Chapter 40

AGENT-BASED MODELS OF FINANCIAL MARKETS

NICHOLAS S. P. TAY, *University of San Francisco, USA*

Abstract

This paper introduces the agent-based modeling methodology and points out the strengths of this method over traditional analytical methods of neo-classical economics. In addition, the various design issues that will be encountered in the design of an agent-based financial market are discussed.

Keywords: agent-based models; computer simulation; bounded rationality; heterogeneous agents; learning; co-evolution; complex adaptive system; artificial intelligence; neural networks; classifiers; genetic algorithms; genetic programming

40.1. Introduction

The sort of phenomena that are interesting in finance and yet difficult to investigate analytically involve the complex interactions among many self-interested heterogeneous boundedly rational agents acting within the constraints imposed by either formal or informal institutions or authorities. To outrival their opponents, each and every agent must continually evolve to adapt to changes that may arise either from exogenous perturbations to the environment or endogenous transitions caused by agents changing their strategies or modifying their behaviors as they learn more about the behaviors of the other agents and the environment they reside in. A good example of such complex adaptive systems is the stock market.

A natural way to study a complex adaptive system like the stock market is to use an agent-based model which entails simulating the stock market on a computer from the bottom up with a large number of interacting heterogeneous boundedly rational artificial agents that are created to mimic the traders in the stock market. Once the environment of the stock market and the behaviors of the agents are specified and the initial state of the model is set, the dynamics of the model from the initial state forward will be driven entirely by agent–agent interactions, and not by some exogenously determined systems of equations. Hence, if any macroscopic regularity emerges from the model, it must be a product of the endogenous repeated local interactions of the autonomous agents and the overall institutional constraints. This is the spirit of the agent-based modeling approach.

What makes the agent-based modeling methodology particularly appealing? To begin with, analytical tractability is not an issue since this approach relies on computer simulations to understand the complex model. Quite the reverse, it is inconceivable how one could obtain closed form solutions of a model as complex as the stock market without first diluting drastically the authenticity of the model. Although analytically tractable heterogeneous agent rational expectations models have been around, the complexity and realism that are captured in agent-based models are beyond the reach of those analytical models.

For instance, consider the problem that a decision maker faces when the outcome is contingent on the decisions to be made by all the participating heterogeneous decision makers, each with their own unique preferences and quirks and private information that are not directly observable by the other decision makers. This decision problem is inherently ill defined and cannot be solved through mathematical deduction or analytical modeling. In real life, when confronted with such an ill-defined situation, decision makers often rely on the rules of thumb that they have distilled from years and years of experience to guide them in their decision-making. This decision making process is formally known as inductive reasoning and it can be captured naturally with the agent-based approach by running computer simulations of a large number of interacting artificial agents who make decisions using rules of thumb that they distill from their repeated interactions with each other.

The ability to build more realistic models with the agent-based method often allows agent-based models to reveal a much richer set of behaviors that are embedded in a system which may otherwise be overlooked by traditional equation-based models. For instance, Parunak et al. (1998) in comparing the differences between equation-based modeling and agent-based modeling of a supply network have found that equation-based model fails to produce many of the rich effects, such as memory effect of backlogged orders, transition effects, or the amplification of order variation, which are observed in an agent-based model of the same supply network. In addition, various agent-based models (Farmer and Joshi, 2000; Johnson et al., 2001; LeBaron et al., 1999; Tay and Linn, 2001) have been successful in accounting for real financial markets phenomena such as market crashes, mean reversion, relatively high level of trading, technical trading, excess volatility, and volatility clustering. These are phenomena that analytical representative agent models of financial markets have tolled to explain without much success.

Another serious shortcoming of analytical representative agent models of financial markets is that by design these models do not specify the dynamic process that will need to happen in order to arrive at the equilibrium or equilibria that are characterized in these models. Consequently, for models that produce multiple equilibria, it is unclear which equilibrium among the multiple equilibria agents would converge on. In contrast, the events that unfold in a computer simulation of an agent-based model are completely transparent, and can be recorded hence providing the modeler a means to go back in the time line of evolution to understand how certain equilibrium or other global regularities came into existence.

The agent-based methodology therefore offers important advantages over the traditional analytical tools of neoclassical economics as it allows a researcher to obtain more germane results. Needless to say, the use of computer simulations as a tool for studying complex models has only became feasible in recent years because of the availability of fast and cheap computing power. Although the agent-based modeling methodology is still in its infancy, there is already a considerable literature on agent-based models. Leigh Tesfatsion at the Iowa State University maintains a website at *http://www.econ.iastate.edu/tesfatsi/ace.htm* to facilitate access to the extensive resources related to the agent-based modeling methodology, and to keep researchers in this field abreast of the latest developments.

