

Agent-based Financial Markets: Matching Stylized Facts with Style

Blake LeBaron *

Brandeis University

April 2004, Revised October 2005, forthcoming in *Post Walrasian Macroeconomics: Beyond the DSGE Model*, edited by D. Colander, Cambridge University Press, forthcoming 2006.

Abstract

Empirical facts from financial data pose some of the most difficult puzzles for equilibrium macroeconomic modeling. Features such as volatility, excess kurtosis, and conditional heteroscedasticity are not easily replicated by any single representative agent model. Most agent-based financial markets are able to match a good subset of these features quite easily. This paper will summarize some of the results from an agent-based model. It will be argued that agent-based approaches also make more sense economically than their representative agent competition. They will also be compared and contrasted with approaches coming from the behavioral finance perspective as well.

*International Business School, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a faculty research fellow at the National Bureau of Economic Research.

1 Introduction

The movements of financial markets, and their connections to the macro economy are one of the most difficult areas for traditional economic theories. This is true both from an empirical and a theoretical perspective. This paper will concentrate on the empirical puzzles from finance that demand new approaches, such as agent-based financial markets. It will be argued that even when traditional modeling approaches fit some subset of empirical features it comes with the cost of moving farther from economic believability and robustness. Agent-based approaches fit more features with frameworks that seem to make more intuitive sense for the functioning of real markets. Also, agent-based frameworks can be used as a testbed for drawing in behavioral results found in both experimental and micro level financial markets. This is crucial for understanding when and where behavioral quirks of economic actors will appear at the macro level.

Many of the most puzzling results from finance deal with problems of behavioral heterogeneity, and the dynamics of heterogeneity. The study of market heterogeneity as a kind of complicated dynamic state variable that needs to be modeled is probably one of the defining features of agent-based models. Empirical features such as trading volume are directly related to the amount of heterogeneity in the market, and demand models that can speak to this issue. Other empirical features are probably indirectly related. Large moves, excess kurtosis, and market crashes all probably stem from some type of strategy correlation that keeps the law of large numbers from functioning well across the market. These changing patterns can only be explored in a framework that allows agent strategies to adapt and adjust over time, and more importantly, to respond to features of the aggregate population around them.

This paper visits some of the empirical puzzles, and shows that many of these can be replicated in an agent-based model. It also compares and contrasts the Post Walrasian approaches of agent-based models with current research in behavioral finance. These two areas have some important overlaps, and also some important differences. It will be argued that agent-based approaches are necessary in understanding which behavioral features will have macro pricing consequences.

The next section goes over some of the empirical finance puzzles. Section 3 looks at some of the empirical results from an agent-based example model. Section 4 looks at some of the important connections between agent-based models and behavioral finance, and also emphasizes that the latter may be a critical tool in sorting out behavioral models. The final section concludes and discusses some policy implications.

2 Empirical Features

Financial time series are arguably some of the most interesting data collected on the functioning of economic relationships. We have relatively long and clean series of financial prices. These prices are generated from well organized financial markets that bring large numbers of investors, and trade goods which can be compared over time. They are also crucial in allocating investment funds, and therefore to the performance of the overall economy. Recently, market microstructure research has utilized data sets that often contain almost every trade on a given market. Although not available for very long periods, these detailed data afford an unprecedented view into the inner workings of large trading institutions. It is surprising that given the amount and quality of data we have on financial markets, many of their features remain somewhat of a frustrating mystery. Before looking into some of the results of agent-based financial markets it will be important to document some of the major puzzles of financial markets.

Volatility is the most obvious and probably the most important puzzle in finance. Why do financial prices and foreign exchange rates move around so much relative to other macro series both on a short term and long term basis? The difficulty of overall financial volatility was first demonstrated in Shiller (1981), and an update is in Shiller (2003).¹ The issue has been that it is difficult to find financial or macro economic fundamentals that move around enough to justify the large swings observed in financial markets. As a potential policy problem, and an issue for long range investors, this might be the most important puzzle faced by financial modelers. This excess volatility could lead to reductions in welfare for savers and potential misallocations of resources for investors. The basic task is therefore to build a market that can take relatively stable fundamentals, and amplify their volatility into a fluctuating price series.

