Can Irrational Investors Survive? A Social-Computing Perspective

Yongjie Zhang, *Tianjin University*

Wei Zhang, *Tianjin University and Tianjin University of Finance & Economics*

The proposed agent-based model accounts for interactions between irrational and rational investors, expanding on existing work and offering hope for irrational-investor survival in artificial stock markets.

Standard financial theory includes a rigorous theoretical system for equilibrium asset pricing. Among the assumptions this system is founded on is the necessity of investors' homogeneity and rationality.[1] In the past few decades, however, researchers have increasingly observed that price dynamics often fail to obey the standard theory's laws.

That is, they've observed financial anomalies in the markets. Various financial economists have exerted tremendous effort to bridge this gap between theory and practice. One of the results is *behavior finance*, which has emerged as a significant school of thought.

One of behavioral finance's intuitive hypotheses is the existence of irrational investors in the markets. Such investors' cognitive bias leads to irrational investment behaviors. These, in turn, create a different mechanism to form price equilibrium than the rational one in standard financial economics. But, if Milton Friedman's induction holds, rational investors will eventually drive irrational investors out of the market,[2] it will shake behavior finance at its foundation.

Behavioral finance literature devotes considerable attention to this critical issue. In general, analytical approaches—such as the BSV (Barberis, Shleifer, and Vishny) model[3]—fail to adequately address investor heterogeneity and irrationality, which led to the increasing use of agent-based social modeling in financial studies. Despite solid work in this area,[4] the literature still lacks a heterogeneity of rational and irrational investors. Therefore, irrational investors' ability to survive in the market remains unknown.

Here, we describe our agent-based computational (ABC) modeling approach, which constructs an interactive model that includes both irrational and rational investor categories. By considering investor heterogeneity endogenously, our approach expands the BSV model to include heterogeneous investors. In so doing, we lay a more "micro" foundation for existing microscopic[5] and analytical[4] models, and thereby reveal the irrational investors' survival status in artificial stock markets (ASMs).

ABC method overview

The ABC method develops a behavioral finance model in five steps, three of which extend and further develop existing behavioral finance theories.

Step 1: Build the conceptual model

In this step, researchers build the conceptual model much as they would in standard behavioral finance methods. However, ABC modeling gives them the freedom to build more complex models using heterogeneous and time-varied investor behaviors. This is possible because the ABC method doesn't require mathematical tractability to solve the model.

Step 2: Design the artificial stock market's architecture

Designing the ASM architecture has the same function as laying equations in the analytical method. In this step, researchers must clearly define the behav-

