Franky Takes on Wall Street:

Adding a Technical Trader to a Limit Order Based Market

Jeremy Barofsky

Special Thanks to Laci Gillemot, Jim Rutt, Paolo Patelli, Eric Smith, Marcus Daniels, Illya Guilo, Doyne Farmer, the Santa Fe Institute, the National Science Foundation, and The General.

August 18, 2002

Introduction

The Walrasian equilibrium paradigm has dominated economics since the turn of the last century. This group of ideas largely defines how businessmen, economists, policymakers, and the general public understand the economy and how to solve its problems. Even with this significant influence, many of the theory's basic assumptions are gross abstractions from empirical reality. Tellingly, most economists themselves do not defend its strictest versions. Nevertheless, they continue to teach this standard model, thereby perpetuating its hold on how we conceptualize, measure, and change our economic lives.

This summer, I worked on a research project that sharply breaks with many of the ideas that underpin this paradigm. Instead, Doyne Farmer led a group of researchers that studied the economy as a complex, evolving system. They investigated the basic institutional structure of modern financial markets- the limit order book- and how random order flow affects price dynamics. I added a contrarian technical trading agent to this simulated market and investigated his profitability. To understand the complexity approach and its context within the sphere of economic ideas, it is imperative that I introduce the established theory.

I must first state that any attempt to describe what economists believe as a group is doomed to failure. Indeed, agreement among economists themselves is rare. As Samuelson notes in the introduction of his classic text, “When a royal commission asked five economists for an opinion, it was said they would get six answers—two from Mr. Keynes,” (Samuelson and Nordhaus, xxiii).

The Context: Walrasian Equilibrium Economics

The assumptions that underlie Walrasian economics are agent optimization, rational expectations, full knowledge of prices, an auctioneer, market clearing, and perfect competition. Agent optimization constitutes the paradigm's theory of behavior. It postulates that agents make decisions by evaluating the expected pay-offs from each possible action and choosing the action that will maximize utility. In this model the concept of utility incorporates all agent values, preferences, and beliefs.

The theory of rational expectations posits that each agent calculates their expected pay-offs by estimating how others are going to act, assuming they also rationally maximize utility. For any situation, all agents make decisions based on the same pay-off and decision-making structure.

For each agent to estimate their pay-offs, they must have full knowledge of prices. Before a transaction can occur, each agent determines her demand curve; what quantity of the given product she is willing to buy or sell at any possible price. Once agent demand curves permit effective pay-off estimation, Walras assumed that an auctioneer took this information and found a price that would equilibrate supply and demand.

Rational agents only make trades when an exchange benefits them. If the quantity that they want to buy or sell is constrained then this clause is not fulfilled. Therefore, Walrasian theory includes the assumption of market clearing to insure that all transactions are, by definition, beneficial.

The assumption of a market in perfect competition means agents are price takers so that they cannot raise prices above their equilibrium levels. This clause also necessitates perfect information because asymmetric information can lead to rent-taking as well. With perfect competition and perfect information, if a product's price strays from its equilibrium value and rent-taking occurs, rational agents enter the market and arbitrage away any profit opportunities, thus returning the market to equilibrium.

Adopting all these assumptions yields the grand conclusion of the Walrasian paradigm: the First Fundamental Theorem of Welfare Economics. Katz and Rosen define the theorem in their intermediate microeconomics text:

"As long as producers and consumers act as price takers and there is a market for every commodity, the equilibrium allocation of resources is Pareto efficient… In other words, an economy composed entirely of price takers—a competitive economy—'automatically' allocates resources efficiently, without any need for centralized direction." (Katz and Rosen, 389).

This is an extremely compelling result because it says that, in its ideal state, the market mechanism generates an optimal resource allocation without the intervention of outside authority. Katz and Rosen continue by explaining that the Fundamental Theorem separates efficiency and equality problems and that policy intervention can be useful, even Pareto-improving, without competitive markets. Nevertheless, this essential conclusion represents a widely held belief for why market economies have outperformed all other types of economic organization.

Efficient Markets Hypothesis

Derived from the above Walrasian analysis or on its own, the Efficient-Markets Hypothesis represents modern finance theory's explanation of how financial markets function. In finance, efficiency usually refers to informational not Pareto efficiency. Sharpe (1995) explains:

"A market is efficient with respect to a particular set of information if it is impossible to make abnormal profits by using this set of information to formulate buying and selling decisions".

