Agent-based Computational Finance

Blake LeBaron *

Brandeis University

July 22, 2004: Preliminary draft, do not quote.

Abstract

This paper surveys research on computational agent-based models used in finance. It will concentrate on models where the use of computational tools is critical in the process of crafting models which give insights into the importance and dynamics of investor heterogeneity in many financial settings.

*International Business School, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a faculty research fellow at the National Bureau of Economic Research. This is a preliminary draft for *The Handbook of Computational Economics, vol. II,*, edited by K. L. Judd and L. Tesfatsion.

1 Introduction

In the mid to later part of the 20th century finance witnessed a revolution. The advent of the efficient markets hypothesis, the capital asset pricing model, and the Black/Scholes options pricing formula put the field on a new, solid scientific foundation. This world was built on the assumption that asset markets were powerful computational engines, and were able to aggregate and process the beliefs and demands of traders, leaving in prices the full set of properly processed information currently available. At the core of asset pricing, efficient market theories give a clean and compelling picture of the world which is as appealing to financial economists as it is potentially unappealing to financial practitioners. [1] It is interesting to note that these foundations came with a very important computational dimension. The early availability of large machine readable data sets, and the computational power to analyze them was the critical foundation for this new financial rigor.[2] In agent-based computational models the computer is once again at the center of a change in thinking about financial markets. This time is it helping to pursue the world in which agents may differ in many ways. Not just in the information they possess, but in their abilities to processes it, attitudes toward risk, and many other dimensions.

Models in the realm of agent-based computational finance view financial markets as interacting groups of learning, boundedly rational agents. The computer may or may not be a necessary tool to understand the dynamics of these markets. This survey will concentrate on the cases where analytic solutions would be impossible, and computational tools are necessary.[3] It is important to distinguish agent-based models from other more general heterogeneous agent models in finance, since the latter have been part of the field for some time. In agent-based financial markets the dynamics of heterogeneity is critical. This heterogeneity is represented by a distribution of agents, or wealth, across either a fixed, or changing set of strategies. In principle, optimizing agents would respond optimally to this distribution of other agent strategies, but in general, this state space is far too complicated to begin to calculate an optimal strategy, forcing some form of bounded rationality on both agents and the modeler. It is important to note that in these worlds bounded rationality is driven by the complexity of the state space more than the perceived limitations of individual agents. It is also important to remember that the simplified rules of thumb used by agents do not suggest that the exercise is forcing some sort of simplified solution on the dynamics of the steady state or

[1] This view is not far off the more general perspective on information dissemination on the economy as a whole put forth in Hayek (1945).

[2] A good early collection of work from this era is Cootner (1964).

[3] The survey by Hommes (forthcoming 2004) covers the more analytic heterogeneous agent models. Also, the recent book by Levy, Levy & Solomon (2000) provides another survey of recent work in the field.

the model, or presupposing that markets are not well represented by equilibrium rational stories. However, it is stressing that these need to be built from a foundation of simple adaptive behavior.

1.1 Why financial markets are good agent-based applications

Financial markets are particularly appealing applications for agent-based methods for several reasons. First, the key debates in finance about market efficiency and rationality are still unresolved. In addition, financial time series yield many curious puzzles which are not well understood. Second, financial markets provide a wealth of pricing and volume data which can be analyzed. Third, when considering evolution, financial markets provide a good approximation to a crude fitness measure through wealth or return performance. Finally, there are strong connections to relevant experimental results which in some cases operate at the same time scales as actual financial markets.

Academic finance has debated the issue of market efficiency for some time. The concept of market efficiency has a strong theoretical and empirical backing which should not be ignored.[4] On the theoretical side, the argument is that traders with less than rational strategies will disappear, and if prices contain any predictable components either in their own past series, or connected to fundamentals, the remaining rational investors will reduce these to zero. This is very close to evolutionary arguments put forth in both Alchian (1950) and Friedman (1953) for the evolution of firms and rational behavior in general. This powerful idea still holds sway in much of the academic financial world, and can be seen in papers such as Rubenstein (2001). As appealing as this idea was, it is interesting to note that there never really was a truly accepted dynamical process for how markets got efficient. The second foundation for efficient market theories was that much of the early empirical work on financial markets suggested that markets were much more unpredictable than the financial practitioner world had suggested.[5]. The random walk model appeared to be a pretty good approximation to the movements of stock prices, and it can be argued that it still is today. We know that markets are probably not completely unpredictable, but they still are very difficult to forecast.[6]

The early ideas of efficient markets were made more formal as modern tools of dynamic optimization were brought to bear on these problems.[7] This led to an even stronger representation for financial markets, the representative agent.[8] This model formally connects asset prices to the beliefs of a single aggregate

[4]The field has been surveyed many places, but the classic surveys remain Fama (1970) and Fama (1991).

[5]Examples are in Cootner (1964) and Fama (1970)

[6]It is also important to note that the radical idea that randomness was a good model for financial prices goes back to the beginning of the 20th century in Bachelier (1900).

[7]See for example, Merton (1971), Breedon (1979), and Lucas (1978).

[8]Constantinides (1989) is a good example describing the assumptions necessary to get a representative consumer in many cases. Also, Kirman (1992) critically assesses the use of representative agents in many economic contexts.

individual who can then be linked to various state variables of the macroeconomy.

The theoretical parts of efficient markets ideas have been attacked for quite some time. One of the most important questions for market efficiency comes from Grossman & Stiglitz (1980). Here, agents have the choice of purchasing an information signal on a financial asset. In a perfectly efficient world with a small cost on the signal, no one would have an incentive to buy the signal. However, if no one bought the signal how did the market get informationally efficient in the first place? It is interesting to note that many of the papers mentioned here, and in Hommes (forthcoming 2004) are based on the paradoxical structure of this model. More recently, the literature on noise trading, DeLong, Shleifer, Summers & Waldmann (1990), has introduced the important idea that risk averse rational types may not be able to "take over" the dynamics from less rational strategies, since they trade less aggressively because they are sensitive to the risk induced by the other traders. We will see that this concept plays an important role in many of the computational models considered here.

The attacks on the empirical side of market efficiency have been more controversial. During the 1980's and 1990's evidence began appearing indicating weaknesses with the efficient market hypothesis, and related equilibrium theories. There was evidence of predictability at long horizons as in Campbell & Shiller (1988), and at shorter horizons as in Lo & MacKinlay (1988). Old prediction methods which had been previously discredited began to appear again.[9] Also, connections between financial markets and macro dynamics were called into question by papers such as Mehra & Prescott (1988) and Hansen & Singleton (1983). Finally, the single factor CAPM model was shown to be insufficient in Fama & French (1992). Predictability alone did not mean the efficient market was dead. Indeed, in his later survey Fama (1991) is well aware that some studies had found some market predictability, but he correctly reminds us that predictability alone does not necessarily mean that markets are inefficient since profitable strategies may be bearing higher risk.[10]

Beyond simple predictability, there is a large range of empirical financial puzzles which remain difficult to explain using traditional asset pricing models. Among these are the overall level of volatility and long swings around fundamentals.[11] Also, the equity premium, which measures the difference between risky and riskless assets is difficult to explain.[12] This feature is directly connected to the failure of macro time series to connect well to financial markets. Series such as consumption are not volatile enough, and do not comove with markets in a way that can justify the magnitudes of risk premia observed in financial series. There have

[9]One example is the use of moving average technical analysis rules as in Brock, Lakonishok & LeBaron (1992).
[10]A good recent survey on this literature is Campbell (2000), and the textbook by Campbell, Lo & MacKinlay (1996).
[11]Shiller (2003) is a good recent survey on this.
[12]This one feature has generated an extensive literature which is surveyed in Kocherlakota (1996) and more recently in Mehra (2003).

been many attempts to address these issues in the academic finance literature, and these wont be surveyed here.[13]

Beyond these puzzles there are a set of facts that are still not well explained by any existing model. Trading volume is probably the most important. Financial markets generally exhibit large amounts of trading volume, and it is difficult to imagine that this can be driven by any situation not involving continuing disagreement between individuals. Beyond the level of volume, there are also some interesting dynamic effects which include persistence and cross correlations with returns and market volatility.[14] Also, recently, volume has also been shown to be a long memory process with persistence extending out many periods, Logato & Velasco (2000). At this time no convincing mechanisms exist for any of these features.

Equally puzzling, but more extensively studied, the persistence of volatility is another major feature that is lacking an explanation. While stock returns are generally unpredictable, their magnitudes are often very predictable.[15] Markets move between periods of relative calm to relative turmoil. This feature remains one of the most robust, and curious in all of finance. Although much is known about the structure of volatility persistence, little is known about its causes. Similar to volume persistence, it is also a potential long memory process.[16] Beyond simple persistence there are some more complicated issues in the dynamics of volume and volatility.[17]

Closely related to the volume and volatility persistence is the issue of fat tails, or excess kurtosis. At frequencies of less than one month the unconditional returns of financial series are not normally distributed. They usually display a distribution with too many observations near the mean, too few in the mid range, and again, too many in the extreme left and right tails. This feature has puzzled financial economists since it was discovered in Mandelbrot (1963). Recently, it has gained more attention since practical problems of risk management critically depend on tail probabilities. Precisely tuned complex derivative portfolios need very good estimates of potential tail losses. Return distributions eventually get close to normal as the time horizon is increased, and at the annual frequency, the normal distribution is not a bad approximation. Fat tails are not entirely independent of volatility persistence. The unconditional distributions of most volatility persistent processes are fat tailed, even when their conditional distributions are Gaussian. Beyond the

[13] Two recent models attempting to address many of these features are Campbell & Cochrane (1999) and Bansal & Yaron (forthcoming 2004).

[14] Many of these are documented in Gallant, Rossi & Tauchen (1992) and Gallant, Rossi & Tauchen (1993).

[15] This has been well known since Mandelbrot (1963), and has led to a large industry of models for fitting and testing volatility dynamics. See Bollerslev, Engle & Nelson (1995) for a survey.

[16] See Ding, Granger & Engle (1993), Andersen, Bollerslev, Diebold & Labys (2003), and also Baillie, Bollerslev & Mikkelsen (1996).

[17] These include connections between volatility and volume to return autocorrelations, LeBaron (1992) and Campbell, Grossman & Wang (1993), and temporal asymmetries in volatility documented in Dacorogna, Gencay, Muller, Olsen & Pictet (2001). Also, there are general indications that volatility tends to lead volume, but not vice versa, Fung & Patterson (1999).

frequency of large moves there is a continuing debate about exactly what the shape of the tails of financial distributions are. It is possible that these may be described by power laws.[18]

One of the reasons for this wide range of puzzles is another reason for finance being a good agent-based test bed. Financial data is generally plentiful, and accurate, and available on many different aspects of financial market functions. Good time series of up to 40 years are available on prices and volume. Longer series are available for lower frequencies, and certain securities. Over the past 20 years, extremely high frequency data has become available. These series often record every trade or every order entering a financial market, and sometimes include some information as to the identity of traders. Therefore, researchers have a detailed picture of exactly how the market is unfolding, and the exact dynamics of trade clearing. Also, series are available which show detailed holdings of institutions such as mutual funds as well as recording the flows of funds coming in and out of these funds. For individuals a few series have been used which reveal the trades of investors' accounts at various brokerage firms. This gives an amazing level of detail about the behavior of individuals which will be useful in the construction and validation of agent-based models. Finally, experimental data is available which can be used to line up and calibrate agent behavior. Several of the models covered here have already done this, and more examples of using experiments are given in Duffy (forthcoming 2004). Finance experiments are particularly appealing since they often can be done at time scales which are reasonable for the real data. It is more credible that you can simulate a day of trading in the laboratory, than to create someone's entire life cycle.

To summarize, financial markets are particularly well suited for agent-based explorations. They are large, well organized, markets trading securities which can be easily compared. Currently, the established theoretical structure of market efficiency and rational expectations is being questioned. There is a long list of empirical features which traditional approaches have not been able to match. Agent-based approaches provide an intriguing possibility for solving some of these puzzles.[19] Finally, financial markets are rich in data sets which can be used for testing and calibrating agent-based models. Data is available on many levels from macro to micro level information.

The remainder of this chapter will summarize recent work. The next section introduces some of the computational tools and design issues that are important in building markets. Section 3 covers models which attempt to recreate an entire market, or artificial market models. Section 4 looks at problems of institutions

[18] A good survey on power laws in finance is Cont (2001).

[19] There are other explanations that may yet prove to be the most important. These come from the area of behavioral finance which allows for deviations from strict rationality, and emphasizes the presence of certain key psychological biases which have been experimentally documented. See Hirshleifer (2001) and Barberis & Thaler (2002) for recent surveys on this literature.

and market design, or market microstructure. Section 5 covers a few other types of markets which don't fit in the earlier categories. Section 6 covers some on going debates and criticisms of agent-based markets, and section 7 concludes and suggests questions for the future.

2 Tools

2.1 Artificial Intelligence Tools

The Genetic Algorithm (or GA), is a key component in many, but not all, agent-based financial markets. It is viewed by some as a power tool for modeling learning and adaptation. It is an alternative to more traditional learning approaches such as Bayesian learning, and adaptive linear models. It is also controversial in that it is not clear that this is a good mechanism for replicating the learning process that goes on inside market participants' heads. This section provides a short introduction with special emphasis on the types of tools that are used in financial settings.

The most common application of the GA is as a simple optimization technique used in various problem solving situations. It is one of several optimization tools that are useful in places where traditional hill climbing may fail such as multi-peaked objectives, or nondifferentiable objective functions, possibly with discrete input varlables. Although in this context, the behavior of the GA is still not completely understood, this is a far simpler setting than the multi-agent models which will be considered in this survey.

Many beginning researchers view the GA as a kind of black box, and simply follow previous work in setup and structure.[20] This approach is probably a mistake. It is important to think more about evolutionary computation in general, than the particular pieces of the GA. The general field of evolutionary computation includes other methods such as evolutionary programming, and evolutionary strategies, and also genetic programming. For the consumer of these techniques distinctions are somewhat unnecessary, and parts of different methods should be used when the problem warrants it.[21]

Setting up an evolutionary learning setting involves first making several choices. First, the mapping from behavioral rules into a genetic structure is important. In some contexts this might involve simply combining real valued parameters into a vector of parameters, or in some instances it might involve coding real values as strings of zeros and ones. It also may involve taking a complex representation such as a neural network and mapping it into some simpler object. One needs to end up with some type of object that both represents

[20] Goldberg (1989) is the classic book for early GA adopters.
[21] A nice balanced overview of all these methods is Fogel (1995).

the behavior, and also can be easily manipulated by evolutionary operators.

In most evolutionary methods there will be a population of the previously mentioned solutions. In the individual optimization setting the information contained in the population is crucial to aiding in the search for solutions. Attached to each solution or rule is a fitness value. This is essentially the objective function for this potential solution. In the traditional optimization setting this isn't a problem since it is most likely a well defined function of the given parameters. This gets more difficult in mutli-agent settings where the question of optimality may be less well defined. Given a fitness value, the population can now be ranked. The computer simulates evolution by removing some set of low fitness solutions. The fraction of the population removed is an important design parameter to be decided. Setting this too high may cause the population to converge too quickly to a suboptimal solution. Setting it too low, may make selection weak, and the GA may converge far too slowly.

2.2 Design questions

In constructing an agent-based financial market the researcher is faced with a large number of basic design questions that must be answered. Unfortunately, there is often little guidance on which direction to follow. This section briefly overviews most of these questions which will be seen again as the setup of different markets is covered in later parts of this survey.

