

A Builder's Guide to Agent Based Financial Markets

Blake LeBaron *

Brandeis University

February 2001

This is a draft of an article that has appeared in:

Quantitative Finance, vol. 1, no. 2, pp 254-261.

Abstract

This paper is intended to guide researchers interested in building their own agent based financial markets. Key design questions are outlined, along with some of the major controversies about which directions to take.

* Graduate School of International Economics and Finance, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a faculty research fellow at the National Bureau of Economic Research, and an external faculty member at the Santa Fe Institute. The author is grateful to J. D. Farmer for useful comments and suggestions.

1 Introduction

Agent based financial markets are an exciting new tool for exploring behavior in financial markets that are far from traditional notions of equilibrium, and involve behavior that is less than fully rational at times. Researchers from several disciplines including economics, computer science, physics, and psychology, are contributing to this very active field. It represents one of the most truly interdisciplinary areas of study operating today.

Even with all this activity new researchers interested in entering the field are faced with a daunting list of design choices. This list can be so large, that some people have become frustrated and returned to more traditional methods for studying finance. This paper is written to try to help those getting started in sorting out some of these questions. It outlines some of what is known about several of the issues, and tries to give new researchers a set of questions that they will need to address.

This paper is not really a survey, but a kind of view from the trenches in terms of building artificial markets.¹ The next section emphasizes what makes financial markets a particularly interesting area to study with agent based methods. The following sections categorize the questions that most market designers will face as they start work in this area. The final section concludes with some ideas about where the field is headed.

2 What makes financial markets special

Agent based methods have been applied in many different economic environments.² However, there are some interesting issues which make financial markets a very exciting place to use these new technologies. They are obviously some of the most dynamic markets in existence, and probably draw more interest from outside economics than any other subfield. They collect very high quality, high frequency data, which gives them an extensive set of facts and features from which to draw. Also, they are well organized, centralized, and trade homogenous products in a generally efficient fashion relative to markets for other goods and services.

Beyond these basic issues are others which are closer to deeper questions in economic theory. The general question of market stability and price formation looms all through economics. It is especially acute in finance

¹Readers interested in recent surveys should consult LeBaron (2000) or Levy, Levy & Solomon (2000). A few early examples of agent based financial markets are: Arifovic (1996), Arthur, Holland, LeBaron, Palmer & Tayler (1997), Bak, Paczuski & Shubik (1997), Beltratti & Margarita (1992), Caldarelli, Marsili & Zhang (1997), Cont & Bouchaud (2000), Farmer (1998), Farmer & Joshi (forthcoming 2001), Kim & Markowitz (1989), Kirman (1993), LeBaron (forthcoming 2001*b*), Levy, Levy & Solomon (1994), Lux (1997), de la Maza & Yuret (1995), Rieck (1994), Steiglitz, Honig & Cohen (1996), Youssefmir & Huberman (1997).

²See the extensive web site of Leigh Tesfatsion at <http://www.econ.iastate.edu/tesfatsi/ace.htm>.

since these markets, more than others, have prices both providing information along with balancing supply and demand. A price increase may induce agents to buy more or less depending on whether they believe there is new information carried in this change. More than any other market, financial markets lay bare the possible behaviors of a large macro system which can act as efficient social coordinating mechanism as suggested in Hayek (1945), or where individual optimizing behavior does not necessary lead to a socially efficient outcome as in Schelling (1978). Financial markets appear to be the perfect place to attack such questions which are at the core of the applicability and importance of complex systems issues in the social sciences.

3 Design issues

3.1 Agents

The most important design question faced in market building comes in the representation and structure of the actual trading agents. Agents can vary from simple budget constrained zero intelligence agents as in Gode & Sunder (1993) to sophisticated genetic programming models as in Chen & Yeh (2001). This variation in design is due to the fact that trading agents must solve a poorly defined task. Their directive is to digest the large amounts of time series information generated during a market simulation, and convert this into portfolio decisions. Given that there are many ways to process this past data, there must be as many ways to construct trading agents.

