategies of traders. By removing strategy from market participants, the researcher may gain an insight into the effect of the market mechanism on the overall market dynamics. The simplicity of these models has the additional benefit, in some cases, of making them analytically tractable. Ladley shows how ZI models may do poorly where there are opportunities for learning (zero-intelligence agents don't learn) and where feedback loops between agents action and the state of the environment they operate generate complex dynamics.

3. TRADITIONAL APPROACHES AND EMPIRICAL EVIDENCE

Much research in financial markets has focused on thinking about fully rational agents (perhaps with some learning) processing information (which may be imperfect) to infer the correct 'fundamental' value of an asset. There is no scope in these models for chartist agents or herding behaviour. The argument is essentially that the predictability of prices should be reduced to zero by rational investors who should earn higher profits and drive less rational traders out of the market. Nonetheless artificial stock market models show that the market does not generally select the rational, fundamental strategy, and that simple technical trading strategies may survive.

The idea of market efficiency is central to much financial research, with both strong theoretical and empirical literatures. In empirical terms, models such as the random walk are arguably pretty good approximations for evidently unpredictable financial markets. A benchmark (theoretical) model for thinking about efficiency is outlined in Grossman and Stiglitz [1980] where agents can purchase a signal about an asset. In an efficient world with a small cost on the signal no one would buy the signal; but how then could it be (informationally) efficient? This paradoxical character of information efficiency is a theme of much criticism of such concepts.

While relating market efficiency to empirical studies is controversial, it is generally accepted that there are many empirical financial phenomena which are difficult to explain using traditional models. As many authors have noted the empirical distributions of returns of many market indices and currencies, over different but relatively short time intervals, shows an asymptotic power law decay (Mandelbrot [1963]; Pagan [1996]; Guillaume et al. [1997]; Gopikrishnan et al. [1999]). A Gaussian, as predicted by the random-walk hypothesis, is recovered only on time scales longer than a month. Moreover, while stock market returns are uncorrelated on lags larger than a single day, the correlation function of the volatility is positive and slowly decaying, indicating long-memory effects. This phenomenon is known in the literature as volatility clustering (Ding et al. [1993]; DeLima and Crato [1994]; Ramsey [1997]; Ramsey and Zhang [1997]). The empirical evidence also points to persistency in trading volume and positive cross-correlation between volume and volatility (Tauchen and Pitts [1983]; Ronalds et al. [1992]; Pagan [1996]). There is also evidence that both the moments of the distribution of returns (Ghashghaie et al. [1996]; Baviera et al. [1998]) and the volatility auto-correlations (Baviera et al. [1998]; Pasquini and Serva [2000]) display multi-scaling.

Recently, the empirical analysis of limit order data has revealed a number of intriguing features in the dynamics of placement and execution of limit orders. In particular, Zovko and Farmer [2002] found a fat-tailed distribution of limit order placement from the current bid/ask. Bouchaud et al. [2002] and Potters and Bouchaud [2003] found a fat-tailed distribution of limit order arrivals and a fat-tailed distribution of the number of orders stored in the order book. The analysis of order book data has also added to the debate on what causes fat tailed fluctuations in asset prices. Gabaix et al. [2003] put forward the proposition that large price movements are caused by large order volumes. A variety of studies have suggested that the mean market impact is an increasing function of the order size. Nonetheless Farmer et al. [2004] have shown that large price changes in response to large orders are very rare. Order submission typically results in a large price change when a large gap is present between the best price and the price at the next best quote (see also Weber [2005] and Gillemot et al. [2006]).

4. MODELLING APPROACHES

In this section we survey the range of models used for financial agent-based modelling where the market mechanism is a major area of interest. In building these kinds of models two key areas are pertinent: understanding the structure of the market and understanding the modelling of behavior. In order to focus on questions of market structure behaviour can be modelled in very simple ways, ranging from leaving it out entirely (having for example market orders placed randomly), to zero intelligence trading (where we have individuals but they trade essentially

randomly, though typically subject to some constraints such as a budget) through to more sophisticated models which include ideas such as bounded rationality, game theoretic principles or approaches from behavioral sciences. Another key element is the way in which heterogeneity is featured in the model: for the simpler models it often arises purely from the random behavior of (statistically) identical individuals, but for the more complex models the heterogeneity often plays a more central role. Approaches from across the range are surveyed below. We focus initially on three categories of models for agents in markets: zero intelligence agents, heterogeneous agents interacting through a market mechanism and heterogeneous agents interacting directly. This approach allows us to think about the level of sophistication in agent behavior and in the structural detail of the interactions in the models.

Direct interactions, or social interactions, are meant to capture how the choice of each agent is influenced by the choices of others. Various alternatives have been considered in the social utility literature: global interaction, where individuals tend to conform to the average behaviour of the entire population, and local interactions, where individuals have an incentive to conform to or information on a specific population sub-group. While interactions could be heterogeneous (with different strengths and signs between pairs) and asymmetric (Iori and Koulovassilopoulos [2004]), the literature has mostly focused on pairwise symmetric spillover, in which case the payoff of a particular choice increases when others behave similarly. Positive social interaction models generate polarised group behaviour even when agents characteristics are uncorrelated. Models allow for the neighborhood composition to evolve over time, possibly in a self-organized way. Typically agents can form new alliances according to some fitness maximisation scheme.

ABMs challenge the neoclassical hypothesis of agents relying on perfect knowledge of the economy and infinite computing capabilities to form rational expectations. Rather, they embrace the bounded rationality paradigm according to which the expectation formation process is driven by adaptive learning or evolutionary selection via genetic algorithms (Chen et al. [2008]). Agents may use technical trading rules or Artificial Neural Networks (ANNs) (Terna [2002]) to forecast market prices. In most of the behavioural models surveyed below agents face discrete choices (submit buy or sell orders, switch between different strategies, sever or form new links, and so on). Bounded rationality in this context may enter via the assumption that, while utility is deterministic, the agents' choice process is stochastic. This formulation captures the difficulty agents face in evaluating different features of the various alternatives and do not necessarily select what is best for them.

4.1. Zero Intelligence Agents. We start with zero intelligence agent models; these models actually vary substantially in what is meant by the term zero intelligence and range from very random behaviors, perhaps constrained by a budget, to those where there may be some kind of strategy specified. However, one can identify two key features in most of such models: lack of (explicit) learning and a minimalist approach to agent behavior. In addition to more obvious agent-based economics models, we consider some approaches from physics which are closely related.

Research in zero intelligence trading for financial markets was started by Gode and Sunder [1993], although there is some related earlier work by Becker [1962] on random agents and aggregate outcomes. In Gode and Sunder [1993] zero intelligence traders are compared to a set of human traders and aggregate outcome compared.

These zero intelligence traders are “not intended as descriptive models of individual behavior”, rather, the point of the model is to think about efficiency arising from the structure. The key result is that for their model, in terms of the aggregate property of allocative efficiency⁴, zero intelligence (“irrational”) traders perform comparably well to human traders. The mechanism studied is that of a double auction in which buyers and sellers submit limit order asks or bids and can accept these asks or bids. Each bid is independently and uniformly drawn from the range $[1, 2, \dots, 200]$. When they match or cross they are accepted. Each buyer has a valuation v_i and each seller a reserve cost c_i . For a unit sold at price p the profit is thus $p - c_i$. Two variants are investigated and compared to results from experiments: one where agents trade randomly over the full range of possible values and one where agents have a budget constraint, that is they must make profit or pay less than their valuation. The latter achieves an average efficiency extremely close to that of human traders.

In Gode and Sunder [1997] this research is continued with an exploration of explanations for allocative efficiency for markets. Again zero intelligence traders are constrained to avoid losses, but otherwise bid randomly. A larger number of markets are investigated, including modifications such as limited collection of bids, a limit of trading to only one round, current bid and ask prices not being made public, and the results suggest that for many market structures the simple rules, rather than complex behaviors, may give rise to most of the efficiency of a market.