Probably the most important question is the design of the economic environment itself. What types of securities will be traded? Will there be some kind of fundamental value, and how does this move? Is there an attempt to model a large subset of the macro economy or just a very specific financial market? As in any economic modeling situation these are not easy questions. In the case of agent-based models they are often more complicated, since the accepted knowledge of how to craft good and interesting worlds of heterogeneous agents is still not something economists are very good at. It is not clear that the knowledge base for building representative agent macro economies will necessarily carry over into the agent-based world. This design question is probably the most important, and the most difficult to give guidance on.

Agent preferences are an important decision that must be made. Questions about preference types are critical. Should they be simple mean/variance preferences, or standard constant relative risk aversion form? Also, myopic versus intertemporal preferences is another issue. The latter brings in more realism at a cost of additional complexity in the learning process. It is also possible that certain behavioral features, such as loss aversion should be included. Finally, there may be an argument in certain cases to avoid preferences all together, and simply concentrate on the evolution of specific behavioral rules. The use of well defined preferences is the most comfortable for most economists. Their use facilitates comparisons with other standard models, and they allow for some welfare comparisons in different situations. Most applications to date have stayed with myopic preferences since the added complexity of moving to an intertemporal framework is significant. It involves learning dynamic policy functions in a world which already may be ill defined.

Beyond preferences the next key decision is on the institution of market clearing. Many of these models are interested in the fundamental problem of price formation, and the method for determining prices is critical. As we will see, many methods are used, but most fall into one of four categories. The first mechanism uses a slow price adjustment process where the market is never really in equilibrium. In this case a market maker announces a price, and agents submit demands to buy and sell at this price. The orders are then summed, and if there is an excess demand the price is increased, and if there is excess supply the price is decreased. The price is often changed as a fixed proportion of the excess demand as in equation 1.

$$p_{t+1} = p_t + \alpha(D(p_t) - S(p_t)) \tag{1}$$

An advantage and disadvantage of this is that the market is never in equilibrium. This might be reasonable for the adaptively evolving situations that are being considered. However, it also may be a problem, since, depending on α, these markets may spend a lot of time far from prices that are close to clearing the market. Another issue is how to handle excess demand. Are excess demanders supplied from some inventory, or is rationing used? Finally, how should the parameter, α, be determined?

A second market mechanism is to clear the market at each period, either numerically, or through some theoretical simplifications that allow for an easy analytic solution to the temporary market clearing price. This method reverses the costs and benefits of the previous method. The benefit is that the prices are clearing markets, and there is no issue of rationing, or market maker inventories that need to be dealt with. There are two critical problems for this type of market. It may impose too much market clearing, and may not represent well the continuous trading situation of a financial market. Also, it is often more difficult to implement. It either involves a computationally costly procedure of numerically clearing the market, or simplifying the demands of agents to yield an analytically tractable price.

These two pricing mechanisms take opposite extremes in terms of market clearing. Two other mechanisms fall somewhere in between. The most realistic mechanism from a market microstructure perspective is to actually simulate a true order book where agents post offers to buy and sell stock. Orders are then crossed using some well defined procedure. This is a great method for market realism. Its only drawback is that many details need to be built in on the trading mechanism, both in the market institution, and also in agent strategies. Any market that hopes to simulate realistic market microstructure behavior should follow this procedure. The final market mechanism that can be used is to assume that agents bump into each other randomly and trade if it benefits them. This is closest to a random field sort of approach as in Albin & Foley (1992). It also may have some connections to floor trading as used in the Chicago futures and options exchanges. It might also be a good representation for informal markets such as foreign exchange trading where, until recently, a lot of trade was conducted over the telephone.

Most agent-based financial markets are evolutionary in nature. In any evolutionary setting fitness evaluation is critical. Agents and strategies must be evaluated according to some criteria. Agents can be evolved using either wealth, or utility based fitness. In the case of wealth, evolution of agents themselves might be unnecessary since agents gaining more wealth will have a larger impact on prices. Utility is another possible fitness measure. Agents can be evaluated based on ex post utility achieved. Rules or trading strategies are often evolved and evaluated. The simplest criterion is to use a forecast based measure such as mean squared error, or mean absolute error, and to promote rules that minimize this. Forecasts are then converted into asset demands using preferences. This is a very transparent route, and it is possible to evaluate and compare agents based on their forecasting performance. This also aligns with the bulk of the learning literature in macroeconomics which often concentrates on forecast evaluation. A second, more direct route, is to ignore forecasts altogether and to deal with asset demands and strategies directly. The strategies are then evolved based on their impact on agents' utilities. This may be more difficult than considering forecast errors, but it eliminates an extra step in converting forecasts to demands, and is a little cleaner from a decision theoretic standpoint. In some cases this also avoids the need to estimate variances, and other higher moments, since risk would be taken into account. Finally, it is important to remember that all these fitness measures will most likely be measured with noise. Furthermore, it is not clear that the time series used to estimate them are stationary. Agents may end up choosing different lengths of history, or memory, in their rule evaluations which can translate into interesting dynamics. In a nonstationary world, there is no a priori argument for any particular history length. This greatly complicates the evolutionary process, and distances these problems from those often considered in the evolutionary computation literature.

One of the biggest problems in market design is how information is presented to the agents, and how they process it. Theoretically, this is the daunting task of converting large amounts of time series information from several series into a concise plan for trading. To handle this researchers are often forced to predefine a set of information variables, and the functional structure used to convert these into trading strategies. A second problem is how information is revealed about securities. Are there special signals visible to only certain agents? Are there costly information variables? How frequent are information releases? Unfortunately, there are no easy answers to these questions.

A final problem is how agents learn from each other. This is often known as "social learning", and has been the subject of much discussion in the agent-based modeling community.[22] At one extreme, agents may operate completely on their own learning rules over time, and only reacting with others through common price and information variables. Other mechanisms try to facilitate some form of communication across agents, or even the transfer of rule based information across individuals from generation to generation. How this information transfer is handled may be critical in market dynamics, since these information correlations cause eventual strategy correlations, which can translate into large price movements, and other features suggestive of a breakdown in the law of large numbers.

One final design issue is the creation of useful benchmark comparisons. It is very important to have a set of parameters for which the dynamics of the market is well understood. This demonstrates certain features in terms of learning dynamics and trading. An important benchmark might be the convergence to a well defined rational expectations equilibrium for certain parameters. The existence of such a benchmark further strengthens the believability of a computational market. The parameter sensitivities also shed which properties of markets lead them to an equilibrium, or far away from one. Finally, the learning dynamic going into the equilibrium may be revealing of exactly how the interactions between traders can operate to enhance the information processing capabilities of the overall markets. In other words, to understand exactly what the dynamics is at or near an efficient market equilibrium.

3 Artificial financial markets

It is easy to get lost in the many different types of models used in agent based financial markets. There are many different approaches that are used, and it is often difficult to distinguish one market from the next. This survey will take an initial stand on trying to categorize the many models that exist in a hope that this will help new researchers to better sort out what is going on the field. At such an early stage in the field it is still possible that some may argue about where markets are being put, or that some belong in multiple categories, or that the categories themselves are wrong. I think the potential help of these categories is well

[22] See Vriend (2000) for a description and examples.

worth these inevitable criticisms.

Most of the earliest models were intended to create an entire functioning financial market. They were often referred to as “artificial financial markets”. The next several subsections deal with different parts of this literature.

3.1 Small type models

Most of the earliest artificial financial markets carefully analyze a small number of strategies which are used by agents to trade a risky asset. The advantage of a small set of strategies comes in tractability, and for many cases these models are more analytic than computational. Many of these models follow the early lead of Frankel & Froot (1988), Kirman (1991), and De Grauwe, Dewachter & Embrechts (1993). In these papers it is assumed that there is a population of traders following two different types of strategies, labeled “technical” and “fundamental”. Technical traders are generally responsive to past moves in prices, while fundamental traders make decisions based on some perceived fundamental value. The relative numbers in the populations usually respond to past performance of the given strategies. The simplicity of these models makes them an important base case for the more complicated computational models which will be discussed later. The extensive set of analytic models will be discussed in Hommes (forthcoming 2004). Several of the small strategy models still require computational techniques to get their dynamics, and these will be discussed here.[23]

One of the earliest “small type” models is Figlewski (1978). This market examines the impact of shifting wealth across heterogeneously informed agents in a simple asset pricing framework. In this market agents make one large informational step that is removed in later markets. It is assumed that they know the wealth level of the other type of agent in the market. This is critical in forming price expectations across the two types. There is an efficient market benchmark, and many of the simulation runs converge to this, but with a slightly larger variance. Certain sets of parameters do not perform well in terms of convergence. Among these is the case where one set of agents has better information in terms of signal variance. In this case the simulated variance in the market is 14 percent larger than they efficient market benchmark. Actually, the simulations show that overall market efficiency might be reduced by the addition of traders with inferior information. Though this paper contains little information on the dynamics of prices and trades, it is still an important early reminder on how wealth dynamics effect the convergence to an efficient market.

[23]Several of the more important early papers in this area which will be discussed in Hommes (forthcoming 2004) are Beja & Goldman (1980) Brock & Hommes (1998), Chiarella (1992), and Day & Huang (1990). .

Kim & Markowitz (1989) are interested in the problem of market instability, and the impact that computerized strategies such as portfolio insurance may have had on the crash of 1987. Portfolio insurance strategies attempt to put a floor on the value of a portfolio through the use of a dynamic trading strategy. As the market falls, investors move holdings to cash to stop their losses. It is obvious that a market with many portfolio insurance related traders can be very unstable. Since the strategy is well defined, this allows for a simple computational test bed on their impact. The authors find that price volatility, trading volume, and the size of extreme price changes is increased as the fraction of portfolio insurance traders increases.

A more recent small type market is Farmer & Joshi (2002). The authors look at strategies which are carefully crafted to replicate a benchmark set of real trading strategies. These include trend following strategies, which increase or decrease their holdings of stock based on recent price trends. There are also value investors who decrease or increase their holdings of stock as the price moves above or below a fundamental value. Traders are allowed to take both long and short positions derived from these strategies. The market is cleared with a market maker who makes up for excess demands and supplies, and adjusts the price accordingly.[24] The authors simulate populations of traders following trend following and value strategies. The population contains a heterogeneous set of parameters for both types of strategies. The resulting dynamics display many features that resemble actual stock return series. They will be discussed further in the section on calibration.

3.2 Model dynamics under learning

The papers described in this section are much more computational than those mentioned previously. In most cases, the small sets of tractable trading rules are replaced with much larger sets of strategies which are usually represented using various computational techniques. This first section concentrates on applications where the economic environment is well understood, and there is often a simple homogeneous rational expectations equilibrium which takes the role of a useful benchmark.

3.2.1 Lettau(97)

Lettau (1997) is a good example of a computational model of this type. It implements a very simple financial market with a learning set of heterogeneous agents. It is simple, transparent, and easily implementable. The model is a portfolio choice environment where investors must decide what fraction of wealth to put in a

[24]This paper contains the best description of the market maker price setting mechanism, and also the best justification for its relevance to real market behavior.

risky asset. There is also a risk free asset paying zero interest. The world is a repeated two period model with myopic preferences based only on wealth in the second period. The price of the risky asset is given exogenously, and it pays a random dividend, d, which follows a normal distribution. The second period wealth of agents is given by,

$$w = s(d - p), \tag{2}$$

and their preferences are assumed to be constant absolute risk aversion which can be parameterized as in,

$$U(w) = -e^{-\gamma w}. \tag{3}$$

This is clearly a very simplified market. No attempt is made to look at the feedback from agents' demands to returns on the risky asset. There is no consumption, and wealth is not linked to agents' impact on asset prices, or evolution. However, it is a very simple test of learning in a financial market.

Given the normally distributed dividend process, there is a well known optimal solution to the portfolio problem given by,

$$s^* = \alpha^*(\bar{d} - p) \tag{4}$$

$$\alpha^* = \frac{1}{\gamma \sigma_d^2} \tag{5}$$

where σ_d^2 is the variance of the random dividend payout. The main exercise in Lettau's paper is to see if and when agents are able to learn this optimal portfolio strategy using a genetic algorithm. In general, agents' policy functions could take the form of

$$s = s(\bar{d}, p), \tag{6}$$

but Lettau simplifies this by using the optimal linear functional form for agent i,

$$s_i = \alpha_i(\bar{d} - p). \tag{7}$$

This gives the agents a head start on the portfolio problem, but they still need to learn the optimal α.[25]

The market is run for S periods with new independent draws of the dividend for each period. Each agent continues to use the portfolio determined by α_i which remains fixed. At the end of each block of S the

[25] A more complicated functional form is tried, but the results only change in that convergence is slowed down by the need to learn more parameters.

genetic algorithm is run, and the set of agent parameters is redrawn. Agents are parameterized with a bit string encoding given by

$$\alpha_i = MIN + (MAX - MIN)\frac{\sum_{j=1}^{L} \mu_{j,i} 2^{j-1}}{2^L - 1} \tag{8}$$

where $\mu_{j,i}$ is the bitstring for the strategy of agent i. The GA first gets a fitness value for each agent estimated over the S periods using

$$V_i = \sum_{s=1}^{S} U(w_{i,s}). \tag{9}$$

This sets the fitness to the ex post estimated expected utility over the sample. A new population is chosen using a technique known as “fitness proportional” selection. Each agent is assigned a probability using

$$p_i = \frac{1/V_i}{\sum_{j=1}^{J}(1/V_j)}. \tag{10}$$

Then a new population of length J is drawn from the old, with probability p_i assigned to each type. This new population is now the basis for the crossover and mutation operators in the GA. Each new rule is crossed with another rule chosen at random according to a fixed crossover probability. Crossover chooses a midpoint in the bit string and uses the first substring of one rule, and the second substring of another rule. This new set of rules is then mutated where each bit is flipped according to a fixed probability. In Lettau's framework the mutation rate is slowly decayed over time, so that eventually mutation probabilities go to zero. This is a form of cooling down the learning rates as time progresses. After mutation, the new population is ready to go back to purchasing the risky asset for another S periods before the GA is run again. This process continues for many time periods with the GA run after each S length sample goes by.

Lettau's results show that in various specifications the GA can learn the optimal parameter in the portfolio policy, however, there are some important caveats. First, the length of S is crucial. For example, for a value of $S = 25$ Lettau reports an average population α of 1.12 which corresponds to an optimal value of $\alpha^* = 1$. For $S = 150$ this goes down to $\alpha = 1.023$. It is not surprising that sample size matters, but this is a fact that can often be forgotten in more complicated setups where this choice is not as transparent. Also, Lettau's estimated α values are all biased above the optimal value. The intuition for this is clear for the case where $S = 1$. Expost it is optimal for s to be 0 or 1 only depending on the draw of d. Lettau sets the mean, $\bar{d}$, to a positive value, so that, on average, it will be better to hold the risky asset. This leads to an upward bias for the smaller values of S. In larger samples this bias dissipates as agents are better able to learn about the advantages of the diversified optimal strategy. This small bias is an important reminder on how learning diversified strategies can be difficult.

This is a very stylized, and simplified agent-based market. There is no attempt to model the price formation process at all. Therefore, this can never be viewed as a real attempt at modeling an actual market where the dependence between today's price and traders' strategies is the critical aspect of the agent-based modeling approach. However, it is a very clean and straight forward setup which is a good learning tool. Also, the biases and sample size issues that it brings up will also pertain to most of the much more complicated models that will be considered later.[26]

3.2.2 Arifovic(1996)

In Arifovic (1996) a much richer and extensive model is constructed. Once again, this paper stays close to a well defined theoretical framework, and is a kind of description of this model in the presence of learning agents. The model that is used is the foreign exchange model of Kareken & Wallace (1981). This is a two country, two period, overlapping generations model. Agents have income and consumption in both periods of their lives. Agents only means for saving income from the first to the second period of their lives is through either country's currency.