Obviously, one needs to make restrictive assumptions to get started, leading researchers in several different directions. The simplest and most direct route is to model agents as well defined dynamic trading rules modeled more or less as strategies used in the real world. This method can lead to very tractable precise results, which gives incites about the interactions between trading rules. The simplicity and analytic tractability of this method comes with a few costs. One of these is the fact that the dynamic interactions are related to these hard wired strategies. What if one were left out, or what if it was in the agents' interest to slightly modify their current strategies? Many markets of this type assume that the trading strategies will continue without modification, although the wealth levels under their control may be diminishing to zero. This leaves some open questions about this coevolutionary dynamic with only a limited amount of new speciation. A second critique is that agents in these markets do not operate with any well defined objective function. There is some usefulness to having well defined objective functions for the agents. If they are endowed with utility functions, or other related objectives, one can measure how well they are doing in terms

of their own limited information processing.³ Utility functions also serve a useful roll in policy evaluation in that they shed light on the relative valuations of certain tradeoffs. An example might be a policy to reduce market volatility which also makes trading less efficient. There may be important tradeoffs involved which only a model with well defined objective functions can answer. A related type of agent operates with well defined objective functions, but simplicity is maintained by limiting information and strategy complexity. This again leads to a fairly tractable market as in Levy et al. (1994).

A second type of agent design moves backward on the evolutionary scale, but has yielded a tremendous amount of information about market mechanisms. This is the “zero intelligence” agent who behaves randomly subject only to a budget constraint.⁴ Zero intelligence traders have been shown to generate very efficient trading outcomes when placed in trading mechanisms modeled after actual markets. Also, their behavior sometimes appears to the observer to look like learning as their budget constraints guide their behavior. These extremely simple agents do not push the frontier of artificial intelligence, or trading agent design, but they are an important reminder of the importance of the overall institutional constraints involved in financial markets. Also, they remind us to extend the question of evolution beyond the trading agents to the actual institutions involved.⁵

Finally, agents may be modeled as learning and adapting strategies as in markets such as the Santa Fe artificial market (SFI market) (Arthur et al. 1997), or Arifovic (1996). These markets use a variety of techniques from artificial intelligence to model continually changing agent strategies. This allows for the possibility of agents learning how to exploit new market inefficiencies. Also, aggregate dynamics in these markets are closest to true emergence in that very little is loaded into the agent ex ante. An interesting feature of many of these markets is that agents are very homogeneous at the start in their abilities, and strategy structures. Differences in behavior and strategy use evolve endogenously as the market runs. Agent heterogeneity becomes a changing feature of the market which can then be studied. The diversity of strategies can often be important in studying traditional finance questions related to market liquidity, or how hard it is to find someone to trade with. Endogenous heterogeneity features also distinguish many of these markets from the large number of heterogeneous agent rational expectations models in economics and finance.

This very general agent modeling approach comes with an increase in computational complexity. Tools

³An example of this is LeBaron (forthcoming 2001 *b*) which tests traditional agent first order conditions subject to the limited information processing of agents.

⁴The origin of the idea that apparent order in economic behavior may be molded as much by the agent’s environment as by their inner desires goes back to Becker (1962), and it is mentioned extensively in Simon (1969). However, the paper by Gode & Sunder (1993) introduced the concept of ZI traders, and Roth (1995) gives a nice comparison to other learning experiments in economics.

⁵See Hodgson (1999) for extensive examples of the institutional side of evolution.

such as artificial neural networks, classifiers, and genetic algorithms become the substrate on which agents are built. These artificial intelligence tools are still poorly understood operating alone, and in groups even less is known. Also, these models are not free of assumptions about trading behavior. One would like to think of the trading rules as evolving out of a pure unformed soup of genetic material, but in actuality assumptions must be made about information, and the structure of trading rules. Rarely are traders able to go out and create totally new indicators from scratch. Usually, by providing various statistical building blocks the designer limits this search.

Another related criticism is that the agents still may be leaving too many obvious trading opportunities on the table. In this case it might be said that the learning agents are simply not smart enough, and a better agent could capitalize on them.⁶ As agent based markets progress this criticism may prove to be one of the best arguments for this approach. One dilemma in this area is where to set the bounds for boundedly rational agents. In the adaptive learning agent world strategies can be allowed to become more complex through evolution.⁷ However, more complicated strategies can become overly specialized and brittle, leading to their eventual demise. This suggests that the bounds in strategy complexity may be driven more by over specialization in a changing environment than to actual limitations in mental capabilities. Hopefully, the this notion of a complexity bound might be easier to simulate and test for than those related to human brain processing power.⁸

An issue that agent design has not accomplished is in understanding how to think about dynamic multi-period preferences and learning. Few papers have tackled the problem of learning intertemporal decisions in the agent finance world. There are no clear routes to learning intertemporal plans in a simple robust fashion while identifying and enumerating state variables in the process. Classifier systems are one possible method for doing this, but they have been shown to have difficulties in solving simple intertemporal consumption plans in economics (Lettau & Uhlig 1999). It is also not clear how well actual people stay with long term plans in financial markets as opposed to simply adapting to current data, and following simple rules of thumb.⁹

A final issue in agent design that has not been explored is whether agents should evolve forecasts of future prices which are then fed into a decision making framework, or if they should evolve their decisions directly. An example of the former would be the SFI artificial market, and Arifovic (1996) is an example of the

⁶An interesting example with agents of varying capabilities is in Beltratti & Margarita (1992).