1541-1672/07/$25.00 © 2007 IEEE
Published by the IEEE Computer Society

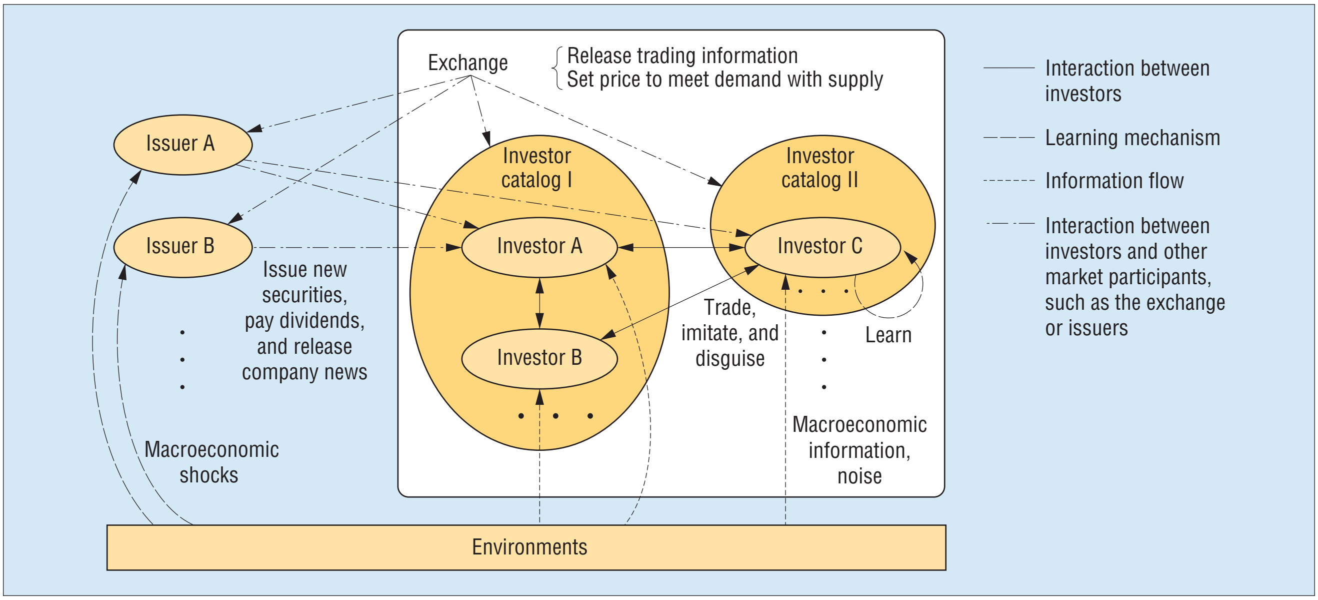

Figure 1. A typical artificial stock market structure. ASMs have four basic categories of actors: issuers, investors, the exchange, and environments.

ior of heterogeneous investors, tradable assets' risk and return characteristics, information flow, and the market trading and equilibrium mechanism.

Step 3: Code and run the ASM program

In this step, researchers choose a proper simulation platform, code and debug the ASM architecture, and design and run computational experiments to generate simulated stock market data.

There are more than a dozen frequently used agent-based simulation platforms and thousands of agent-based simulation applications. Choosing a proper platform can save researchers considerable time and computational cost. Also, selecting an appropriate ASM—such as the Santa Fe Institute's ASM (SFI ASM)[6]—can facilitate coding and debugging work.

Step 4: Examine the ASM data

Typically, the data generated in step 3 is raw information—such as the market price and trading-volume-time series, or agent position and wealth-time series—which researchers can't directly use to develop behavioral finance theories. With the ABC method, researchers can process this raw data using financial econometrical methods to produce meaningful results on the simulated market dynamics. These econometrical results show the ASM's empirical characteristics, which researchers can then use to develop behavioral finance theories in step 5.

Examples of step 4 results might be

- "the market price in this ASM overreacts in the short run and mean-reverts in the long run,"
- "this ASM is not affected by excess volatility," or
- "this ASM's IPO premium for a 10-period window is 45 percent."

Most financial economists can easily understand and use all of these to form new financial economic theories.

If step 4's simulated market dynamics are unexpected or unclear, researchers should return to step 1 and carefully investigate their conceptual models to find the cause. They should also distinguish between new findings and program errors.

Step 5: Presenting new behavioral finance theories

In this step, researchers can extend or develop behavioral finance theories using the ABC method's conceptual model and solution. The aim in utilizing ASMs is to solve the conceptual models, not to simulate real stock markets. Likewise, the conceptual models aren't intended to simulate real stock markets, but rather to enhance insights into how investors interact under given circumstances.

Designing an artificial stock market

An ASM has four basic actor categories: issuers, investors, the exchange, and environments. Although the implementation of these actors might differ in various ASMs, the basic relation among these actors is much like that in figure 1.

Issuers

In the ASM, issuers must try to successfully issue new securities, must pay dividends and interests on their outstanding securities, and must release related company news to the exchange and investors in a timely manner. To some extent, investors' output and profitability are affected by macroeconomic shocks, which in turn will affect the amount of dividends they pay to investors.