In the spirit of Gode and Sunder [1993], Duffy and Ünver [2006] asks whether a simple agent-based model can generate the kind of bubbles and crashes which have been observed in experimental settings; the question here is not a matter of efficiency but of replication and understanding of observed behavior. The experimental setup is a round based trading market with cash and a single asset, at most one unit of which could be sold or bought each round via an order book. The agent-based model used is somewhat more complex than Gode and Sunder [1993], indeed they term it “near-zero-intelligence”, the key difference is rather than having purely random prices (albeit constrained to profitable prices) the average transaction price of the previous round is known. This provides a mechanism for the generation of price bubbles which does not depend on the more sophisticated strategies of many later models.

Another relatively recent example of the kind of work which shows that zero-intelligence trading can give rise to observed market phenomena (in contrast to the experimental results which are the point of comparison for the above research) is Ladley and Schenk-hopp [2007] which looks at a limit-order driven market using a zero-intelligence approach in keeping with the original intuitive style (there is a no-loss constraint on orders) and each buyer/seller has a reservation price as in Gode and Sunder [1993]. Orders are randomly drawn pairs (p, q) drawn from the set of feasible trades (price p must not incur a loss and q units must be available for sale) and when an order is placed the previous order from that trader is removed. In contrast to the original version traders randomly enter and exit the market. The goal is to determine whether characteristics of the order book are a result of the market mechanism or trader strategy. An ‘average’ order book is constructed by looking at best five bid and ask prices and this is compared to empirical findings. The bid-ask spread is found to be about twice the width of either adjacent spread

⁴Allocative efficiency is total profit divided by maximum total profit, or sum of consumer and producer surplus.

and the volume available is almost constant across prices; both of these correspond to empirical findings. As tick size is reduced the volume offered at the best price is reduced, another empirical observation. As the model lacks sophisticated behavior on the part of traders, it suggests these and other properties of such markets may in large part arise from the mechanism.

While the above models take a minimal approach to agent based modelling, it is possible to go even further to the extent of having implicit agents with actions, such as the placing of market orders, occurring randomly at a market rather than at an explicit individual level. A major contribution in this area is Daniels et al. [2003], which is explored in more detail in Farmer et al. [2005a]. In this work order arrival and cancellations are modelled as Poisson random processes (rather than being the explicit actions of agents). Orders arrive in chunks of size σ at rate μ shares per unit time with equal probability of being a buy or sell order. Offers are placed with uniform probability at multiples of a tick size over an infinite interval. At time t the best asks and bids are $a(t)$ and $b(t)$, with spread $s(t) = a(t) - b(t)$. The shape of the order book will be a key consideration: in particular the distribution of stored market orders. Market orders are matched against limit orders, in order of price, and removed. Based on this model predictions can be made for key properties of the market, in particular for the diffusion rate of prices and the spread (difference between best buying and selling prices) and price impact functions. This is possible as the model is simple enough to characterise these properties through dimensional analysis. This kind of model is tested against data from the London Stock Exchange in Farmer et al. [2005b] where it can explain over 95% of the spread and over 75% of the variance of the price diffusion rate with a single free parameter. These, and similar, results suggest there may be simple laws connecting price and properties of the market which do not depend on sophisticated strategies on the part of agents.

4.2. Heterogeneous Agents with Market Mediated Interactions. While zero intelligence models can replicate many stylized facts about financial markets, they cannot address many questions about modelling behavior (as, for good reason, they omit this) and, given we are using an agent-based modelling methodology, are less comprehensive than is necessary. The models we explore below have richer behaviors and agent interactions, though for now we restrict our attention to those with purely market mediated interactions.

An example of early research in building models of financial markets, where traders have strategies and speculate on endogenous price fluctuations, is Caldarelli et al. [1997]. As in many later models, each trader switches between cash and a single stock. All traders start with the same quantity of each and all have access to the complete price history of the stock. At each time period agent's strategy is an amount of stock $S_{i,t}$ to buy/sell. Each agent has a mapping from the history of prices to a fraction of stock $S_{i,t}$ to buy/sell, somewhat in the style of Arthur [1994]. In this case strategies are moving averages of combinations of derivatives of logged prices. This simple model generates a complex price history with the scaling of price variations close to that observed in real financial markets.

The distribution of returns is something poorly captured by traditional financial market models. Research which explores this issue from a behaviourally minimal but structurally detailed way is LiCalzi and Pellizzari [2003]. The model⁵ consists

⁵We go into a little detail about the basic specification as this model is representative at a basic level of many models we look at later.

of an economy with two assets, a bond and a stock. The price of the stock depends on demand and supply of agents. The total supply of cash and stock is constant though not all traders may be active at once. Traders enter market with cash c_i and stock s_i and can buy additional stock or sell stock through an order book process which allows both market orders (which can be fully or partially filled) and limit orders. All agents are fundamentalists (we look at richer models elsewhere in this section) and try to buy low and sell high relative to their individual estimate of the fundamental value, v_i , of the stock. Upon entering the market each agent has an investment horizon h_i , enter market at time t and wish to maximise their gains over time $h_i - t$. Bonds have a risk free return of r , so i requires a sufficient risk premium π_i to invest in the stock, or

$$\frac{v_i}{p} \geq 1 + (r + \pi_i)(h_i - t)$$

and the agent will invest in the bond (sell stock) if

$$\frac{v_i}{p} \leq 1 + r(h_i - t).$$

Based on the above agents buy/sells stock at best prices available in order book when it makes sense (given their valuation) to buy or sell. When such orders are not available they place their own limit orders. This trading takes place in a number of sessions (“days”). Simulations include both the approach above with its risk attitude and knowledge of r and ‘zero intelligence’ trading where only v_i is known. Even in the later simple case we see fat tailed logged returns suggesting this is a result of the structure rather than behavior. While this phenomena may be largely due to structural causes, additional properties such as volatility clustering and short-term correlations cannot be explained by the market structure alone.

In Lux and Marchesi [1999] a model of a financial market with chartists and fundamentalists which gives rise to scaling laws is described. A closely related model is investigated in more detail in Lux and Marchesi [2000]. Here there is a market maker balancing demand and supply of n_c agents with chartist strategies and n_f agents with fundamentalist strategies is explored. The total number of agents is kept as a constant $N = n_c + n_f$. There is further heterogeneity in that within the chartist group agents may be optimistic or pessimistic about the short term future, we have n_+ and n_- of each and an opinion index

$$x = \frac{n_+ - n_-}{n_c}, x \in [-1, 1].$$

The chartists buy or sell (a fixed number of units) if they are optimistic or pessimistic respectively. Fundamentals buy or sell if the market value is below or above the fundamental value. Agents endogenously switch between these groups with the transition probabilities arising from economy wide average profits and parameters for the inertia between groups. The switching probability from positive to negative is

$$\pi_{+-} = \nu\left(\frac{n_c}{N} \exp U\right)$$

and from negative to positive is

$$\pi_{-+} = \nu\left(\frac{n_c}{N} \exp -U\right)$$

where

$$U = \alpha_1 x + \alpha_2 \frac{\dot{p}}{\nu}.$$

and ν is a frequency of opinion revaluation parameter, α_1, α_2 are parameters for the relative importance of majority and price trend. The simulation results give consistent statistical characteristics for the market, including both fat tails and volatility clustering which correspond to empirical observations. Usually the macro behavior of this model is stable, but outbreaks of volatility can occur and the stylised facts of simulated time series data correspond to those observed in real markets. The market behavior is related in Lux and Marchesi [2000] to the concept of “on-off intermittency”: there is an attracting state which become temporarily unstable due to the crossing of some local stability threshold; in this model it is the fraction of traders adopting chartists versus fundamentalist strategies.