Therefore, no rent-taking opportunities are available from asymmetric information. Since all agents rationally and instantaneously incorporate all information about a stock that would determine its value, efficient-markets predicts that a stock's price should equal its underlying value. If price does not equal value then other agents, with perfect information, will arbitrage away profit opportunities until price reflects the stock's equilibrium value.

When prices equal value, only new information will change a stock's underlying value and thus price. Since prices reflect all predictable events and news enters the market randomly (otherwise it would be predictable and, therefore, not news), prices should follow a random walk. This theory predicts that the price autocorrelation between time periods is zero and, essentially, prices are determined by flipping a coin. These conclusions lead to some of finance theory's classic conclusions, like the impossibility of predicting the stock market, the futility of timing the market, and its emphasis on long-term investment.

Some claim that financial markets represent the purest of Walrasian markets because they instantaneously equilibrate supply and demand exhibit. Therefore, they achieve a Pareto-efficient resource allocation more effectively than any other distribution mechanism. These arguments are often cited to support the increasing use of capital markets instead of banks. They have also been used to help explain American productivity growth compared to Japan's or Europe's during the 1990s.

Empirics: The Rhythm and Real Deal, Holyfield

Although the Efficient-Markets Hypothesis generates powerful predictions, many of them are empirically false. Shiller (1997) observed that prices exhibit greater volatility than would be expected with rational expectations. Moreover, Cutler, et al (1989) showed that large price changes occur even without significant news. Indeed, newspapers often cite the market's mood to explain price movements. It is telling that reporters who describe markets on a daily basis anthropomorphize them. With all its mood swings, stock markets might be best described as a teething three-year-old, not a suave calculator of valuation shifts. Prices also exhibit temporal correlation so that volatility tends to cluster (Mandelbrot, 1963).

If prices follow a random walk then prices changes should be normally distributed about zero, whereby the most common changes would be small and large volatility would be increasingly less likely. However, volatility distributions show that extreme instances occur more frequently than normal distributions would predict. This greater probability of large-volatility events leads to price-change distributions with "fat tails" at their extremes and illustrates that price changes are not distributed according to a random walk (Farmer, 2000). In addition, if one assumes that price equals value, trading should occur only infrequently. Yet daily volume in foreign exchange markets exceeds $1 trillion, fifty times more than world GNP per day (Farmer, 2000). That's a whole bunch of new information.

Since I already understood the predictions of efficient-markets theory, it was surprising to learn that financial markets do not trade in equilibrium. Modern markets trade continuously, so at any one time there is always an excess supply or demand because, for example, all the sell-side bond traders rush to the bathroom after their chimichanga fiesta. The mechanism that markets use to store these unfilled orders is called a limit order book.

What about the price-equilibrating auctioneer then? Walras did not just dream up this trading mechanism after too much brie and champagne. In fact, his explanation of how markets reach equilibrium is based on the nineteenth-century Paris Bourse. In this market, trades occurred during a period of tatonnement in which an auctioneer gauged the demand curve of each agent. If the auctioneer could find a market-clearing price that satisfied both buyers and sellers, transactions would occur at given times during the day (Farmer, 2002). Although a few modern markets operate in this manner, the London gold fix trades twice a day and the London Metals Exchange four, virtually all modern markets have discontinued trading with an auctioneer. Today, most modern markets, including the Paris Bourse, transact using some derivation of the limit order book.

Relax- BUT JUST A LITTLE BIT!

Amazingly, the predictions of Walrasian theory cease to hold if even one assumption is relaxed. Rational-expectations theory represents one good candidate for abandonment. Given the quantity of market information and the diversity in human opinion, how could each investor reach the same stock valuation as everyone else? So, let's go wild, let's assume heterogeneous agents. If agents invest according to varied valuation assumptions, then a stock's price cannot equal its objective value- only known by God and maybe Warren Buffet. In fact, when a value investor takes a position, she must implicitedly believe that price and value are not congruent. If they were, she could not gain from correctly forecasting price movements toward value and would not invest in the first place.

Still, as long as investors change their positions until stocks hit their value, don't markets just take their stroll toward equilibrium as Walras says? Not with heterogeneous agents. Say a group of investors with similar stock valuations buys when a stock is at a low price compared to its perceived value. This would induce the price to rise, but as it increased, more investors would decide that no more profits were possible because the stock had hit its value. They would sell and drive the price away from value again. With heterogeneous agents, the stock's perceived value may even become a barrier, not a destination. Indeed, "Evidence from market data suggests that, while prices track values over the very long term, large deviations are the rule rather than the exception," (Farmer 2000, 16).