Investors

Investors buy stock in issuers' public offering and trade stock among themselves. Their trading decisions are made on the basis of both meaningful information and meaningless noise. Investors get noisy information through the environment, as well as from

- issuers providing information on dividends and company revenue;
- the exchange disclosing trading information such as price history, volume, and timely bid and ask; and
- investors who sometimes "create" informa-

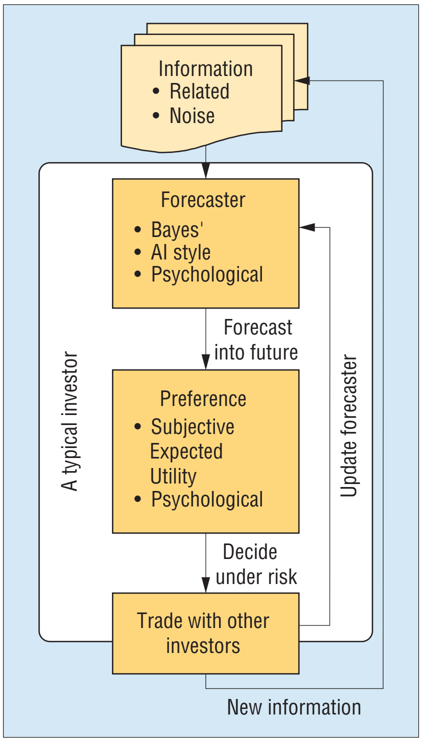

Figure 2. Typical investor behavior in ASMs. The decision-making process consists of a forecast-preference-action sequence. Investors use the results to further train their forecasters, which might be based on Baysian or psychological learning models.

tion to mislead other investors and thereby gain money from them.

Investor behavior is a major topic in behavioral finance research, and, of course, a key factor in ASMs. Figure 2 shows how typical investors make decisions, which we can regard as a forecast-preference-action process. To process noisy information and make future forecasts, the investors use forecasters that depend on conceptual models, which might be Baysian, AI style, or—more typically—psychological, as in our research. Using their forecasts, investors make trading decisions on the basis of their preferences, given the risk conditions, and then trade with other investors. The trading results can then provide investors with useful information for updating their forecasters through either rational or psychological learning.

The exchange and environments

The main role that the exchange and environments play is to serve as "containers" that save and release history information and current shocks. Another key factor of the exchange—also known as specialists or market makers in some ASMs—is the institution of *market clearing*, which is the mechanism used to set the market price to match market demand with supply. Many models are interested in the fundamental problem of price formation and the method for determining prices.

Can BSV investors survive?

As the sidebar "Related Work in Agent-Based Financial Computing" explains, the BSV model is one of behavioral finance's most important explanative models. BSV models how behavioral investors form beliefs, which produces both overreaction and mean-reversion for a wide range of parameter values. However, the BSV model doesn't address market equilibrium problems, which are critical for asset pricing theory.

To extend this model to an equilibrium model with both BSV and rational investors in a risky-asset market, we must address two fundamental questions:

- Will anomalies still prevail if we introduce rational investors into the market?
- Will BSV investors really lose money in the game between BSV investors and rational investors, as per Friedman's paradigm?

BSV forecasting

Investors trade two types of assets on an infinite-period basis:

- the risk-free asset, which pays an unchangeable interest rate and is in perfectly elastic supply, and
- the risky asset, which is available in a limited and constant amount across time. This asset pays a random dividend that follows a random walk, in which the next period's dividends will go up or down with the same probability.

The two investor categories—BSV[3] and rational investors—are both risk-averse utility maximizers. The key difference between rational and BSV investors is in their forecasters. Rational investor forecasters are much like those in the traditional financial theories: they are certain about the distribution and mathematic expectation of future dividends, and use this as a basis for decisions.

BSV investors' forecasters have a cognitive bias that is both representative and conservative. Specifically, BSV investors don't believe that future dividends are generated by a random walk process. Rather, they believe dividends are determined by one of two models—Model 1 or Model 2—depending on the economy's state. Neither model is a random walk; in Model 1, dividends are mean-reverting, while in Model 2 they keep the recent trend.

To form the next period's asset price expectation, BSV investors must forecast dividends into the future. Because BSV investors believe that one of two models generates earnings at any time, they must try to understand which of the two models is currently governing dividends. To do this, they observe earnings each period and use this information to make the best guess as to which model governs the dividends. At time t, having observed the dividends shock, investors calculate the probability that Model 1 generated the current dividend. Then, using the new data observed in each period, BSV investors update their estimate of the expected future price in a Bayesian way.