Agents maximize a two period log utility function subject to their budget constraints as in,

$$\begin{aligned} &\max_{c_{t,t},c_{t,t+1}} \log c_{t,t} + \log c_{t,t+1} \\ \text{st.} \qquad &c_{t,t} \leq w_1 - \frac{m_{1,t}}{p_{1,t}} - \frac{m_{2,t}}{p_{2,t}} \\ &c_{t,t+1} \leq w_2 + \frac{m_{1,t}}{p_{1,t+1}} + \frac{m_{2,t}}{p_{2,t+1}}. \end{aligned}$$

$m_{1,t}$ and $m_{2,t}$ are the money holdings of agents in the two currencies. There is only one consumption good, which has a price in each currency. The exchange rate is given by

$$e_t = \frac{p_{1,t}}{p_{2,t}}. \tag{11}$$

Given this setup, all agents care about in terms of money holdings is the relative returns of the two currencies. In an equilibrium where both currencies are held these returns must be equal.

$$R_t = \frac{p_{1,t}}{p_{1,t+1}} = \frac{p_{2,t}}{p_{2,t+1}}. \tag{12}$$

[26] Another simple example of this can be found in Benink & Bossaerts (2001).

It is also easy to show that the agents' maximization problem yields the following demand for savings.

$$s_t = \frac{m_{1,t}}{p_{1,t}} + \frac{m_{2,t}}{p_{2,t}} = \frac{1}{2}(w_1 - w_2 \frac{1}{R_t}) \tag{13}$$

The model has a fundamental indeterminacy in that when there are price series and an exchange rate which are an equilibrium, there will be infinitely many price/exchange rate series which are also equilibria. One of the interesting issues that Arifovic is exploring is whether the GA learning mechanism will converge to a single exchange rate. Sargent (1993) has explored this same question and found that certain learning algorithms will converge, but the final exchange rate depends on the starting value.

The multi-agent model is set up with a population of agents in each generation. Agents are represented with a bitstring which represents both their first period consumption decision, and the fraction of their savings to put into currency 1. 20 binary bits are used for the first value, and 10 binary bits are used for the second. These two values completely determine a period 1 agent's behavior through life. The price level in this model is determined endogenously. The agent bitstings determine their desired real savings in each currency which gives the aggregate demand for real balances in the two currencies. Nominal currency supplies are given, so this determines the price level in each currency. This setup avoids some of the complexities that appear in other papers in finding prices.

The evolution of strategies is similar to Lettau (1997). The fitness of a strategy is determined by its ex post utility, and a new population is drawn using fitness proportional selection. Agents are paired, and a crossover operator is applied to each pair with a given probability generating two new children. When crossover is not used, the children are direct copies of the parents. These children are then mutated by flipping bits with a certain probability. The fitness of the new rules is then estimated by implementing them on the previous round of prices and returns. At this point all four of the children and parents are grouped together, and the fittest two of this set are put into the next generation's population. This is known as the election operator which was first used in Arifovic (1994). It is designed to make sure that evolution continues to progress to higher fitness levels.

Arifovic analyzes the dynamics of this market for various parameter values. The results show that the first period consumption level converges to a stable value close to the optimum. However, the exchange rate continues to move over time, and never settles to any constant value. There is an interesting interpretation for this dynamic. In the equilibrium the return on the two assets is the same, so the learning agents are indifferent between holding the two currencies. Groups of agents move to holding one currency or another, and move the exchange rate around as they shift demands between currencies. In a a model such as this, it is clear that a constant exchange rate equilibrium can only be maintained through some mechanism which shuts down learning and exploration in the model. Arifovic also shows that similar features are obtained in experimental markets.[27]

3.2.3 Routledge(2001)

Another paper which critically examines the impact of learning on a well known model is Routledge (2001) which implements GA learning in a version of the heterogeneous information model of Grossman & Stiglitz (1980). This is a repeated version of a model where agents can purchase a costly signal about a future dividend payout of a stock. Learning takes place as agents try to convert the noisy signal into a forecast of future dividends. Agents who decide not to purchase the signal must use the current price to infer the future dividend payout. Individual agent representations encode not just the decision on whether to purchase the signal, but also the linear forecast parameters which convert the signal into a conditional expectation of the future dividend payout.

Grossman & Stiglitz (1980) show that there is an equilibrium in which a certain fraction of agents will purchase the signal. Routledge (2001) shows that this can be supported in the GA learning environment. However, there are also sets of parameters for which the original equilibrium proves to be unstable. The dynamics of this instability is very interesting. There is instability and exploration going on around the equilibrium, and by chance a few more informed agents may enter the market. The change in market proportions of informed versus uninformed agents means that the current linear forecast parameters are now wrong. In particular, the uninformed need to learn how to interpret the price with fewer of their type around. Unfortunately, as the number of uniformed falls, the ability of their population to learn decreases due to small sample sizes. These sets of parameters can often converge to the case where everyone is informed.

This is a very interesting learning situation which involves the computational side of the model in a crucial way since there is an issue of small sample sizes in the learning process. It is often nice to think about stylized models where the populations are all taken to infinity, but this may gloss over some of the critical problems in learning imitation. How do agents learn to correctly do something when they do not have many examples around to follow? This may be a question which requires a computational approach to be addressed correctly.

[27] These results are discussed in Duffy (forthcoming 2004).

3.3 Emergence and many agent models

The next set of artificial markets move farther from testing the dynamics of specific models, and more toward understanding which types of strategies will appear in a dynamic trading environment. All have at their core a philosophy of building a kind of dynamic ecology of trading strategies, and examining their coevolution over time. This methodology attempts to answer the basic question of which strategies will survive, and which will fail. Also, one observes which strategies will emerge from a random soup of starting strategies, and which are capable of self-reinforcing themselves, so that survival is possible. They also attempt to perform a very direct exploration into the dynamics of market efficiency. If the market moves into a state where certain inefficiencies appear, then the hope is that the evolutionary process will find new strategies to capitalize on this. The objective is to explore a market that may not be efficient in the textbook sense, but is struggling toward informational efficiency.

3.3.1 The Santa Fe artificial stock market

The Santa Fe Artificial Stock Market is one of the earliest in this set of models. It is described in Arthur, Holland, LeBaron, Palmer & Tayler (1997), and also LeBaron, Arthur & Palmer (1999).[28] It consists of traders evolving strategies using a classifier system which is designed to implement short range forecasts for stock price movements. The basic economic structure of the market draws heavily on existing market setups such as Bray (1982) and Grossman & Stiglitz (1980).

The trading agents have one period myopic preferences of future wealth with constant absolute risk aversion. There are two assets that agents trade in the market, a risky stock, paying a random dividend, d_t, and a risk free bond, paying a constant interest rate, r. The dividend follows an autoregressive process as in,

$$d_t = \bar{d} + \rho(d_{t-1} - \bar{d}) + \epsilon_t, \tag{14}$$

where ϵ_t is gaussian, independent, and identically distributed, and $\rho = 0.95$ for all experiments. It is well known that under CARA utility, and gaussian distributions for dividends and prices, the demand for holding shares of the risky asset by agent i, is given by,

$$s_{t,i} = \frac{E_{t,i}(p_{t+1} + d_{t+1}) - p_t(1 + r)}{\gamma \sigma^2_{t,i,p+d}}, \tag{15}$$

[28] There is also an earlier version of the SFI market which is described in Palmer, Arthur, Holland, LeBaron & Tayler (1994). This market has one crucial difference with the later market in that it implements a excess demand price adjustment mechanism.

where p_t is the price of the risky asset at t, $\sigma^2_{t,i,p+d}$ is the conditional variance of $p+d$ at time t, for agent i, γ is the coefficient of absolute risk aversion, and $E_{t,i}$ is the expectation for agent i at time t.[29] Assuming a fixed number of agents, N, and a number of shares equal to the number of agents gives,

$$N = \sum_{i=1}^{N} s_i \tag{16}$$

which closes the model.

The model includes an important benchmark to compare results with. There exists a linear homogeneous rational expectations equilibrium in which all traders agree on the model for forecasting prices and dividends. In the equilibrium it is easy to show that the price is a linear function of the dividend, which is the only state variable.

$$p_t = b + ad_t. \tag{17}$$

The parameters a and b can be easily derived from the underlying parameters of the model by simply substituting the pricing function back into the demand function, and setting it equal to 1, which is an identity and must hold for all d_t.

The most important part of this market is its implementation of learning and forecasting. This is done with a classifier forecasting system which is a modification of Holland's condition-action classifier,(Holland 1975, Holland, Holyoak, Nisbett & Thagard 1986). It maps current state information into a conditional forecast of future prices and dividends.[30] Current market information is summarized as a bitstring of ones and zeros. This bitstring represents several key variables that many traders use. It is summarized in the following list.

1-7 Price*interest/dividend $> 1/2, 3/4, 7/8, 1, 9/8, 5/4, 3/2$

8 Price > 5-period MA

9 Price > 10-period MA

10 Price > 100-period MA

11 Price > 500-period MA

12 always on

13 always off

[29]It is interesting to note that assumptions behind this specification do not hold in the SFI framework. The expectation is not the true expectation, but is a predictor that comes from the learning framework. This substitution is made often in many of these learning models. Second, the demand specification depends on the distribution of future prices being Gaussian. This may not be the case in the SFI market, so the demand function can only be viewed as an approximation to the utility maximization problem. The accuracy of this approximation has never been tested.

[30]Classifiers are not used extensively in economic modeling. Some other examples are Marimon, McGrattan & Sargent (1990) and Lettau & Uhlig (1999).

The first 7 bits are based on the price/dividend ratio, and are often referred to as “fundamental bits”. Bits 8 through 11 are based on moving average technical trading rules, and are often referred to as “technical bits”. The price is compared to a moving average (MA) of past prices. These are based on standard technical trading methods used by traders. It is important to see that this table preloads a lot of structure into the classifier forecasts before the market has even started.

A forecasting rule consists of two parts. First, there is a string which matches up with the above 13 bits. It consists of 0’,s 1’s, and #’s. A rule is matched when it lines up with the current state vector of the economy. An interesting feature of the classifier is the # symbol, which is a wild card. This matches either a 0 or a 1 in the state vector. For example, the rule 00#11 would match either the string 00111, or 00011. An all # rule would match anything. This allows for the classifier to endogenously determine what the relevant state variables of the economy are, and which can be ignored. In standard classifier systems there is a determination made on which is the strongest rule depending on past performance, and the rule then recommends an action. In the condition-forecast classifier the second part of each rule is a real vector of forecast parameters, $a_{i,j}, b_{i,j}, \sigma^2_{i,j}$ which the agent uses to build a conditional linear forecast as follows,

$$E_{t,i,j}(p_{t+1} + d_{t+1}) = a_{i,j}(p_t + d_t) + b_{i,j}. \tag{18}$$

This expectation along with the variance estimate, $\sigma_{i,j}$ allows the agent to generate a demand function for shares using equation 15. These demands are linear in the current price which allows a market auctioneer to set an equilibrium price each period.

All matched rules are evaluated according their accuracy in predicting price and dividends. Each rule keeps a record of its squared forecast error according to,

$$\sigma^2_{t,i,j} = \beta\sigma^2_{t-1,i,j} + (1 - \beta)((p_{t+1} + d_{t+1}) - E_{t,i,j}(p_{t+1} + d_{t+1}))^2 \tag{19}$$

This estimate of the conditional variance of a forecast is used both for share demand, and to determine the strength of the each rule when they are evolved. Agents individually maintain a set of 100 forecasting rules, and there is no overlap between agents of specific forecasting information.

Learning takes place at random intervals according a given probability in which agents update rules every K periods on average. Learning takes place with a modified genetic algorithm designed to handle both the real and binary components of the rule sets. The worst performing 15 percent of the rules are dropped out of an agent's rule set, and are replaced by new rules. New rules are generated using a genetic algorithm with uniform crossover and mutation. For the bitstring part of the rules, crossover chooses two fit rules as parents, and takes bits from each parent's rule string at random.[31] Mutation involves changing the individual bits at random. Crossing over the real components of the rules is not a commonly performed procedure, and it is done using three different methods chosen at random. First, both parameters, a and b, are taken from one parent. Second, they are each chosen randomly to come from one of the parents. Third, a weighted average is chosen based on strength.

One of the objectives of the SFI market was to examine the dynamics of learning, and to explore its likelihood of convergence to an efficient market equilibrium. Experiments are performed for two values of the learning rate. A fast learning rate sets the average time between runs of the GA to $K = 250$, and a slow learning experiment sets the average time between runs to $K = 1000$. In the latter case, the market converges to the benchmark rational expectations equilibrium, where all agents agree on how to process the fundamental dividend information, and, for the most part, they ignore all other information. Other experiments test a fast learning situation where the GA is run on average every $K = 250$ periods. For this parameter a very different outcome occurs. The market does not appear to converge, and it shows several indications of interesting features in the stock return time series. Among these are nonnormal return distributions, or "fat tails", persistent volatility, and larger amounts of trading volume than for the slow learning case. All of these are elements of the empirical puzzles mentioned in the early sections of this chapter. Though the SFI market does a good job in replicating these facts qualitatively, no attempt is made to quantitatively line them up with actual financial data. Indeed, the SFI market never even clearly states what it considers to be the frequency of the returns series that it generates, or whether the underlying dividend process is realistic.

The SFI market is has formed a platform for other explorations. Joshi, Parker & Bedau (2000) explored the interactions between the technical and fudamental traders. They find that the use of technical trading bits is a dominant strategy in the market. If all other traders are using technical bits, then it would be in the interest of new agents to use them too. Also, if all other agents are using fundamental bits only, then it is optimal for the new agent to add technical bits as well. This strongly suggests that trend following behavior may be difficult to remove from a market. The most sophisticated addition to the SFI classifiers is in Tay & Linn (2001) who replace the classifiers with a fuzzy logic system.

[31] Selection is by tournament selection, which means that for every rule that is needed two are picked at random, and the strongest is taken.

The SFI market has generated much interest since its software is now publicly available. It was originally written in c, and then objective-c, and finally ported to the Swarm system. Johnson (2002) gives an overview and critique of the software from a design perspective, and Badegruber (2003) is an extensive replication and reliability study. It is fair to summarize that the software is not easy to read or use. Much of this stems from its long history on several different platforms. Also, it began before objective languages were popular, and was only adapted to objective form in its later versions. It was not built to be an objective piece of code from the start.

Beyond software critiques there are also important design issues that could be improved in the market. Many of these are covered in LeBaron (forthcoming 2004). Qualitatively, the market and classifier systems have proved to be very complicated and unwieldy in terms of understanding the market dynamics. Many parameters are needed to define the operation of the classifier, and it not clear which of these are important. Also, the information setup of the classifier is often criticized. It can create many rules that will never get matched. To account for this, a system known as a generalizer is used to change zero or ones bits to the # wild card, but this system seems very ad hoc. Ehrentreich (2003) addresses this with a system that does not allow rules that wont get implemented, and also changes the mutation operators. These changes appear to be critical to the dynamics and convergence of the market. Another major question that is left unanswered is how important the definition of the bitstring is to the dynamics of the market. These bit information values are obviously preloaded. Finally, another important critique is that by using CARA utility the market ignores the wealth dynamics of agents. In other words, it is not the case that wealthier agents have a greater impact on prices in the SFI market.