LeBaron [2005] looks at the development of the Santa Fe artificial stock market, one of the first major attempts to build a (somewhat) detailed agent-based model of a financial market. As in many of the models considered elsewhere in this survey, the initial artificial stock market model is one with a risk free asset and a risky stock paying a dividend

$$d_t = d + \rho(d_{t-1} - d) + \mu_t$$

where d, ρ are fixed parameters and $\mu \sim \mathcal{N}(0, \sigma_\mu^2)$. Agents have individual expectations for the price change of the stock and for its variance. The more modern versions of the Santa Fe artificial stock market use a constant relative risk aversion preference for the formation of individual demands for the stock. Agents use a classifier system to estimate the returns. The classifier is based on a the presence of a number of properties, for example “price greater than five period moving average”. These properties are mapped to estimation parameters. Agents each have an *individual* evolving set of 100 rules such that periodically the twenty worst performing rules are removed and replaced with new rules via both crossover and mutation from their existing rules. This kind of individual based selection between rules is similar in style to Arthur [1994]. The model generates many features of real financial data, specifically excess kurtosis in returns, low linear autocorrelation and persistent volatility.

The above models use some kind of switching, either between classes or strategies. An alternative is to think about agents using a mixed strategy, giving different weights to different components. In Chiarella and Iori [2002] a model of a simplified limit order book market is built in order to investigate the effects of differing combinations of strategies on aggregate outcomes. Attention is also given to how structural details of the market (tick size and order lifetime) affect these aggregate outcomes. In the model weights are given to fundamentalist and chartist components of an agent’s strategy and a return is individual estimated via

$$\hat{r}_t = g_1 \frac{(p^f - p_t)}{p_t} + g_2 r_L + n\epsilon.$$

The sign of g_2 indicates a trend chasing (> 0) or contrarian (< 0) chartist component.

Building on Chiarella and Iori [2002], LeBaron and Yamamoto [2007] introduces learning to the order driven market. These parameters (weights) predicated made using equation from Chiarella and Iori [2002], that is g_1, g_2 and n_i are initially assigned randomly and are updated via a genetic algorithm in which the fitness

function f_i is based on the mean squared deviation of realized prices from predicted prices given by

$$f_i = \frac{1}{\sum_{5 \text{ rounds}} (p_t - E^i(p_t))^2}$$

where the $E^i(p_t)$ was based on the individual i 's weighted estimations of return and the probability of a strategy being copied into the next round is

$$P_i = \frac{f_i}{\sum_j^N f_j}.$$

In addition to this copying, there is also a small probability of mutation, where one of the parameters of the price prediction function is replaced by a new value drawn from the original distribution. The key stylized fact they are able to capture is long memory in trading volume, volatility of return and in sign of market orders⁶; this entails that future values of these quantities are (significantly) predictable from past values. A modified version of the rescaled range or R/S statistic is used to test the simulated data and they are able to reject their null hypothesis of short-range dependence (or lack of long memory) in a majority of simulations for all of volume, volatility and sign of market order quantities. When the simulations were carried out without evolution there was insufficient evidence to reject the null hypothesis in most cases.

In the style of Chiarella and Iori [2002], Chiarella et al. [2009] looks at more sophisticated agents with heterogeneous strategies. Again agents have components of each return forecasting strategy (with heterogeneous weights) and different parameters for each component; however, now utility functions are introduced for each agent and further heterogeneity is facilitated by varying risk aversion. The agents maximise their expected utilities via on their individual estimate of stock return, based on their weighted average of the three components. Simplifying a little, this estimate is

$$r = \frac{1}{g_1 + g_2 + n} \left[g_1 \frac{1}{\tau_f} \ln(p_f/p) + g_2 \bar{r} + n\epsilon \right]$$

where these estimates are individual, g_1, g_2, n are the heterogeneous weightings given to each of the fundamentalist, chartist and noise components, τ_f is the time scale for mean reversion to the fundamental price, p_f is the fundamental price and $\bar{r}$ is the chartist component based on estimated return over previous times steps. The agents can place limit or market orders in order book depending on best prices available and their estimate of return. The results suggest that chartist strategies generate longer tails in the distribution of orders (in keeping with empirical findings). The increase in volatility following a large price movement can be explained by the large and opposite contributions to price expectations from the chartist and fundamentalist components.

4.3. Heterogeneous Agents with Direct Interactions. While for many purposes it is ideal to have parsimonious models, as we are adopting a computational approach, we do not have to be limited to those models which can be analytically analysed or even those which attempt to create the simplest possible model of an agent for a particular scenario. Drawing on insights from behavioral sciences and

⁶Whether most orders are buyer or seller initiated.

knowledge of market structures we can build considerably more comprehensive models; these models may include sophisticated learning behaviors and explicit modelling of the direct interactions of agents. Traditional approaches tend to suggest that an irrational departure from market fundamentals should not be sustainable. However, empirical evidence of repeated financial market bubbles (and subsequent crashes) suggests that ruling out such behavior means omitting a major feature of financial markets from models. There is some more traditional work which includes the idea of herding, such as Banerjee [1992] where because agents are making a choice sequentially and basing their choice on those choices already made by others, agents may over-rule their own (better) information (see Hirshleifer and Teoh [2009] for a recent review on herding behaviour in financial markets).

It is possible to identify three stages in the modelling of direct interactions. The first is global interactions, where an agent uniformly randomly interacts with another agent. The second is local interactions on a lattice, where interactions are constrained to a set of neighbors but in a regular way. The final stage is local interactions on a network. With respect to interactions on a network, more recent work takes seriously issues such as how inter-agent structures arise. We can see the evolution from early models where the network is assumed, to later models where the network structure may arise endogenously.

An early piece of work thinking about the interactions of individual agents and the macro consequences is Kirman [1993]; many works in agent-based modelling for financial markets build on this analytic foundation of recruitment. The basic idea here is that we have a system such as an ant hill with N agents (ants) who retrieve food from two sources, “black” and “white”. The state is just the number of ants k using the “black” source. Ants switch between the two resources via an individual process of recruitment: two ants meet and one switches to another’s source with probability $(1 - \delta)$. There is an additional ϵ probability that an ant switches food source without an interaction. So at each time step the system evolves from state k to $k + 1$ with probability

$$P(k, k + 1) = \left(1 - \frac{k}{N}\right) \left(\epsilon + (1 - \delta) \frac{k}{N - 1}\right)$$

and from k to $k - 1$ with probability

$$P(k, k - 1) = \left(\frac{k}{N}\right) \left(\epsilon + (1 - \delta) \frac{N - k}{N - 1}\right).$$

We can characterize the long term behavior of this Markov chain, with particularly interesting results when most of the time the system is at the extreme values. If we think of the choice of source as choice of opinion or strategy for trading, then these states are ones where we may see herding.

Another early approach to crowding behavior, or herding, is Bak et al. [1997] which adopts methods from Physics. There are N traders, each of which can own one share; if they do then they are potential sellers, if they don’t then they are potential buyers. Each has a price $p_s(j)$ or $p_b(j)$, the price they are willing to sell or buy at respectively, which is determined by their individual strategy. Scenarios considered include markets with only fundamental traders and with noise traders (random valuations initially uniformly distributed within a range, then fluctuating randomly). In the most interesting version of the model imitation is introduced for the noise traders, such that when they are selecting a new price they randomly

copy a price from another agent (of either type). What happens as the proportion of rational traders is varied? When there are few rational traders (c. 2%) they can be priced out of the market in a “bubble”. When there are many (c. 20%) the prices are kept within their range.

Vriend [2000] distinguishes between learning at an individual and at a population level, in the former case the agent learns exclusively on the basis of his experience, in the later case the agent also bases his learning on the experience of other players. The example which forms the focus of Vriend [2000] is a Cournot oligopoly game with n firms choosing quantity q_i to be produced and where market price is

$$P(Q) = a + b\left(\sum_i q_i\right)^c$$

and a, b, c are set such that this is a downward slopping curve. Fixed costs of K and marginal costs of k are assumed. Two ways of implementing a genetic algorithm for this kind of scenario are identified. The first is where each individual firm has an output rule (specifying production quantity) and after many periods some kind of crossover and/or mutation is applied based on the relative success of the rules. The second (or individual) learning, as in Arthur [1994] and LeBaron [2005], has a set of rules for each agent and those most successful recently are more likely to be used. For the Cournot oligopoly model these approaches result in completely different aggregate outcomes, with social learning producing a much higher average output than individual learning. In both cases convergence to these distinct levels is relatively fast (it only requires a few genetic algorithm steps, with 100 days of production in between each). The finiteness of the model is important as it allows individual agents some influence over the actions of others. Also this is not really a typical search problem for applying a genetic algorithm (it is far too simple) but this kind of effect may occur for much more complex models.