The sim-ASM

To introduce these BSV strategies, we built the sim-ASM. At the beginning of each period, the exchange announces the period's dividend shock to all investors. BSV investors update their expectations, which are distorted by their cognitive bias. Rational investors—who've known the exact global distribution of future dividends—don't have to update their expectations each period.

Next, the exchange sets the tradable assets' market price, which it discovers using a double-sided auction mechanism (that is, the price is adjusted until total market demand meets the total market supply). Both types of investors update their expectations and try to adjust their positions based on the new expectation and current market price. Given the position adjustment, investors will bid or ask for a specific amount of the risky asset at the current price. The exchange then adjusts the market price to make the total bid amount equal to the total ask amount. When the market is clear—that is, the total bid equals the total ask—the market price is the equilibrium price at period t.

Experiments

Our sim-ASM program is a simplified version of SFI ASM and uses the same market architecture as our earlier agent-based model.[7] We wrote sim-ASM using Objective-C and run it on the open Swarm 2.2

Related Work in Agent-Based Financial Computing

Researchers and financial economists have been investigating artificial stock markets for years. The modeling of agent types is a key element of this work.

Artificial stock markets

It's easy to get lost in the many different ASM types. Researchers use many approaches, and it's often difficult to distinguish one ASM from the next. Here, we introduce a few important ones to illustrate ASM history.

Most of the earliest models were intended to create functioning financial market. Stephen Figlewski, for example, tried to examine the efficient market hypothesis in a "real" simulated market.[1] Later, researchers found that "real" ASMs can't be built. They therefore turned to applications in which the economic environment was well understood and often had a simple, homogeneous rational-expectations equilibrium that played the role of useful benchmark. In the LLS ASM, traders maximized a one-period myopic utility function that was constant relative to risk aversion and used simple forecasting techniques.[2]

Other works concentrated on validating AI algorithms. ASM investors in this category didn't really trade with each other; indeed, in some ASMs, there was only one "representative" investor. Jasmina Arifovic developed an ASM in which investors were trained only to forecast future price as accurately as possible, while the price itself was determined by an exogenous process.[3]

The Santa Fe Institute's ASM (SFI ASM)[4] and the ASM developed by Shu-Heng Chen and Chia-Hsuan Yeh[5] are the most frequently cited ASMs. The core philosophy behind both ASMs is to build a kind of dynamic ecology of trading strategies and then examine their co-evolution over time.

Modeling investor types

After modeling individual investors' cognitive bias, Nicholas Barberis and colleagues asserted that all market participants are "*irrational* investors"—that is, they believe, at time t, that the dividend of a stock at $t + 1$ will obey one of two probability laws and either continue or reverse the recent trend.[6] Based on this and other assumptions, the researchers derived their closed-form analytical BSV (Barberis, Shleifer, and Vishny) model. The BSV model explains anomalies such as overreaction and mean-reversion. It also shows the survival of the irrational investor. However, the BSV model requires investor homogeneity; the difference is that this homogeneity is among *irrational* investors. In contrast, James De Long and colleagues introduced the idea of heterogeneous investors.[7] Their DSSW (De Long, Shleifer, Summers, and Waldmann) model consists of both rational and irrational market players. They define the irrationality by simply adding a normal random bias, with a positive mean, to rational investors' price expectations. DSSW proved that irrational investors can gain a higher expected return than rational ones under some conditions. Their exogenous irrationality definition significantly reduced the complexity of deriving the analytical result. However, it did so at the cost of clarity in the connection between biased individual behaviors and market price equilibrium, while the BSV model increases this clarity.

