

Agent-based Computational Finance

Blake LeBaron *

Brandeis University

April 21st, 2005

Abstract

This chapter surveys research on agent-based models used in finance. It will concentrate on models where the use of computational tools is critical for 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 research associate at the National Bureau of Economic Research. This is a draft for a chapter which appears in *The Handbook of Computational Economics, vol. II*, edited by K. L. Judd and L. Tesfatsion, volume 2, pages 1187-1233, 2006. The author is grateful to many people who have made comments on earlier drafts. These include W. A. Brock, Cars Hommes, Leigh Tesfatsion, Frank Westerhoff, and two anonymous referees.

Contents

1. Introduction
2. Design questions
3. Artificial financial markets
4. Other markets
5. Cautions and criticisms
6. Conclusions

JEL classification: C6, G0, E2, F3, D8

Keywords: learning, evolutionary finance, financial time series, asset pricing, efficient markets, behavioral finance, market microstructure, genetic algorithms, neural networks, artificial financial markets, evolutionary computation

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, laid 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 it is helping to pursue a world view in which agents may differ in many ways, not just in their information, but in their ability to process information, their attitudes toward risk, and in 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.⁴ In agent-based financial markets, dynamic 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 the model,

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

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

³The survey by Hommes (2005) 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.

⁴See Tesfatsion (forthcoming 2005) for more extensive definitions of agent-based approaches in economics.

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

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. Second, financial time series contain many curious puzzles that are not well understood. Third, financial markets provide a wealth of pricing and volume data that can be analyzed. Fourth, 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 that 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.⁵ 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 the 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 is, it is interesting to note that there never really has been a truly accepted dynamical process describing how market efficiency comes about. The second foundation for efficient market theories, supported by much of the early empirical work on financial markets, is that markets are much more unpredictable than the world of the financial practitioner suggests.⁶ In this early literature, the random walk model appeared to be a pretty good approximation for the movements of stock prices, and it can be argued that the same holds true today. We know that markets are probably not completely unpredictable, but they still are very difficult to forecast.⁷

The early ideas of efficient markets were made more formal as modern tools of dynamic optimization were brought to bear on these problems.⁸ This led to an even stronger representation for financial markets, the representative agent.⁹ This model formally connects asset prices to the beliefs of a single aggregate individual who can then be linked to various state variables of the macroeconomy.

⁵The field has been surveyed many places, but the classic surveys remain Fama (1970) and Fama (1991).

⁶Examples are in Cootner (1964) and Fama (1970).

⁷It 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).

⁸See, for example, Merton (1971), Breedon (1979), and Lucas (1978).

⁹Constantinides (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.

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 (2005), are based on the paradoxical structure of this model. More recently, the literature on noise trading, [e. g., DeLong, Shleifer, Summers & Waldmann (1990)], 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. An example of this was the use of moving average technical analysis rules as in Brock, Lakonishok & LeBaron (1992). 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 the real return on 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

¹⁰A good recent survey on this literature is Campbell (2000); see also, the textbook by Campbell, Lo & MacKinlay (1996).

¹¹Shiller (2003) is a good recent survey on this.

¹²This one feature has generated an extensive literature which is surveyed in Kocherlakota (1996) and more recently in Mehra (January/February 2003).

do not comove with markets in a way that can justify the magnitudes of risk premia observed in financial series. There have 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, volume has recently been shown to be a long-memory process with persistence extending out many periods [see, e. g., 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 lacks an accepted explanation. While the direction of stock returns is generally unpredictable, their magnitudes are often very predictable.¹⁵ Stock markets repeatedly switch between periods of relative calm and periods of 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.¹⁶ Similar to volume persistence, it is also a potential long-memory process.¹⁷ Beyond simple persistence there are some more complicated issues in the dynamics of volume and volatility.¹⁸

Closely related to 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 by 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. At the annual frequency, the normal distribution is not a bad approximation. Fat tails are not

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

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

¹⁵This 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.

¹⁶One of the few examples of theoretical models generating persistent volatility is McQueen & Vorkink (2004).

¹⁷See Ding, Granger & Engle (1993), Andersen, Bollerslev, Diebold & Labys (2003), and also Baillie, Bollerslev & Mikkelsen (1996).

¹⁸These 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).

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 frequency of large moves there is a continuing debate about the exact shape of the tails of return distributions. It is possible that these may be described by power laws.¹⁹

One of the reasons for this wide range of puzzles is another justification for finance being a good agent-based test bed. Financial data are generally plentiful, accurate, and available on many different aspects of financial market functions. Good time series of up to forty years are available on prices and volume. Series of lengths up to one hundred years are available for lower frequencies, and for certain securities. Over the past twenty 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 that show detailed holdings of institutions, such as mutual funds, and that record the flows coming in and out of these funds. For individuals a few series have been used that reveal the trades of investors' accounts at various brokerage firms.²⁰ 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 are available that 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 (2005). Finance experiments are particularly appealing since they often can be done at time scales that are reasonable for the real data. It is more credible that you can simulate a day of trading in the laboratory, than to simulate someone's entire life cycle.

To summarize, financial markets are particularly well suited for agent-based explorations. They are large well-organized markets for 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 that traditional approaches have not been able to match. Agent-based approaches provide an intriguing possibility for solving some of these puzzles.²¹ Finally, financial markets are rich in data sets that can be used for testing and calibrating agent-based models. High quality data are available at many frequencies, and in many different forms.

¹⁹Good surveys on power laws in finance are Cont (2001), Dacorogna et al. (2001), Mantegna & Stanley (1999), and Lux (2002). Power laws are difficult to formally test empirically. However, Solow, Costello & Ward (2003) is one framework for attempting to build a test of power law behavior. LeBaron (2001c) provides an example showing how visual tests of power laws can be deceiving.

²⁰Barber & Odean (2000) is example of this research.

²¹There are other explanations that may yet prove to be 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.

The remainder of this chapter will summarize recent work on agent-based computational models in finance. The next section introduces some of the computational tools and design issues that are important in building markets. Section 3 covers artificial market models that attempt to recreate an entire market. Section 4 covers a few other types of markets which do not fit into the earlier categories. Section 5 covers some on-going debates and criticisms of agent-based markets, and section 6 concludes and suggests questions for the future.

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.

2.1 Preferences

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 altogether, and to concentrate simply 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.

2.2 Price determination

Many models considered here focus on 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. An early example of this is Day & Huang (1990). 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; if there is an excess demand the price is increased, and if there is an 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 is excess demand handled? Are excess demanders supplied from some inventory, or is rationing used?

A second market mechanism is to clear the market in each period either numerically, or through some theoretical simplifications that allow for an easy analytic solution to the temporary market clearing price. Two examples of this method are Brock & Hommes (1998) and Arthur, Holland, LeBaron, Palmer & Tayler (1997). 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 it may not well represent 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 a simplification of the demands of agents to yield an analytically tractable price.²²

²²A close relation to this method is to assume that prices are a function of the aggregation of expectations as in Kirman (1991) and De Grauwe, Dewachter & Embrechts (1993). Although trades don't actually take place, these papers do provide a clean mechanism for determining the current period price, and they can concentrate on agent expectation formation.