Building on the approaches in papers such as Kirman [1993] and Lux and Marchesi [2000], Westerhoff [2009] builds a model with strategy switching between fundamental and technical trading. In contrast with Kirman [1993], above, the switching of opinions is now more sophisticated: the probability of adopting the rule used by another trader now depends on past profitability of the rule and no longer has the same symmetric, random specification. This is accomplished via fitness variables A^C and A^F for the chartist and fundamentalist strategies respectively, each a discounted sum of the past returns. The Kirman [1993] style dynamics are modified by including an weighting on transition probabilities given by $0.5 + s\lambda$ where $s \in \{-1, 1\}$ and the sign of s reflects the relative fitness of the strategy to be changed to. This switching leads to periods dominated by a fundamentalist rule, but with major shifts towards technical rules which increase volatility and may result in bubbles and crashes.

A major area of research for agent-based models of financial markets is the Minority Game. In it an odd number of agents choose between two options independently and want to be in the minority. There have been hundreds of research papers on this topic as it is seen to be a good model for thinking about financial market issues. It developed from Arthur [1994] which considered a model of individual inductive reasoning about aggregate outcomes for attendance at the ‘El Farol’ bar. This kind of model, where every agent wishes to be in the minority is seen to encapsulate key qualities of financial markets. Challet et al. [2005] includes both an introduction to and a comprehensive collection of many major papers on the Minority Game.

In Challet et al. [2000] the basic formulation of the minority game is built on in multiple ways such as the introduction of heterogeneity (in the form of classes of agent), an increase in memory length, the possibility of having more strategies and agents getting information of other agents. The idea is to think about various financial market issues in the context of this well understood model. The initial formulation is standard with agents $i = 1, \dots, N$ and actions $a_i(t) = \pm 1$. The gain of agent i at time t is $g_i(t) = -a_i(t)A(t)$ where

$$A(t) = \sum_{j=1}^n a_j(t).$$

A particularly interesting section is that which looks at what happens if an agent knows ahead of time the actions of a subset of other agents. This agent can then adopt a different strategy depending on the aggregate decision of the subset. This agent always gains at least as much as the average. If there are more agents with this extra information then the gain is reduced.

A model with interacting agents which can give rise to volatility clustering is presented in Iori [2002]. A modified random field Ising model is used to model the behavior of agents in a financial market. There is an $L \times L$ lattice, with each node i being an agent, connected to his four nearest neighbors. Initially each agent owns the same amount of capital with $M_i(0)$ units of cash and $N_i(0)$ units of stock. At each time step three actions, $S_i(t)$, are possible: -1 if they sell a unit of stock, 0 if they do nothing and 1 if they buy. A market maker clears orders and adjusts prices. Agents make decisions based on an idiosyncratic signal $\nu_i(t)$ (a shock to personal preferences) and through exchanges of information between neighbors. The aggregate signal is:

$$Y_i(\tilde{t}) = \sum_{i,j} J_{ij} S_j(\tilde{t}) + A \nu_i(t)$$

where J_{ij} captures influence. For simple cases of J_{ij} this model is well understood in Statistical Physics, for example if $J_{ij} = 1$ then this is the Ising model and traders would all agree (with large resultant fluctuations in price). In addition to the above formulation, friction is introduced otherwise agents would sell given any positive or negative signal, however small. Synchronization effects (which generate large fluctuations in returns) are shown to arise purely from imitation among these simple traders. These fluctuations exhibit the mutliscaling phenomena observed empirically (Pasquini and Serva [2000]).

Traditional models of financial markets have particular difficulties in explaining the presence of bubbles. Föllmer et al. [2005] looks at a model of a financial market where the demand of agents for assets is determined by their forecasts of prices. The agents switch between rules in a way which is driven by the success of the rules and influenced by other traders. Expectations of prices can be heterogeneous, though agents are not, in contrast to related approaches, systematically wrong. The prices can move far from the fundamentals, but the fundamentals do determine the long run behavior. These rules are supplied by the recommendations R_t^i of a guru or financial expert i . Agents choose from the available experts randomly, with choices weighted by the discounted average of past profits for those recommendations. This model allows clear investigation of the effects of different kinds of rules (or gurus). In particular it is seen that the switching in forecasting method can actually be

self-fulfilling and may result in bubbles. Chartist experts increase both variance and kurtosis of the limiting empirical distribution of logarithmic prices; they cause (temporary) bubbles and crashes in the model.

In Stauffer and Sornette [1999] clusters of agents aggregate and shatter via variation of a parameter p for connectivity. These clusters act together and the idea is that there may be times when traders act very individualistically and times when herding is strong. This is developed from a model in Cont and Bouchaud [2000] which is simple enough (in contrast to say Bak et al. [1997]) to allow for some analytical results. There is a market with N agents and a single asset with price at time t of $x(t)$. Demand for agent i is a random variable $\phi_i \in [-1, 0, 1]$ where positive represents a bullish agent (wanting to buy) and negative a bearish agent (wanting to sell). If $\phi_i = 0$ then the agent doesn't trade in that period. So the excess demand for the market is

$$D(t) = \sum_{i=1}^N \phi_i(t)$$

and as demand is assumed to be symmetric, or

$$P(\phi_i = +1) = P(\phi_i = -1) = a \text{ and } P(\phi_i = 0) = 1 - 2a$$

then average excess demand is 0. Price change is assumed to be proportionate to the excess demand with a parameter λ for market depth (controlling how sensitive the price is to excess demand). The key element of the model is communication between agents, which is modelled here by a set of clusters which coordinate individual demand; so the excess demand is now the weighted sum over cluster demands, or

$$\frac{1}{\lambda} \sum_{\alpha=1}^k W_{\alpha} \phi_{\alpha}(t)$$

Modelling clustering of agents via the a random graph model for links between agents allows us to characterize the distribution of cluster sizes, depending on a single parameter for overall willingness of agents to coordinate their demand. Once we have distribution of cluster sizes it is possible to characterize the distribution of aggregate demand and hence price changes. In particular two key results are derived: heavy tailed density of price changes and the heaviness of tails (kurtosis of price change) is inversely proportional to order flow; these results hold for a range of parameter values. In Stauffer and Sornette [1999] the stylized facts of interest arise not from a parameter value being within a certain range, but from the variation of 'herding strength'.

The kind of herding models we have seen above (in the style of Kirman [1993]) have been connected to network structure. A review of the mean field approximation approach to determining transition rates and resulting equilibrium distribution is provided in Alfarano [2008]. Transition probabilities are reformulated for individuals based on their neighbors, now the probability of switching for an individual is

$$p_i = \frac{a + \lambda n(i, j)}{a + \lambda N}$$

where a is the idiosyncratic parameter, λ is the global herding intensity and $n(i, j)$ is the number of i 's neighbors in the opposite state. The probability of not switching

is simply given by $1 - p_i$. Results for regular, scale free, random⁷ and small world networks are compared to the mean field results which are a reasonable approximation for the simulated results on networks. Only the random network captures the stylized fact of constant variance in proportion in particular states for any system size. When heterogeneity in behavior is introduced it has little effect on outcomes, in contrast to introducing a heterogeneous network structure which can have a major effect on the aggregate outcomes.

Endogenous network formation for financial markets is considered in Tedeschi et al. [2012]. Unlike work such as Föllmer et al. [2005], where the idea of ‘gurus’ is something built into the market via the availability of a set of rules, the idea of a ‘guru’ here arises as an endogenous result of an information network. Each agent has an outward connection to another agent and the ‘guru’ is the agent with the most incoming links.