Because the analytical approaches have difficulties dealing with investor heterogeneity and irrationality, new approaches were needed. Due to groundbreaking work by W. Brian Arthur and colleagues,[4] agent-based social modeling—a significant component of social computing[8]—has become more popular in financial studies. Apart from modeling financial markets, a few researchers have reported computational experiments aimed at studying financial markets' economic phenomena. Among these, Moshe Levy and colleagues[9] used an agent-based "microscopic" stock market model to test how investor behavior influenced market dynamics. The researchers defined heterogeneous investors in their LLS (Levy, Levy, and Solomon) models and found evidence of a relationship between investor heterogeneity and asset prices. But, as with DSSW, they failed to give a reasonable explanation of the cause of investors' behavioral heterogeneity, and had to define it exogenously. Meanwhile, among their heterogeneous investor catalogs, the rational investor as defined by Robert Lucas is still missing. Blake LeBaron[10] used SFI ASM to run similar experiments, comparing the memory lengths of two rational investor groups during genetic algorithm learning.

References

1. S. Figlewski, "Market 'Efficiency' in a Market with Heterogeneous Information," *J. Political Economy*, vol. 86, no. 4, 1978, pp. 581–597.
2. M. Levy, H. Levy, and S. Solomon, "A Microscopic Model of the Stock Market: Cycles, Booms, and Crashes," *Economics Letters*, vol. 45, no. 1, 1994, pp. 103–111.
3. J. Arifovic, "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies," *J. Political Economy*, vol. 104, no. 3, 1996, pp. 510–541.
4. W. Arthur et al., "Asset Pricing under Endogenous Expectations in an Artificial Stock Market," *The Economy as an Evolving Complex System II*, W. Arthur, S. Durlauf, and D. Lane, eds., Addison-Wesley, 1997, pp. 15–44.
5. S. Chen and C. Yeh, "Evolving Traders and the Business School with Genetic Programming: A New Architecture of the Agent-Based Artificial Stock Market," *J. Economic Dynamics and Control*, vol. 25, no. 3, 2001, pp. 363–394.
6. N. Barberis, A. Shleifer, and R. Vishny, "A Model of Investor Sentiment," *J. Financial Economics*, vol. 49, no. 3, 1998, pp. 307–343.
7. J. De Long et al., "Noise Trader Risk in Financial Markets," *J. Political Economy*, vol. 98, no. 4, 1990, pp. 703–738.
8. F. Wang et al., "Social Computing: From Social Informatics to Social Intelligence," *IEEE Intelligent Systems*, vol. 22, no. 2, 2007, pp. 79–83.
9. M. Levy, H. Levy, and S. Solomon, *Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena*, Academic Press, 2000.
10. B. LeBaron, "Empirical Regularities from Interacting Long and Short Memory Investors in an Agent-Based Stock Market," *IEEE Trans. Evolutionary Computation*, vol.5, no. 2, 2001, pp. 442–455.