⁷For an example of this from game theory see Lindgren (1992).

⁸See Bookstaber & Langsam (1985) and Bookstaber (1999) for some examples of this in both finance and biology.

⁹See O'Donoghue & Rabin (1999) for examples of how well people hold to intertemporal optimization plans.

latter. Simply evolving decision rules directly would appear to be a much simpler and cleaner approach, but there can be advantages to thinking about forecasts first, and decisions later. This can make a simple least squares objective appropriate, and various tools from time series econometrics can be applied. The forecasting structure of these methods can allow for the use of more detailed analytics.¹⁰

3.2 Trading

The second most important part of agent based markets is the actual mechanism that governs the trading of assets. Once one leaves the relatively simple world of equilibrium modeling, it is necessary to think about the actual details of trading. This can be both a curse and a blessing to market designers. On the bad side it opens up another poorly understood set of design questions. However, it may have the beneficial effect of allowing one to study the impact of different trading mechanisms, all of which would be inconsequential in an equilibrium world. Most agent based markets have solved this problem in one of three ways: assume a simple price response to excess demand, build the market such that a kind of local equilibrium price can be found easily, explicitly model the dynamics of trading to look like the continuous trading in an actual market.

Most of the earliest agent based markets used the first method to model price movements. Most markets of this type poll traders for their current demands, sum the market demands, and if there is an excess demand, increase the price. If there is an excess supply they decrease the price.¹¹ A simple form of this rule would be

$$p_{t+1} - p_t = \alpha(D_t - S_t)$$

where D_t and S_t are the demand and supply at time t respectively. This method has several important strengths. It is fast. It emphasizes a market continuously in disequilibrium, and it allows some amount of analytics depending on the agent structures used. However, it comes with some significant costs as well. In the early SFI market experiments the dynamics were found to be very sensitive to the α parameter. Setting α too large made the price overreact and caused prices to thrash back and forth. Setting α too low caused the market to react too sluggishly to excess demands. The price would slowly rise while investors were never getting the shares they demanded. These trends would also be picked up by the learning agents, and magnified as more traders adopt trend following strategies.

¹⁰Much of the traditional learning literature in economics is based on least-squares time series learning as in Evans & Honkapohja (1995), Grandmont (1998) and Sargent (1993).

¹¹It is not well known, but the earliest version of the SFI artificial stock market in Palmer, Arthur, Holland, LeBaron & Tayler (1994), used a form of this price clearing method.

One interesting extension of this price setting method is to assume that the excess supply or demand for shares gets filled by a market maker.¹² The price adjustment mechanism is similar to the previous one, but excess demands and supplies are filled by a market maker. This eliminates the problem of asking what happens to agents who place orders that never get satisfied, but opens new questions about the inventory behavior of the market maker, and whether it is at least sensible, if not optimal. Beyond parameter sensitivity this price adjustment mechanism has another problem in that it may sidestep some important issues in market dynamics by assuming a constant market depth, or liquidity. This is essentially the amount the price is moved due to order flow imbalances. In actual markets this depth is changing, and should be thought of as an emergent property of the current trading agent features as opposed to a fixed constant. It is not clear how well the fixed price adjustment can handle this issue.

A second price setting method is to make several market structure assumptions so that a kind of temporary equilibrium price can be found.¹³ In all of these cases the structure of the model yields a well defined demand function for the agents. This can then be cleared either analytically, or in some cases computationally. While this method settles some of the parameter questions of the first method, and can allow for changing market depth, it does require putting more economic structure on the market. Specifically, the requirement that demand functions are well behaved enough for temporary price determination may be very restrictive. Furthermore, it represents a market that is at least temporarily in balance. This feature may not be a good representation for continuous trading in high frequency financial markets, but might not be a bad assumption for lower frequency price dynamics.