3.3.2 Chen and Yeh(2001)

If the general goal of the markets in this section is to see strategies form out of a general set of functional building blocks with little structure entered initially by the designer, then then the model of Chen & Yeh (2001) is probably the best example of this. These authors use a computational methodology known as genetic programming. This allows the agents to evolve actual predictor functions for financial forecasting.[32]

The economic set up in Chen & Yeh (2001) is similar to the SFI market except for the fact that the price adjustment is made using excess demands as in Palmer et al. (1994). Also, demands are based on the forecast of future prices and dividends. This is where genetic programming is implemented. The forecast

[32] There have been some implementations of this technology on actual data as in Neely, Weller & Dittmar (1997) and Allen & Karjalainen (1998). The origins of genetic programming go back to Koza (1992).

takes the form of

$$E_{i,t}(p_{t+1} + d_{t+1}) = (p_t + d_t)(1 + \theta_1 \tanh(\theta_2 f_{i,t})), \tag{20}$$

where $f_{i,t}$ is evolved using genetic programming. It takes as inputs $p_{t-j} + d_{t-j}$ for $j = 1, 2, \ldots, 10$.

A second important innovation is the use of a common pool of rules which the authors refer to as a "business school". This allows for some strategy learning to occur across agents in a very natural way.[33] The rules in the common pool are evolved according to forecast accuracy. Agents then decide whether to update their own strategies based on current performance. When they decide to update they begin drawing rules from the common pool, and compare the performance with their current rules. If the new rule is better they switch, but if they are unsuccessful in finding a better rule after several tries they quit, and stay with their current rule.

Simulations of the this market display some features of actual return time series. They exhibit fat tails, and visually they do not settle down to any price level. However, there are several features that disagree with the actual data. Among these are a large level of positive skew, and also the linearly filtered return series are independent, which indicates there may be no persistent effects in volatility. Another interesting feature that the authors test for is a unit root in the price series. The standard tests cannot reject a unit root. This is a little curious since the dividend process is stationary. It is probably sensible that in the long run prices should not diverge too far from the fundamental, and therefore should also be stationary.

3.3.3 Beltratti and Margarita(1992)

The market model presented in Beltratti & Margarita (1992) and in Beltratti, Margarita & Terna (1996) is quite different from the other markets described here. It is again searching for an emergent pattern in the trading behavior of adaptive agents. However, unlike the previous markets, this one has no organized central trading institution. Agents trade in a completely disaggregated fashion in a market where they randomly bump into potential trading partners. This is similar to structures such as Albin & Foley (1992).

The traders build a forecast of what they think the stock is worth using past information and an artificial neural network. The network builds a forecast of the following form,

$$E_{i,t}(p + t + 1) = f(p_{i,j,t-1}, \Delta p_{i,j,t-1}, \pi_{t-1}, \Delta \pi_{t-1}) \tag{21}$$

where π_{t-1} is the average transaction price at time $t-1$ across all agents, $p_{i,j,t-1}$ is the agents last price

[33] This is the recurring theme of social learning from Vriend (2000).

execution that it received, and Δx refers to the one period change in x. This is an interesting function because the traders are using both local and global information When traders meet they compare price forecasts. The one with the larger forecast price will purchase 1 share from the trader with the smaller price. The trade is executed at the simple average of the two prices. The market keeps track of the average execution price across the random pairings, and this is used in the information sets of the traders. After a day of trading, agents are allowed to update the weights of the neural networks in a direction which improves forecast accuracy.

Beltratti et al. (1996) present many experiments with this basic structure. One of the more interesting explorations tackles the problem of heterogeneous agents with differing levels of complexity. This is covered in Beltratti & Margarita (1992). The population consists of different neural network structures. Trader sophistication is represented by more complicated neural networks. The more complicated structure comes at a given higher complexity cost that is paid directly by the traders, c. The simulations show the eventual heterogeneous population depends critically on the value of c. For low levels of c, agents purchase the extra network complexity, and for high levels of c, they eventually only use the simple networks. There is an interesting mid range where both types of strategies are able to coexist.[34]

3.3.4 Youssefmir and Huberman(1997)

Although not technically a financial market per se, Youssefmir & Huberman (1997) add important insights to the question of emergent phenomenon in financial markets, and in particular the issue of volatility clustering. They analyze a simple resource allocation model where agents are able to choose between various resources. The payoff of each resource depends on the number of agents using it. In the competitive environment that they analyze, agents payoffs are decreasing in the number of agents using the resource. The connection between this multi-agent system and financial markets is that agents are allowed to use various strategies to optimize their behavior. They chose strategies which have done well over the recent past, in terms of payments from the resource extraction problem. The choices are determined with some amount of noise, so fluctuations between strategies with similar payoffs are possible.

The important result in Youssefmir & Huberman (1997) is that they provide another mechanism for generating volatility clustering. Also important is the fact that their model is simple enough for them to do a very complete job of analyzing just what is causing the persistence in volatility. The model has a stable Nash equilibrium in the resource choices. When the system is close to this equilibrium, the payoffs for large

[34]There are interesting connection between this model and Brock & Hommes (1998).

subsets of the strategies are the same. In the random choice setup of the model, when an agent is faced with new strategies equivalent to their current ones, they are just as likely to switch as to stay put. This causes a random walk in strategy space as the population moves around through different strategy choices. For certain parameters, this movement can be large enough that the system can be drawn out of the equilibrium region, and can restart large fluctuations in the agent switching decisions.

The clustered volatility in this model is reminiscent of financial markets. It also seems similar to the dynamics in Arifovic (1996). In both cases there is an equilibrium to which the model converges, and in this state of the world various strategies become equivalent, and agents begin exploring this space. As the population spreads out enough the dynamics of the system are changed, and the local stability breaks down. Youssefmir & Huberman (1997) can be criticized for not being a financial market, but the authors are cautious about this, and it is hoped that the social dynamics they generate will carry over into more detailed financial settings. A more important critique might be to question the decision making process of agents. There is a need for agents to continually update their choices, and they will switch to an equivalent rule with some positive probability when it is offered to them. An interesting modification is to consider what happens when you only switch when some improvement in performance is possible. This would stop the random fluctuations in strategies at or near the equilibrium, and might lead to a more general convergence result.

3.3.5 Chakrabarti and Roll(1999)

All of the papers reviewed in this survey, and almost all agent-based markets in general, model individual traders as small. They view themselves has having no price impact, and feel that there is little information to be gained by observing other individual's trades. Chakrabarti & Roll (1999) is an interesting exception to this. These authors model an information acquisition process where agents observe other large traders in market, and adjust their own beliefs based on the observed actions of others. This is in the spirit of other sequential trading models such as Welch (1992).

The individual agents receive a signal each period, and also observe the trades of others as they go by in the market. Their own trading strategies are based on optimally forecasting the final payment of the security using Bayesian updating on their initial priors. Though the individual strategies are analytically defined, the final dynamics of the market as a whole requires a computational experiment. To explore the impact of many different parameters the authors employ a novel approach. They run many simulations at randomly chosen parameter values, and record various results. To analyze all this data, they run multiple linear regressions on the parameters, to observe their impact. This may seem like a lengthy and indirect method to understand parameter sensitivity, but it may be important when there are many parameters, and the interactions between parameters are not well understood.

The authors analyze many properties of the market including price volatility, and price prediction error or tracking. An interesting result is that when signal diversity increases, price volatility increases, but also the price is a better forecast of future value. This implies that some markets can have more trading activity which leads both to greater movements in the price, but also better learning and information sharing of signals. This should remind policy makers that simple measures of volatility alone may not always be a good measure of market quality. Other interesting results include the fact that a more diffuse prior on the value of the stock can lead to better learning in the market. This is because when the traders believe less in their initial information, they have greater incentive to try to glean knowledge from the better informed market as a whole. Their model allows for another interesting experiment. The size in which a trade gets noticed by others is a parameter of the model. Trades which are smaller than this value go unnoticed, but the larger ones are observed. The authors find that reducing this threshold reduces price volatility, and increases forecast accuracy. This is again very suggestive that in the end the learning processes in this sequential market is effective, if not perfect.

3.4 Calibration and validation

The markets discussed in this section emphasize the replication of many of the empirical puzzles that were mentioned at the beginning of this chapter. In each case the agent based model itself is less important than the replication of various empirical results from financial market time series.

3.4.1 Levy, Levy, and Solomon (1994)

Levy, Levy & Solomon (1994) presents a market where the outcome of agents strategies is also emergent.[35] Similar to the earlier markets the outcome depends on the presence of many different heterogeneous agent types. However, it differs in not having them form complicated strategies and predictors. Traders maximize a one period myopic utility function which is now constant relative risk aversion. This technical change is important in that now agents' impact on prices depend on their relative wealth levels.

The economic foundations of the model are similar to other artificial financial markets. There is a risk

[35] This model is presented in the book, Levy et al. (2000) which also contains useful summaries of many other agent based markets.

free asset which pays a constant interest rate. There is a risky stock paying a random dividend which follows a multiplicative random walk,

$$d_{t+1} = d_t(1 + z_{t+1}). \tag{22}$$

where z_t is drawn from a well defined distribution designed to roughly replicate actual dividend growth.

The market consists of several types of traders. There are rational traders who posses a model for pricing the stock based on the dividend fundamental. They use this to predict the future price, and to then set their optimal portfolio fraction accordingly. A second, and more important, type for this model uses past information only to determine its current portfolio. This trader looks at the past m periods of returns, and finds what fraction of stock and bond holdings would have been optimal over this period. This is a kind of memory length for traders. It allows for some to believe that only a short period of the past is necessary for forecasting, and others to a believe a much longer series is necessary. The short memory types represent a kind of short term trader who is only interested in the latest fads, and believes the older returns data are irrelevant. The memory length history of past returns is used to make a portfolio recommendation for the next period. There is often a population of these traders with many different memory lengths, m_i.

The authors progressively add richer sets of the heterogeneous memory traders who trade along side the fundamental traders. For groups of only one, or two memory types, the stock price dynamics clearly reflects the memory length, in that it displays distinct cycles. However, when a full spectrum of these traders is added, the prices show no perceptible cycles, and display very realistic features. The returns show relatively large positive autocorrelations at shorter horizons, and negative autocorrelations at longer horizons. The authors suggest that this is representative of actual markets where it has been shown that at the short horizon stock returns demonstrate some positive autocorrelations, which shift to negative at longer horizons.[36] Many empirical aspects of the model are explored including large amounts of trading volume, and its positive correlation with volatility. The market also is capable of endogenously generating market crashes. The authors are also very concerned with the coexistence of both the rational strategy, and the finite memory strategies. They give some examples showing the two types can coexist with neither one evolutionarily driving the other out.

The model has been criticized recently by Zschischang & Lux (2001). These authors claim that some of the original results are sensitive to the initial conditions in the simulation. They further indicate that the results may be sensitive to the number of agents in the simulation. This critique is interesting, but it was

[36]See Hirshleifer (2001) for summaries of these empirical results.

done for a set of only three different memory lengths of traders, $10, 141, 256$. It remains to be seen if it has implications over more general distributions of memory length.

3.4.2 Volatility Long Memory

One of the most interesting features that various artificial markets try to replicate is the persistence of volatility. While stock returns themselves are relatively uncorrelated, the squares of absolute values of returns are highly correlated. Positive correlations continue out a year or more, and decay at much more of a hyperbolic rate, rather than an exponential rate. This is a possible indication of a long memory of fractionally integrated time series process for volatility. The mere fact that volatility is persistent is puzzling enough, but the fact that it also possesses long memory like properties presents a high hurdle for agent-based financial markets to hit in terms of empirical replications.

The model of Iori (2002) is interesting both in its structure, and in its ability to fit the facts. The model has an Ising model like structure in the dissemination of information and strategies across agents.[37] In this model traders receive a signal that combines information on choices of local neighbors as in

$$Y_{i,t} = \sum_{(i,jj)} J_{i,j} S_{j,t} + A\nu_{i,t}. \tag{23}$$

$S_{j,t}$ are the decisions of other traders in the neighborhood of i, and $J_{i,j}$ controls the weighting and the neighborhood size. $\nu_{i,t}$ is a noise term. This signal is an input into a trader's final decision to purchase or sell one share of the stock. The interesting part of this decision is that agents have a range of inaction on the signal. For $-w_t < Y_{i,t} < w_t$ there is no trade by agent i, and $S_{i,t} = 0$. When the signal is less than $-w_t$, the agent sells one unit, $S_{i,t} = -1$, and when the signal is greater than w_t the agent buys one unit, $S_{i,t} = 1$.

It is clear that the decisions of agent i again feed into the signals of others. The belief formation and demand part of this model is iterated several times until there is convergence. Then the demands to buy and sell shares are calculated as the number of positive and negative values of $S_{i,t}$ respectively, and are recorded as D_t and Z_t. There is a market maker who covers the order imbalance, and adjusts the price using

$$p_{t+1} = p_t(\frac{D_t}{Z_t})^{\alpha}. \tag{24}$$

Stock returns are measured as the log difference of this price series, and the volatility is estimated with

[37] Other examples of this will be discussed in Hommes (forthcoming 2004). A few of these are Cont & Bouchaud (2000) and Stauffer & Sornette (1999).

the absolute values of these returns. The model generates returns which are nearly uncorrelated, but the volatility series generates close to long memory like behavior with a long hyperbolic decay pattern in the autocorrelations. Further, the model is also able to display the strong positive correlation between trading volume and volatility which is observed in the data. It also appears that the threshold part of the model is critical for volatility clustering to occur.

Kirman & Teyssiere (2001) is another model capable of generating long memory like effects in return volatility. It is a modified version of Kirman (1991) which is described more extensively in Hommes (forthcoming 2004). It is a form of the earlier mentioned small rule models where agents follow a finite set of well defined portfolio rules. These are defined as technical and fundamental, and the traders shift back and forth between these according to an epidemiological process of contagion. The authors perform extensive tests on the long memory properties of the model volatility as compared to actual volatility in several foreign exchange series, and find a good quantitative alignment.

3.4.3 Calibrating to macro fundamentals

Several papers have taken the step of trying to tie markets to actual market fundamentals. In **?** the authors use U.S. aggregate real dividends interpolated to daily frequencies as a fundamental input into market with heterogeneous value investors and trend followers. Their market generates reasonable long swings away from the dividend process along with uncorrelated daily returns. It also generates most of the important features described in the first section including, fat tails, volatility persistence, and trading volume persistence. As mentioned earlier the model offers interesting tractability since it is built from a foundation of realistic trading strategies.

LeBaron (2001*a*) and LeBaron (2002) perform some extensive calibration exercises. They are based on an agent-based framework presented in LeBaron (2001*b*). This model combines several features of the models mentioned previously. It is uses a neural network structure to represent agent portfolio strategies. In this model agents do not build forecasts. The neural network maps past information into a recommended portfolio holding directly, and avoids the intermediate step of mapping a forecast into a portfolio policy. It also avoids having to estimate the return variance using a separate volatility equation. Agents are defined by heterogeneous memory lengths as in Levy et al. (1994). Some agents evaluate strategies using a small past history of returns, while others use longer histories. Also, the preferences for the agents are constant relative risk aversion, so agents with more wealth control a larger fraction of the market. The strategy population is separate from the agents, and they chose strategies which are optimal based on their history length. This has some similarities to the social learning mechanisms in Chen & Yeh (2001). The strategies are evolved using a modified genetic algorithm designed to respect the neural network architecture. Finally, the economic structure is similar to many previous markets, in that there is a risky and risk free asset that are traded. The risky asset pays a well defined stochastic dividend following a geometric random walk which with drift and volatility set equal to the values from real post war aggregate U.S. dividends. The time period in the model is set to 1 week.