In the introductory remarks on her website, Tesfatsion observes that agent-based research may generally be organized according to one of the following four research objectives: (1) empirical understanding, (2) normative understanding, (3) qualitative insight and theory generation, and (4) methodological advancement. The first objective focuses on seeking answers that are established on the repeated interactions of agents to explain the emergence of global regularities in agent-based models. Some examples of global regularities in financial markets are mean reversion and volatility clustering. Researchers in this group are interested in understanding if certain types of observed global regularities can be attributed to certain types of agent-based worlds. The second objective concerns using agent-based models as laboratories to aid in the discovery and design of good economic policies or good institutional structures. Researchers with this objective in mind are interested in using agent-based models to evaluate whether certain economic policies or institutional designs and processes will promote socially desirable outcomes over time among agents that are driven solely by their self interests. Tesfatsion phrased the third objective as "How can the full potentiality of economic systems be better understood through a better understanding of their complete phase portraits (equilibria plus basins of attraction)?" Unlike analytical models, the causal mechanisms in agent-based models are not direct and are very difficult to discern because of the complex nature of the interactions among the agents and between the agents and the environment. The goal here is to use the phase portraits as a means to enrich our understanding of the causal mechanism in these systems. The fourth objective addresses issues related to improving the methods and tools used by agent-based researchers.

For someone who is just starting out in this line of research, it is worthwhile to begin by reading "A Guide for Newcomers to Agent-based Modeling in the Social Sciences" by Axelrod and Tesfatsion which is available on the homepage of Tesfatsion's website. In addition, it is beneficial to read the survey articles written by Hommes (2004), Duffy (2004), LeBaron et al. (1999), LeBaron (2000, 2004a), and Tesfatsion (2002) and a book by Batten (2000) that provides an overview of agent-based models and offers some historical perspectives of this methodology.

The next section discusses the design issues that will be encountered in the design of an agent-based model. This discussion benefited greatly from the insights that LeBaron has provided in his excellent overviews of the various design issues (LeBaron, 2000, 2001c, 2004a).

40.2. Design Considerations

A typical agent-based model is made up of a set of autonomous agents that encapsulate the behaviors of the various individuals in a system we are interested in studying and the investigation involves simulating on a computer the interactions of these agents over time. Accordingly, there are two important design considerations in the development of an agent-based model – the design of the agents and the design of the environment.

How naive or sophisticated the agents should be modeled really depends on the objective of the research. For instance, if the research objective is to understand how certain market structures affect the allocative efficiency of a market independent of the intelligence of the agents as in Gode and Sunder (1993), then one can simply model the agents as naive "zero intelligence" agents. Zero intelligence agents are agents that are not capable of formulating strategies or learning from their experience; hence their behaviors will be completely random. Gode and Sunder populated their double auction market with zero intelligence agents that are designed to submit their bids and asks at random over a predefined range and remarkably they discover that zero intelligence agents when subjected to a budget constraint are able to allocate the assets in the market at over 97 percent efficiency. The lesson to be learned here is that not all macroscopic regularities that emerge from agent-based models are necessarily consequences of the actions taken by the agents as they evolve and learn from their interactions. In this case, the high level of allocative efficiency that is attained in a double auction market is due to the unique structure of the market itself.

However, in many agent-based models, the objective is to investigate the outcome of the interactions among many heterogeneous agents that are designed to mimic their counterparts in the real world. In these models, the key design issues related to the design of the agents are the agents' preferences and their decision-making behaviors.

Agents could have either myopic or intertemporal preferences. The latter is more realistic but will make the model much more complex. As we have alluded to earlier, the decision problem that the agents face is usually ill defined, and thus cannot be solved by deductive reasoning. A reasonable solution is to assume that the agents rely on inductive reasoning to arrive at a decision (see Arthur, 1994, 1999; Rescher, 1980). Inductive reasoning or induction is a means for finding the best available answers to questions that transcend the information at hand. In real life, we often have to draw conclusions based upon incomplete information. In these instances, logical deduction fails because the information we have in hand leaves gaps in our reasoning. In order to complete our reasoning, we fill those gaps in the least risky, minimally problematic way, as determined by plausible best-fit considerations. Consequently, the conclusions we draw using induction are suggested by the data at hand rather than logically deduced from them.