The accounting scandals that have rocked corporate America underline the fact that complete information is an empirical impossibility. Insider trading probably constitutes an unsolvable problem since violations are difficult to differentiate from normal transactions and even harder to prove in court. In addition, these cases require clear winners and losers, more likely in a bear market. During the go-go 90s, even though insider trading may have occurred, there were enough winnings to make everyone happy.

Abandoning perfect information leads to a concept developed by George Soros (1987) called "price reciprocity". Since there is always unknown information, an investor could logically believe that price changes reflect value shifts that they don't know about yet. Imagine if a stock's price declines, an investor might surmise that the stock's value had declined based on information unknown to him. He would act on this lower valuation by selling his shares and driving the price lower still. If the price changes are large enough, they might induce other investors to react similarly and cause a cascade that leaves the stock ravaged.

Price reciprocity and the positive feedback loops that it creates might help explain clustered volatility and investor herd-behavior often exhibited during financial crises. The Walrasian belief that price changes reflect fundamental value inadequately explains the crises that have swept from Bangkok to Buenos Aires in the past decade. Some of the IMF's policy recommendations, which are based on the conventional approach, arguably worsened the crises, most notably in Indonesia.

If Walrasian assumptions yield incorrect predictions, don't reflect empirical reality, and their abandonment could help explain some mysteries of modern financial markets, why do they persist as the cornerstones of economic theory?

Why do these assumptions persist?

I believe there are at least three reasons: the cultural, the political, and the theoretical. Culturally, these assumptions persist because the mathematical tools that economists have mastered are nearly useless once one abandons Walrasian assumptions. Without them constraining the sphere of agent action, prediction becomes nearly impossible. Differential equations cannot be solved for much more than a few iterations when heterogeneous agents interact in a non-linear environment. Still, only in the last two decades have computers existed with enough processing power to investigate how non-linear economies function.

Politically, Walrasian conclusions provided an answer to Marxist ideology and its emphasis on state intervention: Free markets, left to their own devices, allocate resources in a manner most beneficial to society. It is unsurprising that, just as the Red Menace extended its slimy tentacles around the globe, governments began employing economists as advisers. Yet, it was even dangerous for mainstream practitioners to deviate from their own doctrine. “A conservative alumnus of MIT warned the university’s president …that Paul Samuelson would jeopardize his scholarly reputation if he were allowed to publish his apologetics for the ‘mixed economy’,” (Samuelson and Norhaus, xxi).

On the theoretical level, these assumptions render unnecessary a theory describing out-of-equilibrium dynamics. Once at equilibrium, a market enters a timeless and asymptotic state that can persist indefinitely as the market’s past becomes irrelevant. The explanation for how it got there reads something like, ‘Demand increases, which raises price, and leads to an increase in supply and a new price that equilibrates supply and demand until the next shock.’ But this account reifies the process, ignoring the agency of market participants. Who actually raises prices in response to increased demand, why, how long does it take, how much does price move? All these questions can be disregarded with equilibrium, even though few economists have compelling answers. As Mas-Colell, et al (1995) note “economists are good…at recognizing a state of equilibrium but are poor at predicting precisely how an economy in disequilibrium will evolve.” This admission is especially worrying if one believes that economies rarely, if ever, achieve equilibrium.

The More Things Change the More They Stay the Same

To be fair, mainstream economics has moved on all fronts to confront these limitations. Indeed, most economists would probably acquiesce if asked whether they believed in some or most of the dominant paradigm. Especially in the last twenty years, movements like bounded rationality (Simon), other-regarding behavior (Gintis and Bowles), and behavioral economics (Kahneman and Tversky), among others, have moved the mainstream toward the systematic relaxation of Walrasian assumptions.

Even with these significant changes, the fundamental model remains remarkably stagnant. One could argue- and some have- that attacking these conventional assumptions is like battling straw men. Why bother, they ask because no one really believes them anyway. Sadly, even though many economists disagree with large portions of the theory, they continue to teach it in their undergraduate classes without hesitation. This is not simply an issue of ivory-tower inertia. One must consider that those taking economics courses grow up to become the politicians and business people who lead our nation. The standard paradigm forms the basis of their economic thinking and their collective decisions determine how our economy functions. Without replacing the canonical model this process will continue. I can attest to it myself after three years of undergraduate studies.