There are several puzzles related to financial volatility that may be at least as interesting as the simple level of volatility. First, financial volatility is extremely persistent. The persistence of volatility in many financial markets has led to an entire industry of models, and is an area of intense interest both in academic and commercial areas. However, although there is a lot of empirical activity, the underlying microeconomic motives for volatility persistence are still not well understood. There are very few models which have even tackled this problem. This is probably due to the fact that in a homogeneous agent framework this is simply a very difficult problem. There also may be something very interesting about the nature of this phenomenon. Volatility appears to be very persistent. It has autocorrelations which decay slowly out to almost a year or

¹A related paper is Hansen & Singleton (1982), which uses a very different methodology, but comes to similar conclusions. A simple macro economic variable, consumption, which should be able to explain the movements of returns in an intertemporal representative agent framework, simply doesn't move around enough to account for changes in returns. It also doesn't move in a way that is correlated with returns either, and this is also part of the problem.

more. This is an indication of a long memory process. This may be a very interesting signature for return dynamics, and it sets a pretty high hurdle for any financial model to meet.

The second feature is excess kurtosis. Financial returns at relatively high frequencies (less than one month) are not normally distributed. There is not much of a strong theoretical reason that they need to be, but the hope has often been that some form of the central limit theorem should drive returns close to normality when aggregated over time. Recently, a new field, Econophysics, has appeared which stresses that returns also have additional structure that can be described using power laws.² The determination and testing of power laws remains a somewhat open area, and the set of processes that generate acceptable power law pictures is also not well understood.³ It is also important to realize that most persistent volatility processes generate excess kurtosis, so there is a connection between these two phenomena. The connection between large moves and endogenous correlations, or reductions in heterogeneity in the population makes these a crucial fact for agent-based modelers. They are also an important practical fact to deal with for finance professionals in the risk management business.

Another broad empirical puzzle that has captured an enormous amount of attention is the equity premium puzzle, which considers the returns on stocks versus bonds. In most representative agent optimizing models it is difficult to explain this return spread.⁴ This remains a critical puzzle, though many explanations have been put forth. It is important to realize that in many ways this can't be separated from the volatility puzzle mentioned previously. If there are mechanisms that magnify market volatility that we don't quite understand, then the premium on equity might be reasonable to justify this extra risk component.

Finally, trading volume is a major issue that can only be tackled from a heterogeneous agent perspective. Most traditional financial models remain completely silent on this issue. This is troubling since it often appears that trading volume is too large to be explained with standard models. Daily foreign exchange trading volume is often well over a trillion dollars, which represents a large fraction of total U.S. annual GDP moving through the market each day. Beyond this, trading volume is very persistent. It is another time series which also might be a long memory process. This is more difficult to determine since many trading volume series also have persistent time trends in them, but there is some early evidence that volume is also a long memory process (Lobato & Velasco (2000)). There is also interesting, and possibly related, evidence for long memory in internet packet traffic (Leland, Taqqu, Willinger & Wilson (1994)). It would be fascinating if the internet and trading volume mechanisms were similar.

²See Mantegna & Stanley (1999) for an introduction and examples.

³LeBaron (2001) gives some examples of this issue.

⁴This result goes back to Mehra & Prescott (1988), and was surveyed more recently in Kocherlakota (1996).

3 Empirical Examples

This section presents some short examples of an agent-based model replicating many of the empirical features mentioned so far. The basic model structure is taken from LeBaron (2002*a*), and LeBaron (2002*b*). Details are presented in these other papers, but a very brief sketch will be presented here.

The model is based on a single risky asset available in fixed supply, paying a dividend that follows a random walk growth process aligned with aggregate growth and volatility in U.S. dividends after World War II. It is simulated at weekly frequencies. There is a risk-free asset paying zero interest available in infinite supply. Agents in this model all maximize intertemporal constant relative risk aversion (CRRA) utility which is restricted to log preferences. This locks down the consumption as a constant fraction of agents' wealth and concentrates learning on determining optimal dynamic portfolio strategies.

