Financial-market stability in the presence of heterogeneous adaptive agents

Ziqiang Wu

College of Management and Economics, Tianjin University, China

Goal

To investigate the interplay between financial-market dynamics and emerging investor personalities of agents with the ability of self-adaptation and social learning.

Background and motivation

The financial system is the core of our current economy. Financial markets, as the most important part of the financial system, provide an efficient way to trade assets of various kinds. Asset prices determine investor decisions and depend on them in turn, leading to complex dynamics that are prone to drastic fluctuations in both asset price and investor wealth; an issue that, although not new, has recently attracted mounting public attention. Traditional financial research has usually assumed that investors are homogeneous; in the real world, however, they are heterogeneous and adaptive. This means that they can adjust, over and over again, their beliefs and strategies according to the market situation and their own performance histories. Such adaptability is in fact necessary to achieve market efficiency, but it is also the force at the origin of great market tremors and financial crises.

Agent-based models allow analyzing the interplay between investor personalities and financial market dynamics without a need for artificially constraining the heterogeneity or adaptability of investors. To establish the proper context for such modeling efforts, we elaborate on these two key points in turn.

Investors are heterogeneous. In a real financial market, there are many participants, like fund managers, investment bankers, speculators, and individual investors, and these differ, e.g., with respect to their beliefs, motives, risk preferences, time horizons, and trading frequency. Wolf et al. (2007) showed that competition may result in stable coexistence of different explorative and risk-taking strategies in animals, leading to characteristic “animal personalities”. According to Simon (1956) and to Kahneman and Tversky (1979), economics and finance are witnessing an important (if actually slow) paradigm shift, from a representative, rational-agent approach towards a heterogeneous, bounded-rational approach.

Investors are adaptive. Based on empirical studies, Ippolito (1992), Del Guercio and Tkac (2002), Benartzi and Thaler (2007), and De Jong et al. (2009) found that investors change their behavior, which I turn affects market prices. Investors do this either through individual learning or through social learning. Chen and Yeh (2001) introduced a “business school” mechanism to implement social learning. Chang (2007) described a learning process via various individual and social learning factors. Ryuichi (2010) allowed agents to learn other agents’ decision parameters directly, which may be considered unrealistic, since such parameters are usually not observable as such and can be inferred from observable actions only with great uncertainty.

Social learning through imitation tends to reduce investor heterogeneity, creating a trend to “follow the herd”. This can often be a successful strategy, but it may also create market situations that greatly reward minority decisions. This is reminiscent of the El Farol Bar Problem of Arthur (1994) and the Minority Game of Challet and Zhang (1997), in which congestion entails low efficiency. This suggests that there is a collectively optimal level of social learning that preserves investor heterogeneity while still fostering market efficiency.

The Santa Fe Institute (SFI) artificial stock market model (Arthur et al. 1997) is an early and widely studied agent-based market model, designed to explore asset-pricing theory by employing heterogeneous agents. Such agent-based methods have recently become popular, because they enable a repeatable and controllable approach to otherwise intractably complex dynamics. Agents in the SFI artificial stock market achieve individual learning though updating their strategies via a genetic algorithm (GA). LeBaron (2001) revised the SFI model to include social learning via a strategy pool, which is not only more efficient but also better reflects reality. In this project, we will thus follow the latter approach.

Research questions

Modeling a group of agents with different risk attitudes and the ability to learn from history and each other, I will address the following questions:

- Which population structures emerge, and how do they affect asset prices and wealth distributions?
- How does the introduction of social learning influence market dynamics?
- Can we identify population structures that promote or threaten market stability?
- What is the relationship between trading frequency and wealth distribution?

Methods and work plan

Model overview

Our model is agent-based and uses discrete time. Agents can divide their wealth over two assets: firstly, there is a risk-free bond (not modeled explicitly), paying a known interest rate; secondly, there is a risky stock, paying an uncertain dividend that fluctuates according to an AR(1) model. At the end of each period, the agents may reallocate their wealth; by selecting the number of stocks and bonds in their *portfolio*, they attempt to maximize the expected utility of their wealth. To assess this expected utility, the agents apply *forecasting* strategies to estimate the value of the stock in the next period; the number of stocks they want in their portfolio depends on the current and expected future *price*. To satisfy the *demand* thus created, the market offers a *trading* scheme that adjusts prices and demands accordingly. By continually comparing predicted and actual stock values, agents assess the usefulness of their forecasting strategies and *learn* to use good ones. All agent decisions are influenced by four evolvable traits characterizing an investor's *personality*.

Portfolio

Agents divide their wealth between a bond and a stock. With x_t^i denoting the number of shares held by agent i at time t, p_t the price of a share, d_t the dividend, r_t the risk-free interest rate (initially assumed to be fixed for all times, $r_t = r_{\mathrm{f}}$), the wealth W_{t+1}^i of agent i at time $t+1$ is given by

$$W_{t+1}^i = x_t^i(p_{t+1} + d_{t+1}) + (1 + r_t)(W_t^i - p_t x_t^i)\,.$$

The expectation of the share value $v_{t+1} = p_{t+1} + d_{t+1}$ (and therefore of W^i_{t+1}) depends on an agent's forecasting strategies (see below). Investment into the bond is modeled implicitly: wealth not allocated to the stock is implicitly invested in the bond. Agents attempt to maximize their utility

$$U(W^i_{t+1}) = -\exp(-\gamma^i W^i_{t+1}),$$

which represents constant absolute risk aversion (CARA) (von Neumann and Morgenstern 1944), with γ^i denoting the coefficient of absolute risk aversion of agent i.