Agents have cash and stocks and a wealth relative to the wealthiest agent which is used as a measure of fitness,

$$f_t^i = \frac{W_t^i}{W_t^{\max}}.$$

Agents may randomly rewire (choose another agent to be connected to) and they do this based on the fitness of agents. The probability of i rewiring to agent j from k is

$$p_r^i = \frac{1}{1 + e^{-\beta^i(f_t^j - f_t^k)}}.$$

Agent’s expectations are a combination of their own individual expectations and those of the agent they are connected to. In the model gurus emerge endogenously, rise and fall in popularity over time, and are possibly replaced by new gurus. Traders have an incentive to imitate and a desire to be imitated since herding turns out to be profitable. The assumption that noise traders quickly go bankrupt and are eliminated from the market is unrealistic in presence of herding and positive feedback. We show that more sophisticated strategies underperform the guru and his followers and positive intelligence agents can not invade a market populated by noise traders when herding is high.

5. CALIBRATION OF AGENT-BASED MODELS OF FINANCIAL MARKETS

Relating agent-based modelling and empirical knowledge is a key challenge of agent-based modelling in general. In the case of financial markets we typically have large volumes of high resolution data which should be helpful both for calibrating and evaluating models. The typical approach taken for ABM, as with most of the work surveyed above, is to replicate stylized facts from financial markets. However, in addition to this kind of qualitative replication, attempts have been made to more fully calibrate models from empirical data of particular financial markets. We consider some general issues then look at specific examples of calibration of financial market models.

A general guide to empirical validation of agent-based models can be found in Fagiolo et al. [2007]. It highlights key issues facing modellers attempting empirical validation, attempts to classify models and identifies unresolved issues. Problems

⁷Uniform random probability of any particular link existing.

and solutions are split into three categories: (i) relating theory and empirical research, (ii) relating models and real world systems and (iii) how the empirical validation deals with the these first two issues. Agent-based modelling is summarised. A key aspect of the approach adopted here is to think of there being a ‘real world data generating process’ (rwDGP) and ‘model data generating process’ (mDGP). The later must be simpler than the former and its ‘goodness’ is to be evaluated by comparing simulated outputs and real-world observations. The lack on consensus about validation is remarked upon and four categories of heterogeneity in approaches identified: those relating to nature of object studied, in the goal of the analysis, in the modelling assumptions and in the method sensitivity analysis. Three methods of validation with particular relevance to modelling markets are examined: the indirect approach the Werker-Brenner approach and the ‘history-friendly’ approach. Indirect calibration is where stylised facts are identified and a model build with reference to known microeconomic description, then the stylised facts are used to restrict parameters. The Werker-Brenner approach combines Bayesian inference, retaining only those parameters associated with highest likelihood of empirical data, and an attempt to identify structure from remaining models. The ‘history-friendly’ approach uses case studies (for example of particular financial markets) and, for it, a good model is one which generates stylised facts for those studies.

Richiardi [2012] is a recent introduction to agent-based computational economics with a emphasis on interpretation of results and estimation. In addressing estimation the necessary approach is contrasted with that for an analytical model. One must compare artificial data with real data and should change the structural parameters of the model such that these two sets of data become as close as possible. There are various ways in which one might measure this closeness and form an objective function for the optimisation algorithm. One method suggested is that *the method of simulated moments* different orders of moments are weighted by their uncertainty. For real data this can be estimated and for simulated data this can be reduced by repeated simulation.

Judd [2006] offers a general overview of methodological computational issues related to agent-based economic modelling. The main appeal of computational approaches is outlined, namely that the elements of economic investigations previously sacrificed for simplicity can be investigated. Two common objections to numerical approaches are examined. Firstly, the lack of generality to which it is argued that theories look at a “continuum of examples”, but perhaps a measure zero set of plausible or interesting examples. Viewed this way, Judd argues that the relevance and robustness of examples is more important than the number. A second common objection, errors, is dismissed, as when handled carefully these can be negligible. The main question for Judd is how we systematically do computational (economic) research? We can’t prove theorems using computers (in the conventional sense) but we can search for counter examples to a proposition, use Monte Carlo sampling methods (which can be clearly expressed in terms of classical or Bayesian statistics), use regression methods to obtain “shape” of some distribution and we can perhaps straightforwardly adapt a computer model to a new case (something which is often not at all straightforward for a theorem).

In Chen et al. [2012] the development of agent-based (computational) modelling is described from an econometrics viewpoint. Of particular relevance to this section

are the accounts of (i) ABM and stylized facts and (ii) the use of econometric methods for estimation for ABM. The first of these gives a comprehensive description of thirty stylized facts observed in the literature (we looked at research focusing on many of these in section 4) and how these relate to the number of agent types in the model. The later offers a clear account of the major options in estimating an agent-based model. Essentially one can carry out direct or indirect estimation. The former case may be possible for simpler agent-based models. One uses statistical techniques to estimate the probability of parameters. The later case will typically be necessary for more complex models. We see examples of both approaches below.

One of the first examples of validation/estimation of an agent-based model of financial markets is Gilli and Winker [2003]. This uses a stochastic approximation of an objective function for estimating the parameters of a foreign exchange model. Bianchi et al. [2007] looks at a case study of validating the “Complex Adaptive Trivial System” model of Mauro et al. [2005]. This model has reproduced many stylized facts of financial markets with ad hoc parameter values. The calibration process here takes a sample of Italian firms and estimates parameters for the model using this real world data. The model is modified, mostly making it more realistic (introducing realistic heterogeneity for firms), though the new model uses a homogeneous market interest rate (as micro-level data is not available to do otherwise). The process used is one of indirect inference, minimizing the distance between the actual and simulated distributions of the model, this allows for a close match of the simulated results to the empirical data. A ‘simple’ agent-based model of order flow is validated in Mike and Farmer [2008] where we have a random order placement process. The original model is from Daniels et al. [2003] which we introduce above, though now additional empirical regularities are modelled, specifically the order signs, the order price and order cancellations all now have empirically motivated models, allowing for a more realistic model of order placement than the previous approach. The model is constructed based on a single stock and tested on 24 others. For those with small tick sizes and low volatility the constructed model works particularly well.

In Ghonghadze and Lux [2010] a framework for collective opinion formation is created and compared to two more standard time series models, when applied to EU business and consumer survey data. Specifically the model’s performance in out-of-sample forecasting is compared to $\text{ARMA}(p, q)$ and $\text{ARFIMA}(p, d, q)$ univariate time series models. In the model we have two opinion states, positive and negative, with n^+ and n^- of agents holding each view; let $n_t = (n_t^+ - n_t^-)/2$ be the configuration and assuming we have $2N$ agents the aggregate expectation is the ratio $x_t = n_t/N$. A Master equation can be formulated for this system and a continuous approximation can be numerically solved allowing us to calibrate the system from the EU data. It typically does better than the ARMA models and performs similarly to the ARFIMA models for individual series (though looking at the performance across all the data, it does better than the ARFIMA models in a majority of cases).

Housing bubbles had an important role to play in the recent financial crisis, though this is an area that has historically been of little interest for macroeconomics. In Geanakoplos et al. [2012] a retrospective model of housing which includes large amounts of actual data from Washington, DC. As it is a detailed agent-based model it can include large amounts of heterogeneous, individual level data which many

models would have to omit or aggregate. This includes information on race, income (from detailed IRS income data for area), wealth, age and household position. In addition demographic trends such as population size, death rates and migration patterns can be included along with economically relevant parameters such as loan-to-value ratios.

6. POLICY AND ABMS

Below we look at several applications of agent-based modelling of financial markets to policy. Agent-based modelling seems to be particularly useful when experimenting with policy rules. Agent-based models may capture details an analytical model cannot and may be more acceptable to policy makers as they are less abstract (Dawid and Neugart [2011]). As we mentioned above building large scale forecasting models is a goal of many agent-based modellers, though a difficult task.

A publication by Agentlink, Luck [2005], attractively presents numerous uses of agent-based modelling for commercial purposes, including their success for practical use in financial markets. It notes, for example, the ability of agent-based traders to outperform human traders by 7% and the use of agent-based auctioning systems for the decentralized allocation of resources in many substantial real world settings. Even at the time of its publication a large proportion of trades on many financial markets are carried out by some kind of automated trader (potentially exactly corresponding to a trading agent in an financial market ABM).

