New ways to understand financial markets

Edward Tsang
Working Paper WP046-10
Centre for Computational Finance and Economic Agents (CCFEA), University of Essex
21 October 2010

1. Fundamental vs technical analysis

Traditionally, researches in financial markets were broadly divided into two camps: fundamental and technical analysis. Fundamental analysts assume that assets, such as shares, have intrinsic values, which they attempt to assess. The value of a company's shares is determined by all the factors that could affect the company's future profitability. These include the company's operation as well as the environment in which the company operates. Examples of relevant factors are: the company's asset structure, operations, revenues, expenses, personnel, etc. and global factors include interest rates, inflation rates, etc.

Technical analysts, who are sometimes referred to as chartists, assume that there are recurring patterns in share prices. They attempt to find such patterns from the history of the prices. For example, moving averages of a share may be computed. A simplistic rule may be: "if the price has risen above the 20-days moving average (when it was below before), then it is time to buy". Another simplistic rule may be: "if the price is above the maximum price in the past 250 days, then it is time to sell". In practice, more complicated patterns were used.

Fundament analysts dismiss technical analysts' approach for lack of theoretical foundations. There is no guarantee that patterns found earlier will repeat themselves later. Technical analysts believe that prices are driven by traders, and traders do response to price changes in particular ways, which can be charted. It is not within the scope of this paper to join this heated debate. Instead, this paper suggests that there are new, sensible ways to study financial markets, which could offer us something that neither fundamental analysts nor chartists can.

2. Classical economics built on unrealistic assumptions

Fundamental analysts are based on research in classical economics. Classical economics use mathematics to assess asset prices or model market behaviour. Mathematical models can be rock solid, as long as the assumptions are sound. Unfortunately, some of the assumptions used in theoretical economics are seriously flawed. It is very dangerous to use models built on flawed assumptions to explain what we observe in real markets. It is also very dangerous to rely on such models to set government policies.

One of the most fundamental assumptions in classical economics is rationality. The assumption is that decision makers, whether they are individuals or institutions, are perfectly rational. That means whenever an investor has clear objectives, he/she will be able to make the right, optimal, decision. We shall argue below that this is not true. Another often made assumption is homogeneity in markets. That means every decision maker has the same objectives, and will arrive at the same valuation on the same asset. We shall argue below that this assumption is unrealistic as well.

Classical economics assume that investors are all rational. This means that they will make the optimal decision, given their objectives. For example, as soon as company X announces a 250 Million Pounds profit, investors will be able to reassess the share price of X, and decide whether buying, selling or holding shares of X is the optimal decision. In fact, all rational investors will be able to assess the impact of the announcement thoroughly, and come up with the "correct" decision.

For simple decisions, decision making may be trivial. For example, if a producer is given two offers for its product, £10 and £20, it may be easy to decide which offer to accept. But when more information is taken into consideration, the decision may not be trivial. For example, these two customers may have different potentials for more trades in the future. Furthermore, from a computational point of view, many problems are intractable by nature. Chess is a good example, in which we have yet to find the optimal moves from start. To find the optimal moves, all combinations of moves must be considered. The combinations grow exponentially. Clever algorithms may help us to eliminate some combinations. Clever heuristics may help us to eliminate more combinations, but risk missing some optimal moves.

Combinatorial explosion prevents one from finding optimal solutions for even very small problems within days, or indeed, years. If combinatorial explosion can be contained easily, passwords are useless. Combinatorial explosion drives a large research agenda in computer science. Being perfectly rational requires a decision maker to contain combinatorial explosion, which requires breakthroughs in computer science which are yet to come.

Economists generally accept that human decision makers are not perfectly rational. They are "bounded rational", which means they can make nearly optimal decisions. If the perfect rationality assumption is replaced by bounded rationality, then all the text economic theories that were based on the former must be revised. Effectively, all economics textbooks must be rewritten. Unfortunately, there is no consensus on exactly what bounded rational means. In order to revise quantitative economics theories, one must be able to study the behaviour of bounded rational decision making mathematically, or at least experimentally.

3. Technical analysis scratches the surface of matters

Many technical analysts believe that useful patterns have been found, which led to profitable trades. Some argued that price changes reflect collective behaviour by investors in the market; while the behaviour of each individual may be unpredictable, collective behaviour form patterns, which can be analysed. For example, one may not be able to predict the behaviour of each wildebeest, but their herding behaviour can be analysed. Although one cannot predict where an individual will stand before he enters a lift, it is a pretty good guess that if two strangers enter a lift, they will stand in opposite corners.