Forecasting

Each agents i has a list of strategies $(a_j, b_j, c_j)^i$ to forecast the expected stock value as

$$E^i_t(v_{t+1}) = a^i_j(v_t) + b^i_j(v_t - v_{t-1}) + c^i_j,$$

which amounts to estimating stock values as an AR(2) process. In each period, agents reorder their strategies according to the squared deviation of their recently predicted stock value from the actual current stock value. Each strategy is thus assigned an updated weight index

$$w^i_{t,j} = (1-\mu^i)w^i_{t-1,j} + \mu^i(v_t - E^i_{t-1}(v_t))^2,$$

where μ^i describes the agent's propensity to accept new information. Agents pick a strategy for making the next forecast with a probability proportional to $(w^i_{t,j})^{-\sigma^i}$, where σ^i describes the relevance the agent assigns to a success-based ordering of its strategies.

Price and demand

It can be shown (von Neumann and Morgenstern 1944) that agents maximize their expected CARA utility by holding

$$x^i_t = \frac{E^i(v_{t+1}) - p_t(1+r_t)}{\gamma w^i_{t-1,*}}$$

shares, where $w^i_{t-1,*}$ is the weight index of the forecasting strategy the agent currently uses (which roughly reflects the historical variation in this rule's forecasting success, see its definition above). The demand thus generated depends on the current price p_t.

Trading

The agents' respective demands are collected and trading prices are adjusted until supply meets demand. While trading is in this sense optimal, we allow for the possibility that agents do not have equal access to the market: in an attempt to mimic the situation in real markets, we thus assume that wealthier agents enter the market more frequently. Agents that do not enter the market within a given period do not participate in the price-setting scheme and cannot reallocate their wealth during that period.

Learning

If the weight index of a strategy exceeds an agent's weight threshold τ^i, the agent replaces this strategy by a new strategy drawn randomly from a set of centrally published strategies. This strategy pool contains n strategies, m of which are the m most successful strategies in the agent population and the rest are randomly generated strategies. Large values of m correspond to a high degree of social learning.

Personalities

For each agent, the vector $(\gamma, \mu, \sigma, \tau)^i$ describes its “investor personality”. In each period, these personality traits are subject to mutations with small probability, upon which new trait values are drawn from Gaussian distributions of given widths centered on the old trait values.

Model implementation

The algorithm sketched above will be implemented in C, while the model output will be analyzed in Matlab.

Work plan

The work plan is as follows:

- Using a set of “reasonable” personality trait values, motivated by common sense, run the model with a monomorphic, non-evolving investor population, and visualize asset prices, wealth distribution, and strategies.
- Identify, if possible, model parameters that produce asset prices and wealth distributions with “realistic” qualitative features. Assess whether there are accompanying strategy patterns and how these can be interpreted.
- Let investor traits evolve and observe the resultant trait distributions. Vary model parameters to establish stable polymorphisms through evolutionary branching.
- Add social learning and observe its effects. Look for a critical level of social learning at which its costs (loss of heterogeneity, leading to an inability to react to market fluctuations) outweigh its benefits (ability to follow and stabilize trends).
- Document parameter choices and describe emerging patterns that appear to be correlated with stable/unstable market phases.
- Observe the effects of trading frequency on wealth distribution.

Relevance and link to EEP’s research plan

Poverty and Equity is one of the three global problem areas on which IIASA research is planned to focus. In real financial markets, big investors often have a systematic advantage over small investors. With this project, we hope contribute to an increased understanding of the conditions required for the emergence of fair and equitable markets. By interfacing research on financial markets with recent advances in understanding the evolution of animal personalities, this interdisciplinary work is also related to EEP’s research project on *Evolving Biodiversity*.

Expected output and publications

This research is intended to be published as a coauthored article in an international scientific journal.

Reference

Arthur WB (1994). Inductive reasoning and bounded rationality. *American Economic Review* 84: 406–411

Arthur WB, Holland JH & LeBaron B (1997). Asset pricing under endogenous expectations in an artificial stock market. Pages 15–44 in Arthur WB, Lane D & Durlauf S eds. *The Economy as an Evolving Complex System II,* Addison-Wesley, Boston, USA

Benartzi S & Thaler RH (2007). Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives* 21: 81–104

Challet D & Zhang YC (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A* 246: 407–418

Chang SK (2007). A simple asset pricing model with social interactions and heterogeneous beliefs. *Journal of Economic Dynamics and Control* 31: 1300–1325

Chen SH & Yeh CH (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: 363–393

Del Guercio DD & Tkac PA (2002). The determinants of the flow of funds of managed portfolios: Mutual funds vs. pension funds. *Journal of Financial and Quantitative Analysis* 37: 523–557

Ippolito RA (1992). Consumer reaction to measures of poor quality: Evidence from the mutual fund industry. *Journal of Law and Economics* 35: 45–70

Kahneman D & Tversky A (1979). Prospect theory: An analysis of decisions under risk. *Econometrica* 47: 263–291

LeBaron B (2001). Evolution and time horizons in an agent-based stock market. *Macroeconomic Dynamics* 5: 225–254

Ryuichi Y (2011). Volatility clustering and herding agents: Does it matter what they observe? *Journal of Economic Interaction and Coordination* 6: 1–19

Simon HA (1956). Surrogates for uncertain decision problems. Unpublished ONR Research Memorandum #38, GSIA, Carnegie Institute of Technology, USA. Pages 235–244 in Simon HA (1982) *Models of Bounded Rationality, Volume 1, Economic Analysis and Public Policy*. MIT Press, Cambridge, USA

von Neumann J & Morgenstern O (1944). *Theory of Games and Economic Behavior.* Princeton University Press, Princeton, USA

Wolf M, van Doorn GS, Leimar O & Weissing FJ (2007). Life-history trade-offs favour the evolution of animal personalities. *Nature* 447: 581–584