One early and successful commercial application of ABM was developed by Bios Group for the National Association of Security Dealers Automated Quotation (NASDAQ) Stock Market ⁸. In 1997, the NASDAQ Stock Market was about to implement a sequence of apparently small changes: reduction in tick size, from 1/8th to 1/16th and so on down to pennies. In the agent-based NASDAQ model, market maker and investor agents (institutional investors, pension funds, day traders, and casual investors) buy and sell shares by using various strategies. The agents' access to price and volume information approximates that in the real-world market, and their behaviors range from very simple to complicated learning strategies. Neural networks, reinforcement learning, and other artificial intelligence techniques were used to generate strategies for agents. The model produced some unexpected results. Specifically, the simulation suggests a reduction in the market's tick size can reduce the market's ability to perform price discovery, leading to an increase in the bid-ask spread. A spread increase in response to tick-size reduction is counterintuitive because tick size is a lower bound on the spread. In the remainder of this section we look at some recent research drawing on ABMs for policy work.

The impact of Tobin style transaction taxes on an artificial financial market is explored in Mannaro et al. [2008]. The motivation for this tax comes from the proposal by James Tobin to charge a small tax or 0.1% on all foreign exchange transactions which should discourage short term speculation while leaving longer term investors relatively unaffected; this, it is widely believed, would reduce market volatility⁹. A model similar to many of those above, that is with one stock, cash and various classes of traders, is considered under a transaction tax regime. A number of computational experiments are carried out. It seems that in this model

⁸www.cbi.cgey.com/journal/issue4/features/future/future.pdf

⁹This is a controversial view and many have argued the opposing viewpoint. This kind of transaction tax has, at the time of writing, not been fully implemented in practice.

transaction taxes increase volatility and that when both a taxed and un-taxed market is available the volume traded in the taxed market reduces which may increase volatility. Another examination of transaction taxes in an agent-based model is Pellizzari and Westerhoff [2009] in which two micro-structures are considered: a continuous double auction and a central dealership. In the former case, while volume decreases with the transaction tax so too does liquidity, eliminating gains in stability from reduced volume of trading. In the later case as liquidity is provided by the dealership then the volatility of the market can be significantly reduced via the imposition of a transaction tax.

Hommes and Wagener [2009b] study the effects of financial innovation upon price volatility and welfare. They introduce hedging instruments in an asset pricing model with heterogeneous beliefs and show that more hedging instruments may destabilize markets and decrease welfare when agents are boundedly rational and choose investment strategies based on reinforcement learning.

Gsell [2008] incorporates algorithmic trading into the ABM of Chiarella et al. [2009]. Two strategies of order splitting have been implemented: (1) a simple static execution strategy, where the overall volume is executed linearly over time, (2) a dynamic execution strategy whose aggressiveness varies over time depending on the current market situation and the algorithms previously achieved performance. The results of the simulation show that Algorithmic Trading has an impact on market outcome in terms of price impact and market volatility.

In more recent work Anand et al. [2010], related to the understanding of current economic issues, a rule based approach is adopted with a focus on modelling credit derivative markets. When considering the purchase of an asset backed security (ABS) agents can choose whether to rely on a signal from a rating agency or they can carry out independent risk analysis. If many other agents also believe the rating agency then it is rational to believe that the ABS is liquid, irrespective of its underlying quality. In this model this simple but rational approach can result in a highly fragile state of the market as rules spread through the economy.

Unlike many agent-based models which take a fairly minimal approach the Eurace model is a large macroeconomic model which includes a number of interacting markets and agents whose balance sheets are rigorously modelled. Cincotti et al. [2010] gives an account of how the Eurace model was used to look at the provision of credit. The model includes detailed financial markets, credit markets and a central bank which can pursue quantitative easing. Two policy options, quantitative easing and fiscal tightening, are explored across multiple runs of the model and the results suggest that while quantitative easing increases inflation in both short and long run it leads to a better macroeconomic results (higher output).

In Thurner et al. [2010] leverage is connected to systemic financial risk. Here the kind of heavy tailed fluctuations which have arisen in some of the above models, through strategies such as trend following, are shown to arise from the effects of leverage. The focus is on collateralized loans with margin calls; a type of loan where a loan to value (of collateral) ratio must be maintained alongside interest payments by repaying rather than rolling over debt. This has a feedback effect (selling collateral reduces value of collateral which demands further sale of collateral and so on) which increases as the level of leverage increases. In this context the policy of leverage restriction may have unintended consequences, causing a local failure to become systemic.

7. CONCLUSIONS

As this survey has shown, much ABM research has concentrated on proof of concepts rather than the development of robust tools to control and forecast complex real-world financial markets. Nevertheless, stylized models are extremely useful for understanding how complex macro-scale phenomena emerge from micro-rules. This kind of exercise has allowed for the testing of existing economic theories and their refinement toward greater realism.

The ABM surveyed in this review show a clear evolution towards better micro funded behavioural approaches to modelling the agents' decision process. Still in many situations, agents, while behaving purposefully, use rules of thumbs and inductive reasoning to make decisions. While the fully rational utility maximiser of classical economics does not represent real people, more sophisticated models of behavior may be important for a fuller understanding of financial market dynamics. In Brenner [2006] various learning processes are surveyed that could guide modeling of economic agents behaviour as closely as possible to that of humans. Choosing between the various approaches can be difficult and there is not yet a consensus about which approaches are best in which situations. Evolutionary approaches are good for population level results though not individual dynamics. Fictitious play is both simple and supported by evidence. Where more information about beliefs is available stochastic belief learning may be a good approach. In short there are a rich and increasing set of behavioral models which could be applied to financial market ABMs.

When it come to modelling interactions, much work has been done both in Physics (Newman [2010]) and in Economics (Goyal [2009]; Jackson [2010]) on the subject of networks. Financial systems are networks, which have become increasingly more complex and interlinked. Nonetheless the literature on financial networks is still at an early stage (Allen and Babus [2009]), with most of the research concentrating on financial stability and contagion. The focus of ABM is typically on understanding systems dynamics on given network structure. From a regulatory perspective, questions such as optimal networks design or optimal design of incentives that lead to the formation of networks with desirable characteristic offer interesting opportunity for ABM research.

The calibration and validation of ABM is challenging. One key advantage for ABM with its explicit modelling of heterogeneous individuals is the possibility of calibration using fine-grained microeconomic data (see for example Geanakoplos et al. [2012]) and evidence from laboratory experiments with human subjects (Duffy [2006], Hommes and Lux [2009], Heckbert [2009]). Using experimental techniques, well-defined decision scenarios can be reproduced, and strategies that humans actually use in dealing with complex situations may be revealed. This approach offers a way to capture the heuristics of decision making in a model which is grounded in empirical data.

In section 6 we give multiple examples of the use of ABM for financial market related policy. However, the application of ABM to financial market policy, and indeed, macroeconomic policy matters in general is clearly still in its infancy. Drawing on increasingly sophisticated modelling techniques, detailed structural modelling and better calibration methods, such as the kind of individual based calibration made possible by agent-level modelling, has great promise for the future.