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. Examples of this are Chiarella & Iori (2002) and Farmer, Patelli & Zovko (2005). This method is very realistic and allows detailed analysis of trading mechanisms. Its only drawback is that these same institutional details need to be built into both the market architecture, and the learning specifications of agents. 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). A finance example of this is Beltratti & Margarita (1992). This mechanism 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. It would appear realistic for situations where no formal trading markets have been established. However, it may not be very natural in places where trading institutions are well defined, and function to help buyers meet sellers in a less-than-random fashion.

2.3 Evolution and learning

Much of the agent-based literature has used tools taken from the artificial intelligence literature to model learning. One of these is the Genetic Algorithm (or GA), which is a key component in many, but not all, agent-based financial markets.²³ 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.

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 situations where traditional hill climbing methods can fail, such as multi-peaked objectives, or nondifferentiable objective functions, possibly with discrete input variables. 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 that will be considered in this survey.

²³More information on genetic algorithms along with many other learning algorithms is presented in Brenner (2005) and Duffy (2005).

Many beginning researchers view the GA as a kind of black box, and simply follow previous work in setup and structure.²⁴ This approach is probably a mistake. It is important to think more about evolutionary computation in general than about the particular pieces of the GA. The general field of evolutionary computation includes other methods such as evolutionary programming, and evolutionary strategies, and 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.²⁵

Setting up an evolutionary learning framework requires several preliminary steps. 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 represents the behavior and that 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 multi-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.

In financial settings agents and strategies can be evolved using either wealth, or utility-based fitness. In the case of wealth, evolution of the 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

²⁴Goldberg (1989) is the classic book for early GA adopters.

²⁵A nice balanced overview of all these methods is Fogel (1995).

performance. This also aligns with the bulk of the learning literature in macroeconomics, which often concentrates on forecast evaluation.

A second route is to ignore forecasts altogether and to deal directly with asset demands and strategies. 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.

2.4 Information representation

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 as well as 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 only to certain agents? Are there costly information variables? How frequent are information releases? Unfortunately, there are no easy answers to these questions.

This is another area where technology is often taken from the artificial intelligence literature. In Arthur et al. (1997) a method known as a classifier system is used, which will be described later in this chapter and in Brenner (2005) and Duffy (2005). In Beltratti & Margarita (1992) and LeBaron (2001*b*) neural networks are used to represent trading strategies. However, strategies can be as simple as a vector of parameters as in Lettau (1997).

2.5 Social learning

How agents learn from each other is another important design question. This is often known as “social learning”, and has been the subject of much discussion in the agent-based modeling community.²⁶ At one

²⁶See Vriend (2000) for a description and examples.

extreme, agents may operate completely on their own, learning rules over time, and only reacting with others through common price and information variables. However, in financial settings it may be useful to try to implement some form of communication across agents, or even to transfer rule-based information across individuals from generation to generation. How this information transfer is handled may be critical in market dynamics; 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.

2.6 Benchmarks

The 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. Parameter sensitivities can reveal critical factors in a simulation that lead a market towards or away from an equilibrium. Finally, the dynamics of the learning process may be just as interesting in a neighborhood of an equilibrium as far away from an equilibrium. To make this distinction the definition of a benchmark is essential.

3 Artificial financial markets

It is easy to get lost in the many different types of models used in agent-based financial markets. Several approaches are used, and it is often difficult to distinguish one model 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 in the field. At such an early stage, it is still possible that some may argue about how markets are being categorized, or that some markets belong in multiple categories, or that the categories themselves are wrong. 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 Few-type models

Most of the earliest artificial financial markets carefully analyze a small number of strategies that are used by agents to trade a risky asset. The advantage of a small set of strategies comes in tractability, and in 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 et al. (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. Most of these models are analytic, but several with small strategy sets still require computational techniques to get their dynamics. These will be discussed here.²⁷

One of the earliest few-type financial market models was developed by Figlewski (1978). This market model examines the impact of shifting wealth across differentially-informed agents in a simple asset pricing framework. In this market agents possess a critical piece of information which might be unrealistic when considering real financial 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. 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 affect the convergence to an efficient market.

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 traders using portfolio insurance strategies can be very unstable. Since the strategy is well defined, this allows for a simple computational test bed to assess 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.

²⁷Other important early papers in this area which are discussed in Hommes (2005) are Beja & Goldman (1980), Brock & Hommes (1998), Chiarella (1992), Cont & Bouchaud (2000), Day & Huang (1990), Lux (1997), and Zeeman (1974).

3.2 Model dynamics under learning

The papers described in this section are more computational than those mentioned previously. In most cases the small sets of tractable trading rules are replaced with larger sets of strategies, which are usually represented using various computational techniques. These will be referred to as many-type models. This first section concentrates on applications where the economic environment is well understood and where there is often a simple homogeneous rational expectations equilibrium which gives a useful benchmark comparison.

Lettau (1997) provides a good example of a computational model of this type. He implements a financial market model with a set of heterogeneous learning agents, that is simple, transparent, and easy to implement. The model is a portfolio choice environment where investors must decide what fraction of wealth to put in a 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 risky asset has an exogenously given price and pays a random dividend, d , which follows a normal distribution. The second period wealth of agents is given by,

$$w = s(d - p), \quad (2)$$

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

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

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 straightforward 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,

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

$$\alpha^* = \frac{1}{\gamma\sigma_d^2} \quad (5)$$

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

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

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

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

This gives the agents a head start on the portfolio problem, but they still need to learn the optimal α .²⁸

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 genetic algorithm (GA) is run, and the set of agent parameters is redrawn. Agents are parameterized with a bitstring encoding given by

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

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

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

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

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

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 each of the two bitstrings, and then combines the first substring of one rule, with 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.

Lettau’s results show that in various specifications the GA can learn the optimal parameter for the

²⁸A more complicated functional form is tried, but the only change is that convergence is slowed down by the need to learn more parameters.

portfolio policy, nevertheless, there are some important caveats. First, the specification of S is crucial. For example, Lettau ran experiments for which the optimal value of α was $\alpha^* = 1.0$. With $S = 150$, he found that the experimentally-determined value of alpha in his agent population was 1.023. However, for $S = 25$ this average population alpha increased to 1.12, substantially different from the optimal value. 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 $S = 1$. Ex post it is optimal for S to be 0 or 1 depending only 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 that 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 cannot be viewed as an attempt to model an actual financial market, in which the dependence between today's price and traders' strategies is the most critical aspect of the agent-based modeling approach. However, it is a very clean and straightforward setup and hence a good learning tool. Also, the biases and sample size issues that it brings up will also pertain to many of the much more complicated models that will be considered later.²⁹

In Arifovic (1996) a much richer more extensive model is constructed. Once again, the model stays close to a well-defined theoretical framework while extending the framework to include 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,

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

$m_{1,t}$ and $m_{2,t}$ denote the money holdings of agents in the two currencies. There is only one consumption

²⁹Another simple example of this can be found in Benink & Bossaerts (2001).

good, which has a price in each currency. The exchange rate is given by

$$e_t = \frac{p_{1,t}}{p_{2,t}}. \quad (11)$$

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

$$R_t = \frac{p_{1,t}}{p_{1,t+1}} = \frac{p_{2,t}}{p_{2,t+1}}. \quad (12)$$

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

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

The model has a fundamental indeterminacy in that, if there exists one price series and an exchange rate paring that constitutes an equilibrium, then there will exist infinitely many such 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) explored this same question; he found that certain learning algorithms 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. A bitstring of length 30 is divided as 20 binary bits for consumption in period 1, and 10 for the fraction of savings put into currency 1. 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, never settling to any constant value. There is an interesting interpretation for this dynamic price process. 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, the exchange rate moves around as they shift demands between currencies. In a model such as this, it is clear that a constant exchange rate equilibrium can only be maintained through some mechanism that shuts down learning and exploration in the model. Arifovic also shows that similar features are obtained in experimental markets.³⁰

Routledge (2001) also critically examines what happens when leaning agents are introduced into a well-known model. He 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 are 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 uninformed agents falls, the ability of their population to learn decreases due to small sample size. Typically the end result is convergence to a situation in which all agents

³⁰These results are discussed in Duffy (2005).

are informed.³¹

3.3 Emergence and many-type models

The next set of artificial market models moves farther from testing specific models and more towards 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 of examining their coevolution over time. This methodology attempts to determine 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.