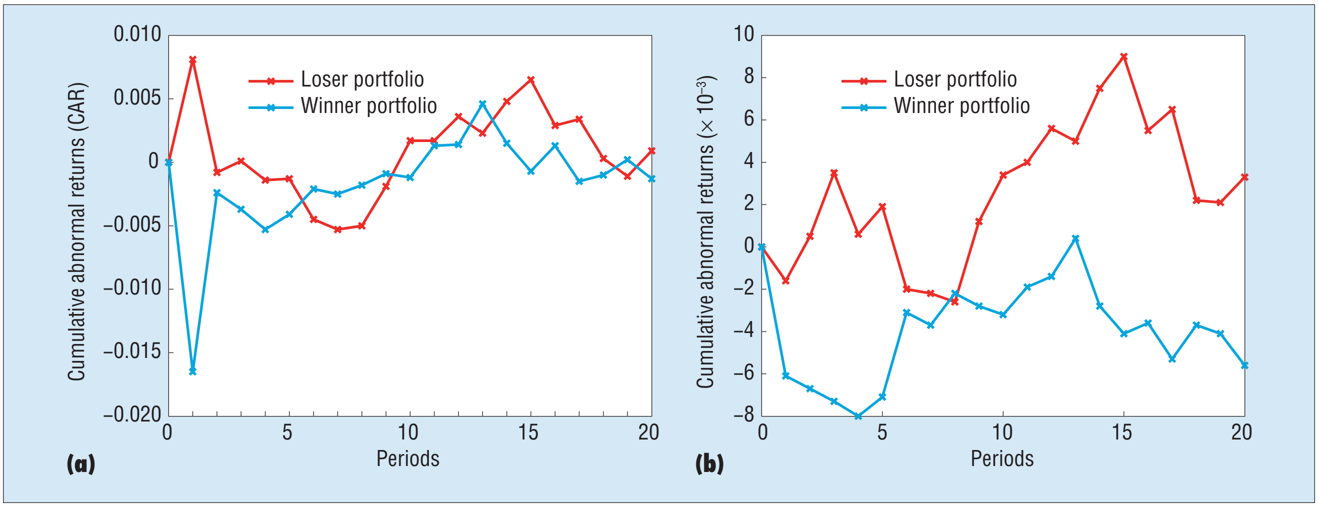

Figure 3. Experimental results. These representative results show overreaction in (a) a sample market and (b) a comparison market.

platform in Linux. We used sim-ASM in two experiments: a sample market experiment and a comparison market experiment. The only difference between the two is that in the sample market experiment, we had 10 BSV investors versus 10 rational investors, while in the comparison market, we had 20 rational investors and no BSV investors. We designed the comparison market to provide benchmarking data to use in our event study method to examine the sample market's overreaction anomaly.

Each experiment consisted of 5,000 periods. We derived the experiment parameters from BSV and SFI ASM. At each period's end, we wrote the risk asset's price, the dividends, the dividends' shock, and the wealth of a representative investor of each type into a data file. This constituted our source data, which we used in our empirical analysis (described in the next section). Figure 3 shows representative results from one experiment. Although we ran several experiments, we found no significant differences among the results.

Results and findings

The aim of our experiments was to discover whether BSV investors can survive—and anomalies prevail—in a market that has both BSV and rational investors. To accomplish this, we compared the empirical results of examining the overreaction and mean-reversion phenomena in the sample market and comparison market.

As figure 3a and our ANOVA (analysis of variance) testing show,[8] in the sample market, the cumulative average residual (CAR) of the winner and loser portfolios differ significantly for one period length. This means that, in the sample market, the price overreacts for one period after both positive and negative dividend shocks, and then mean-reverses. In the comparison market (see figure 3b), the CAR of both portfolios moves randomly, so there's no overreaction phenomenon.

So, obviously, BSV investors can survive and anomalies will prevail in a market with both investor types. That is, the market still shows the overreaction and mean-reversion phenomenon even when rational investors are added to the irrational groups. This is because rational investors won't hold more risky assets unless they are compensated for bearing the risk that BSV investors will change their current models and thus adversely change the risk asset's price. At each period, both BSV and rational investors believe that the risky asset is mispriced in the equilibrium market, but because the price is more uncertain than it is with no BSV investors, neither group is willing to bet too much on this mispricing. Thus, BSV investors "create their own space." In so doing, they also create space for the existing BSV model to maintain its explanative power when extended into a mixed-investor market.

Combining BSV and DSSW to find survival dynamics

The DSSW (De Long, Shleifer, Summers, and Waldmann) model introduced the idea of heterogeneous investors by including both rational and irrational market players.[4] The model's concept of "noise trader risk" sheds light on several financial anomalies. DSSW holds that the unpredictability of noise traders' beliefs creates a risk in the asset's price that deters rational arbitrageurs from aggressively betting against them. As a result, prices can diverge significantly from fundamental values, even in the absence of fundamental risk. Moreover, bearing a disproportionate amount of risk—that they themselves create—lets noise traders earn a higher expected return than rational investors.

DSSW successfully achieves analytical equilibrium between heterogeneous investors, but the model fails to offer a reasonable explanation of the cause of investors' behavioral heterogeneity and must predefine the heterogeneity of investors exogenously. While BSV investigates investors' expectation-formation process and models it endogenously, it's not an equilibrium model. Therefore, combining the two models might help us learn more about market dynamics.