The model is compared with values drawn from the S&P 500, and it is able to replicate a large range of features quantitatively. These range from simple statistics such as means and variances of returns, to the more complicated dynamic features of long memory volatility persistence, and volatility/volume cross correlations.[38]

Bullard & Duffy (2001) introduce learning into a more traditional macroeconomic framework for asset prices. The model is a multiperiod overlapping generations setup with a constant returns to scale aggregate production technology. Also, important is the fact that there is a government issuing money in the economy at a constant growth rate which is greater than the growth rate of the economy. Therefore, the forecasting of inflation and real returns becomes an important problem for agents in this economy. They forecast future price levels using a recursive regression framework. This learning mechanism yields excess volatility in the asset market. The authors perform a search over their parameter space using a genetic algorithm to find parameters generating results similar to actual data. They find parameters which are able to give them reasonable volatility in asset returns along with the low volatility in per capita consumption growth. For the most part, the parameters that generate these results are consistent with U.S. macro economic data.

3.4.4 Foreign exchange forecasts

Some of the most puzzling facts in finance come from foreign exchange markets. Empirical research on forecast surveys suggest that professional foreign exchange forecasts are biased estimates of the future spot rate.[39] In a recent paper, Marey (2004), tries to calibrate heterogeneous agent models to these forecast survey results. Using a stylized foreign exchange market with an analytic market clearing which is similar to Kirman (1993) the author explores the dynamics of markets with heterogeneous forecasts built on specific functional forms corresponding to different rule of thumb forecast types. Many of these types populate the earlier mentioned "small type" models as well.

[38] An interesting feature is that the model replicates the tendency for volatility to generally lead trading volume. This is consistent with results in Gallant et al. (1993).

[39] Frankel & Froot (1987) is an early example.

The general results in terms of empirical replication are bad for markets containing many types of forecasting agents. Models with extrapolative forecasts and distributed lag forecast fare poorly, as do bandwagon types of expectations. The simulated markets are evaluated in terms of how well they replicate various macro features in foreign exchange markets. In the end, the best performance comes from markets with both extrapolative and regressive forecasts operating simultaneously. Regressive forecast predict a return to some kind of fundamental value. It is interesting that this is often the set of forecasts which is implemented in many of the "small type" models. This paper is an important empirical defense of this conjectured picture of the world.

3.4.5 Other empirical approaches

Several other papers look at empirical features of some of the simpler agent-based markets. Arifovic & Gencay (2000) perform extensive time series tests on data generated from the model in Arifovic (1996). They present strong evidence that the generated returns series are not just nonlinearly dependent, but are most likely chaotic. The latter conclusion is based both on Lyapunov exponent estimates along with phase diagrams on the dynamics. From the standpoint of understanding the model these results are very interesting. However, from the standpoint of actual data they do diverge from actual data where it is difficult to see enough structure for convincing evidence for chaos, Brock, Hsieh & LeBaron (1991). A similar set of time series diagnostics is performed in Chen, Lux & Marchesi (2001) on data generated in the model of Lux (1997). They obtain similar, but slightly more complicated results. There is strong evidence for volatility persistence in all the series. However, tests for independence depend on the different subsamples used. These results are actually inconsistent with each other, and suggest differences in the power of the tests for the sample sizes used.

All of the results cited so far have been concerned with empirical replication of stylized features. Few of these have attempted to formally estimate parameters. A recent exception to this is Winker & Gilli (2001) where the authors estimate parameters in the Kirman (1991) model. They search over two parameters in the model with an objective of fitting two features of actual financial returns, kurtosis, and the first order volatility coefficient in an ARCH(1) specification. Since the search space is relatively simple as well as the objective, this paper provides the most detailed view into the sensitivity of the results to various parameter specifications.

3.4.6 Currency crises

Arifovic & Masson (1999) are interested in a very particular form of emergence, the emergence of currency crises. Their model is designed to replicate and better understand periodic capital crises in developing countries. Recent history has shown developing markets to be susceptible to periodic crises, and large capital outflows, leading to significant impacts on overall macroeconomic well being. As in many other cases in finance, the exact mechanism leading to these dynamics is not completely understood. Arifovic & Masson (1999) is the first paper to tackle this problem in a completely agent-based setup. This approach is justified in that models of financial crises, and financial contagion are often characterized by multiple equilibria. Also, they can be sensitive to how the model is set up, and to the information process that is used.

Building on a model by Masson (1999), Arifovic and Masson examine the behavior of a developing country's capital market. Agents are modeled as real valued pairs $(\pi_t^i, \delta_t^{e,i})$ which represent the probability of a devaluation, and the amount of this devaluation respectively. They then make simple risk neutral portfolio decisions between developed and developing country debt. They compare the returns of developed country debt, $1 + r^*$ with the expected returns of developing country debt which is given by

$$\frac{1 + r_t}{1 + \pi_t^i \delta_t^{e,i}}. \tag{25}$$

Given that each agent is risk neutral, they invest 100 percent of their wealth in the higher expected return asset using their own probability assessments. The investors are assumed to have a constant wealth level, $\bar{w}$, so the amount invested in the developing country is given by

$$D_t = \sum_{i=1}^{n} \lambda_t^i \bar{w} \tag{26}$$

where λ_t^i is the fraction of wealth invested in the developing country which is zero or one. The key question in this model is how to clear the market for developing country debt. It is done with a kind of interest parity equation which sets the return on developing country debt equal to a weighted average of the agents' beliefs.

$$1 + r_t = (1 + r^*)\Pi_{i=1}^n (1 + \pi_t^i \delta_t^{e,i})^{1/n} \tag{27}$$

Once r_t is set, then agents' demands for developing country debt can be determined along with aggregate demands. The dynamics of the country's reserves are then determined by,

$$R_t = R_{t-1} + T_t + D_t - (1 + r_{t-1})D_{t-1}, \tag{28}$$

where R_t are the reserves at t, and T_t are exogenous trade flows which are assumed to follow and autoregressive process.

While R_t stays positive there is now devaluation, but when $R_t < 0$ there is a crises, and the currency devalues. The devaluation amount is determined to reduce the amount of debt so that the reserves can be brought back to zero. This sets it at

$$1 + \delta_t = \frac{-R_t}{D_t} = \frac{(1 + r_{t-1})D_{t-1} - R_{t-1} - T_t - D_t}{D_t}. \tag{29}$$

Once it is determined whether there was a devaluation, and by how much, agents' performance can be determined. Their fitness is simply $\mu_t^i = r^*$ if they invested in the developed country, and

$$\mu_t^i = \frac{1 + r_t}{1 + \delta_t} - 1 \tag{30}$$

for the developing country which is the expost return. At this point the agents are evolved. Evolution takes place using a very simple mechanism involving imitation and mutation only. Agents first compare their rule with another drawn from the population according to the following fitness determined probability

$$p_i = \frac{\mu_t^i}{\sum_{i=1}^n \mu_t^i}. \tag{31}$$

If the fitness of this other agent exceeds their own fitness, then the agent switches to the other's beliefs. Otherwise, it keeps its beliefs. The final step is a kind of mutation like step called experimentation. At random, agents will experiment, and change their values of π_t^i or δ_t^i to some new value drawn from a well defined distribution.

Initial values in the simulation are chosen from an actual currency crises, Argentina at the end of 1996. The dynamics of the market begins with no devaluations. As agents continue to get feedback on no devaluations, the strategies with lower π_t^i survive and thrive, more capital flows into the country. These flows drive r_t down. Eventually, r_t falls so far that some investors begin moving funds back out of the country. At this point, one of two things can occur. If R_t is low then as investors shift out of the country, D_t can fall fast enough to trigger a crisis. After the crises, $\delta_t^{e,i}$ and π_t^i both increase in the population driving r_t, and restoring capital inflows. In the second case, the R_t can be large enough, so the country avoids a crisis, and the high π_t^i investors were wrong. The higher returns in the developing country cause π_t^i to fall again as the population slowly decides that it is safe to invest.

The model is an interesting agent-based version of a foreign exchange crises. The bandwagon effect of beliefs is modeled in an evolutionary framework. Finally, similar to other markets in this section, the model is aligned with some actual data.

3.5 Model Validation

Fitting agent-based models to empirical results has often come under attack. It is claimed that the large set of parameters available for tweaking can be used to fit almost any feature that is desired. This may or may not be exactly true. The problem is that while there may be many parameters, the actual effective degrees of freedom is not clear in many of these models. A second issue is that in finance these models are clearly trying to break new ground in terms of empirical features. As was mentioned in the introduction, they are trying to go places where traditional models have remained silent. An agent-based model replicating the equity premium alone would not be very interesting. However, most experiments attempt to replicate features such as volatility persistence, fat tails, and trading volume, which up to now have remained off the mainstream model agenda. They are trying to also estimate features at many different time horizons. Finally, some results, such as long memory, are specifically concerned with the interconnections of results across many different time scales. These features may be very powerful in helping to pin down viable agent-based representations.

There are are also ways to lessen some of these problems. One is try to endogenize some of the parameters, or to put them under evolutionary control. An example of this is LeBaron (2002) which takes the crucial parameter of memory lengths, and turns it into a problem in evolution. Agents are allowed to exist with many different parameter values, and evolution determines which will end up dominating. This is in contrast to presetting a single, or group of memory lengths ex ante. Another possibility, is to concentrate on periods when the market is subject to stress as in high volatility, or near a crash. It is possible that the dynamics of these periods in both the real and artificial markets may be much more revealing of the actual underlying structures at work. It is also possible that much of the available price record may not be all that informative about what is going on, and should be down weighted.[40]

[40] There are very few studies that attempt to do this. One example is LeBaron (2001*c*).

Experiments along with micro level trading information also provide a useful means to model validation. A powerful result would be to calibrate agents in experiments, and then to fit various macro features using these micro agents. Calibrating to experiments has already been done in many contexts.[41]

Finally, older model testing methods such as out of sample predictability can be used. Agent-based markets can be used as out of sample predictors, although the implementation is not as obvious, and as straight forward as in more traditional models. One method is to replace the endogenous price series with actual prices during the learning phase of the market. Then the price series can be turned off, and the artificial market continues to run, generating a kind of out of sample forecast. Some early attempts at this will be described in future sections using the minority game structure.

4 Trading institutions and market microstructure

Most of the markets considered up to now have abstracted away from actual trading institutions. This is somewhat of a puzzle in the agent-based finance world, since a bottom up approach would appear to call for starting from the basics of how trades are executed.[42] This survey has mentioned the issue of the trading mechanism as being an important design question many times. Most models build stylized economic structures that avoid the institutional details of trading. However, research has begun appearing which implements more realistic trading systems.[43] Market design and market microstructure questions appear to be well suited for agent-based approaches. First, there is a large amount of data available. Second, there are critical policy questions which clearly need to be tested in an environment with heterogeneous, adaptive strategies.

It is interesting that one of the earliest agent-based markets modeled the trading process in detail. Rieck (1994) looks at the evolution of trading strategies with a simple order book trading mechanism. It has many similarities to some of the emergence papers mentioned in the previous sections, in that the coevolution of strategies is the crucial issue that the author is interested in. Also, strategies are evolved using evolutionary techniques, but these are applied to functional forms which are designed to replicate actual trading strategies. The results show that fundamental strategies are not able to take over the market and drive the price to the fundamental value. In markets with fundamentalists and other strategies competing the price moves away

[41] In finance Arifovic (1996) is one of the earliest examples. These are surveyed in Duffy (forthcoming 2004).

[42] Beltratti & Margarita (1992) is an interesting exception in these models, since trades are made in a random matching process.

[43] This area of financial research overlaps with work on market design which is covered more extensively in Mackie-Mason & Wellman (forthcoming 2004).

from fundamentals, but eventually returns. Although the methods are different, results in Rieck (1994) are suggestive that many agent-based financial market results could be replicated with more microstructure inspired trading mechanisms.[44]

4.1 Comparing trading institutions

In Audet, Gravelle & Yang (2001) the authors explore the question of order book versus dealer markets. Order book markets allow customers to place orders in an order book. The book is completely transparent, allowing all market participants to know the current depth and liquidity in the market. It is representative of most modern electronic trading systems. In dealer markets traders make trades through a set of intermediary dealers in the market. Dealer markets give customers much less information on the state of the market, but they also allow them to reveal less information concerning their current demands. The question that the authors are interested in is which type of market is preferred, and if there is a general optimal market structure. This latter question is a very interesting one, given that many financial markets have both trading systems operating in parallel.

Their setup is an agent-based approach which uses neural networks to represent agent strategies. The order book mechanism brings customers together directly in an electronic trading book, while the dealer based market involves two stages. In the first, orders are submitted to the dealers, and in the second, dealers adjust their inventories by trading with other dealers only. Before any trading starts, a Nash equilibrium in the strategies is found using a numerical search procedure. Once the strategies are determined, the agents are allowed to trade, generating market data, and utility payoffs for the various agents. The average utility for the order book market is always increasing in the number of customers entering the market, so for a large enough market, in terms of customers, the order book is preferred. The key number recorded is the break even value for which customers would be indifferent between the order book, and the dealer market. They use this number to quantify which type of market would be preferred in more general situations. If the break even number of customers is small, then order books are preferred, and when it is large dealer markets are preferred.

Changes on the parameters include changing the structure of the order flow, the number of dealers, and the risk aversion of the dealers. When order flow becomes larger and more correlated across customers dealer markets become more desirable. This result suggests that in situations when customers need to move large orders through a trading system the anonymity of the dealer system dominates the transparency of the order

[44] Yang (2002) does replicate many of the Santa Fe market results with a microstructure foundation.

based system. The authors also find that increasing the number of dealers and reducing their risk aversion also makes dealer markets more desirable. In both these situations the dealer market is able to absorb more risk.

This agent-based approach presents a detailed analysis of a critical policy question, and analyzes actual trading mechanisms in use. While the model uses an agent-based approach, it is important to indicate some distinctions from many of the other papers. The strategies for agents are not adapting. They are found as a fixed Nash equilbrium in strategy space, and then simulated. It will be interesting to know how well these results hold up in a world of learning and adaptation, even though this may come with the cost of making things more difficult to interpret.

Bottazzi, Dosi & Rebesco (2003) is another exploration of various financial trading institutions. They are interested in comparing the Walrasian trading protocol with two more realistic mechanisms, a batch auction procedure, and a limit order book. Traders are modeled with constant absolute risk aversion, and share demands similar to markets such as the Santa Fe Market structure. The agents perform simple forecasts on future prices and volatility as weighted averages over the recent past. They then pick a point on their demand schedule which gives a price quantity pair which they would like to execute in the market. The agent then randomly decides between executing a limit order which will get stored in the order book, or a market order which will be implemented immediately. They measure the efficiency of the different market structures while changing critical parameters. The most of important of these turns out to be the probability of submitting a market order. When the probability of a market order is large, the batch market is preferred, but when this probability is low, the order book market is preferred. This would appear to suggest that in markets where traders do not wish to reveal any information in the limit book, a batch auction is necessary to aggregate the sparse amount of information coming in through order flow. It does also repeat the message of Audet et al. (2001) that optimal trading structures may depend critically on the type of order flow coming into a market.