The third and final trading method involves modeling the actual trading mechanism in markets. This means actually building a trading mechanism that replicates that used in actual markets, including possible limit orders, and order crossing rules. This method is most appealing for modeling high frequency data, and is not sensitive to some of the previous critiques. However, it is only easy to implement when the market clearing mechanism is a mechanical rule. When it is actually done through human intervention, as in the New York Stock Exchange specialist system, it requires the modeling of another learning and adapting agent which could get complicated.¹⁴ One area where this approach can be easily used is in replicating experimental results from very simple markets.¹⁵ In these cases the agent based market faithfully follows the market design used in human experimental markets. Since these markets are carefully controlled experiments

¹²An example of this is Farmer & Joshi (forthcoming 2001).

¹³There are many different versions of this. These include Arifovic (1996), Brock & Hommes (1998), Levy et al. (1994), Arthur et al. (1997), and LeBaron (forthcoming 2001*b*).

¹⁴Two examples of this modeled after the foreign exchange market are Chakrabarti (1999) and Yang (2000).

¹⁵Examples of this would be Gode & Sunder (1993) and Chan, LeBaron, Lo & Poggio (1999).

this gives the artificial market builder a lot of structure to work with. Furthermore, replicating results from experimental markets may turn out to be useful for validation in the future.

Another related solution is to study stylized models that might approximate market behavior, but are not “market like” in their construction.¹⁶ The minority game presents a situation in which agents must choose from two options. Those who end up in the group with the smaller number of people get a payoff. The game is repeated over time, and agents try to make predictions and design strategies over time. While this shares some features with markets in terms of forecasting, there are some market features which are missing.

A key missing feature is the absence of a price which could help stabilize the market through the balancing of well behaved supply and demand functions. Almost by definition in economics a market is a situation where the rate of exchange, or price, is trying to balance out supply and demand. In absence of this basic mechanism the connections between minority games and real markets is somewhat stretched.

Another feature is that the game has no symmetric pure strategy Nash equilibrium. In other words there is no nonrandom strategy that if assigned to all agents would be optimal for everyone given everyone else holds to this strategy. This becomes a problem for many minority game simulations in that they only include pure, or nonrandom, strategies in the agent’s strategy sets. In these cases the dynamics may be interesting, but the fact that there is no convergence to a stable equilibrium is a property of the structure of the model/strategy choice, and is not surprising.¹⁷ Other agent based markets often include a stable Nash equilibrium in the set of learnable strategies.¹⁸

The fact that the dynamics of the minority game has to be unstable, along with its distance from traditional Walrasian market setups makes it a less useful tool for studying important questions about financial markets. The finance setting calls for a tool to explain what it is about agents’ information processing that makes a market unstable. As mentioned previously financial markets are probably special in the way information comes into play in interacting with prices. If all markets shared the bad dynamics of the minority game then it is not clear that economies would even function. The key is to capture the special nature of the financial information problem and to layer it on top of a more traditional model for market supply and demand. This is not to say that the minority game is not useful. It deserves a place next to the prisoner’s dilemma, and many other simple games. Bell et al. (1999) argue that it is a good model for several congestion problems.¹⁹

¹⁶An example of this is the minority game developed in Challet & Zhang (1997) which is based on the “El Farol” problem of Arthur (1994). This is closely related to models in Schelling (1978).

¹⁷See Bell, Sethares & Bucklew (1999) for models where agents are able to successfully learn mixed (random) strategies. Also, these authors find a curious result where more information can actually make learning less efficient in this model.

¹⁸See for example Arthur et al. (1997) and LeBaron (forthcoming 2001*b*) for examples.

¹⁹Anyone who has driven on a multi-lane interstate will have played a similar game concerned with the choice of the fastest

3.3 Securities

The traded securities are another important part of any agent based market just as they are in more traditional models of financial markets. The agent based modeling world often is quite simple in the types of traded securities because the complexity introduced through the heterogeneous agents usually pushes researchers to streamline the assets available for trading. Many agent based models are inspired by more traditional economic models which can make comparisons useful. Security prices often are connected to some type of fundamental value such as a dividend, or earnings, and these can then be calibrated to actual values from data. To most people not associated with building markets, the traded securities can appear very primitive. Obviously, this yields greater tractability at this very early stage in this field, but extensions will have to be made in the future.

The first change involves the revelation of new information. Often the dividend, or fundamental, is revealed each period. This is a luxury that few investors in the real world enjoy. This will have to be modified in a reasonable way such that traders need to infer a fundamental value from a noisy signal. The uncertainty surrounding a stock's valuation should be slowly diminishing over time if ever. Other extensions along these lines might be to build in commonly used data such as earnings forecasts, and other fundamental information.