The Santa Fe Artificial Stock Market, SF-ASM, is one of the earliest in this set of models. It is described in Arthur et al. (1997), and also in LeBaron, Arthur & Palmer (1999).³² The basic objective of the SF-ASM is to understand the behavior of an environment of evolving trader behavior, where prediction strategies compete against each other. Part of this objective is to find if and when the market converges to a tractable rational expectations equilibrium. A second part is to explore the dynamics of the computational model for the cases in which convergence does not occur, and to compare these to results from real financial time series.

The basic economic structure of the market draws heavily on existing market setups such as Bray (1982) and Grossman & Stiglitz (1980). The traders have one-period myopic preferences of future wealth with constant absolute risk aversion (CARA) utility functions. 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,

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

where ϵ_t is gaussian, independent, and identically distributed, and $\rho = 0.95$ for all experiments. It is well

³¹Routledge (1999) presents similar results in an analytic framework.

³²There 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 an excess demand price adjustment mechanism. The later version uses a form of market clearing.

known that, assuming CARA utility functions, and Gaussian distributions for dividends and prices, the demand for holding shares of the risky asset by agent i is given by,

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

where p_t is the price of the risky asset at t , $\sigma_{t,i,p+d}^2$ 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 . Assuming a fixed number of agents, N , and a number of shares equal to the number of agents gives,

$$N = \sum_{i=1}^N s_i, \quad (16)$$

which closes the model.

The SF-ASM includes an important benchmark for comparison. 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,

$$p_t = b + ad_t, \quad (17)$$

where d_t is the only state variable. 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 the equilibrium holding of shares for each agent.

The most important part of the SF-ASM 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.³³ Traders build their own individual forecasts of future prices and dividends by matching specific forecasting rules to current market conditions. In the classifier system traders can use, or ignore, any part of a predefined set of current information in their forecasts. In the SF-ASM classifiers are used to select between different forecasts that are conditioned on certain pieces of market information. Information is coded into bitstrings, and each bit is connected to different ranges for various indicators. The information bits are classified either as fundamental or technical. Fundamental bits refer

³³Classifiers are not used extensively in economic modeling. Examples of other studies using classifiers are Marimon, McGrattan & Sargent (1990) and Lettau & Uhlig (1999). See Brenner (2005) and Duffy (2005) for more discussion.

to the current price relative to the current dividend level. Technical bits are trend following indicators that refer to the current price relative to a moving average of past prices.³⁴ A classifier forecasting rule is matched to a specified vector of these conditions, and corresponds to a linear price-dividend forecast of the form

$$E_{t,i}(p_{t+1} + d_{t+1}) = a_j(p_t + d_t) + b_j. \quad (18)$$

The classifier selects the appropriate real-valued pair, (a_j, b_j) . Therefore, the classifier selects a piecewise linear forecasting rule which is then used in the demand relationship (15). It is important to note that given the linear structure of the forecasting rule and the rational expectations equilibrium in (17), neither fundamental nor technical bits would provide additional information if the market were in the equilibrium.

At the end of each period, each trader with probability p engages in a learning process to update his current set of forecasting rules for the next period and with probability $(1 - p)$ leaves his current set of forecasting rules unchanged. The probability, p , is an important model parameter that determines the average number of periods between learning for each trader as a function $K = K(p)$. This K is referred to as the “learning rate.” Learning takes place with a modified genetic algorithm (GA) 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.³⁵ Mutation involves changing the individual bits at random. Crossover also is implemented on the real components of the forecasting rules too. This is one of the earlier applications of a real-valued crossover operator in finance.

One of the objectives of the SF-ASM 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 slow-learning experiment sets the average time between runs of the GA to $K = 1000$, and a fast-learning experiment sets the average time between runs to $K = 250$. In the first case, the market converges to the benchmark rational expectations equilibrium, where all agents agree on how to process the fundamental dividend information. They also ignore all other information. In the fast-learning experiments, $K = 250$, a very different outcome occurs. The market does not appear to converge, and it shows several indications of

³⁴The bits code these based on conditions. An individual bit would refer to the test $p_t/ma_t > 1$. If this is true the bit is set to 1, and if it is false it is set to 0.

³⁵Selection is by tournament selection. This means that, for every rule that is needed, two are picked at random and the strongest is taken.

interesting features in the stock return time series.³⁶ 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 SF-ASM does a good job in replicating these facts qualitatively, no attempt is made to quantitatively line them up with actual financial data. Indeed, the SF-ASM 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 SF-ASM is has formed a platform for other explorations. Joshi, Parker & Bedau (2000) explore the interactions between the technical and fundamental 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.

The SF-ASM market has generated much interest since its software is now publicly available. It was originally written in the programming language C, 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) provides 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.

Another important software replication issue arising from work with the SF-ASM is presented in Polhill, Izquierdo & Gotts (2005). These authors show that the precise trajectory dynamics in the SF-ASM can be altered by making mathematically irrelevant changes in the code. For example one might change,

$$d = \frac{a + b}{c} \tag{19}$$

to

$$d = \frac{a}{c} + \frac{b}{c}. \tag{20}$$

Although these two equations are the same, they generate different code in the compiler. This change appears

³⁶This parameter sensitivity is closely related to the changes observed in Brock & Hommes (1998) as the intensity of choice parameter is changed.

to have no impact on the general results, but it does impact the exact replication of trajectories. Runs using the two different forms will eventually diverge. This is an interesting reminder about the importance of nonlinearities inside these large systems, and on the difficulties in replicating exact trajectories across different computing platforms. While general statistical properties and features should be maintained, exact replications may be an elusive goal.

In addition to software critiques of the SF-ASM, there are also important design issues to consider. Many of these are covered in LeBaron (forthcoming 2005). Qualitatively, the classifier system has proved to be a very complicated and unwieldy way to model and understand the market dynamics. Many parameters are needed to define the operation of the classifier, and it not clear which of these is important. Also, the implementation of the classifier is often criticized. Ehrentreich (forthcoming 2005) addresses the GA and its impact on the classifier bitstrings. His claim is that the original SF-ASM GA mutation operator was biased, and he implements a new operator that he claims is unbiased. In his modified market bitstrings contains fewer 1's and 0's which connect forecasts to information bits. Also, the emergence of technical trading rules does not occur. This is an interesting modification, but the entire classifier system makes it difficult to judge what is unbiased in terms of mutation. There is a generalizer system which periodically removes bits in the classifier from rules that haven't been used recently. This puts a downward pressure on bit-setting in the Ehrentreich (forthcoming 2005) system. Second, it is not clear whether one has to have an unbiased mutation operator in terms of bitstrings. One could view a biased operator as putting many possible conditional rules out in public view, and then it is the agents' choice to ignore them. The traders are designed to ignore useless rules since the forecast performance of these rules will be inferior to the others. Disagreements about the "right" mechanism here indicate why the classifier system is a difficult to implement and completely understand. One major question about the classifier 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 pre-loaded. Finally, another important critique is that by assuming CARA utility functions, the SF-ASM 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 SF-ASM.