The sv-ASM

Our research investigates investors' rates of return on total assets, rather than on market price. We do this both for simplicity and clarity. To aid our investigation, we developed an ASM based on SFI ASM to clear out the market, including heterogeneous investors. We wrote this sv-ASM program according to the conceptual model on the open Swarm 2.2 platform in Linux (the sv-ASM code is available at http://sv-asm.cvs.sourceforge.net).

The sv-ASM's conceptual model is much like that of sim-ASM, except that it has four categories of investors in the risky-assets market. In addition to BSV investors, sv-ASM includes

- noise traders, who are totally random in how they buy and sell;
- rational-expectation investors, who are smart arbitrageurs who always adopt the genetic algorithm to find and use any opportunity to earn more money; and
- passive investors, who follow the "buy and hold" strategy, and never buy or sell their risky-asset holdings.

In sv-ASM, artificial agents use a forecaster that—according to SFI ASM—can generate rational-expectation equilibrium behavior under certain technical circumstances. (The sv-ASM's genetic algorithm is the same as that used in SFI ASM, which lets agents form rational expectations even when the underlying pricing model is mismatched.[6])

Roughly speaking, the genetic algorithm realizes "rationality" in the following way: At each period, once each agent has selected a price and dividend forecasting rule, the agents make the predictions and then figure out their demand for risky assets. It's now a trivial exercise for an auctioneer to find a price that clears the market by balancing the demand to the fixed supply of risky assets. The sv-ASM uses the same market clearance mechanism as sim-ASM.

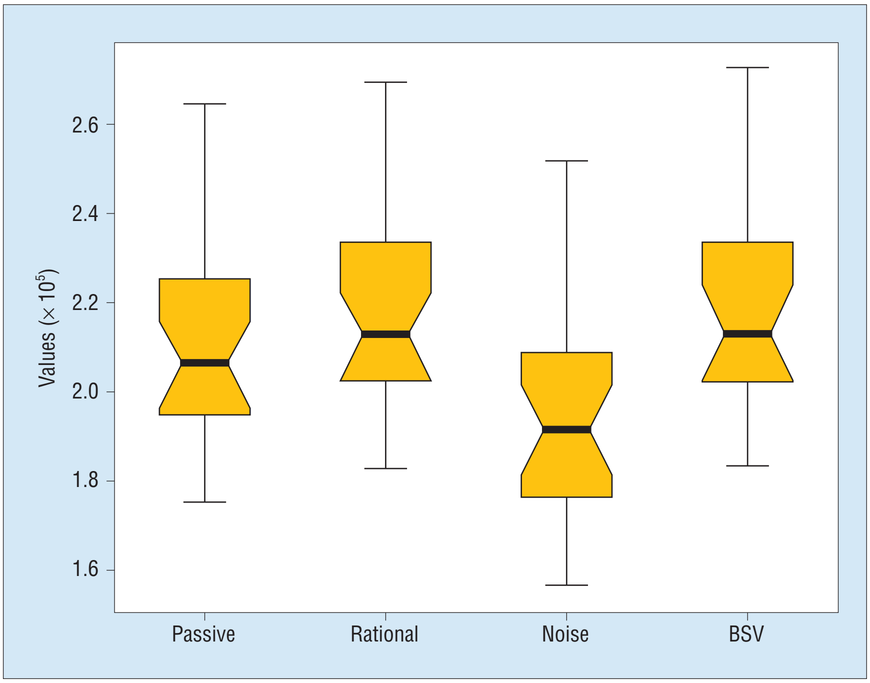

Figure 4. The sv-ASM experiment results. A comparison of the final wealth for each investor category.

Experiments

We ran 24 experiments with different random seeds of dividend generation. We then used an ANOVA process to calculate and compare the different investors' rates of return on total assets in each experiment. In addition, we calculated the frequency of extreme loss for the different investor categories to get a risk indication for the various investor types.

Each experiment consisted of 300,000 periods. (Because sv-ASM's agents use the genetic algorithm to form expectations, its forecasters require a long training time.) At each period's end, we wrote the risk asset's price, the dividends, the dividends shock, and each investor type's wealth into a data file and used the data for empirical analysis.