Another major change is even more important in connecting agent based markets to real markets, the addition of more tradeable securities. Few agent based markets operate with many more than two assets which is clearly a limitation in thinking about the real world. Issues such as trading volume, diversification, and derivative trading cannot be studied in these single risky asset setups. Adding more securities is certainly desirable, but the complexity and computational burden this would bring to the entire market building process is a daunting task.

3.4 Evolution

Evolution is crucial in most agent based markets. In many ways it is the core dynamic at work both practically and philosophically. Early arguments about evolution and irrational speculation go back as far as Friedman (1953) where it was suggested that irrational speculators would be driven out of the market since their performance would not stand up to the better rational traders. This argument does not hold up when thinking along more coevolutionary terms. The argument is loaded with notions of absolute measures of trader performance, which don't exist in the endogenous population of a financial market. A trader's

performance depends critically on the behavior of others. Therefore it is impossible to judge any strategy ex-ante as being irrational without knowing the environment in which it will be placed.

In finance this question was reopened by the early literature on noise trading and limits to arbitrage.²⁰ These models stressed that rational traders might hold back some of their positions fearing irrational traders might move the price in a direction they weren't predicting. This is equivalent to trying to hold onto a short position in what one knows is a rising bubble. Even though you are doing the right thing, it gets progressively harder to stay with the position. This suggests an approach which looks at performance which can only be fully evaluated relative to others in the market. Agent based markets provide a perfect setting for testing these ideas, but this opens the question of how and where to implement evolution.

Evolution can appear in many different ways, and in many social settings. What is referred to as evolution, could often be called learning. One of the practical appeals of evolution is that it can be used to take parameters for which one's prior beliefs are weak, and put them under evolutionary control. In other words, give up on trying to set these to some arbitrary value, and let evolution decide their levels.²¹ This approach is very much in the spirit of evolutionary game theory in trying to get the evolutionary dynamic to select eventual stable equilibria.

Applying evolution in the social sciences opens up the major issue of what should be used as a fitness criteria. Finance may be particularly lucky in that measures of wealth seem like a good first choice, but this is not entirely clear. Several studies have used utility measures as a proxy for fitness, but it is not obvious why this should lead to overall agent survival. In Blume & Easley (1990) the authors show that utility maximization alone is not synonymous with wealth maximization. In their model a wealth selection mechanism does not necessarily chose the most economically rational agents for survival. Beyond being an interesting philosophical question about tastes and preferences, this issue is an important one in finance because of the problem of risk aversion. Risk averse agents may make optimal and cautious portfolio plans ex-ante, but in the long run, these may not make them the leaders in terms of wealth maximization. It is quite possible that evolution will carry a market toward the boldly overconfident in the long run.²²

Evolution also greatly affects the evaluation of rule fitness, and the learning mechanism that governs their changes over time. Often a mechanism such as a genetic algorithm is used as the learning engine for new rule generation and testing. This also relies on a fitness measure which is crucial in the direction that

²⁰See DeLong, Schleifer, Summers & Waldmann (1991) for early evolutionary ideas on noise trading, and Axelrod (1984) for the important early work on coevolution in a social setting.

²¹An example of this using the length of agent memory is LeBaron (forthcoming 2001 *b*).

²²See Kyle & Wang (1997).

the market may take. Should this be a rule that maximizes expected returns, wealth, utility, forecast error, or some estimate of survivability? Given that the fitness measure will be extremely noisy, how should agents adjust to this in their rule selection? In some fitness/selection mechanisms populations often narrow down very quickly to a small subset of potential strategies. About the only recommendation that can be made here is to stay with relatively robust fitness measures which make it difficult to eliminate strategies.

3.5 Benchmarks/Calibration

One of the key unanswered questions in building agent based markets is that of validation. How do you know you have a market that has any connection to the real world? This validation question appears constantly in the agent based economics community. There is a side of this issue which is not really different from more traditional theoretical modeling. These are all only toy models that represent a complicated social situation in a highly stylized fashion. What sets agent based markets apart is the large number of parameters for which our priors are extremely diffuse. Therefore, they do give the researcher many degrees of freedom from which to align to any interesting empirical feature of the data that one wishes to match. In this situation success may be difficult to judge.

Though a complete solution to this issue is still out of reach, there are a few things the agent based market builder can try to test and validate various structures. The first is to think about useful benchmark cases where the behavior of a market is well defined. An example of this is to think about which parameter values may lead to convergence into a well defined homogeneous agent equilibrium. Set the market to these parameters, and show that it does indeed converge. Such behavior gives results outside of this parameter realm greater potency. Furthermore, understanding exactly where the parameter boundaries are between simple and complex behaviors is crucial to understanding the mechanisms that drive agent based markets. Researchers should be strongly encouraged to “tweak the dials” in their models, and not to simply report that a price series looks similar to a real market.