If the general goal of the financial 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 the model of Chen & Yeh (2001) is probably the best model directly addressing this problem. These authors use a computational methodology known as genetic programming to model agent learning. The authors allow the traders to

evolve actual predictor functions for financial forecasting.³⁷

The economic setup in Chen & Yeh (2001) is similar to the SF-ASM except that the price adjustment occurs in response to 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 learning is implemented. The forecast takes the form of

$$E_{i,t}(p_{t+1} + d_{t+1}) = (p_t + d_t)(1 + \theta_1 \tanh(\theta_2 f_{i,t})), \quad (21)$$

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

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.³⁸ The rules in the common pool are evolved according to forecast accuracy. Traders then decide to update their own strategies based on current performance. They draw rules from the common pool, comparing their performance with their current rules. If the new rule is better they switch, but if they are unsuccessful after several tries, they quit and stay with their current rule.

Simulations of this financial 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. For example, there is a large level of positive skew. 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 should therefore also be stationary.

Another financial market model is the Genoa artificial market, Raberto, Cincotti, Focardi & Marchesi (2001). In the original version of their market model the authors used random-order selection, meaning that buy and sell limit orders are generated at random by traders. Traders first determine whether they are a buyer or seller at random, and then place a limit buy or sell order determined by their budget constraints. These limit prices in each case are generated as random variables. In contrast to the previous markets, these traders are generally fairly unsophisticated. As in the study by Cont & Bouchaud (2000) they exhibit a kind of herding behavior. Buyers and sellers group into larger dependent sets, which then move together.

The Genoa artificial stock market has an interesting market-clearing property. The limit orders are all

³⁷Genetic programming is discussed in Brenner (2005) and Duffy (2005). 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).

³⁸This is the recurring theme of individual versus social learning; see Vriend (2000).

collected after a short period has gone by, and then the market is cleared by crossing the supply and demand curves given by the current limit orders. The market-clearing price is then used to clear all of the trades that can be executed on the limit order book. This interesting batch-order book market is very simple and direct. Similar to the other models discussed earlier, the Genoa market generates uncorrelated returns, fat tailed return distributions, and persistent price volatility.

The financial 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 models, 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,

$$E_{i,t}(p + t + 1) = f(p_{i,j,t-1}, \Delta p_{i,j,t-1}, \pi_{t-1}, \Delta \pi_{t-1}), \quad (22)$$

where π_{t-1} is the average transaction price at time $t - 1$ across all traders, $p_{i,j,t-1}$ is the last price execution that the trader received, and Δx refers to the one-period change in x . This is an interesting function because it implies the traders are using both local and global information. When two traders meet, they compare their price forecasts. The trader with the larger forecasted price then purchases 1 share from the trader with the smaller forecasted 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 included in the information sets of the traders. After a day of trading, traders are allowed to update the weights of their neural networks in a direction that they perceive will improve 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, c , that is paid directly by the traders. The simulations show the eventual heterogeneous population depends critically on the value of c . For low levels of c , traders purchase the extra network complexity, and for high levels of c , they eventually only use the simple networks. There is an

interesting mid- of c values where both types of strategies are able to coexist.

In all of the papers reviewed so far, the traders are assumed to behave competitively. That is, they view themselves as having no price impact, and they believe 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 the 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 traders receive a signal each period, and they also observe the trades of others. Their own trading strategies are based on optimally forecasting the final payment of the security using Bayesian updating from their initial priors. Though the individual strategies are analytically defined, the final dynamics of the market as a whole requires a computational experiment. The authors employ a novel approach to explore the impact of many different parameters. 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 parameter values, to observe their impact on empirical market outcomes. This may seem like a lengthy and indirect method to understand parameter sensitivity, but it may be important when there are many parameters, and when 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 the price is also a better forecast of future value. This implies that increased trading activity can lead both to greater price movements and to better learning and information-sharing through price 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 have less belief in their initial information, they have a greater incentive to glean knowledge from the better informed market as a whole. The authors' model allows for another interesting experiment. One of the parameters of their model is the threshold level at which a trade between two agents is noticed by other traders. Trades which are smaller than this threshold level go unnoticed, but the larger trades are observed. The authors find that reducing this threshold reduces price volatility and increases forecast accuracy. This is again suggests that, in the end, the learning processes in this sequential market are effective although not perfect.

3.4 Calibration

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 Memory and return autocorrelations

Levy, Levy & Solomon (1994) presents a financial market model with outcomes emerging from agent strategiest.³⁹ Similar to the market models covered above, these outcomes depend on the presence of many different heterogeneous agent types. However, the traders in Levy et al. (1994) do not form complicated strategies and predictors. Traders maximize a one-period myopic utility function exhibiting constant relative risk aversion rather than constant absolute 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 agent-based financial markets. There is a risk-free asset which pays a constant interest rate. There is a risky stock paying a random dividend that follows a multiplicative random walk,

$$d_{t+1} = d_t(1 + z_{t+1}), \quad (23)$$

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 fundamental traders who possess 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 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 who believes the older returns data are irrelevant.⁴⁰ 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.

The authors progressively add richer sets of the heterogeneous memory traders who trade along side the

³⁹This model is presented in the book, Levy et al. (2000), which also contains useful summaries of many other agent-based markets.

⁴⁰See Mitra (2005) and Sargent (1999) for analysis of short memory, mis-specified forecasting models.

fundamental traders. For sets with only one, or two memory types, the stock price dynamics clearly reflect the memory length, in that distinct cycles are observed. 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 stock returns demonstrate small positive autocorrelation over short horizons, but negative autocorrelation over longer horizons.⁴¹ 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 fundamental strategy, and the finite-memory strategies. They give some examples showing the two types can coexist with neither one evolutionarily driving the other out.⁴²

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 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

One of the most interesting empirical features that various financial market models try to replicate is the persistence of asset price volatility. While stock returns themselves are relatively uncorrelated, the squares or absolute values of returns are autocorrelated, reflecting a tendency for markets to move from periods of relative quiet to more turbulent periods. Significant positive autocorrelations for absolute stock returns continue out a year or more, and decay at a rate which is slower than exponential. This slow decay rate cannot be captured by traditional time series models, and may indicate the presence of fractional integration in volatility.⁴³ The mere fact that volatility is persistent is puzzling enough, but the fact that it may be fractionally integrated 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 these facts. The model

⁴¹See Hirshleifer (2001) for summaries of these empirical results.

⁴²In Levy et al. (2000) begin to explore some simple multi asset models. Their goal is to begin to understand how well the predictions of the Capital Asset Pricing Model hold up in heterogeneous agent situations. Their early findings are supportive of the CAPM, but the model only allows heterogeneity to enter in a limited way, through mean expectations.