The Complexity Approach

The project that I did this summer takes a fundamentally different approach to economic research by viewing the economy as a complex system. The group, led by Doyne Farmer, aimed to produce the simplest possible model of how a financial market works but which still included its salient features. They threw out the conventional, a priori assumptions and created an agent-based computer simulation of the limit order book. This allowed them to investigate the model's dynamics through a wide scope of parameter values and add features to slowly increase its realism. One could view this as modeling from the bottom-up. In contrast, the Walrasian top-down approach begins its models with myriad assumptions that are later relaxed for empirical realism.

The model replaces agent optimization with agent heuristics (simple decisions rules based on local information) to determine behavior. Even without rationality, investors still exhibit behavioral regularities that can be simulated. Indeed, most investor behaviors can be separated into groups of technical traders, value investors, and market makers. Each strategy differs based on the information input used.

Technical traders use past prices to forecast future changes. One type of technical trader- trend followers- assumes that prices are inertial; therefore they buy if trends are up and sell if trends are down. On the other hand, contrarian technical traders believe the market will reverse itself. Therefore, when prices move higher they sell and when prices decline they buy to take advantage of the reversal. I added this type of trading strategy to the limit order book model. The reason I used this strategy will become clear presently.

In contrast, value investors use some objective measure of value, usually price to earnings ratios, dividends, or management quality, to make their investment decisions. Like contrarians, they tend to induce counter-cyclical price movements, while trend followers tend to reinforce recent trends. Market makers are simply liquidity providers that buy at the best sell price and sell at the best buy price, profiting from the spread (the difference between the two, see Figure 1) while making it narrower.

By building models from the ground up, the new dynamics that each investor strategy creates can be investigated alone or together with others. One agent might induce dynamics that another agent can exploit, or the competition of two agents could create a niche for a new tactic. Strategies can development competitive, predator-prey, or symbiotic relationships just like biological species. In effect, this model generates a market ecology, but with investor strategies instead of species and profits not food. The most successful strategies gain market power and propagate themselves, while the bad ones lose money and die out. In this way, the feedback between price dynamics and investor strategies causes constant market evolution. Indeed, Farmer (2002) models the interaction of investor strategies using versions of the Lotka-Volterra equations, which are standard for describing biological population dynamics.

This constant evolution tests claims of informational efficiency because as each new strategy develops, prices should become increasingly less patterned and "abnormal profits" more difficult to achieve. Moreover, since these models simulate the agents that move supply and demand, they can help investigate out-of-equilibrium market evolution and determine whether equilibrium is achieved or not.

Using the limit order book as the basis for simulation also sets this approach apart from Walrasian models that assume transactions occur in the same way in all markets. As we will see, this institutional regularity induces unexpected price dynamics itself. These types of simulations may provide clues as to how institutions affect specific markets and how policies should change to accommodate for these variations. As the former Soviet republics have followed divergent and mostly unsuccessful paths toward market economies, institutions have become an especially important field of analysis.

How a Limit Order Book Works

Limit order books permit continuous, double-auction trading, meaning that each book has a buy and sell side transacting simultaneously. There are essentially two types of orders that make up a book: a limit order and a market order. For the former, an investor indicates the number of shares she wants to buy or sell and the order is filled immediately. The price she gets is based on the best bid (buy) price or best ask (sell) price within the book. In contrast, an investor placing a limit order indicates the price at which he wants to buy or sell, but the order execution time is indefinite and if the order is not filled within a certain time period he can cancel it or it will expire (See figure 1).

In Figure 1 the length of each block represents the volume of shares being offered at each price. Imagine that an investor places a buy market order. The order is matched with the best available sell price (best ask). If the order volume is greater than the volume at the best ask, then the next best ask price is also used to fill the market order. Therefore, the investor would pay different prices to buy the shares he desires. By eating up volume on the sell side, the best ask price would increase and so would the midpoint price (best bid + best ask / 2). When referring to the model, price means the mid-point price. Since the best bid price remains unchanged, this buy market order also widens the spread.

The Model

Laszlo Gillemot programmed the limit order book model that I worked with all summer using the Objective C language. In his implementation, orders enter the book randomly. The model could also be seen as a market with one investor trading randomly. Specifically, limit orders enter the book in a Poisson process whereby orders are uniformly distributed among prices. Buy limit orders can fall anywhere below the best ask price and sell limit orders can fall anywhere above the best bid price. Market orders enter the book at a certain default rate (See Appendix 1 for model parameters).