A final approach to validation and calibration follows closely in the spirit of the calibration literature in economics.²³ It would recommend using parameters estimated from actual experimental markets in the simulated agent markets. In other words, try to learn about learning in the laboratory, and then take what you know into situations that would be impossible to simulate in the laboratory into an agent based approach. This is subject to the usual criticisms of calibration in that the parameters may not be relevant in different situations, but this would appear to be a fruitful way to get agent models onto a stronger footing.

²³See for example Kydland & Prescott (1996).

3.6 Time

Almost all adaptive systems including financial markets encounter several important issues in dealing with time. This is often ignored in many papers since the problems are generally not well understood even though their impact could be quite large.

The first issue of time concerns the past, and how agents intend to deal with it. In almost all learning mechanisms there needs to be a feedback from past performance into current rule selection. How distant a past should this be? In many settings it is fixed to be some lagged performance measure, or set weighting of the past. Several papers have begun to explore this dimension and show that its impact on dynamics may be quite rich.²⁴ The interesting experiment is to construct a world populated with agents of varying memory lengths. Some believe only the past 6 months of data are relevant, while others might be looking at 20 years. This may be related to important behavioral heterogeneity observed in the real world. Also, it forms a key test of how well agents can learn about whether the world they live in is stationary. If the economic environment truly is stationary then more data should be preferred to less and longer memory agents should survive. In early experiments it appears that long memory agents have a difficult time driving out short memory types LeBaron (forthcoming 2001*b*), and therefore stationarity itself may be a difficult concept to learn in a multi-agent environment.

Another aspect of time is more concerned with relative rates of change than memory lengths. In the SFI stock market it was found that the speed at which agents were updating behavioral rules had an enormous impact on the outcome. In this case it appears that when the coevolution of the market is slowed sufficiently the changing fitness landscapes can be simple enough for agents to jointly find a socially optimal Nash equilibrium corresponding to an efficient market. However, when traders are changing strategies relatively frequently this search becomes futile, and probably never ends.²⁵ These sensitivities are also a little troubling in that it may be that observed outcomes from agent based markets depend critically on a few relative speeds in evolution. This could lead to a criticism pointing out that these are free parameters for which little is known, but are capable of moving the results anywhere you want. In the worst case it would mean that the shell of the agent based market is not a very restrictive theory at all, and these markets would be unfalsifiable.

²⁴See Levy et al. (1994) and LeBaron (forthcoming 2001*b*).

²⁵LeBaron (forthcoming 2001*b*) shows that directly slowing learning can lead to convergence to an efficient market equilibrium. It also shows that an indirect, but possibly more economically relevant mechanism will do the same thing. The introduction of a cost to changing rules (similar to, but not exactly a transaction tax) will have the same effect. This friction slows the agent adaptation in a similar fashion to slowing learning directly. Features such as these suggest a connection to Bell et al. (1999) where the impact of reducing information available to agents had the curious effect of improving social welfare in a minority game setting. In other words there may be times with less optimization is preferred to more.

Hopefully, the situation is not this bad.²⁶

A final aspect of time that most current agent based markets ignore is that of synchronicity. This is an issue of great interest to agent based modeling in all of the social sciences, and has led to some important controversies.²⁷ In most financial models trading is synchronized by the designer. However, in the real world there are no fixed periods in which all trades must take place, and the price gets set. Agents arrive asynchronously and must find others or specified dealers to trade with. There may be information events that tend to bring them to the market at similar times, but these do not guarantee an exactly similar arrival time. Future agent based financial models will have to better address this issue along with the problem of more realistic trading mechanisms. Unfortunately, there is no easy answer here on how to proceed. Even agent based software continues to have some difficulties with this problem since this is concerned with modeling inherently parallel activities on serial machines. It is possible that the best answer to this will have to involve implementation on large scale parallel computer networks to make sure that no synchronization artifacts are seeping into the results.

4 Future

Ten years ago this field was populated with only a handful of people. Now it is booming with many different styles of market simulations appearing all the time. Similar to the “dot com” world there will probably be some sort of convergence over time back down to a smaller set of platforms, but with extensive experimentation done with each one.