4.2 Tick sizes

One of the more interesting policy questions in financial market design has been the issue of decimalization. Many financial markets have a long history of trading in fractions such as 1/8 of a dollar, but many have recently shifted to decimal trading. Several agent-based markets have explored this question. Darley, Outkin, Plate & Gao (2000) was the result of a study done for NASDAQ by the Bios Group. They simulated realistic market trading mechanism and explored the impact of changing the tick size with various populations of trading agents. They measured performance of the market based on two criteria: its ability to closely track fundamentals, and overall price volatility. The trader population contained varying numbers of parasite trading strategies. These strategies were modeled after the "SOES Bandits", which is the nickname given to day trading strategies using the Small Order Execution System. With few parasites the impact of changing the tick size was minimal, but when there were many parasites around, reducing tick size increased the impact of the parasite strategies and actually reduced the market's ability to track fundamentals. This is interesting, since it goes against the perceived wisdom that any reduction in the tick size must help the efficient operation of the market. The impact on price volatility was a little more complicated. At first when the tick size is reduced the volatility falls, but eventually volatility changes direction and increases. In terms of minimizing price volatility there is an actual optimal tick size.

Yeh (2003) is another analysis of market performance under varying tick sizes. This model uses a genetic programming methodology similar to Chen & Yeh (2001). Agents again predict future prices and dividends, and use this along with their current portfolio information to determine a reservation price. Agents will buy a fixed unit of stock if they can get it for less than this price, and will sell if they can find a buyer. Orders proceed using a random ordering of agents, and an order book. When an agent's reservation does not cross the current book, the agent places a limit order using a simple algorithm that includes the reservation price along with the inside quotes. Simulation results for different tick sizes show that it does impact market performance. The results suggest that reducing the tick size generally helps the final stock customers, but hurts liquidity providers as the affected spreads are reduced when tick sizes fall.

4.3 Market makers

Most financial markets depend critically on the behavior of a broker dealer who brings liquidity and continuity to a real time trading market. Chan & Shelton (2001) concentrate specifically on the behavior of this dealer, and its ability to learn optimal strategies. The dealer is situated in a market with random order flow coming in from both informed, and uninformed agents. Informed agents receive a signal about the true value of the stock, while uniformed agents trade randomly. The dealer declares a fixed price at which he or she is willing to buy or sell (the bid/ask spread is zero), and accepts some amount of order flow. Then based on the dealer's inventories, recent order flow, and several other market state variables, the dealer revises the trading price for the next round. The policy function maps these state variables into the the price adjustment decision. Dealers seek to maximize their expected profits, and to keep inventory risk low. The first part of the paper tests various learning mechanisms to see if an optimal policy function can be learned using several different reinforcement learning algorithms. In most cases they are able to find the optimal policy function for the market maker.

The model is extended to the case where the bid and ask price might differ. Extending the policy function in this way eliminates the analytic comparison, but takes the market into a more realistic domain for comparison with actual markets. The prices and spreads track the actual value well, and dealers adjust the spreads to keep inventories small. Constructing realistic market makers is an important question for all agent-based financial markets. This paper shows that it can be tackled directly, and that reinforcement learning methods can find reasonable dealer strategies.

4.4 Empirical features

One of the most well documented features in intra-day data is the U-shaped pattern in bid ask spreads which are wide at the opening of trading, narrow during the day, and again widen before the close. There are also similar patterns in the volatility of spreads, return volatility, and trading volume as well. Chakrabarti (1999) seeks to replicate these features in a microstructure trading model which is built using an agent-based framework for foreign exchange dealers. Dealers receive random order flow through the day which gives them information about the aggregate order flow, and the eventual value of the foreign currency they are dealing in. This information is augmented by quotes they receive from other dealers during the intra-day trading period. Dealers are risk averse and are concerned about variances of their positions during the day, and also on the positions they hold over night. The reservation prices for these dealers are determined in a Bayesian learning framework. Each dealer agent determines an optimal selling (ask) and buying (bid) price at each time step. The spread between these is the return that compensates the dealer for risk in the inventory position. Trade takes place in a random matching process of the dealers. They trade when a calling dealer's bid is greater than the responding dealer's ask, or when the calling dealer's ask is less than the responding dealer's bid. Dealers use information from the other dealer spreads to update their own beliefs about order flow as they move through the day. As trading proceeds all traders' information improves, and order flow uncertainty falls. This leads to smaller spreads from the morning into the day. As the day reaches the close, the impact of overnight risk takes over for the dealers and spreads rise again.

The model is simulated for a wide range of parameters. The methodology chooses 729 unique parameter combinations, performs simulation runs for each set of parameters and records their results as separate observations. Parameter sensitivity is then determined using least squared regressions. The results show a general presence of the U-shaped spreads, and return volatility over simulated trading days, and the results are robust across many parameters. An interesting general result is that there is more unexplained variation in the afternoon variables. The author conjectures that this indicates the importance of path dependence of the prices and trades executed through the day. Finally, the nonlinear impact of the parameters on the results is explored. In most cases there are significant nonlinear effects in both quadratic and cross terms. This suggests a very complex relationship connecting the underlying information and preference parameters to final market outcomes.

One of the simplest and most direct approaches to replicating features in high frequency data is Farmer, Patelli & Zovko (2003). This is not exactly an agent-based market since order flow is completely random, and there is no individual trading agent per se. However, it forms a useful benchmark that agent-based modelers should be aware of. It also has very strong empirical foundations. Random order flow is calibrated to the actual order flow for several different stocks on the London Stock Exchange. This flow is then fed into a market clearing mechanism with a standard electronic order book. The types of incoming orders, limit or market, are determined randomly, and calibrated to the actual order flow from the data. The authors are able to show that their trading mechanism with no intelligent behavior in the agents is able to replicate many features from their price history data sets. This experiment is an interesting comparison with the ZI traders in Gode & Sunder (1993). These experiments help us to better understand which empirical regularities are the result of learning behavior, and which are simply a feature of the trading institutions in existence.

Although not designed to replicate empirical features directly, the Penn-Lehman Trading Project, Kearns & Ortiz (November/December 2003), is a large research operation that seeks to blend market making and high frequency trading behavior with direct market feeds from various Electronic Communication Networks, ECN's. Their system updates an off line limit order book either in real time or using historical data feeds. The order book is merged with the live orders coming into the ECN, and is also impacted by trade executions performed by the off line experimental agent traders.[45] This gives them a trading test bed to put different trading strategies into. Initially, this has been used in a trading competition with different submitted trading algorithms trading amongst themselves. The results of this experiment are probably only the beginning of much larger studies. However, there already are some critical insights in their strategy sets. Certain strategies turned out to make large trades, and demand a lot of liquidity from the market. In doing so, they had to move deeply into the order book, and paid large costs for trading. This market with live data penalized these strategies for their large trade making behavior in a very natural and realistic fashion. It will be interesting

[45] This data merge between the artificial market and the real order flow is nontrivial, and is described in www.cis.upen.edu/∼mkearns/papers/plat.pdf.

see what other results come from this large scale research collaboration.

4.5 Policy

Beyond the institutional design questions, the trading process opens up several interesting policy questions. If one accepts the fact that market volatility is too big relative to some rational market metric, then it opens the question of policies designed to inhibit volatility and trading activity. These often involve interfering with trades, and or intervening in markets directly. In three recent papers Frank Westerhoff tackles three of the most common policies, price limits, central bank intervention, and trading taxes, Westerhoff (2003*b*), Westerhoff (2003*c*), and Westerhoff (2003*a*).

The question of price limits considered in Westerhoff (2003*c*) is probably one of the most interesting for market policy makers. Price limits restrict the movement in prices to a certain percentage up or down. If the price moves to this limit, no trades can occur above or below the limit. Trading can continue, but trades must stay at or within the limit boundary. This is slightly different from a market circuit breaker where trade is actually halted for a fixed amount of time when the boundary price is reached.[46] The agent-based model is a small type model with technical and fundamental traders. The technical trader's demand depends on recent price trends, and the fundamental traders base their holdings on where the price is relative to a fundamental value. The fraction of fundamentalists in the population increases as the distance between the price and the fundamental increases. This generates a natural fundamental reverting dynamic which can be locally unstable when the price is close to the fundamental. The price is adjusted as a function of excess demand as is done in many of the previous markets.

The author implements a price limit by restricting the percentage change in the price in the excess demand price adjustment mechanism. Since the markets were not necessarily clearing in this market with no price restriction, the price limit is very easy to implement. Price volatility shows the obvious expected connection by increasing as the price limits are widened, and flattening out with a price limit in the range of 3 to 4 percent. A second measure of market performance is also recorded. This is the distortion, or distance of the traded price from the fundamental. The results for this are less obvious, and more interesting. As the price limit is increased from zero, the distortion initially falls. Strong restrictions on price adjustments keep the price far from the fundamental. However, the distortion begins to increase after the price limit passes about 0.5 percent. Increasing limits after this point actually reduces fundamental tracking by increasing the level of technical trading activity.

[46] Part of the reason for this difference is that many circuit breakers are implemented on indices rather than individual stocks.

This model makes an important point that choosing an optimal price limit may be a difficult task, and depends critically on the weighting of the volatility and distortion objectives. Though this paper is impressively aggressive in attacking an important policy question, two important caveats should be considered. First, the result depends critically on the price adjustment parameter, α. In other models the price dynamics depend on how quickly the price is adjusted to excess demand. Here, the entire question of whether, and how much of a price limit is needed can also depend critically on α. A second, and possibly more important question, is whether the behavior of traders might change after the limits are imposed. Traders may trade more aggressively as the price nears the limit since they fear that they might not be able to get in or out of the market. This kind of behavior could conceivably increase volatility when a limit is imposed. It might be more interesting to consider some form of adaptive behavior, where agents are able to change their strategies in response to the policy change.[47]

5 Other related markets

5.1 Corporate Finance

Most of the papers considered in this survey could loosely be considered part of the investments side of finance. There is no consideration for the issuance of securities by firms, or the design and evolution of securities themselves. A recent exception is Noe, Rebello & Wang (2003) which represents the first paper to consider corporate finance related issues in an agent-based framework. The authors are interested in the problem of which securities firms will issue to raise investment capital, along with the simultaneous learning of investors on how to price these securities. Firms need to issue securities that maximize their profits, but cannot do this independent of investors' pricing strategies. On the other hand, investors must learn how to price and evaluate the securities issued by firms, but they can only do this for securities they have seen in the past. The importance of this coevolutionary process of firm and investor learning turns out to be critical in the authors' model.

Their model is composed of a firm with an investment project that needs to be financed, and two potential investors. The firm can chose from a fixed set of 6 different securities that it can issue. These include standard debt and equity like securities, along with some more complex ones. The latter includes convertible and subordinated debt, and something known as a "Do-or-die" security. In each case the security

[47]This criticism is a form of the Lucas critique, Lucas (1976), and is related to the meta-trader question considered in this survey.

represents a well defined contract for splitting the payout to the firm's risky project between investors and the firm. Both the firm and investors encode their strategies as bit strings for use with a GA.[48] The firm maintains a pool of 80 potential security issue decisions which is vector of numbers (binary coded) between 1 and 6 corresponding to the 6 types of securities. The firm will chose one of these at random each period. The fitness of a strategy is updated with the realized cash flow received by the firm after the investment project has been completed, and the investors have been paid. Evolution takes place by replacing all the rules which encode the least profitable strategies with rules that encode the most profitable strategy. Then a mutation operator is applied to all rules.

The investors are also encoded as bitstrings as well. The two investors maintain a price table that indicates the price that they will pay for each possible security. The investor has a table of 80 possible pricing strategies for each security. In each round, each investor choses a pricing strategy at random from the appropriate security table after the firm has decided on the security that it will issue.[49] The security goes to the highest bidder in each round. The profitability of the strategy from the investor's perspective is recorded, and the populations are adjusted with a selection procedure where the 10 worst strategies are replaced by the 10 best. At this point the GA is applied to the population with crossover, mutation, and the election operator.

The authors then run this simulation for many rounds and in many different design situations. One of the most interesting results comes from the choice of securities. Experiments are performed that try to separate out the joint learning processes. Firms play against a fixed set of investors who know the appropriate pricing functions. In this situation equity and subordinated debt dominate the market, and straight debt is rarely used in stark contrast to the real world. When learning is allowed for both parties, debt moves to becoming the most commonly used security with subordinated debt next, and equity third. This shows the importance of the coevolutionary learning dynamic. In this world the preponderance of debt may have more to do with the ability of agents to learn how to price this relatively simple security, and the ensuing positive feedback this has on the issuance decision. Several other results from the model are also interesting. Investors tend to systematically underprice the securities in all cases. Also, the cases where the firm is not able to raise sufficient investment funds actually falls in the case with two sided learning relative to investor only learning.

The results in this paper will eventually need to be explored under different learning specification, and investment structures, but it is an interesting first step of agent-based models into the field of corporate

[48]The structure is very similar to Arifovic (1996).

[49]As in the earlier GA papers there is a binary to real mapping that occurs to get the real valued price.

finance. Furthermore, the importance of the coevolution of agent behavior along with security design is interesting both for finance, and economics in general where behaviors evolve along with the institutions that guide this behavior.

5.2 Minority Game

The minority game is an application of agent-based techniques which is on the fringe of finance. Simulations based on it have attracted a lot of attention in the Physics community, and some, but not as much, in the economics and finance worlds. This dichotomy is related to questions about whether it is actually a financial market. Whether it is exactly a financial market or not, the minority game has generated interesting results which should be of interest to financial researchers. This section presents a very short introduction to this large and fast growing literature.[50]

The first version of the minority game was presented in Challet & Zhang (1997). Its origins go back to a game known as the "El Farol" problem in Arthur (1994). There are also related resource games in Schelling (1978), so the game itself has a pretty long history. What is new in the recent research are the precise codings and dynamics of how agents adapt their behavior in an endogenously changing environment.

The game itself is very simple. There are two doors, labeled 0 and 1. There are an odd number of players who will chose a door in each round of the game. After choices are made, the door with the fewest number of players (the minority choice) wins. The game is then repeated. The objective is to be a contrarian in this world, and not to follow everyone else. Most papers on the minority game encode the strategies for players as a table of bit strings. The table has $m+1$ columns, and 2^m rows. Each row is a unique match for each of the 2^m possible histories plus a final bit which is the recommended play in the current round. The entire table is one strategy, and players are assumed to hold s different strategies which are usually generated at random. These different strategies will be continually monitored and compared according to the performance.

A general finding of the game is that it doesn't stabilize to any pattern or limit cycle. The agents continually shift back and forth between the doors in a fashion, which in the aggregate appears nearly random. Early studies explored the impact of changing the history, m, and other key parameters in the model on the time series features of the game. In particular, they studied the variability of the number of agents going to each door. A very interesting feature was documented in Savit, Manuca & Riolo (1999). When m is increased the volatility of the market initially falls. It reaches a minimum, and then eventually rises. The volatility is compared to the benchmark level where agents are choosing their strategies randomly.

[50] Interested readers should go to the website for the minority game at http://www.unifr.ch/econophysics/minority/.

For the very small m the volatility is larger than the random benchmark suggesting that in these cases the agents adapt against each other in a manner which generates excess volatility. For very large m the volatility asymptotes to the random value, suggesting that long history strategies eventually become equivalent to a kind of random number generator in strategy choices.