⁴³See Baillie et al. (1996) for an example of a fractionally integrated volatility process. Also, see LeBaron (2001c) for further discussion of fractional integration in stock return series.

is based on the spatial spread of information across traders.⁴⁴ In this model each trader i in each period t receives a signal $Y_{i,t}$ that combines information about the decisions of this trader's local neighbors. For example,

$$Y_{i,t} = \sum_{(i,j)} J_{i,j} S_{j,t} + A\nu_{i,t}, \quad (24)$$

where $S_{j,t}$ are the decisions of other traders in the neighborhood of i , and $J_{i,j}$ controls the weighting and the neighborhood size, and $\nu_{i,t}$ is a noise term. $J_{i,j}$ declines as the distance between i and j increases. This signal is an input into a trader i 's final decision to purchase or sell one share of the stock. The interesting part of this decision is that agents are assumed to 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 trader i in turn feed into the signals of other traders. The traders' belief formation and demand processes are 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 who adjusts the price using

$$p_{t+1} = p_t \left(\frac{D_t}{Z_t} \right)^\alpha. \quad (25)$$

Stock returns are measured as the log difference of this price series, and the volatility is estimated with the absolute values of these returns. The model generates returns that are nearly uncorrelated, but the volatility series generates a very persistent autocorrelation pattern which is similar to actual asset return data. Further, the model is also able to display the strong positive correlation between trading volume and volatility that is observed in the data. It also appears that the thresholding of the signals is critical for volatility clustering to occur.

Kirman & Teyssiere (2001) develop another model capable of generating very persistent return volatility. It is a modified version of Kirman (1991) which is described more extensively in Hommes (2005). This model is a form of the earlier-mentioned few-type models in which 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 properties of returns generated by the model, and show good qualitative agreement with actual foreign exchange series

⁴⁴Other examples of this type of model are discussed by Hommes (2005). These examples include Cont & Bouchaud (2000) and Stauffer & Sornette (1999).

in terms of long range persistence in volatility.

3.4.3 Macro fundamentals

Several papers have taken the step of trying to tie markets to actual market fundamentals. In Farmer & Joshi (2002) 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 financial market model generates reasonable long swings away from the fundamental pricing as well as uncorrelated daily returns. It also generates most of the important empirical features described in previous sections, including, fat tails, volatility persistence, and trading volume persistence. The model also offers interesting tractability since it is built from a foundation of realistic trading strategies.

LeBaron (2001*a*) and LeBaron (2002*a*) perform some extensive calibration exercises. These exercises are based on an agent-based model presented in LeBaron (2001*b*). This model combines several features of the models mentioned previously. It uses a neural network structure to represent trader portfolio strategies. In this model traders do not build forecasts. The neural network maps past information directly into a recommended portfolio holding directly, and thus 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. Traders are defined by heterogeneous memory lengths as in Levy et al. (1994). Some traders evaluate strategies using a short 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 evolves separately from the traders; the traders choose strategies perceived to be optimal based on time series with lengths corresponding to the traders' memory lengths. 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 of the financial market models reviewed above in that there are only two traded assets, a risky asset and a risk-free asset. The risky asset pays a well-defined stochastic dividend following a geometric random walk with drift and volatility calibrated to match 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 volatility persistence, and volatility/volume cross correlations.⁴⁵

⁴⁵An interesting feature is that the model replicates the tendency for volatility to lead trading volume. This is consistent

These results appear to be connected to the presence of short-memory traders. Eliminating the latter group leads the market to converge to a well-defined rational expectations equilibrium. Other modifications are shown to improve learning and to induce the market to converge. Among these are slowing down the rate at which agents switch rules, and having them switch strategies only when a new strategy beats the current one by a certain threshold [LeBaron (2002b)]. Both of these operate to slow down the learning process, which one would think would make things worse.

The strategies used in this market are emergent in that they are not prewired into the model. It is interesting to note that the learning process does evolve as the market progresses. LeBaron (2001a) shows that in the early stages of the market, predictability is quite high. Regressions of returns on simple lagged returns can yield R-squared values as high as 0.7. These patterns are quickly learned by agents, however, and this unrealistically high predictability is greatly reduced. It is also interesting that the dividend-price ratio remains a consistently good predictor in many different time periods, which is consistent with results from real financial data.

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 the government issues money in the economy at a constant growth rate that 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 parameter values that are able to give them reasonable volatility in asset returns along with a low volatility in per capita consumption growth. For the most part, the parameter values that generate these results are consistent with U.S. macroeconomic data.

3.4.4 Other calibration examples

This section briefly summarizes several other calibration examples which try to line up with interesting data sets, and scrutinize time series generated by agent-based financial markets. Arifovic & Masson (1999) implement an agent-based model of foreign exchange currency crises which is aligned with empirical results from foreign exchange crises periods. Another model examining foreign exchange markets is Marey (2004) which uses foreign exchange survey forecasts to calibrate agent behavior. Finally, several papers such as Chen,

with results in Gallant et al. (1993).

Lux & Marchesi (2001) and Arifovic & Gencay (2000) perform detailed tests on the nonlinear properties of the time series output from various agent-based financial markets. They general find evidence similar to that from actual markets.

3.5 Estimation and validation

While many market models have been calibrated to financial time series, very few computational models have attempted to actually fit parameters to data in a direct estimation procedure. Obviously, in most computational models this will be a costly procedure in terms of computer time. 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 and objective are relatively simple, this paper provides the most detailed view into the sensitivity of the results to various parameter specifications.

Estimation of few-type models has been a much more common activity, and has already yielded some interesting early results. The simpler structure of these models permits their conversion into tractable, albeit nonlinear, time series structures. Vigfusson (1997) is one of the first papers to estimate one of these models. The framework is based on a model by Frankel & Froot (1988) for studying exchange rate movements, which was mentioned in section 3.1 and is also covered by Hommes (2005). The model is implemented empirically as a Markov switching model, as in Engel & Hamilton (1990), where the two states correspond to fundamental and chartist regimes. Exchange rate predictions are generated as a weighted average of the two different regimes, where the weights are given by conditional probabilities of the two states. Some general support is given to the different conditional forecasts in different states of the world, but some of the results are mixed. Ahrens & Reitz (2005) is a more recent test of a Markov switching model. They find better evidence in favor of the model with chartists and fundamentalists, and they also test several different specifications for the chartist regime. They see an interesting connection between volatility and the types of traders, and many of these results appear robust across different subsamples. They also document an interesting result that volatility is larger during the fundamental regime. This result is interesting, but a little difficult to explain.

Westerhoff & Reitz (2003) and Reitz & Westerhoff (2004) fit a nonlinear threshold model of a financial market to various time series. This model is also inspired by the few-type models with chartists and fundamentalists trading in a market. The model results are generally supportive of the transition between the two different types of trading strategies. They find different market dynamics depending on how close the

price is to fundamental value. These are an interesting first test of heterogeneous trader behavior. More extensive tests will be necessary to judge the general robustness of this modeling framework.

A common concern about all agent-based computational modeling is validation. The usual criticism leveled at agent-based financial markets is that there are too many degrees of freedom. Researchers are able not just to move freely through large parameter spaces, but can also change entire internal mechanisms at their discretion in the attempt to fit sets of stylized facts. Anyone using agent-based financial markets must acknowledge that there is some truth to these criticisms. However, these comments should not stop all experimentation. Furthermore, there are directions in which the field is moving that will give these markets a more solid foundation.