Even though orders enter the market randomly, I was surprised to learn that the book generates patterned prices. Indeed, as Figure 2 shows, there is a negative autocorrelation in prices that persists for roughly one hundred time steps even averaged over one thousand runs. Notice that the blue line never reaches zero over the length of the graph. Negative autocorrelation means that when prices increase they will tend to decrease in the future and vice versa. These statistics show that even over one hundred time lags, prices still are more likely to change direction than continue their trend. As the graph shows, the predicted reversal gets weaker as time steps increase, yet the pattern remains for a significant period of time. My agent attempts to use this predicted price reversal to make profits. This reversal pattern also explains why he is a contrarian trader. The red line shows how much autocorrelation variability can exist for just one run.

Figure 2 also shows the sum of average autocorrelation, which is a measure of the total pattern in prices. Since the technical trader uses these patterns to predict price change and make profits, this sum measures the maximum profits my agent can generate from the random order flow model if he harvested the entire pattern.

Figure 3 illustrates why this negative price autocorrelation occurs. The red bid strength and blue ask strength represent another way to conceptualize limit order book volume. They show the average volume over time at each book price. Although limit orders enter the book in a uniform distribution, average book strength slopes down because the entrance of market orders eats up volume around the bid and the ask price. Since orders enter randomly, the book can look very different at any one time, but on average it conforms to this overall shape.

Now imagine that a buy market order enters and chews up a large portion of the ask strength volume (the black region). Then the best ask and mid-point price both would increase. The best ask price now holds a lot of volume, so it would take many random market orders to move the price up more. Yet, the best bid price holds a much smaller order volume. It is likely that enough sell market orders will enter the book to move the best bid price and shift the price back down. So, the manner in which the limit order book holds excess orders also happens to induce negative price autocorrelations.

Franky the Agent

On top of this random order flow model, I implemented my contrarian technical trading agent. The questions that I attempted to answer were: Can a technical trading strategy be profitable in this model, and if so at what parameter values? And, how does the entrance of a technical trader affect market dynamics like volatility and price autocorrelation? People name their dogs, their cars, even their furniture. Since I spent long hours slaving to create my trader, I decided to name him Franky…Franky the Agent.

Since Franky is a technical trader he uses past prices to make investment decisions. He calculates a signal from past prices using this exponential moving average equation:

signal = past returns / Tau + (1- 1 / Tau) * past signal

Past returns is calculated by the difference between price two time steps ago and price one time step ago. Past signal is just the signal value from the previous time step, which incorporates all the past signals before it. Tau represents a weighting variable between past returns and past signal. If Tau equals one, then Franky determines his signal and his investments based only on past returns. As Tau increases, Franky uses past signal more while calculating his present signal. So, as past signal is weighted more heavily, he takes into account the price changes from an increasingly longer past period.

To take advantage of that negative price autocorrelation, he must do the reverse of what his signal suggests. He calculates his desired position by the equation:

desired position = -(volume * signal)

In my implementation, Franky trades with volume equal to one, so that desired position = -signal. If price increases, he predicts it will decline in the future, so his signal increases and his desired position decreases. If he holds a negative position by selling shares and price does decrease, then he could buy back his position at the lower price and pocket the difference.

Franky calculates his signal and desired position each time step. His buying or selling at that time is determined by the equation:

position change = desired position – past position

Franky uses only market orders to reach his desired position. If he ends a simulation holding shares, he simply buys or sells enough so that his final position equals zero.

Figure 4 illustrates how Franky determines his signal and position changes based on prices. Notice that as prices drop before time 3000 (left panel), the signal (red line) also drops. Since price declines, Franky believes they will increase in the future. So his position change (green line) shoots up, as he buys shares for this profit opportunity. Also notice that after the signal plummets, it moves toward zero but at an exponentially declining rate. After the initial price change, both his signal and his desired position move toward zero. So, after buying a large amount of shares, his position change goes slightly negative until the signal reaches zero again.

Any investor will tell you that trades incur transaction costs and Franky is no exception. He buys shares at the best ask price and sells at the best bid. Consequently, he loses the spread on each trade. He can only hope to make profits if the price pattern is strong enough to overcome this spread loss.

First, I investigated Franky's profitability over different Tau values using this initial implementation. I generated 50 runs at each Tau value with 105,000 time steps and calculated the mean profits per Tau. Figure 5 shows the results: Franky's profits never move above zero, but he loses money at an exponentially decreasing rate as Tau increases. Since a high Tau means that he uses more price history to calculate his signal, the signal becomes weaker as Tau increases. A weaker signal means that Franky makes fewer trades and, consequently, loses less money. Still, his price-movement predictions never overcome his spread losses.