As mentioned earlier, the minority game has been expanded and extensively analyzed in a large literature.[51] Some interesting additions have included adding a price like mechanism on top of the game, and replicating several of the time series features in actual financial markets. Another interesting extension is for some traders to sit out a round of the game if they are not confident about their signal.[52]

A very interesting application is to use the minority game in a forecasting context. In Johnson, Lamper, Jefferies, Hart & Howison (2001) a financial series is converted into a binary string depending on whether the price went up or down. The agents are then allowed to continue playing the game after the price series is shut off, and the continued model dynamics are used in a kind of out of sample forecasting context. They are able to produce some small forecasting gains in some high frequency data. It remains to be seen how robust and reliable these numbers are, but this is an interesting test of the minority game model on real data. Another interesting application is to recommend possible policies to try and mitigate bubbles and crashes and stabilize markets. Hart, Lamper & Johnson (2003) is an example of this. Using their minority game framework the authors show that by trying to move the market in the direction of a crash before it really gets going can relieve the crash pressure, and the dynamics of the market are stabilized. This is kind of like a release valve in the dynamics of the market which lets off steam before the final big crash occurs. It is an interesting policy recommendation, and it will be interesting to see how robust this is across many different models.

The most important question for finance is whether the minority game is actually a good representation for a financial market. The minority game community holds that the contrarian nature of the game is analogous to optimal trading strategies, because it is best to be slightly out of step with the rest of the world. It is not clear whether this feature alone makes it a good model for finance. Also, it may not be the case that traders should be completely going against the tide. It may be optimal to ride a bubble for a while, but to get out just in time. Whether this is a minority strategy or not is not an easy question. Also, the minority game has been criticized for not really being a market in that there are no prices, and no trade actually occurs. This is changing as the game continues to be augmented, but the modifications often have

[51] A recent survey is Jefferies, Hart, Hui & Johnson (2000), and the previously cited website is also very useful.

[52] This is known as the Grand Canonical Minority Game, Jefferies et al. (2000) and references therein. It shares some similarities to the model mentioned earlier by Iori (2002).

the flavor of a kind of last minute addition. Trading itself will probably never be at the core of the minority game. Even with these criticisms this is an important body of work because it is a fairly simple agent based world in which the aggregate dynamics are determined by interacting forecast behaviors.

6 Cautions and criticisms

Agent-based markets are criticized from many different angles. The most common is that the models have far too many parameters, and the impact of many of these parameters is not well understood or tractable. This issue has already been discussed in the section on calibration. Another important issue that is brought up is the stability of a given market's results to the addition of new trading strategies. Specifically, are there strategies which would smoke out obvious patterns in the data, and change the dynamics. The computational models are trying to continuously defend against this with the constantly learning agents, but something outside the learning structure is possible. An initial defense of this is that most markets generate very little autocorrelation, and therefore yield no obvious trading strategies. However, there is a possibility of more complex nonlinear strategies. Arifovic (2001) is an example testing this sort of issue, and finds that the more complicated agents do not do better in the simulated market environment. This problem is still one of the most important for agent-based modelers to worry about, and no one should feel immune to this criticism.

Another very common, and very difficult criticism is that most of the markets deal with a rather small number of assets. Often agents trade only one risky, and one risk free asset alone.[53] It is certainly true, that with all of the new technology that is in use in these models it was important to start with the simplifying case of a single risky and and riskless asset. However, this simplification may eliminate many interesting features. The criticisms of the the single representative agent may carry over equally well to the representative risky asset. Questions naturally arise about calibrating to aggregate dividends, and exactly what these mean since they are not paid by any single stock. Also, recent events such as the technology bubble of the 1990's remind us that often bubbles are very sector dependent, leaving much of the market unaffected. Finally, when thinking about trading volume, it is really necessary to have a multi-asset world where traders are allowed to move back and forth between stocks. The single asset market puts an extreme restriction on the amount of trading volume that can be generated in a simulated market. Another related problem is that even though most artificial markets have two assets they actually shut down pricing in one market. In many cases the risk free rate is fixed, and the market is not a general equilibrium representation of a market for

[53] A recent exception to this is Westerhoff (forthcoming 2004).

risky and risk free assets. This is problematic in that the level and volatility of the risk free asset itself has been another asset pricing puzzle. Getting the risk free rate to be as low and stable as it is in the macro data is not easy, and most agent-based models simply avoid this problem completely. Endogenously opening multiple markets for trading is still a difficult problem, but it needs to be addressed at some point. Probably once more of the basic technologies have settled down, the problem of multi-asset markets will be tackled more frequently.

Egenter, Lux & Stauffer (1999) address another interesting question for agent-based modelers to consider. What happens as the number of agents is increased? They have performed some tests on models which can be studied analytically, and find that the dynamics can change dramatically as the number of agents becomes large. What initially looks like random behavior for small numbers of agents can become more or less predictable as the numbers become very large. Is it possible that many of the nice features that many models display are artifacts of the small numbers of traders that can be handled in the computer. This is a very important question. One caveat to it is that setting the number of agents to infinity might not be realistic in some settings. There may be places where the small sample issues in agent populations are important, and critical to the observed dynamics in real markets. This issue will definitely be an important one for the field to tackle in the future.

Almost all of the agents that are modeled and discussed in this survey operate inductively. They adapt and adjust to rules and forecasts which have performed well in the recent past. The early spirit of agent-based models is clearly to push from more traditional deductive styles of learning toward induction. However, it is often asked if there still may be a role for some form of deductive reasoning. Is it going too far to think of agents simply looking for patterns in the past, and using behaviors that work? Can they be allowed to do some form of deductive reasoning? Can they learn commonly held theories in finance, such as present value analysis, or the Black-Scholes option pricing formula? An interesting question is whether an agent-based model can be constructed which allows for a little deductive reasoning while keeping the general inductive spirit of simple rules of thumb.

A final problem, that is often ignored is that of timing. Almost all agent-based models need to make explicit assumptions about the timing of decisions, information and trade. Of course, any asset pricing model needs to make these choices, but in more analytic settings more events can be assumed to take place simultaneously. In the computer this sequence of events often needs to be spelled out. The amount to which results depend on arbitrary timing decisions is definitely important. One example that has been discussed here is the delayed price adjustment approach, where prices are adjusted based on current excess demand in the market. It is important to note that in a world of evolving strategies, this timing may have a large impact in that the strategies themselves adapt to the specific timing and trading structures. There are no perfect solutions to this problem beyond moving to more models where asynchronous behavior plays a bigger role.

7 Conclusions

This paper has given an overview of the current state of research in agent-based computational finance along with some ideas concerning the design and construction of working simulations. It is important to note that this is a very young field, and it still shows the kind of open ended exploratory nature of such an endeavor. However, several crucial trends are starting to appear.

First, the models are beginning to divide into several different types. These range from the small type models which are covered extensively in Hommes (forthcoming 2004) to the models with very open ended strategies determined through genetic programming operators. The small type models offer an important dimension of tractability relative to the the larger models, and they often provide definitive connections between parameters and results which might not be seen or noticed in the more complex frameworks, so it is easy to see their appeal. However, the reasons for going to the computer, and the use of more sophisticated artificial intelligence algorithms in the more complicated models with many traders needs to be understood as well. In most cases the reasons are twofold. First, they take emergence very seriously in that they do not prebias toward any particular strategy loaded ex ante by the researcher. The strategies that end up being used are those that could appear, and possibly remain stable inside a learning structure. They therefore partially answer a criticism of the small type models being somewhat ad hoc in their choice of trading strategies. Second, they try to use the computer and the learning algorithms to continuously answer the meta-trader question. In principle, these algorithms are continuously searching the time series record to smoke out new trading opportunities. This is something that is not present in the small type models. The obvious limitation here is that their ability to seek out and take at advantage of any inefficiencies that may appear depend critically on the data representations and implementations of the learning algorithms. Both methods clearly have strengths and weaknesses, and researchers using both should be well aware of these.

The field is also starting to separate from the more stylized earlier artificial markets toward more market microstructure oriented explorations. The latter try to model very explicitly the actual mechanisms of trade that are being used in the market as opposed to building a stylized trading framework. These microstructure oriented models are well designed to answer questions concerning the construction and design of these same trading mechanisms. In some of these markets it is the institutions that are at the center of the investigation, and the agents are just a mechanism for testing their behavior. Some of the policy questions addressed in this work are much more sharply defined than in other agent-based models. An example of this would be the explorations into decimalization on markets, or the implementation of price limits. From a policy perspective this would seem like a very natural place for the field to move as it matures.

Up to this point very little reference has been made to the growing literature on behavioral finance. It is important to define where agent-based financial markets sit relative to this larger field. First, they are clearly behavioral models themselves, since the agents are boundedly rational, and follow simple rules of thumb. This is a key characteristic of any behavioral model, and these are no exception. Where these approaches diverge is that they often consider agents who maximize relatively standard preferences. No attempt is made to model common behavioral biases such as loss aversion or hyperbolic discounting. This is not because they cannot handle these problems, but it seems sensible in the early stages to not add too many more complications to models which are already very complicated. It is important to note that agent-based technologies are well suited for testing behavioral theories. They can answer two key questions that should be asked of any behavioral structure. First, how well do behavioral biases hold up under aggregation, and second, which types of biases will survive in a coevolutionary struggle against others. Therefore, the connections between agent-based approaches and behavioral approaches will probably become more intertwined as both fields progress.

Whether computational or not, all of the models mentioned in this survey share a common tie to ideas from nonlinear dynamics and chaos. The relationship between model structure and noise in nonlinear systems can be very complicated, and these markets are no exception. In many cases the markets operate as noise magnifiers, taking a small amount of input noise, or underlying fundamental risk, and increasing its level to a much larger observed macro value. Noise can also help to stabilize a nonlinear system by keeping it off unstable trajectories. As is well known, nonlinear systems can also be difficult to forecast, and most of the markets described here share this feature. Unfortunately, this may also make them difficult to estimate using traditional econometric tools. Agent-based modelers should be aware of these nonlinear issues, and take them into account when evaluating market simulations.

Financial markets are an important challenge for agent-based computational models. They may be one of the important early test areas where these methods start to show their worth. This is both because of the many open questions, and the current problems of accepted theories, and because of the large amount of data available for testing. It will be interesting to see if sometime in the future we eventually replace the stylized theories of equilibrium market dynamics, with a more realistic picture of the continuing struggle of learning and adapting agents who push markets in the direction of efficiency, even though they never quite reach this goal.

References

Albin, P. S. & Foley, D. K. (1992), 'Decentralized, dispersed exchange without and auctioneer: A simulation study', *Journal of Economic Behavior and Organization* **18**(1), 27–52.

Alchian, A. (1950), 'Uncertainty, evolution, and economic theory', *Journal of Political Economy* **58**, 211–221.

Allen, F. & Karjalainen, R. (1998), 'Evolution of trading rules in financial markets', *Journal of Financial Economics* **51**, 245–271.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), 'Modeling and forecasting realized volatility', *Econometrica* **71**, 529–626.

Arifovic, J. (1994), 'Genetic algorithm learning and the cobweb model', *Journal of Economic Dynamics and Control* **18**, 3–28.

Arifovic, J. (1996), 'The behavior of the exchange rate in the genetic algorithm and experimental economies', *Journal of Political Economy* **104**, 510–541.

Arifovic, J. (2001), 'Performance of rational and boundedly rational agents in a model with persistent exchange rate volatility', *Macroeconomic Dynamics* **5**, 204–224.

Arifovic, J. & Gencay, R. (2000), 'Statistical properties of genetic learning in a model of exchange rate', *Journal of Economic Dynamics and Control* **24**, 981–1005.

Arifovic, J. & Masson, P. (1999), Heterogeneity and evolution of expectations in a model of currency crisis, Technical report, Simon Fraser University, Vancouver, BC, Canada.

Arthur, W. B. (1994), 'Inductive reasoning and bounded rationality', *American Economic Review* **84**, 406–411.

Arthur, W. B., Holland, J., LeBaron, B., Palmer, R. & Tayler, P. (1997), Asset pricing under endogenous expectations in an artificial stock market, *in* W. B. Arthur, S. Durlauf & D. Lane, eds, 'The Economy as an Evolving Complex System II', Addison-Wesley, Reading, MA, pp. 15–44.

Audet, N., Gravelle, T. & Yang, J. (2001), Optimal market structure: Does one shoe fit all?, Technical report, Bank of Canada, Ottawa, CA.

Bachelier, L. (1900), Theorie de la Speculation, PhD thesis, Ecole Normale Superieure, Paris, France.

Badegruber, T. (2003), Agent-based Computational Economics: New aspects in learning speed and convergence in the Santa Fe Artificial Stock Market, PhD thesis, Universitat Graz, Graz, Austria.

Baillie, R. T., Bollerslev, T. & Mikkelsen, H.-O. (1996), 'Fractionally integrated generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **74**, 3–30.

Bansal, R. & Yaron, A. (forthcoming 2004), 'Risks for the long run: A potential resolution of asset pricing puzzles', *Journal of Finance* .

Barberis, N. & Thaler, R. (2002), A survey of behavioral finance, *in* G. Constantinides, M. Harris & R. Stulz, eds, 'Handbook of Economics and Finance', North-Holland.

Beja, A. & Goldman, M. B. (1980), 'On the dynamic behavior of prices in disequilibrium', *Journal of Finance* **35**, 235–248.

Beltratti, A. & Margarita, S. (1992), Evolution of trading strategies among heterogeneous artificial economic agents, *in* J. A. Meyer, H. L. Roitblat & S. W. Wilson, eds, 'From Animals to Animats 2', MIT Press, Cambridge, MA.

Beltratti, A., Margarita, S. & Terna, P. (1996), *Neural Networks for economic and financial modeling*, International Thomson Computer Press, London, UK.

Benink, H. & Bossaerts, P. (2001), 'An exploration of neo-austrian theory applied to financial markets', *Journal of Finance* **56**(3), 1011–1027.

Bollerslev, T., Engle, R. F. & Nelson, D. B. (1995), ARCH models, *in* 'Handbook of Econometrics', Vol. 4, North-Holland, New York, NY.

Bottazzi, G., Dosi, G. & Rebesco, I. (2003), Institutional architectures and behavioral econologies in the dynamics of financial markets, Technical report, S. Anna School of Advanced Studies, Pisa, Italy.

Bray, M. (1982), 'Learning, estimation, and the stability of rational expectations', *Journal of Economic Theory* **26**, 318–339.

Breedon, D. (1979), 'An intertemporal asset pricing model with stochastic consumption and investment', *Journal of Financial Economics* **7**, 265–296. %K %X %F.

Brock, W. A. & Hommes, C. H. (1998), 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model', *Journal of Economic Dynamics and Control* **22**(8-9), 1235–1274.

Brock, W. A., Hsieh, D. & LeBaron, B. (1991), *Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence*, MIT Press, Cambridge, MA.

Brock, W. A., Lakonishok, J. & LeBaron, B. (1992), 'Simple technical trading rules and the stochastic properties of stock returns', *Journal of Finance* **47**, 1731–1764.

Bullard, J. & Duffy, J. (2001), 'Learning and excess volatility', *Macroeconomic Dynamics* **5**, 272–302.

Campbell, J., Grossman, S. J. & Wang, J. (1993), 'Trading volume and serial correlation in stock returns', *Quarterly Journal of Economics* **108**, 905–940.

Campbell, J. Y. (2000), 'Asset pricing at the millennium', *Journal of Finance* **55**, 1515–1568.

Campbell, J. Y. & Cochrane, J. H. (1999), 'By force of habit: A consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy* **107**, 205–251.

Campbell, J. Y., Lo, A. W. & MacKinlay, A. C. (1996), *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.

Campbell, J. Y. & Shiller, R. (1988), 'The dividend-price ratio and expectations of future dividends and discount factors', *Review of Financial Studies* **1**, 195–227.

Chakrabarti, R. (1999), 'Just another day in the inter-bank foreign exchange market', *Journal of Financial Economics* **56**(1), 29–64.