Inductive reasoning follows a two-step process: possibility elaboration and possibility reduction. The first step involves creating a spectrum of plausible alternatives based on our experience and the information available. In the second step, these alternatives are tested to see how well they answer the question at hand or how well they connect the existing incomplete premises to explain the data observed. The alternative offering the ‘‘best fit’’ is then accepted as a viable explanation. Subsequently, when new information becomes available or when the underlying premises change, the fit of the current alternative may degrade. When this happens a better alternative will take over.

How can inductive reasoning be implemented in an agent-based financial market model? Arthur (1994, 1999) envisions inductive reasoning in a financial market, taking place as follows. Initially, each agent in the market creates a multitude of decision-making rules (this corresponds to the possibility elaboration step discussed above). Next, the decision-making rules are simultaneously tested for their effectiveness based on some criteria. Finally, effective decision-making rules are retained and acted upon in buying and selling decisions. Conversely, unreliable rules are dropped (this corresponds to the possibility reduction step). The rules that are dropped are then replaced with new ones in the first step and the process is carried out repeatedly to model how individuals learn inductively in a constantly evolving financial market.

Some examples of criteria that have been used for appraising the effectiveness of the decision rules includes utility maximization, wealth maximization, and forecast errors minimization. Once a decision has been made on a criterion for evaluating the decision-making rules, the next task is to decide the length of historical data to be used in computing the criterion. Although many agent-based models tend to allow the agents to adopt identical history length, this is not necessary. It is in fact more realistic to permit agents in the same model to adopt different history length as in LeBaron (2001a,b).

To take the modeling to the next step, decision will have to be made concerning what the decision making rules look like and how they are to be generated in the models? One possibility is to model the decision-making rules after actual trading strategies used in real financial markets. The benefit of this approach is that the results are likely to be tractable and precise and it will also shed light on the interaction among these actual trading strategies. However, this approach does not allow the agents any flexibility in modifying the strategies or developing new strategies. This could impose ad hoc restrictions on the model's dynamics. Some common tools that have been employed to allow the agents more degrees of freedom in structuring and manipulating the decision making rules as they learn are artificial neural networks (LeBaron, 2001a), genetic programming (Chen and Yeh, 2001), and classifiers that are evolved with genetic algorithms (LeBaron et al., 1999). Even with these artificial intelligence tools, the modeler will need to predefine a set of information variables and functional forms to be used in the strategies or decision-making rules. Although these tools can successfully mimic the inductive reasoning process described earlier, it is not known if any of these tools indeed faithfully represent the inductive reasoning used by actual human traders. It is also unclear at this juncture whether this issue matters. Another related decision that has to be made by the modeler concerns whether the agents should be allowed to learn only from their own experiences or from the collective experience of all the agents in the model. The latter is known as "social learning."

We will turn our attention next to the design of the financial market environment. Most agent-based models simplify the environment to a market with one risky asset and one risk-free asset. Clearly, this is an oversimplification of actual financial markets, but there are good reasons for doing so. Given that the agent-based methodology is new and researchers barely comprehend the implications of this methodology, it is prudent for them to begin by exploring what the new method can reveal about the dynamics in a fairly simple market environment. Moreover, doing so also facilitates comparisons with results from well-known neoclassical models of a market with one risky asset and one risk-free asset.