Some steps that researchers can take to ameliorate these problems include replicating difficult empirical features, putting parameters under evolutionary control, and using results from experimental markets. The first of these suggestions involves making sure that an agent-based financial market fits facts which are not well replicated by standard models. Examples of this would be empirical features such as the long range persistence of volume and volatility in financial time series. This requirement sets a higher standard for empirical replication, and also pushes the envelope in terms of our understanding about which mechanisms may be at work in financial markets.⁴⁶ The second suggestion is to put as many parameters as possible in the market under evolutionary control. An example of this change is reflected in the differences between markets such as the SF-ASM, and LeBaron (2001*b*). In the first case fixed learning rates for all traders implicitly give them a common perspective on how much past data is allowed into fitness evaluation. It turns out that this parameter is crucial to the behavior of the market. In the second case, traders with different perspectives on the past compete against each other. If there were an optimal memory length of past data to use, this would dominate the market in terms of wealth. Thus the setting of this effective value of this parameter is reflected in the wealth distribution of traders in the model, and is part of the evolutionary dynamics. The final suggestion would be to use results from experimental economics to build better, and more realistic learning dynamics in the artificial financial markets. This seems like a promising procedure, but as yet there are not that many examples of it.⁴⁷

⁴⁶An extension of this is to concentrate model fitting on extreme event periods in the market as in Ecemis, Bonabeau & Ashburn (2005).

⁴⁷Duffy (2005) surveys agent-based models and experimental work. There are several examples given there of agent-based learning mechanisms which fit experimental data. Whether such mechanisms could be taken into settings that are more complicated than the human experimental settings is an interesting and open question.

4 Other markets

This section covers several financial market models which are different from those considered above. Among these are markets which consider detailed trading institutions and learning market makers, and also models which consider the coevolution of strategies and financial securities.

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. Most models build stylized economic structures that avoid the institutional details of trading. However, research has begun appearing which implements more realistic trading systems.⁴⁸ Market design and market microstructure questions appear to be well suited for agent-based approaches. First, large amount of data are available. Second, there are critical policy questions which clearly need to be tested in an environment with heterogeneous, adaptive strategies. Since some of this area is covered in other handbook chapters, the descriptions of models here will be relatively brief.

It is interesting that Rieck (1994), one of the earliest agent-based financial market studies, specifies the trading process in detail. Rieck (1994) looks at the evolution of trading strategies with a simple order-book trading mechanism. His model has many similarities to some of the emergence papers mentioned in the previous sections in that the coevolution of strategies is the crucial issue of interest. Also, strategies are evolved using evolutionary techniques, but these are applied to functional forms that are designed to replicate actual trading strategies. His results show that fundamental strategies are not able to take over the market and drive the price to the fundamental value. Rieck (1994)'s findings suggest that many results obtained in agent-based financial market models without detailed specifications for trading strategies could be replicated using more empirically-based micro trading mechanisms.⁴⁹

Much simpler models have been implemented using an order-book trading mechanism with the goal of replicating empirical features of actual market data. One example of this is Chiarella & Iori (2002). This is a few-type model with technical, fundamental, and noise traders placing orders in an electronic order book system. They show that the interaction of these three types generates realistic price series, and trading activity. Just how much agent complication is necessary to replicate high frequency features in financial time series is a question which is addressed in Farmer et al. (2005). This is not exactly an agent-based market since order flow is completely random and there is no individual trading agent per se. Random order flow

⁴⁸This area of financial research overlaps with work on market design which is covered more extensively in Mackie-Mason & Wellman (2005) and Marks (2005).

⁴⁹Yang (2002) replicates many of the SF-ASM results with a microstructure foundation.

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. Even though traders in this model have no learning capabilities, the authors are able to show that the outcomes of their model replicate many features from actual price history data sets. As in the work by Gode & Sunder (1993) on “zero intelligence traders,” these results help us to understand which empirical regularities are the result of learning behavior, and which are simply a feature of the trading institution.⁵⁰

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 that uses 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 overnight. The reservation prices for these dealers are determined in a Bayesian learning framework. Each dealer 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 values. The author chooses 729 unique parameter value combinations, performs simulation runs for each combination of parameters values 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 of return volatility over the simulated trading

⁵⁰A very early example of this type of research on random order flow is Cohen, Maier, Schwartz & Whitcomb (1983). Another recent research direction has been to link electronic trading agents to live data feeds coming off of actual markets. In the Penn-Lehman Trading Project, Kearns & Ortiz (November/December 2003) the survival of different strategies can be monitored as they interact with live market data.

days, and these results are robust across many of the tested parameter combinations. 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 parameter values 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.

Most financial markets depend critically on the behavior of a market maker, or dealer, who brings liquidity and continuity to a real time trading market. Several agent-based models explore the behavior of dealers. Gu (1995) takes the model of Day & Huang (1990) and explores changing the market maker behavior built into the model. This analysis includes estimating the market maker profitability under different parameters. The results show that a profit-maximizing specialist may be interested in generating some amount of extra market churning. The specialists objectives will not align with price variance minimization which could be construed as the maintenance of an orderly market. Westerhoff (2003*c*) also explores the impact of inventory restrictions in a setup with an implied market maker. The market maker price adjustment reactions differ depending on the current inventory position along with current excess demands. The market maker is assumed to make greater price adjustments when these two variables are of the same sign. Increasing this adjustment level leads to increased volatility. Although interesting, this result does depend critically on very specific behavioral assumptions made for the market maker.⁵¹

Most of the papers considered in this survey could loosely be considered part of the investment 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 issues in an agent-based framework. The authors are interested in the problem of which securities firms will issue to raise investment capital, and how investors learn 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.

In Noe et al. (2003) a firm has an investment project that needs to be financed, and there are two potential

⁵¹A related paper is Chan & Shelton (2001) which models a dealer learning optimal behavior when faced with a random order flow. Further research in the area of market design includes papers examining tick sizes (Darley, Outkin, Plate & Gao (2000) and Yeh (2003)), order book versus dealer markets (Audet, Gravelle & Yang (2001), and price limits, trading taxes, and central bank intervention (Westerhoff (2003*b*), Westerhoff (2003*d*), and Westerhoff (2003*a*).

investors. The firm can choose from a fixed set of six different securities that it can issue. These include standard debt and equity securities, along with some more complex ones. The latter include convertible and subordinated debt, as well as something known as a "do-or-die" security. In each case the security represents a well-defined contract for splitting the payout to the firm's risky project between the firm and the two investors. Both the firm and investors encode their strategies as bitstrings for use with a GA. The firm maintains a pool of 80 potential security issue decisions which is a vector of numbers (binary coded) between one and six corresponding to the six types of securities. The firm will choose 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 that 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 encoded as bitstrings. 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 chooses a pricing strategy at random from the appropriate security table after the firm has decided on the security that it will issue.⁵² 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 in which 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 firms 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 situation where the firm is not able to raise sufficient investment funds actually occurs more often with two-sided learning than investor-only learning.

The results in this paper will eventually need to be explored under different learning specifications and

⁵²As in the earlier GA papers there is a binary-to-real mapping that determines the real valued price.

investment structures, but it is an interesting first attempt to use agent-based models in the field of corporate finance. The coevolution of agent behavior along with the institutions that guide this behavior is interesting both for finance and for economics in general.