This shake out will push artificial markets in several directions. One will be towards realism and validation. Markets in this area will have to concentrate on getting serious about the data. Tools from the extensive literature on calibration in economics will clearly be useful. They will have to face all the questions on validation which are continually asked. Several initial responses to these criticisms are possible. One is to perform calibration on a grand scale, by matching lots of features in the data. Traditionally, one might try to get a small set of expected returns, or variances. Pushing to match many more features including trading volume, and various cross correlations will be important.²⁸

The second direction involves better understanding of what the generic properties of multi-agent systems are. There appears to be a growing set of features that traditional financial models find difficult to generate,

²⁶Ellison (2000) shows some early results on this from evolutionary game theory.

²⁷See Axtell (2000) for some interesting experiments on this subject.

²⁸See LeBaron (forthcoming 2001a) for an example of aligning to many empirical features in a calibration exercise.

but agent based models readily provide. Among these are fat tailed return distributions, persistent volatility, and widely fluctuating trading volume. There may be some very interesting generic issues at work here, and it is possible that broad classes of agent based markets will yield these results. Part of the agent based modeling world will probably head in the direction of simpler market models, and tools which can reveal their generic properties. This line of research may help uncover some basic features common in the multiagent/complex systems world returning to our original Hayek/Schelling distinctions for markets.

Tantalizingly out of range remain questions about what to do with these markets. Most practitioners are interested, but perplexed as to how to use the results. Can they be forecasting tools? Can they indicate that a financial crisis is eminent? Applications along these lines have still not been taken on by the agent based community. One reason for this is these models do not lend themselves to more traditional econometric testing. This is not to say that they can't be used. Methods such as "forward testing" where a market is run using real data as the price input up to the current date, and then allowed to continue on into its future may end up being useful. On the other hand it is possible that forecasting related questions will not be well served by agent based markets. They may play a very different role in evaluating the stability and efficiency of different trading mechanisms in normal and crisis periods. This would give more robust policy answers to various current questions about trading mechanics, liquidity and depth, and transparency in all types of markets as new technologies continue to impact their landscape.

One final large expectation is that the research from artificial markets has a positive spillover into other areas of economic and social science. Hopefully, the technologies that are tried and tested here, will prove useful in other places. These spillovers may lead to new insights about human behavior in many different situations far beyond the boundaries of finance.

References

- Arifovic, J. (1996), 'The behavior of the exchange rate in the genetic algorithm and experimental economies', *Journal of Political Economy* **104**, 510–541.
- 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.
- Axelrod, R. (1984), *The Evolution of Cooperation*, Basic Books, New York.
- Axtell, R. (2000), Effects of interaction topology and activation regime in several multi-agent systems, Technical report, Brookings Institution, Washington, DC.
- Bak, P., Paczuski, M. & Shubik, M. (1997), 'Experts, noise traders, and fat tail distributions', *Economic Notes* **26**, 251–290.
- Becker, G. S. (1962), 'Irrational behavior and economic theory', *Journal of Political Economy* **70**, 1–13.
- Bell, A. M., Sethares, W. A. & Bucklew, J. A. (1999), Coordination failure as a source of congestion in information networks, Technical report, NASA Ames Research Center, Moffett Field, CA.
- 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.
- Blume, L. & Easley, D. (1990), 'Evolution and market behavior', *Journal of Economic Theory* **58**, 9–40.
- Bookstaber, R. (1999), 'Risk management in complex organizations', *Financial Analysts Journal* **55**(2), 18–20.
- Bookstaber, R. & Langsam, J. (1985), 'On the optimality of course decision rules', *Journal of Theoretical Biology* **116**, 161–193.
- 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.

- Caldarelli, G., Marsili, M. & Zhang, Y.-C. (1997), 'A prototype model of stock exchange', *Europhysics Letters* **40**(5), 479–484.
- Chakrabarti, R. (1999), 'Just another day in the inter-bank foreign exchange market', *Journal of Financial Economics* **56**(1), 29–64.
- Challet, D. & Zhang, Y. C. (1997), 'Emergence of cooperation and organization in an evolutionary game', *Physica A* **246**, 407.
- Chan, N., LeBaron, B., Lo, A. W. & Poggio, T. (1999), Agent-based models of financial markets: A comparison with experimental markets, Technical report, Massachusetts Institute of Technology, 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.
- Cont, R. & Bouchaud, J. P. (2000), 'Herd behavior and aggregate fluctuations in financial markets', *Macroeconomic Dynamics* **4**, 170–196.
- de la Maza, M. & Yuret, D. (1995), A model of stock market participants, in J. Biethahn & V. Nissen, eds, 'Evolutionary Algorithms in Management Applications', Springer Verlag, Heidelberg, pp. 290–304.
- DeLong, J. B., Schleifer, A., Summers, L. H. & Waldmann, R. (1991), 'The survival of noise traders in financial markets', *Journal of Business* **64**, 1–19.
- Ellison, G. (2000), 'Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution', *Review of Economic Studies* **67**, 17–45.
- Evans, G. W. & Honkapohja, S. (1995), 'Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models', *Econometrica* **63**, 195–206.
- Farmer, J. D. (1998), Market force ecology and evolution, Technical report, Santa Fe Institute, Santa Fe, NM.
- Farmer, J. D. & Joshi, S. (forthcoming 2001), 'The price dynamics of common trading strategies', *Journal of Economic Behavior and Organization* .