Technical analysts build their faith on the assumption that future prices are driven by what happened in the past. There is certainly some truth in it. However, past share prices are only the results of interactions among all the traders. They do not reflect all the factors that lead to price movements. Some of these factors are the order book, positions and margins, which are explained below.

Prices are changed by buy and sell orders by the traders. A trader may instruct his broker to buy or sell a specific quantity of shares at market prices. This is called a market order. Another trader may instruct her trader to buy or sell a specific quantity of shares at a specific price. This is called a limit order. Limit orders are stored in an order book, ordered by the prices specified. A limit order to buy at price X will be executed when there are sellers who are willing to sell at price X. The order book gives significant information about market liquidity, i.e. how easy it is to buy or sell a certain number of shares. Therefore, if price movements are predictable at all, as technical analysts believe, then information about order books should be included in technical analysis. Unfortunately, this is rarely the case as the majority of analysts do not have access to order books. An individual broker will have access to its own order book, but it may not reflect the full picture.

Another set of factors that affect price movements are the margins that individual traders trade with. That means each trader deposits a certain amount of money with her broker. The broker allows them to trade more than their deposit. For example, with a deposit of £1,000, she may be allowed to buy £20,000 worth of stocks. In that case, she has a margin of 5%. However, if, at any point, her position makes a potential loss beyond £1,000, the broker will force her to cut her loss and close her positions (i.e. selling the shares that she is holding or buying the shares that she has shorted). The positions of the traders, including the number of shares they hold or short, the prices that they bought or sold the shares for, and the traders' margins, will all affect the future changes in prices. Unfortunately, due to unavailability (like order books), such information is not widely included in technical analysis. This seriously limits technical analysts' ability to predict price movements.

To summarize, even if one accepts technical analyst's assumption that future prices are affected by what happened in the past, past prices and indicators derived by most technical analysts today are pretty poor in predicting price changes in the future. Past prices do not reflect information such as the order book and traders' positions and margins. Even if (a) observable patterns can be identified, and (b) these patterns repeat themselves, technical analysis is only scratching the surface of the markets.

4. Agent-based artificial markets explore possible worlds

As a result of advances in computing, a new approach to study markets was proposed. The idea is to model the market mechanism and the participants, and observe the simulations. Simulation results could show us what would happen under what situations. This allows us to gain insight in the possible outcomes of a market. This may help to design new market rules (which is called mechanism design), or to design and evaluate strategies (to beat the market). This idea was popularized by the Santa Fe Institute (e.g. see Ehrentreich 2007).

Why should one be interested in artificial market? The attractiveness of models is that they allow us to study multiple futures. Life cannot be rewound. From time to time, the Central Bank has to decide whether to change the interest rate, and by how much. Once it decides to raise interest rate by, say, 0.5%, no one would know what would have happened had interest rate been raised by 1% instead. With a model, we can simulate the effects of different scenarios. Aeroplane and car designers would use wind tunnels to test their designs before they go into production. It would be inconceivable to launch a policy before analysing it thoroughly. Models and simulations could be used to study alternative policies, hence play the role of economic wind tunnels (Chen 2005).

We know that "*all models are wrong, but some are useful*" (Box & Draper 1987). Nevertheless, models enable us to study a policy or a strategy scientifically. A good model could help us to identify components that are most relevant to the subject of our study. Wrong models can be refined. Therefore, models enable us to advance knowledge collectively. "*More calculation is better than less, Some calculation is better than none*" (Sun Zi 600BC)

Models allow us to detect adverse situations. Some of the simulations could reveal paths to undesirable situations. For example, simulations could by chance encounter conditions under which a large number of banks in an economy could fail. Safety mechanism can be built in to guard against such conditions, e.g. through the setting of reserve requirements.

Models also help us to study causal relations. When we observe a phenomenon in the market, it is often very difficult to determine what has caused it to happen. Through model simulations, the experimenter may keep everything constant, and change only one or a set of variables. If the observed results of the simulation deviate from the results from the original model, deviation is accounted for by the variables changed. This allows the experimenters to establish causal relations.