Results and findings

Figure 4 shows a comparison of the final wealth for the different investor types. According to Friedman,[2] investors who aren't fully rational will gradually lose money when trading with rational investors, and will thus eventually be eliminated from the market. So, in real markets, there won't be any more behavioral investors. That is, the BSV model is useless because the BSV investors are eliminated with the probability of one.

It's therefore important that we examine Friedman's hypothesis that behavioral investors will lose money in their trading with rational investors. Such empirical testing can hardly be done in real markets due to the unavailability of data. However, our ABC approach gives researchers the freedom to use the data of individual wealth fluctuation.

To calculate and compare rates of return on the different investor categories' total wealth from our experiments, we again used an ANOVA process. The results show that the rational and BSV investors achieved the same expected rates of return on total wealth, and also faced the same bankruptcy probability—which means that rational arbitrageurs cannot "eliminate" BSV investors, even in the long run.

Although still in an early stage, our ABC modeling approach could be refined to help solve complex financial economic models in which investors' behavior is heterogeneous and affected by various kinds of cognitive bias. In our examples here, for instance, ASM investors don't learn from each other. Such social learning has been the subject of much discussion in the agent-based modeling community. At one extreme, investors might operate completely on their own, learning rules over time and only interacting with others through common price and information variables. Other mechanisms might try to facilitate some form of communication across investors, or even transfer rule-based information across individuals from generation to generation. How such an information transfer is handled could be critical to market dynamics.

Acknowledgments

The authors are grateful to the Natural Science Funds of China for financial support (Project 70471062), to Gu Zhi-hui, Huang Yi-ping, Liu Hong-hai, Yu Ying, Zhang Qing-yuan, and Zhao Shuai-te, and to participants in the Seminar on Behavioral and Agent-Based Computational Finance, jointly supported by Ma Jun-lu, the director of Department of Finance, Nankai University, and Wei Zhang.

References

1. R. Lucas, "Asset Prices in an Exchange Economy," *Econometrica*, vol. 46, no. 6, 1978, pp. 1426–1445.

The Authors

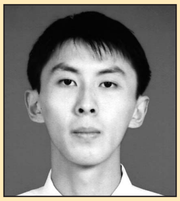

Yongjie Zhang is an assistant professor at the Institute of System Engineering, Tianjin University. His research interests include behavioral finance, agent-based modeling, financial engineering, and empirical studies on financial-market data. He has a PhD in management science from Tianjin University. Contact him at congoz@sina.com.

Wei Zhang is a professor at the Institute of System Engineering at Tianjin University and a professor of management science and vice president at Tianjin University of Finance and Economics. His research interests include economic system engineering theory, agent-based modeling, behavioral finance, financial engineering, venture capital, and financial econometrics. He has a PhD in engineering from Tianjin University. Contact him at zhangwei@nsfc.gov.cn.

2. M. Friedman, "The Case for Flexible Exchange Rates," *Essays in Positive Economics*, M. Friedman ed., Univ. Chicago Press, 1953, pp. 157–203.

3. N. Barberis, A. Shleifer, and R. Vishny, "A Model of Investor Sentiment," *J. Financial Economics*, vol. 49, no. 3, 1998, pp. 307–343.

4. J. De Long et al., "Noise Trader Risk in Financial Markets," *J. Political Economy*, vol. 98, no. 4, 1990, pp. 703–738.

5. M. Levy, H. Levy, and S. Solomon, *Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena*, Academic Press, 2000.

6. W. Arthur et al., "Asset Pricing under Endogenous Expectations in an Artificial Stock Market," *The Economy as an Evolving Complex System II*, W. Arthur, S. Durlauf, and D. Lane, eds., Addison-Wesley, 1997, pp. 15–44.

7. W. Zhang et al., "BSV Investors Versus Rational Investors: An Agent-Based Computational Finance Model," *Int'l J. Information Technology & Decision Making*, vol. 5, no. 4, 2006, pp. 455–467.

8. W. De Bondt and R. Thaler, "Does the Stock Market Overreact?" *J. Finance*, vol. 40, no. 3, 1985, pp. 793–805.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.