Another key design issue concerns the design of the trading mechanism that has a direct influence on how prices are determined in the market and how the market is cleared. LeBaron (2004a) observes that most of the agent-based models employ one of the following four designs for trading mechanism. The simplest trading mechanism is one that allows mutually beneficial trades to be consummated between agents that meet at random. Though this trading mechanism is quite simple, it bears some resemblance to the trades conducted on the floor of the Chicago futures and options exchanges and over the telephone in the foreign exchange markets. But for markets where the market makers play an important role in filling the buy and sell orders, this mechanism would not be an adequate representation. A second trading mechanism, which is more sophisticated than the previous one is an analytical market-clearing device akin to one espoused in Grossman (1976). This device provides a closed form solution for the market-clearing price hence enabling the agent-based markets to be cleared analytically each period. A critical advantage of this design is that it avoids having to deal with the difficult issue of explicitly modeling the decision-making behaviors of the risk-adverse market maker. Unfortunately, this advantage is also a serious shortcoming in that this is not a realistic picture of what is happening in real markets that trade continuously and are rarely in equilibrium. The third trading mechanism attempts to address this issue. It assumes that agents submit trade orders to buy (D_t) and sell (S_t) at a price, p_t, which is announced beforehand by a market maker. The market maker then appraises the aggregate of the orders submitted by the agents, and adjusts next period price by a fixed fraction of the excess demand or supply according to $p_{t+1} = p_t + \alpha(D(p_t) - S(p_t))$. Granted that this adaptive price process may be a more reasonable model of how prices adjust in real markets, the problem with this mechanism is that it does not address how the market maker manages the imbalance between demand and supply in the market. Moreover, there is no guidance on how the parameter value for α should be determined and certain α values may in fact cause the market to deviate far from the market clearing price for a substantial period. The most sophisticated and also the most realistic trading mechanism is one that either models the market maker explicitly or implements an order book system that can accept and cross out the buy and sell orders from agents according to some defined procedure (Audet et al., 2001). The only downside of this approach is that the design of the agent-based model is much more complicated as many details at the institutional as well as the agent level will need to be clearly specified. But, this is inevitable if the objective is to simulate realistic market microstructure behavior.

To sum up, there are many design questions that need to be addressed in the development of an agent-based model and there is yet no clear guidance on how best to address these questions. Inevitably, design decisions will have to be made however arbitrary these decisions may be but it is important to keep in mind that the choices made by the designer may ultimately have important consequences on the results.

REFERENCES

Arthur, W.B. (1994). "Inductive reasoning and bounded rationality." *American Economic Review*, 84(2): 406.

Arthur, W.B. (1999). "Complexity and the economy." *Science*, 284(5411): 107.

Audet, N., Gravelle, T., and Yang, J. (2001). "Optimal market structure: does one shoe fit all?" *Technical Report*. Ottawa, CA: Bank of Canada.

Batten, D.F. (2000). *Discovering artificial economics: How agents learn and economies evolve*. Boulder, CO: Westview Press.

Chen, S.-H and 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.

Duffy, J. (2006). "Agent-based models and human subject experiments," in K.L. Judd and L. Tesfatsion, (eds.), *Handbook of Computational Economics Volume 2: Agent-Based Computational Economics, Handbooks in Economics Series*. Amsterdam: Elsevier North-Holland.

Farmer, J.D. and Joshi, S. (2000). "The price dynamics of common trading strategies." Santa Fe Institute Working Paper 00-12-069, Santa Fe, New Mexico.

Gode, D.K. and Sunder, S. (1993). "Allocative efficiency of markets with zero intelligence traders." *Journal of Political Economy*, 101: 119–137.

Grossman, S.J. (1976). "On the efficiency of competitive stock markets where traders have diverse information." *Journal of Finance*, 31(2), 573.

Hommes, C.H. (2006). "Heterogeneous agent models in economics and finance." In K.L. Judd. and L. Tesfatsion (eds.) *Handbook of Computational Economics Volume 2: Agent-Based Computational Economics, Handbooks in Economics Series*. Amsterdam: Elsevier North-Holland.

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

LeBaron, B., Arthur, W.B., and Palmer, R. (1999). "Time series properties of an artificial stock market." *Journal of Economic Dynamics and Control*, 23:(9,10), 1487–1516.

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. (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, Blake (2001c). *A builder's guide to agent-based financial markets*. Technical Report, International Business School, Brandeis University, Waltham, MA.

LeBaron, B. (2006). "Agent-based computational finance," in K.L. Judd and L. Tesfatsion, (eds.), *Handbook of Computational Economics Volume 2: Agent-Based Computational Economics, Handbooks in Economics Series*. Amsterdam: Elsevier North-Holland.

Parunak, H., Savit, R., and Riolo, R.L. (1998). "Agent-based modeling vs. equation-based modeling: a case study and users' guide." In MABS, Sichman, C. and Gilbert (eds.) *Multi-agent systems and Agent-based Simulation*. Berlin: Springer-Verlag.

Rescher, N. (1980). *Induction: An essay on the Justification of Inductive Reasoning*. Pittsburgh: University of Pittsburgh Press.

Tay, N.S.P. and Linn, S.C (2001). "Fuzzy inductive reasoning, expectation formation and the behavior of security prices." *Journal of Economic Dynamics and Control*, 25(3,4): 321–361.

Tesfatsion, L. (2002). "Agent-based computational economics: growing economies from the bottom up." *Artificial Life*, 8(1): 55–82.