- Friedman, M. (1953), The case for flexible exchange rates, in 'Essays in positive economics', University of Chicago Press, Chicago, IL.
- Gode, D. K. & Sunder, S. (1993), 'Allocative efficiency of markets with zero intelligence traders', *Journal of Political Economy* **101**, 119–37.
- Grandmont, J.-M. (1998), 'Expectations formation and stability of large socioeconomic systems', *Econometrica* **66**, 741–782.
- Hayek, F. A. (1945), 'The use of knowledge in society', *American Economic Review* **35**(4), 519–30.
- Hodgson, G. M. (1999), *Evolution and Institutions*, Edward Elgar, Northampton, MA.
- Kim, G. & Markowitz, H. (1989), 'Investment rules, margin, and market volatility', *Journal of Portfolio Management* **16**(1), 45–52.
- Kirman, A. P. (1993), 'Ants, rationality, and recruitment', *Quarterly Journal of Economics* **108**, 137–56.
- Kydland, F. E. & Prescott, E. G. (1996), 'The computational experiment: An econometric tool', *Journal of Economic Perspectives* **10**, 69–85.
- Kyle, A. & Wang, F. A. (1997), 'Speculation duopoly with agreement to disagree: Can overconfidence survive the market test?', *Journal of Finance* **52**, 2073–2090.
- LeBaron, B. (2000), 'Agent based computational finance: Suggested readings and early research', *Journal of Economic Dynamics and Control* **24**(5-7), 679–702.
- LeBaron, B. (forthcoming 2001a), 'Empirical regularities from interacting long and short memory investors in an agent based stock market', *IEEE Transactions on Evolutionary Computation* .
- LeBaron, B. (forthcoming 2001b), 'Evolution and time horizons in an agent based stock market', *Macroeconomic Dynamics* **5**(2).
- 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.
- Lindgren, K. (1992), Evolutionary phenomena for social dynamics, in C. G. Langton, C. Taylor, J. D. Farmer & S. Rasmussen, eds, 'Artificial Life II', Addison-Wesley, Reading, MA, pp. 295–312.
- Lux, T. (1997), 'Time variation of second moments from a noise trader/infection model', *Journal of Economic Dynamics and Control* **22**, 1–38.
- O'Donoghue, T. & Rabin, M. (1999), 'Doing it now or later', *American Economic Review* **89**, 103–24.
- Palmer, R., Arthur, W. B., Holland, J. H., LeBaron, B. & Tayler, P. (1994), 'Artificial economic life: A simple model of a stock market', *Physica D* **75**, 264–274.
- Rieck, C. (1994), Evolutionary simulation of asset trading strategies, in E. Hillebrand & J. Stender, eds, 'Many-Agent Simulation and Artificial Life', IOS Press.
- Roth, A. E. (1995), Introduction to experimental economics, in 'The Handbook of Experimental Economics', Princeton University Press, Princeton, NJ, pp. 3–109.
- Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.
- Schelling, T. (1978), *Micromotives and Macrobehavior*, Norton, New York, NY.
- Simon, H. (1969), *The Sciences of the Artificial*, MIT Press, Cambridge, MA.
- Steiglitz, K., Honig, M. L. & Cohen, L. M. (1996), A computational market model based on individual action, in S. H. Clearwater, ed., 'Market-Based Control: A paradigm for distributed resource allocation', World Scientific, New Jersey, pp. 1–27.
- Yang, J. (2000), Price efficiency and inventory control in two inter-dealer markets, Technical report, Bank of Canada, Ottawa, Canada.
- Youssefmir, M. & Huberman, B. A. (1997), 'Clustered volatility in multiagent dynamics', *Journal of Economic Behavior and Organization* **32**, 101–118.