REFERENCES

- Simone Alfarano. Should network structure matter in agent-based finance? 2008.
- Franklin Allen and Ana Babus. *The Network Challenge*, chapter 21, pages 367–382. Wharton School Publishing, 2009.
- Kartik Anand, Alan Kirman, and Matteo Marsili. Epidemics of rules , information aggregation failure and market crashes. (1999):1–13, 2010.
- W. Brian Arthur. Inductive reasoning and bounded rationality. *American Economic Review*, 84:406–411, 1994.
- P. Bak, M. Paczuski, and M. Shubik. Price variations in a stock market with many agents. *Physica A*, 246:430–453, 1997.
- Abhijit V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- R. Baviera, M. Pasquini, M. Serva, D. Vergni, and A. Vulpiani. Efficiency in foreign exchange markets. 1998. URL <http://xxx.lanl.gov/cond-mat/9811066>.
- G. Becker. Irrational behavior and economic theory. *The Journal of Political Economy*, 70:1–13, 1962.
- Carlo Luigi Bianchi, Pasquale Cirillo, Mauro Gallegati, and Pietro Vagliasindi. Validating and calibrating agent-based models: A case study. *Computational Economics*, 30(3):245–264, 2007.
- J.-P. Bouchaud, M. Mézard, and M. Potters. Statistical properties of stock order books: empirical results and models. *Quantitative Finance*, 2:251–256, 2002.
- Jean-philippe Bouchaud. Economics needs a scientific revolution. *Nature*, 455 (October):1181, 2008.
- Thomas Brenner. Agent learning representation: Advice on modelling economic learning,. *Handbook of Computational Economics*, 2:895–947, 2006.
- Mark Buchanan. Meltdown modelling. *Nature*, 460(August):680–682, 2009.
- G. Caldarelli, M. Marsili, and Y.-C. Zhang. A prototype model of stock exchange. *Europhysics Letters*, 40:479–484, 1997.
- Anirban Chakraborti, Ioane Muni Toke, Marco Patriarca, and Frédéric Abergel. Econophysics review: II. Agent-based models. *Quantitative Finance*, 11(7):1013–1041, July 2011. ISSN 1469-7688.
- D. Challet, M. Marsili, and Y.-C. Zhang. *Minority Games: Interacting Agents in Financial Markets*. Oxford University Press, 2005.
- Damien Challet, Matteo Marsili, and Yi-Cheng Zhang. Modeling market mechanism with minority game. *Physica A*, 276:284–315, 2000.
- Shu-Heng Chen. Varieties of agents in agent-based computational economics: A historical and an interdisciplinary perspective. *Journal of Economic Dynamics and Control*, 36(1):1–25, 2012. ISSN 01651889.
- Shu-Heng Chen, Tzu-Wen Kuo, and Kong-Mui Hoi. Genetic programming and financial trading: How much about what we know? In C. Zopounidis, M. Doumpos, and P. Pardalos, editors, *Handbook of Financial Engineering*, pages 99–154. Springer, 2008.
- Shu-Heng Chen, Chia-Ling Chang, and Ye-Rong Du. Agent-based economic models and econometrics. *The Knowledge Engineering Review*, 27(Special Issue 02):187–219, 2012.
- Carl Chiarella and Giulia Iori. A simulation analysis of the microstructure of double auction markets. *Quantitative Finance*, 2:346–353, 2002.

- Carl Chiarella, Giulia Iori, and Josep Perelló. The impact of heterogeneous trading rules on the limit order book and order flows. *Journal of Economic Dynamics and Control*, 33:525–537, 2009.
- Silvano Cincotti, Marco Raberto, and Andrea Teglio. Credit money and macroeconomic instability in the agent-based model and simulator eurace. *Economics: The Open-Access, Open-Assessment E-Journal*, 3:0–33, 2010. ISSN 1864-6042.
- Rama Cont and Jean-Philippe Bouchaud. Herd behavior and aggregate fluctuations in financial markets. *Macroeconomic Dynamics*, 4:170–196, 2000.
- M Cristelli, L Pietronero, and A Zaccaria. Critical Overview of Agent-Based Models for Economics. pages 1–51, 2011.
- Marcus G. Daniels, J. Doyne Farmer, László Gillemot, Giulia Iori, and Eric Smith. Quantitative model of price diffusion and market friction based on trading as a mechanistic random process. *Physical Review Letters*, 90(10), 2003.
- Herbert Dawid and Michael Neugart. Agent-based models for economic policy design. *Eastern Economic Journal*, 37:44–50, 2011.
- Christophe Deissenberg, Sander van der Hoog, and Herbert Dawid. EURACE: A massively parallel agent-based model of the European economy. *Applied Mathematics and Computation*, 204(2):541–552, October 2008. ISSN 00963003.
- P. DeLima and N. Crato. Long range dependence in the conditional variance of stock returns. *Economic Letters*, 45:281, 1994.
- Z. Ding, C.W.J. Granger, and R.F Engle. A long-memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1:83106, 1993.
- John Duffy. Agent-based models and human subject experiments. volume 2 of *Handbook of Computational Economics*, pages 949 – 1011. Elsevier, 2006.
- John Duffy and M. Utku Ünver. Asset price bubbles and crashes with near-zero-intelligence traders. *Economic Theory*, 27:537–563, 2006.
- The Economist. Agents of change, July 2010. URL <http://www.economist.com/node/16636121>. Retrieved 16 February 2011.
- Giorgio Fagiolo, Alessio Moneta, and Paul Windrum. A critical guide to empirical validation of agent-based models in economics: Methodologies, procedures, and open problems. *Computational Economics*, 30(3):195–226, October 2007. ISSN 0927-7099.
- J. D. Farmer, S. Krishnamurthy D. E. Smith I. Gillemot G. Iori, and M. G. Daniels. A random order placement model of price formation in the continuous double auction. *The Economy as an Evolving Complex System III*, pages 133–173, 2005a.
- J. Doyne Farmer. Agent-based modeling. *Institutional Investor*, September 2012. URL <http://www.institutionalinvestor.com/Article/3094958/People/J-Doyne-Farmer-on-Agent-Based-Modeling.html>.
- J. Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, pages 685–686, 2009.
- J. Doyne Farmer, L. Gillemot, F. Lillo, S. Mike, and A. Sen. What really causes large price changes? *Quantitative Finance*, 4:383–397, 2004.
- J. Doyne Farmer, Paolo Patelli, and Ilija I Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences of the United States of America*, 102(6):2254–9, 2005b. ISSN 0027-8424.
- Hans Föllmer, Ulrich Horst, and Alan Kirman. *Equilibria in financial markets with heterogeneous agents: a probabilistic perspective*, volume 41. 2005.

- X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley. A theory of power-low distributions in financial market fluctuations. *Nature*, 43:267–270, 2003.
- John Geanakoplos, Robert Axtell, Doyne J. Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M. Palmer, and Chun-Yi Yang. Getting at systemic risk via an agent-based model of the housing markets. *American Economic Review*, 102:53–58, 2012.
- S. Ghashghaie, W. Breymann, J. Peinke, P. Talkner, and Y. Dodge. Turbulent cascades in foreign exchange markets. *Nature*, 381:767–770, 1996.
- Jaba Ghonghadze and Thomas Lux. Modeling the dynamics of eu economic sentiment indicators: An interaction-based approach. 2010.
- L. Gillemot, J.D. Farmer, and F. Lillo. There's more to volatility than volume. *Quantitative Finance*, 6:371–384, 2006.
- M. Gilli and P. Winker. A global optimization heuristic for estimating agent based models. *Computational Statistics & Data Analysis*, 42(3):299 – 312, 2003. ISSN 0167-9473. [jce:title;Computational Econometrics;ce:title;.](#)
- Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):pp. 119–137, 1993. ISSN 00223808.
- Dhananjay K. Gode and Shyam Sunder. What makes markets allocationally efficient? *Quarterly Journal of Economics*, 112:603–630, 1997.
- P. Gopikrishnan, V. Plerou, M. Meyer, L.A.M. Amaral, and H.E. Stanley. Scaling of the distribution of fluctuations of financial market indices. 1999. URL <http://xxx.lanl.gov/cond-mat/9905305>.
- Sanjeev Goyal. *Connections: an introduction to the economics of networks*. Princeton University Press, 2009.
- Stanford E. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.
- Markus Gsell. Assessing the impact of algorithmic trading on markets: A simulation approach. 2008. URL <http://ssrn.com/abstract=1134834>.
- D.M. Guillaume, M.M. Dacorogna, R.R. Davé, U.A. Müller, and R.B. Olsen. From the birds eye to the microscope: a survey of new stylized facts of the intra-day foreign exchange markets. *Finance and Stochastics*, 1:95–130, 1997.
- Andrew G. Haldane. Rethinking the financial network. Speech delivered at the Financial Student Association, April 2009.
- S. Heckbert. Experimental economics and agent-based models. In *18th World IMACS / MODSIM Congress*, 2009.
- Thorsten Hens and Klaus Reiner Schenk-Hopp, editors. *Handbook of Financial Markets: Dynamics and Evolution*. Elsevier, 2009.
- David A. Hirshleifer and Siew Hong Teoh. Thought and behavior contagion in capital markets. In Thorsten Hens and Klaus Reiner Schenk-Hoppé, editors, *Handbook of Financial Markets: Dynamics and Evolution*. Elsevier, 2009.
- Cars Hommes. Heterogeneous agent models in economics and finance. volume 2 of *Handbook of Computational Economics*, pages 1109–1186. Elsevier, 2006.
- Cars Hommes and Thomas Lux. Individual expectations and aggregate behavior in learning to forecast experiments. *CeNDEF Working Papers*, 2009.
- Cars Hommes and Florian Wagener. Complex evolutionary systems in behavioral finance. In T. Hens and K.R. Schenk-Hoppé, editors, *Handbook of Financial Markets: Dynamics and Evolution*. Elsevier, 2009a.