One final agent-based model which is often compared to financial markets is the minority game.⁵³ This is a repeated game in which agents must choose one of two doors, left or right. If the minority of agents chooses left this group wins, and if the minority chooses right this group wins. The connection to finance is through the notion of contrarian strategies, where it is best to move against the herd. Connecting this model to finance is a controversial subject since its basic version does not have a natural role for prices. Also, it would appear that the contrary nature of the minority game is somewhat forced, and in real financial markets it may be better to follow the herd for a short period of time. An interesting application of the minority game to financial data is Johnson, Lamper, Jefferies, Hart & Howison (2001). In this model the authors convert a financial series into a binary string depending on whether the price rises or falls. The agents play the game for many periods watching the real financial time series as the input into their rule selection process. 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.⁵⁴

5 Cautions and criticisms

Agent-based markets have been criticized from many different angles. The most common criticism is that the models have far too many parameters, and the impact of many of these parameters is not well understood. This issue has already been discussed in the section on calibration. However, beyond simple parameter questions, these models have made use of a wide selection of the available computational tools and methods. Table 1 gives a short overview of the design structures of some of the agent-based financial market models described in this paper. This is far from being an all inclusive list, since many of the models described in this chapter would not fit well into the criteria for the list. This emphasizes what should have become clear from the earlier sections: agent-based financial models have been built using many different features and designs. This is natural for a field at this early stage, but it has made comparisons across market platforms difficult.

⁵³There are several early implementations of this model. These include Arthur (1994), and Challet & Zhang (1997). However, early versions of similar models can be found in Schelling (1978). See Jefferies, Hart, Hui & Johnson (2000) for a recent survey. Interested readers should go to the website for the minority game at <http://www.unifr.ch/econophysics/minority/>.

⁵⁴Another agent-based model indirectly related to finance is the resource allocation setup in Youssefmir & Huberman (1997).

Unlike analytic models, there are still relatively few general principles that one can confidently apply to the construction of different agent-based market models. This is a problem, but the situation should improve as the field evolves.

(Insert table 1 about here.)

Another important issue that is brought up is the stability of a given agent-based model's results to the addition of new trading strategies. Specifically, are there strategies that would smoke out obvious patterns in the data and change the dynamics? Agent-based models are trying to continuously defend against this with the continuously 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 arbitrage opportunities for new trading strategies to exploit. However, there is a possibility that more complex nonlinear strategies could detect such opportunities. Arifovic (2001) is an example testing this sort of issue, and finds that the more complicated agents do not do better in her 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 pertinent criticism is that most agent-based financial market models assume a small number of assets. Often agents trade only one risky asset, and one risk-free asset alone.⁵⁵ It is certainly true that, with all of the new methodological tools in use in these models, it was important to start with the simplifying case of one risky and one risk-free asset. However, this simplification may eliminate many interesting features. The criticisms of models with a single representative agent may carry over equally well to models with a single representative risky asset. Questions naturally arise about calibrating to aggregate dividends, and exactly what this calibration means, since aggregate dividends are not paid by any single stock. Also, recent events such as the technology bubble of the 1990s remind us that bubbles are often very sector dependent. 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 agent-based markets have two assets, they actually shut down pricing in one market. In many cases the risk-free rate is fixed, hence the market is not a general equilibrium model. This is problematic in that explaining 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 actual macro

⁵⁵Two recent exceptions to this are Chiarella, Dieci & I. Gardini (2004) and Westerhoff (forthcoming 2004). Also, Levy et al. (2000) perform some experiments in multi-asset settings with options.

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. Once researchers are more confident they have mastered agent-based modeling tools, they will probably tackle multi-asset market modeling 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 that can be studied analytically, and they find that the dynamics can change dramatically as the number of agents becomes large. What initially looks like random behavior for a small numbers of agents can become increasingly predictable as the number of agents becomes very large. Is it possible that many of the nice features that many models display are artifacts of the limitation to relatively small numbers of traders imposed by computer modeling? This is a very important question. One response to this question is that assuming an infinite number of agents might not be realistic in some settings. There may be real-world market situations in which the thinness of the market is an important and critical issue for the determination of the market's dynamics. 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 adopt rules and forecasts which have performed well in the recent past, and they adjust these rules and forecasts to perform better in the future. The early spirit of agent-based models is clearly to push away from more traditional deductive styles of learning and towards more inductive styles of learning. 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 have worked in the past? 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 that allows for a little deductive reasoning while keeping the general inductive spirit of simple rules of thumb.

A final problem, often ignored, is 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 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 degree 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 since the

strategies themselves adapt to the specific timing and trading structures. It will be interesting to see if agent-based financial models start permitting actions to take place more asynchronously, and if this has an impact on any of the early results.

6 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 few-type models covered in section 3.1, in which traders are assumed to choose from among relatively small fixed sets of trading strategies, to the many-type models covered in sections 3.2 and 3.3 in which traders choose from among large and possibly evolving sets of trading strategies. The few-type models offer an important dimension of tractability relative to the many-type 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, a key reason for doing computer modeling is that the use of more sophisticated trading strategies in many-type models needs to be understood as well. There are two basic reasons for this. First, many-type models take emergence very seriously in that they do not bias toward any particular strategy loaded *ex ante* by the researcher. The strategies that end up being used are those that appear and persist inside a learning structure. They therefore partially answer a criticism of the few-type models that their specification of trading strategies is *ad hoc*. Second, they use the computer and the learning algorithms to continuously search the time series record to smoke out new trading opportunities. This is something that is not present in the few-type models. The obvious limitation is that their ability to seek out and take advantage of any inefficiencies that may appear depends critically on the data representations and implementations of the learning algorithms. Few-type and many-type models clearly each have both strengths and weaknesses that users should take into account.

Agent-based modelers are also starting to move from the more stylized earlier financial market models toward more models incorporating explicit market microstructure. 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 to be 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 agent-based models have this characteristic. Where agent-based financial market models have diverged to date from behavioral finance models is their typical presumption that agent preferences have relatively standard representations. Typically, no attempt is made to model common behavioral biases such as loss aversion or hyperbolic discounting. This is not because agent-based models cannot handle these behavioral aspects. Rather, it has just seemed sensible in this early stage of the field to refrain from adding 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 share this feature. 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 modelers. Financial markets may be one of the important early areas where agent-based methods show their worth, for two basic reasons.

First, the area has many open questions that more standard modeling approaches have not been able to resolve. Second there is a large amount of financial data available for testing. It will be interesting to see if, sometime in the future, financial economists 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

- Ahrens, R. & Reitz, S. (2005), 'Heterogeneous expectations in the foreign exchange market: Evidence from daily DM/US dollar exchange rates', *Journal of Evolutionary Economics* **15**, 65–82.
- 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. (2004), ‘Risks for the long run: A potential resolution of asset pricing puzzles’, *Journal of Finance* **59**, 1481–1509.
- Barber, B. & Odean, T. (2000), ‘Trading is hazardous to your wealth: The common stock investment performance of individual investors’, *Journal of Finance* **55**, 773–806.
- 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.
- 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.
- Brenner, T. (2005), Agent learning representation: Advice on modeling economic learning, in K. L. Judd & L. Tesfatsion, eds, ‘Handbook of Computational Economics’, Elsevier.
- 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., 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.
- Chiarella, C., Dieci, R. & I. Gardini (2004), Diversification and dynamics of asset prices under heterogeneous beliefs, Technical report, University of Technology, Sydney, Australia.
- Chiarella, C. & Iori, G. (2002), 'A simulation analysis of the microstructure of double auction markets', *Quantitative Finance* **2**, 346–253.
- Cohen, K. J., Maier, S. F., Schwartz, R. A. & Whitcomb, D. K. (1983), 'A simulation model of stock exchange trading', *Simulation* **41**, 181–191.
- 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. (2005), Agent-based models and human subject experiments, in K. L. Judd & L. Tesfatsion, eds, 'Handbook of Computational Economics', Elsevier.
- Ecemis, I., Bonabeau, E. & Ashburn, T. (2005), Interactive estimation of agent-based financial markets models:Modularity and learning, in 'Proceedings Genetic Evolutionary Computation Conference 2005', ACM Press.
- 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. (forthcoming 2005), 'The Santa Fe artificial stock market re-examined - suggested corrections', *Journal of Economic Behavior and Organization* .
- Engel, C. & Hamilton, J. D. (1990), 'Long swings in the dollar: Are they in the data and do markets know it?', *American Economic Review* **80**, 689–713.
- 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. (2005), 'The predictive power of zero intelligence models in financial markets', *Proceedings of the National Academy of Sciences of the United States of America* **102**, 2254–2259.