Chakrabarti, R. & Roll, R. (1999), 'Learning from others, reacting and market quality', *Journal of Financial Markets* **2**, 153–178.

Challet, D. & Zhang, Y. C. (1997), 'Emergence of cooperation and organization in an evolutionary game', *Physica A* **246**, 407.

Chan, N. T. & Shelton, C. (2001), An electronic market-maker, Technical report, MIT, Cambridge, MA.

Chen, S. H. & Yeh, C. H. (2001), 'Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market', *Journal of Economic Dynamics and Control* **25**, 363–394.

Chen, S., Lux, T. & Marchesi, M. (2001), 'Testing for non-linear structure in an artificial financial market', *Journal of Economic Behavior and Organization* **46**, 327–342.

Chiarella, C. (1992), 'The dynamics of speculative behaviour', *Annals of Operations Research* **37**, 101–123.

Constantinides, G. M. (1989), Theory of valuation: Overview and recent developments, *in* S. Bhattacharya & G. M. Constantinides, eds, 'Theory of Valuation: Frontiers of Modern Financial Theory', Rowman and Littlefield, Totowa, New Jersey, pp. 1–24.

Cont, R. (2001), 'Empirical properties of asset returns: stylized facts and statistical issues', *Quantitative Finance* **1**, 223–236.

Cont, R. & Bouchaud, J. P. (2000), 'Herd behavior and aggregate fluctuations in financial markets', *Macroeconomic Dynamics* **4**, 170–196.

Cootner, P., ed. (1964), *The Random Character of Stock Market Prices*, MIT Press, Cambridge.

Dacorogna, M. M., Gencay, R., Muller, U. A., Olsen, R. B. & Pictet, O. V. (2001), *An Introduction to High-Frequency Finance*, Academic Press, San Diego, CA.

Darley, V., Outkin, A., Plate, T. & Gao, F. (2000), Sixteenths or pennies? observations from a simulation of the Nasdaq stock market, Technical report, Santa Fe Institute.

Day, R. H. & Huang, W. H. (1990), 'Bulls, bears, and market sheep', *Journal of Economic Behavior and Organization* **14**, 299–330.

De Grauwe, P., Dewachter, H. & Embrechts, M. (1993), *Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets*, Blackwell, Oxford.

DeLong, J. B., Shleifer, A., Summers, L. H. & Waldmann, R. (1990), 'Noise trader risk in financial markets', *Journal of Political Economy* **98**, 703–738.

Ding, Z., Granger, C. & Engle, R. F. (1993), 'A long memory property of stock market returns and a new model', *Journal of Empirical Finance* **1**, 83–106.

Duffy, J. (forthcoming 2004), Agent-based models and human subject experiments, *in* K. L. Judd & L. Tesfatsion, eds, 'Handbook of Computational Economics', Elsevier.

Egenter, E., Lux, T. & Stauffer, D. (1999), 'Finite-size effects in monte-carlo simulations of two stock market models', *Physica A* **268**, 250–256.

Ehrentreich, N. (2003), The santa fe artificial stock market re-examined - suggested corrections, Technical report, Martin Luther University of Halle-Wittenberg Dept. of Banking and Finance, Halle/Saale, Germany.

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

Fama, E. F. (1991), 'Efficient capital markets: II', *Journal of Finance* **45**, 1575–1617.

Fama, E. F. & French, K. R. (1992), 'The cross-section of expected stock returns', *Journal of Finance* **47**, 427–465.

Farmer, J. D. & Joshi, S. (2002), 'The price dynamics of common trading strategies', *Journal of Economic Behavior and Organization* **49**, 149–171.

Farmer, J. D., Patelli, P. & Zovko, I. (2003), The predictive power of zero intelligence models in financial markets, Technical report, Santa Fe Institute, Santa Fe, NM.

Figlewski, S. (1978), 'Market "efficiency" in a market with heterogeneous information', *Journal of Political Economy* **86**(4), 581–597.

Fogel, D. B. (1995), *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*, IEEE Press, Piscataway, NJ.

Frankel, J. A. & Froot, K. A. (1987), 'Using survey data to test standard propositions regarding exchange rate expectations', *American Economic Review* **77**(1), 133–153.

Frankel, J. A. & Froot, K. A. (1988), Explaining the demand for dollars: International rates of return and the expectations of chartists and fundamentalists, *in* R. Chambers & P. Paarlberg, eds, 'Agriculture, Macroeconomics, and the Exchange Rate', Westview Press, Boulder, CO.

Friedman, M. (1953), The case for flexible exchange rates, *in* 'Essays in positive economics', University of Chicago Press, Chicago, IL.

Fung, H. G. & Patterson, G. A. (1999), 'The dynamic relationship of volatility, volume, and market depth in currency futures markets', *Journal of International Financial Markets, Institutions and Money* **17**, 33–59.

Gallant, A. R., Rossi, P. E. & Tauchen, G. (1992), 'Stock prices and volume', *The Review of Financial Studies* **5**, 199–242.

Gallant, A. R., Rossi, P. E. & Tauchen, G. (1993), 'Nonlinear dynamic structures', *Econometrica* **61**, 871–908.

Gode, D. K. & Sunder, S. (1993), 'Allocative efficiency of markets with zero intelligence traders', *Journal of Political Economy* **101**, 119–37.

Goldberg, D. E. (1989), *Genetic Algorithms in search, optimization and machine learning*, Addison Wesley, Reading, MA.

Grossman, S. & Stiglitz, J. (1980), 'On the impossibility of informationally efficient markets', *American Economic Review* **70**, 393–408.

Hansen, L. & Singleton, K. (1983), 'Stochastic consumption, risk aversion, and the temporal behavior of asset returns', *Journal of Political Economy* **91**, 249–265.

Hart, M. L., Lamper, D. & Johnson, N. F. (2003), An investigation of crash avoidance in a complex system, Technical report, Physics Department, Oxford University, Oxford, UK.

Hayek, F. A. (1945), 'The use of knowledge in society', *American Economic Review* **35**(4), 519–30.

Hirshleifer, D. (2001), 'Investor psychology and asset pricing', *Journal of Finance* **56**(4), 1533–1597.

Holland, J. H. (1975), *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, MI.

Holland, J. H., Holyoak, K. J., Nisbett, R. E. & Thagard, P. R. (1986), *Induction*, MIT Press, Cambridge, MA.

Hommes, C. H. (forthcoming 2004), Heterogeneous agent models in economics and finance, *in* K. L. Judd & L. Tesfatsion, eds, 'Handbook of Computational Economics', Elsevier.

Iori, G. (2002), 'A microsimulation of traders activities in the stock market: The role of heterogeneity, agents' interactions and trade frictions', *Journal of Economic Behavior and Organization* **49**, 269–285.

Jefferies, P., Hart, M., Hui, P. M. & Johnson, N. F. (2000), From market games to real-world markets, Technical report, Physics Deptartment, Oxford University, Oxford, UK.

Johnson, N. F., Lamper, D., Jefferies, P., Hart, M. L. & Howison, S. (2001), 'Application of multi-agent games to the prediction of financial time-series', *Physica A* **299**, 222–227.

Johnson, P. (2002), 'What I learned from the artificial stock market', *Social Science Computer Review* **20**(2), 174–196.

Joshi, S., Parker, J. & Bedau, M. A. (2000), Technical trading creates a prisoner's dilema: Results from an agent-based model, *in* 'Computational Finance 99', MIT Press, Cambridge, MA, pp. 465–479.

Kareken, J. & Wallace, N. (1981), 'On the indeterminacy of equilibrium exchange rates', *Quarterly Journal of Economics* **96**, 207–222.

Kearns, M. & Ortiz, L. (November/December 2003), 'The penn-Lehman automated trading project', *IEEE Intelligent Systems* pp. 22–31.

Kim, G. & Markowitz, H. (1989), 'Investment rules, margin, and market volatility', *Journal of Portfolio Management* **16**(1), 45–52.

Kirman, A. P. (1991), Epidemics of opinion and speculative bubbles in financial markets, *in* M. Taylor, ed., 'Money and Financial Markets', Macmillan, London.

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

Kirman, A. P. (1993), 'Ants, rationality, and recruitment', *Quarterly Journal of Economics* **108**, 137–56.

Kirman, A. P. & Teyssiere, G. (2001), 'Microeconomic models for long-memory in the volatility of financial time series', *Studies in Nonlinear Dynamics and Econometrics* **5**, 281–302.

Kocherlakota, N. (1996), 'The equity premium: It's still a puzzle', *Journal of Economic Literature* **34**(1), 42–71.

Koza, J. R. (1992), *Genetic Programming: On the programming of computers by natural selection*, MIT Press, Cambridge, MA.

LeBaron, B. (1992), 'Some relations between volatility and serial correlations in stock market returns', *Journal of Business* **65**(2), 199–219.

LeBaron, B. (2001*a*), 'Empirical regularities from interacting long and short memory investors in an agent based stock market', *IEEE Transactions on Evolutionary Computation* **5**, 442–455.

LeBaron, B. (2001*b*), 'Evolution and time horizons in an agent based stock market', *Macroeconomic Dynamics* **5**(2), 225–254.

LeBaron, B. (2001*c*), Volatility magnification and persistence in an agent based financial market, Technical report, International Business School, Brandeis University, Waltham, MA.

LeBaron, B. (2002), Calibrating an agent-based financial market, Technical report, International Business School, Brandeis University, Waltham, MA.

LeBaron, B. (forthcoming 2004), Building the santa fe artificial stock market, *in* F. Luna & A. Perrone, eds, 'Agent-based Theory, Languages, and Experiments', Routledge Publishing.

LeBaron, B., Arthur, W. B. & Palmer, R. (1999), 'Time series properties of an artificial stock market', *Journal of Economic Dynamics and Control* **23**, 1487–1516.

Lettau, M. (1997), 'Explaining the facts with adaptive agents: The case of mutual fund flows', *Journal of Economic Dynamics and Control* **21**, 1117–1148.

Lettau, M. & Uhlig, H. (1999), 'Rules of thumb versus dynamic programming', *American Economic Review* **89**, 148–174.

Levy, M., Levy, H. & Solomon, S. (1994), 'A microscopic model of the stock market: cycles, booms, and crashes', *Economics Letters* **45**, 103–111.

Levy, M., Levy, H. & Solomon, S. (2000), *Microscopic Simulation of Financial Markets*, Academic Press, New York, NY.

Lo, A. W. & MacKinlay, A. C. (1988), 'Stock prices do not follow random walks: Evidence from a simple specification test', *Review of Financial Studies* **1**, 41–66.

Logato, I. & Velasco, C. (2000), 'Long memory in stock-market trading volume', *Journal of Business and Economic Statistics* **18**, 410–426.

Lucas, R. E. J. (1976), 'Econometric policy evaluation: A critique', *Carnegie-Rochester Conference Series on Public Policy* **1**, 19–46.

Lucas, R. E. J. (1978), 'Asset prices in an exchange economy', *Econometrica* **46**, 1429–45.

Lux, T. (1997), 'Time variation of second moments from a noise trader/infection model', *Journal of Economic Dynamics and Control* **22**, 1–38.

Mackie-Mason, J. K. & Wellman, M. P. (forthcoming 2004), Automated markets and trading agents, *in* K. L. Judd & L. Tesfatsion, eds, 'Handbook of Computational Economics', Elsevier.

Mandelbrot, B. B. (1963), 'The variation of certain speculative prices', *Journal of Business* **36**, 394–419.

Marey, P. S. (2004), 'Exchange rate expectations: controlled experiments with artificial traders', *Journal of International Money and Finance* **23**, 283–304.

Marimon, R., McGrattan, E. & Sargent, T. J. (1990), 'Money as a medium of exchange in an economy with artificially intelligent agents', *Journal of Economic Dynamics and Control* **14**, 329–373.

Masson, P. (1999), 'Contagion: Macroeconomic models with multiple equilibria', *Journal of International Money and Finance* **18**, 587–602.

Mehra, R. (2003), The equity premium: Why is it a puzzle?, Technical Report 9512, National Bureau of Economic Research.

Mehra, R. & Prescott, E. C. (1988), 'The equity risk premium: A solution?', *Journal of Monetary Economics* **22**, 133–136.

Merton, R. (1971), 'Optimum consumption and portfolio rules in a continuous-time model', *Journal of Economic Theory* **3**, 373–413.

Neely, C., Weller, P. & Dittmar, R. (1997), 'Is technical analysis in the foreign exchange market profitable? a genetic programming approach', *Journal of Financial and Quantitative Analysis* **32**, 405–26.

Noe, T., Rebello, M. & Wang, J. (2003), 'Corporate financing: An artificial agent-based analysis', *Journal of Finance* **63**, 943–973.

Palmer, R., Arthur, W. B., Holland, J. H., LeBaron, B. & Tayler, P. (1994), 'Artificial economic life: A simple model of a stock market', *Physica D* **75**, 264–274.

Rieck, C. (1994), Evolutionary simulation of asset trading strategies, *in* E. Hillebrand & J. Stender, eds, 'Many-Agent Simulation and Artificial Life', IOS Press.

Routledge, B. R. (2001), 'Genetic algorithm learning to choose and use information', *Macroeconomic Dynamics* **5**, 303–325.

Rubenstein, M. (2001), 'Rational markets: Yes or no? The affirmative case', *Financial Analyst Journal* **17**, 15–29.

Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.

Savit, R., Manuca, R. & Riolo, R. (1999), 'Adaptive competition, market efficiency, and phase transitions', *Physical Review Letters* **82**, 2203–2206.

Schelling, T. (1978), *Micromotives and Macrobehavior*, Norton, New York, NY.

Shiller, R. J. (2003), 'From efficient market theory to behavioral finance', *Journal of Economic Perspectives* **17**, 83–104.

Stauffer, D. & Sornette, D. (1999), 'Self-organized percolation model for stock market fluctuations', *Physica A* **271**, 496–506.

Tay, N. S. P. & Linn, S. C. (2001), 'Fuzzy inductive reasoning, expectation formation and the behavior of security prices', *Journal of Economic Dynamics and Control* **25**, 321–362.

Vriend, N. (2000), 'An illustration of the essential difference between individual and social learning, and its consequences for computational analysis', *Journal of Economic Dynamics and Control* **24**, 1–19.

Welch, I. (1992), 'Sequential sales, learning, and cascades', *The Journal of Finance* **47**, 695–732.

Westerhoff, F. (2003*a*), 'Central bank intervention and feedback traders', *International Financial Market, Institutions, and Money* **13**, 419–427.

Westerhoff, F. (2003*b*), 'Heterogeneous traders and the Tobin tax', *Journal of Evolutionary Economics* **13**, 53–70.

Westerhoff, F. (2003*c*), 'Speculative markets and the effectiveness of price limits', *Journal of Economic Dynamics and Control* **28**, 493–508.

Westerhoff, F. (forthcoming 2004), 'Multi-asset market dynamics', *Macroeconomic Dynamics* .

Winker, P. & Gilli, M. (2001), Indirect estimation of the parameters of agent based models of financial markets, Technical Report 38, FAME.

Yang, J. (2002), Agent-based Modelling and Market Microstructure, PhD thesis, Concordia University, Montreal, Canada.

Yeh, C. H. (2003), Tick size and market performance, Technical report, Yuan Ze University, Chungli, Taiwan.

Youssefmir, M. & Huberman, B. A. (1997), 'Clustered volatility in multiagent dynamics', *Journal of Economic Behavior and Organization* **32**, 101–118.

Zschischang, E. & Lux, T. (2001), 'Some new results on the Levy, Levy, and Solomon microscopic stock market model', *Physica A* **291**, 563–573.