- William Brock Cars Hommes and Florian Wagener. More hedging instruments may destabilize markets. *CeNDEF Working Paper*, 2009b.
- Giulia Iori. A microsimulation of traders activity in the stock market: the role of heterogeneity, agents' interactions and trade frictions. *Journal of Economic Behavior & Organization*, 49:269–285, 2002.
- Giulia Iori and Vassilis Koulovassilopoulos. Patterns of consumption in discrete choice models with asymmetric interactions. In William Barnett, Christophe Deissenberg, and Gustav Feichtinger, editors, *Economic Complexity: Non-linear Dynamics, Multi-agents Economies, and Learning*, volume 14. Elsevier, 2004.
- Matthew O. Jackson. *The Missing Links: Formation and Decay of Economic Networks*, chapter 2, pages 19–43. Russel Sage Foundation, 2007.
- Matthew O. Jackson. *Economic and Social Networks*. Princeton University Press, 2010.
- Kenneth L Judd. Computationally intensive analyses in economics. volume 2 of *Handbook of Computational Economics*. Elsevier, 2006.
- Alan Kirman. Ants, rationality, and recruitment. *Quarterly Journal of Economics*, pages 137–156, 1993.
- Alan Kirman. The economic crisis is a crisis for economic theory. *CESifo Economic Studies*, 56(4):498–535, 2010.
- Alan P. Kirman. Reflections on interaction and markets. *Quantitative Finance*, 2(5):322–326, 2002.
- Alan P. Kirman. *Complex Economics*. Routledge, 2011.
- Dan Ladley. Zero intelligence in economics and finance. *The Knowledge Engineering Review*, 27(02):273–286, 2012. ISSN 0269-8889.
- Dan Ladley and Klaus Reiner Schenk-hopp. Do stylised facts of order book markets need strategic behaviour? pages 1–24, 2007.
- Blake LeBaron. Building the santa fe artificial stock market. In F. Luna and A. Perrone, editors, *Agent-Based Theory, Languages, and Experiments*, number June. Routledge, 2005.
- Blake LeBaron. Agent-based computational finance. volume 2 of *Handbook of Computational Economics*. Elsevier, 2006.
- Blake LeBaron and Ryuichi Yamamoto. Long-memory in an order-driven market. *Physica A*, 383:85–89, 2007.
- Marco LiCalzi and Paul Pellizzari. Fundamentalists clashing over the order book: a study of order-driven markets. *Quantitative Finance*, 3:1–11, 2003.
- Mike Luck. 50 facts about agent-based computing. 2005. URL <http://www2.econ.iastate.edu/tesfatsi/AgentLink.50CommercialApplic.MLuck.pdf>.
- T. Lux and M. Marchesi. Scaling and criticality in a stochastic multi-agent model of a nancial market. *Nature*, 397:498–500, 1999.
- T. Lux and M. Marchesi. Volatility clustering in financial markets. int. *International Journal of Theoretical and Applied Finance*, 3:675–702, 2000.
- Benoit Mandelbrot. The variation of certain speculative prices. *Journal of Business*, 36:394–419, 1963.
- Katiuscia Mannaro, Michele Marchesi, and Alessio Setzu. Using an artificial financial market for assessing the impact of tobin-like transaction taxes. *Journal of Economic Behavior & Organization*, 67(2):445–462, August 2008. ISSN 01672681.

- Gallegati Mauro, Domenico Delli Gatti, Corrado Di Guilmi, Edoardo Gaffeo, Gianfranco Giulioni, and Antonio Palestrini. A new approach to business fluctuations: heterogeneous interacting agents, scaling laws and financial fragility. *Journal of Economic Behavior & Organization*, 56(4):489–512, 2005.
- Szabolcs Mike and J. Doyne Farmer. An empirical behavioral model of liquidity and volatility. *Journal of Economic Dynamics and Control*, 32:200–234, 2008.
- Mark Newman. *Networks: An Introduction*. Oxford University Press, 2010.
- A. Pagan. The econometrics of financial markets. *Journal of Empirical Finance*, 3, 1996.
- M. Pasquini and M. Serva. Clustering of volatility as a multiscale phenomenon. *The European Physical Journal B - Condensed Matter and Complex Systems*, 16: 195–201, 2000. ISSN 1434-6028.
- Paolo Pellizzari and Frank Westerhoff. Some effects of transaction taxes under different microstructures. *Journal of Economic Behavior & Organization*, 72(3): 850–863, December 2009.
- M. Potters and J.-P. Bouchaud. More statistical properties of stock order books and price impact. *Physica A*, 324:133–140, 2003.
- J.B. Ramsey. On the existence of macrovariables and of macrorelationships. *Journal of Economic Behavior & Organization*, 30:275–299, 1997.
- J.B. Ramsey and Z. Zhang. The analysis of foreign exchange rates using waveform dictionaries. *Journal of Empirical Finance*, 4:341–372, 1997.
- Matteo G. Richiardi. Agent-based computational economics: a short introduction. *The Knowledge Engineering Review*, 27(02):137–149, April 2012. ISSN 0269-8889.
- G.A. Ronalds, P.E. Rossi, and E. Tauchen. Stock prices and volume. *Review of Financial Studies*, 5:199242, 1992.
- E Samanidou, E Zschischang, D Stauffer, and T Lux. Agent-based models of financial markets. *Reports on Progress in Physics*, 70(3):409–450, 2007. ISSN 0034-4885.
- Dietrich Stauffer and Didier Sornette. Self-organized percolation model for stock market fluctuations. *Physica A*, 271:496–506, 1999.
- G.E. Tauchen and M Pitts. The price variability-volume relationship on speculative markets. *Econometrica*, 51:485–505, 1983.
- Gabriele Tedeschi, Giulia Iori, and Mauro Gallegati. Herding effects in order driven markets: The rise and fall of gurus. *Journal: Journal of Economic Behaviour and Organization*, 81:82–96, 2012.
- Pietro Terna. Cognitive agents behaving in a simple stock market structure. In F. Luna and A. Perrone, editors, *Agent-Based Methods in Economics and Finance: Simulations in Swarm*, pages 188–227. Kluwer Academic, 2002.
- Stefan Thurner, J Doyne Farmer, and John Geanakoplos. Cowles foundation discussion paper no. 1745. *Cowles Foundation for research in economics*, (1745), 2010.
- Jean-Claude Trichet. Speech of jean-claude trichet, president of the european central bank, on november 18, 2010. 2010.
- Nicolaas J. Vriend. An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control*, 24:1–19, 2000.
- Rosenow B. Weber, P. Order book approach to price impact. *Quantitative Finance*, 5:357–364, 2005.

- Frank Westerhoff. A simple agent-based financial market model: direct interactions and comparisons of trading profits. *BERG Working Paper Series*, 2009.
- I.I. Zovko and J.D. Farmer. The power of patience: A behavioral regularity in limit order placement. *Quantitative Finance*, 2:387–392, 2002.