Artificial markets have been used to study rules for markets and strategies for succeeding in markets. Following are some examples. (It is certainly not exhaustive.) Nicolaisen et al (2001) used artificial markets to design and evaluate market mechanism in the electricity market. Alexandrova-Kabadjova (2007) used artificial markets to study bank strategies in the credit cards market. Jin and Tsang explored game strategies in bargaining. Gosling and Tsang (2006) investigated trading strategies for middlemen in a supply chain. Martinez-Jaramillo and Tsang (2009) discovered conditions under which artificial markets exhibit behaviour similar to real markets. Chen and Yeh used artificial markets to study the role of learning by traders (2001). Sunders (2004) and Cliff (2009) used "zero intelligent agents" to examine how much is the market behaviour due to the rules in the market, and how much of it is due to traders' intelligence.

5. Market as a hard science

Physics is studied by conducting experiments, collecting data and generalizing the findings. Biology started from observations, experimentation and data collection, before generalization were made. There is no reason why markets should not be studied similarly.

Olsen (2009 and Olsenworld web link) suggests that one should study markets by looking at its microstructures. That means one should look at every single movement in prices (these are called high-frequency data), as opposed to sampling the price once every n seconds, minutes or hours, as technical analysts typically do. One should also look at every single observable action by every trader (placing of each order), in conjunction with his/her positions and margins. Olsen started their research by looking at the price movements in the foreign exchange market.

One difficulty with having so much data on hand is knowing where to pay attention to. Olsen et al defined a concept called "directional change", which characterizes price movements. In an upward run, a directional change with threshold t% is defined as the event of the price dropping from its highest point by t%. This directional change starts a downward run.

Similarly, in a downward run, a directional change of threshold t% happens when the price going up from its lowest point by t%, which starts an upward run. Directional changes help us to focus on significant events, where significance is defined by t. The observer may choose to use any t that suits him/her.

Directional changes capture the movements in the market. Time is intrinsically defined by directional change events. Physical time can be, and is often independent of intrinsic time. In one minute, there could be no directional changes. In the minute after, 100 directional changes could be observed.

Directional change is interesting because it is easy to understand and easy to observe. They are made more attractive by the regularities discovered by Glattfelder et al (2008). These regularities follow the pattern of power laws, and are independent of the threshold t. For example, if it take x minutes for the prices to drop t% from a peak, then it takes twice as much time (i.e. 2 times x minutes) for the price to reach the bottom. This is a rather striking regularity, given the seemingly irregularity of trades in the foreign exchange market. It is also interesting because it covers all values of t. Such regularities show the usefulness of directional change as an observable event.

Another approach to studying microstructures of markets is worth mentioning. It is based on the following observation: Every mechanism in a market is by design. Therefore, we should know exactly how the market clearing program handles orders placed by the traders. That means it should be possible to write down the clearing process concretely using a formal language. Writing down these processes provides us with a chance to study their properties mathematically, statistically or algorithmically. This approach is taken by Tsang and Olsen (2010), through the use of Event Calculus, which is summarized below.

Stock markets are mainly driven by the participants' positions, margins and orders, which were explained earlier. For simplicity, we only consider in this paper double-auction markets, where the only participants buyers and sellers, plus a mechanism that handles the orders. (The other major alternative is to have a market-maker, who buys and sells from the traders.) At any point in time, the consequence of an order (to buy or sell a specific quantity at a specific price) is determined by the market participants' positions, margin and orders (unless the clearing program introduces randomness in the process). That means if all these data are available, one should be able to tell how much the market is going to drop if someone sells 10 million shares. This will help us to know how fragile the market is. For example, if a market will drop by, say, 15% as a result of a sell order of 10,000 shares, and selling 10,000 shares was not that uncommon in the past, then we know that the market is in a fragile situation. With proper management, an early warning system can be built to help preventing future crises in financial markets.

6. Computational intelligence in economics research

Computational intelligence plays a major role in economics research, whichever approach one chooses to take.

In fundamental analysis, the perfect rationality assumption must be relaxed. Bounded rationality can be defined by the algorithms and heuristics that one uses. Once bounded rationality is defined, one can define rational decisions and equilibriums. This was demonstrated in bargaining theory (Jin 2007) (Jin et al 2009).

In technical analysis, computational intelligence algorithms can find complex patterns, e.g. see (Bauer 1994), (Neely & Weller 1997), Li & Tsang (1999, 2000) and Garcia-Almanza & Tsang (2007). These works show that learning methods can find complex patterns that are impossible to find manually.

Artificial market and wind tunnel testing build on computational techniques. It also benefits from advances in hardware (which is no credit to computational intelligence).

7. Concluding Summary

Classical economics relies on some unrealistic assumptions. Technical trader is only scratching the surface of matters. They should not be the only approaches to studying markets. Markets should be studied more like hard science, where (i) experiments are performed and repeated, and (ii) theories are driven by observations, with massive data collected. Artificial markets and market science are two new approaches to studying markets. The former facilitates repeatable experiments. The latter is data driven and the latter builds on data collection, data analysis, and process description.