Since he loses the spread for each transaction, my next implementation added that Franky would only trade if the spread dropped below a certain value: a spread trigger. If he only traded with a small spread he might lose less on each transaction and get into the black. I tested different combinations of signal triggers and Tau values and calculated mean profits over 20 runs for each parameter duo. As Figure 6 shows, the spread trigger made no significant difference, as his profitability never rose above zero. This is an unexpected result and it may have been due to a program bug. I must investigate this further.

Next, I tried a signal trigger. A signal far from zero indicates that random orders have cut deeply into one side of the limit order book, making a price reversal more likely and probably larger. In this implementation, Franky only trades if his signal moves above or below a set value. Again, I swept through parameter values of Tau and the signal trigger and calculated mean profits for each combination. This addition had the effect of lowering Franky's transactions just like the spread trigger. As Figure 7 illustrates, I found a very small ridge of profits as Tau moves from 2.5 to 3.5 and the signal trigger moves from 0.03 to 0.04 (left panel, red line = Tau 2.5, blue line = Tau 3.0, green line = Tau 3.5). The right panel shows average profits, Tau, and signal trigger values in three-dimensional space. The region of profits shown in the left panel can be seen as Tau moves from 2.5 to 3.5 and average profits rise slightly above zero

Future Directions

There is, of course, much more to understand about Franky's impact on the limit order book. Initially, the most important task would be to fully investigate his region of profitability. If Franky's contrarian technical trading strategy cannot make consistent profits in this simple model, then it is unlikely to be successful in the harsh competition on Wall Street. Other than increasing the granularity of my parameter sweeps, I might find profits by implementing both a spread and a signal trigger together.

In the future, I would like to investigate Franky's strategy within the ecology of agents that I have outlined. This would also allow observation of how agent behavior and market dynamics co-evolve. Finally, Farmer's group is attempting to determine actual market parameters using data from the London Stock Exchange. Exploring Franky's behavior under these circumstances would provide at least some guidance to his performance in a real market.

As these models become more complex, their ability to test hypotheses and public policies increases. An interesting direction would be to create a society that invests in the stock market through institutional investors. A fictious economy could determine the stock market's general direction, but the economy's movements would be partially determined by a wealth effect from the stock market. I would like to explore the dynamics generated by feedback between the economy and the market. One could even feed in past data and gauge whether the model generated qualitative market behavior.

Conclusions

Coming into this summer I knew that the Walrasian economic paradigm did not coincide with my intuition of how an economy works nor, for that matter, empirical reality. I also believed that understanding the economy as a complex system might solve some of these problems. Yet, I only understood how the ideas of complexity could be applied to economics as an analogy. I had no sense of how this approach might be utilized in economic research or help explain a market's dynamics.

This project made concrete my intuitions and barebones ideas. It is now clear that nonlinearity constitutes a useful tool for economists to advance their knowledge. Agent-based modeling provides a powerful medium to explore how agents interact and evolve in complex systems. The limit order book simulation and Franky the agent represents an example of how this approach can produce tangible insight into the workings of financial markets. I even learned many of the tools necessary to do the nitty-gritty simulation that this research requires.

My hope is that the complexity approach to economics will generate a new theory that subsumes Walras. Instead of it being taught as the standard example, equilibrium could be presented as a special case whose assumptions are nearly never achieved in practice. Further research using complexity could generate greater understanding of and more effective solutions to our economic problems. This limit order book model is a small, but instructive example of this new paradigm's promise.

Bibliography:

Farmer, Doyne. 2000. Market force, ecology, and evolution. Submitted to Journal of Economic Behavior and Organization.

Farmer, Doyne. 2002. How market making supports technical trading. Santa Fe Institute working paper.

Katz, Michael, Rosen, Harvey. 1998. Microeconomics. Boston: McGraw-Hill, Inc.
Mas-Colell, Andreu, Whinston, Green. 1995. Microeconomic Theory. New York: Oxford University Press.

Samuelson, Paul, Nordhaus, William. 1995. Macroeconomics. Boston: McGraw-Hill, Inc.

Sharpe, W., G. Alexander, J. Bailey. 1995. Investments. New York: Prentice Hall, 106.

Shiller , Robert. 1997. Market Volatility. Cambridge, MA: MIT Press.

Soros, Geroge. 1987. The alchemy of finance. New York: John Wiley and Sons.