- 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. (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.
- Gu, M. (1995), 'Market mediating behavior: An economic analysis of the security exchange specialists', *Journal of Economic Behavior and Organization* **27**, 237–256.
- Hansen, L. & Singleton, K. (1983), 'Stochastic consumption, risk aversion, and the temporal behavior of asset returns', *Journal of Political Economy* **91**, 249–265.
- 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. (2005), 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 Department, 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 dilemma: 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. & 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. (2001a), ‘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. (2001b), ‘Evolution and time horizons in an agent based stock market’, *Macroeconomic Dynamics* **5**(2), 225–254.
- LeBaron, B. (2001c), ‘Stochastic volatility as a simple generator of apparent financial power laws and long memory’, *Quantitative Finance* **1**, 621–631.
- LeBaron, B. (2002a), Calibrating an agent-based financial market, Technical report, International Business School, Brandeis University, Waltham, MA.
- LeBaron, B. (2002b), ‘Short-memory traders and their impact on group learning in financial markets’, *Proceedings of the National Academy of Science: Colloquium* **99**(Supplement 3), 7201–7206.
- LeBaron, B. (forthcoming 2005), 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. (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.
- Lux, T. (2002), Financial power laws: Empirical evidence, models and mechanisms, Technical report, University of Kiel, Kiel, Germany.
- Mackie-Mason, J. K. & Wellman, M. P. (2005), 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.
- Mantegna, R. N. & Stanley, H. E. (1999), *An introduction to econophysics: Correlations and complexity in Finance*, Cambridge University Press, Cambridge, UK.
- 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.
- Marks, R. (2005), Market design, in K. L. Judd & L. Tesfatsion, eds, 'Handbook of Computational Economics', Elsevier.

- McQueen, G. & Vorkink, K. (2004), 'Whence GARCH? A preference-based explanation for conditional volatility', *Review of Financial Studies* **17**, 915–949.
- Mehra, R. (January/February 2003), 'The equity premium: Why is it a puzzle?', *Financial Analysts Journal* pp. 54–69.
- 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.
- Mitra, K. (2005), 'Is more data better?', *Journal of Economic Behavior and Organization* **56**, 263–272.
- 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.
- Polhill, J. G., Izquierdo, L. R. & Gotts, N. M. (2005), 'The ghost in the Model (and other effects of floating point arithmetic)', *Journal of Artificial Societies and Social Simulation* **8**.
- Raberto, M., Cincotti, S., Focardi, S. M. & Marchesi, M. (2001), 'Agent-based simulation of a financial market', *Physica A* **299**, 319–327.
- Reitz, S. & Westerhoff, F. (2004), Commodity price cycles and heterogeneous speculators: A STAR-GARCH model, Technical report, University of Osnabruck, Dept. of Economics, Osnabruck, Germany.
- 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. (1999), 'Adaptive learning in financial markets', *Review of Financial Studies* **12**, 1165–1202.
- 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.
- Sargent, T. (1999), *The Conquest of American Inflation*, Princeton University Press, Princeton, NJ.
- 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.
- Solow, A., Costello, C. & Ward, M. (2003), ‘Testing the power law model for discrete size data’, *The American Naturalist* **162**(5), 685–689.
- 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.
- Tesfatsion, L. (forthcoming 2005), Agent-based computational economics: A constructive approach to economic theory, in K. L. Judd & L. Tesfatsion, eds, ‘Handbook of Computational Economics’, Elsevier.
- Vigfusson, R. (1997), ‘Switching between chartists and fundamentalists: A Markov regime-switching approach’, *International Journal of Finance and Economics* **2**, 291–305.
- 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. (2003a), ‘Central bank intervention and feedback traders’, *International Financial Market, Institutions, and Money* **13**, 419–427.
- Westerhoff, F. (2003b), ‘Heterogeneous traders and the Tobin tax’, *Journal of Evolutionary Economics* **13**, 53–70.
- Westerhoff, F. (2003c), ‘Market-maker, inventory control and foreign exchange dynamics’, *Quantitative Finance* **3**, 363–369.

- Westerhoff, F. (2003*d*), ‘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* .
- Westerhoff, F. H. & Reitz, S. (2003), ‘Nonlinearities and cyclical behavior: The role of chartists and fundamentalists’, *Studies in Nonlinear Dynamics and Econometrics* **7**.
- 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.
- Zeeman, E. (1974), ‘On the unstable behavior of stock exchanges’, *Journal of Mathematical Economics* **1**, 39–49.
- Zschischang, E. & Lux, T. (2001), ‘Some new results on the Levy, Levy, and Solomon microscopic stock market model’, *Physica A* **291**, 563–573.

Table 1: *Model Structures*

Authors	Preferences	Price determination	Evolution, Fitness	Strategy representation
Arifovic (1996)	CRRA	Market clearing	GA, utility	Real parameters
Arthur et al. (1997)	CARA	Market clearing	GA, forecast	Classifier
Beltratti & Margarita (1992)	CRRA	Random matching	Hill climbing, forecast	Neural network
Bullard & Duffy (2001)	CRRA	Market clearing	OLS, forecast	Real parameters
Chen & Yeh (2001)	CARA	Price adjustment	GP, forecast	GP functions
Chiarella & Iori (2002)	None	Order book	None, none	Real parameters
Farmer & Joshi (2002)	None	Price adjustment	None, none	Real parameters
LeBaron (2001 <i>b</i>)	CRRA	Market clearing	GA, utility	Neural network
Lettau (1997)	CARA	Exogenous	GA, utility	Real parameters
Levy et al. (1994)	CRRA	Market clearing	None, utility	Real parameters
Raberto et al. (2001)	None	Order book	None, none	Real parameters
Routledge (2001)	CARA	Market clearing	GA, utility	Real parameters
Tay & Linn (2001)	CARA	Market clearing	GA, forecast	Fuzzy logic

This is a short description of some of the multi-agent computational models considered here along with their design structures described in section 2. Preferences describe the types of preferences used by agents. Price determination describes the method for determining asset prices. Evolution refers to which computational evolution mechanisms, if any, are used. Fitness is the fitness measure used to evolve strategies, and to determine agent strategy choices. Strategy representation is the way strategies are stored in the computer. Often this is a predefined functional form, and the representation is simply a vector of real parameters. GA stands for the genetic algorithm. CARA and CRRA are constant absolute risk aversion, and constant relative risk aversion, respectively.