References

[1] Alexandrova-Kabadjova, B., Artificial payment card market - an agent based approach, PhD Thesis, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, 2007

[2] Bauer, R.J.Jr., Genetic algorithms and investment strategies, John Wiley & Sons, Inc., 1994

[3] Box, G. & Draper, N., Empirical model-building and response surfaces, Wiley, 1987

[4] Chen, K-Y, An Economics Wind Tunnel: The Science of Business Engineering, in Morgan, J., Experimental and Behavioral Economics - Advances in Applied Microeconomics Volume 13, Elsevier Press, 2005, Chapter 6, 133-167

[5] Chen, S-H. & Yeh, C-H., 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, Vol.25, 2001, 363-393

[6] Cliff, D., ZIP60: Further Explorations in the Evolutionary Design of Trader Agents and Online Auction-Market Mechanisms, IEEE Transactions on Evolutionary Computation, Vol.13, No.1, February 2009, 3-18

[7] Dacorogna, M.M., Gencay, R., Muller, U., Olsen, R.B. & Picktet, O.V., An introduction to high frequency finance, Academic Press 2001

[8] Ehrentreich, N., Agent-Based Modeling: The Santa Fe Institute Artificial Stock Market Model Revisited, Springer, 2007

[9] García-Almanza, A.L. & Tsang, E.P.K., Detection of stock price movements using chance discovery and genetic programming, International Journal of Knowledge-based and Intelligent Engineering Systems, Vol.11, No.5, December 2007, 329-344

[10] Glattfelder, J.B., Dupuis, A. & Olsen, R., Patterns in high-frequency FX data: discovery of 12 empirical scaling laws, Quantitative Finance, forthcoming (earlier

version: An extensive set of scaling laws and the FX coastline, Working Paper WP025-08, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, September 2008)
[11] Gosling, T. & Tsang, E.P.K., Tackling the simple supply chain model, Proceedings, 2006 Congress on Evolutionary Computation, Vancouver, Canada, 16-21 July 2006, 7943-7950
[12] Jin, N., Constraint-based co-evolutionary genetic programming for bargaining problems, PhD Thesis, Department of Computer Science, University of Essex, UK, 2007
[13] Jin, N., Tsang, E. & Li, J., A constraint-guided method with evolutionary algorithms for economic problems, Applied Soft Computing, Vol.9, Iss.3, June 2009, 924-935
[14] Li, J. & Tsang, E.P.K, Investment decision making using FGP: a case study, Proceedings of Congress on Evolutionary Computation (CEC'99), Washington DC, USA, July 6-9 1999.
[15] Li, J. & Tsang, E.P.K, Reducing failure in investment recommendations using genetic programming, Computing in Economics and Finance Conference, Barcelona, July 2000
[16] Martinez-Jaramillo, S. & Tsang, E.P.K., An heterogeneous, endogenous and co-evolutionary GP-based financial market, IEEE Transactions on Evolutionary Computation, Vol.13, No.1, 2009, 33-55
[17] Neely, C., Weller, P. & Ditmar, R., Is technical analysis in the foreign exchange market profitable? A genetic programming approach, Journal of Financial and Quantitative Analysis, 32, 1997, 405-26
[18] Nicolaisen, J., Petrov, V. & Tesfatsion, L., Market power and efficiency in a computational electricity market with discriminatory double-auction pricing, IEEE Transactions on Evolutionary Computation, Volume 5, Number 5, October 2001, 504-523
[19] Olsen, R. & Cookson, C., How science can prevent the next bubble, Financial Times, 16th February 2009
[20] Olsenworld http://www.olsen.ch/
[21] Sun Zi, The art of war, 600BC, http://en.wikipedia.org/wiki/The_Art_of_War
[22] Sunder, S., Market as an artefact aggregate efficiency from zero intelligence traders, In Augier, M.E. & March, J.G. (ed.), Models of a Man: Essays in memory of Herbert A. Simon, Cambridge, MA: MIT Press, 2004, 501-519
[23] Tsang, E.P.K., Computational intelligence determines effective rationality, International Journal on Automation and Control, Vol.5, No.1, January 2008, 63-66
[24] Tsang, E.P.K., Olsen, R. & Masry, S., Event Calculus on high frequency finance, Working Paper WP038-10, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, (revised) October 2010