Agents chose over a set of portfolio strategies that map current asset market information into a recommended portfolio fraction of wealth in the risky asset. This fraction can vary from zero to one since short selling and borrowing are not allowed. Information includes lagged returns, dividend price ratios, and several trend indicators. Agents must evaluate rules using past performance, and it is in this dimension where they are assumed to be heterogeneous. Agents use differing amounts of past information to evaluate rules. In other words, they have different memory lengths when it comes to evaluating strategies. Some agents use 30 years worth of data, while others might use only 6 months.⁵ In this way this model implements behavioral features. First, agents are clearly boundedly rational in that they do not attempt to determine the entire state space of the economy, which would be unwieldy if they attempted this. Also, they are assumed to have "small sample bias" since they don't all choose to use as much data as possible. If it actually is better to use longer sample sizes, then wealth will shift to the longer memory agents, and the shorter memory types will steadily control less and less wealth, and will steadily be evolved to zero.⁶ The set of trading rules is represented by a neural network, and is evolved over time using a genetic algorithm. The diverse set of strategies of the different agents are used to numerically form a market excess demand function. The equity is assumed to be in constant supply of one share, and the price is then determined by numerically clearing this market. It is a form of temporary Walrasian equilibrium.

Figure 1 compares the actual market price for a set of agents with many different memory lengths ranging from 6 months to 30 years. This will be referred to as the all memory case. The actual price is compared to a

⁵Formally, they try to find the rule that maximizes the log of their portfolio return following a dynamic strategy going into the past.

⁶This is also a type of constant gain learning as used in Sargent (1999). It is also related to a form of "learning stationarity," in that agents need to know whether the time series they are looking at are stationary or not.

theoretical price derived by assuming the market has converged to a homogeneous equilibrium where agents hold all equity, and consume the dividends paid. The actual price takes large swings from the equilibrium and exhibits what appear to be large crashes. The bottom panel of the figure examines the dividend price ratio, which varies significantly over time.

A similar picture corresponding to the S&P 500 is shown in figure 2. This picture is drawn from the Shiller annual data, and uses Shiller’s constant dividend discount price, P^* , as a comparison.⁷ Obviously, there are many choices for a “rational price” in the actual data, but P^* is a good first comparison. The figure shows swings around P^* that are similar to the all memory simulation. The frequency of large price swings is a little smaller than in the actual data, but the patterns are similar. The lower panel displays the dividend yield again. This also shows large swings, which are slightly less frequent than in the simulation.

In the last price figure agents are required to be of relatively long memory. They vary between 28 and 30 years in the data sets they use for decision making. This will be referred to as long memory. Figure 3 displays these results, which show a dramatic convergence to the equilibrium price, where the dividend price ratio is constant. This benchmark demonstrates the importance of short memory strategies, and also confirms that without these strategies the learning mechanism is able to find the homogeneous equilibrium.

Table 1 presents summary statistics for the generated weekly stock returns, and it compares the two different simulations with data from the S&P 500. The table shows that the heterogeneous memory framework amplifies volatility, and also generates leptokurtic returns the same as in the actual market data. The long memory returns are close to Gaussian in terms of kurtosis. The column labeled VaR presents the value-at-risk, or 1 percent quantile level on returns, and again shows that the all memory runs correspond well to actual return distributions in terms of this simple property of the left tail. Finally, the column labeled ARCH performs Engle’s test for the presence of conditional heteroskedasticity. This is not found for the long memory case, but is present for all the other series. The table also shows the impact of short memory traders on trading volume, which is larger in the all memory case.

This simple agent-based model has met several difficult hurdles presented in actual financial series. First, it has boosted volatility, and generates price series which go through persistent deviations from equilibrium prices, and violently crash. More precisely, the variance of returns is boosted to the level of actual variances, and returns show strong evidence for leptokurtosis and GARCH. This is far from the only model to show these features. Models such as Arifovic & Gencay (2000), Brock & Hommes (1998), Lux (1998), and Levy, Levy & Solomon (2000) display similar features.

⁷See Shiller (1981), and Shiller (2003) for a recent analysis.

Figure 4 shows the persistence of volatility in more detail where it is represented by the autocorrelations of absolute values of returns. For the all memory case, and the actual S&P returns, the figure shows a large amount of persistence going out for more than 52 weeks. This kind of slow decay in autocorrelations is a characteristic of a long memory process, and the agent-based model can replicate this feature well.⁸ The long memory only case generates no persistence.⁹

Figure 5 turns to the dynamics of trading volume. The actual level of trading volume in this model is not a reasonable number to compare to actual data, since volume must be affected by the fact that there is only one stock available. However, the dynamics of volume is equally as interesting as the dynamics of volatility. The autocorrelation patterns of trading volume in the various simulations are compared with trading volume from IBM from 1990-2000 in figure 5. The stock and the all memory simulation again show a very strong persistence pattern, which is similar to that presented for volatility. The persistence goes out for a long range, and the decay rate is very slow. Figure 6 looks at the connections between volume and volatility by displaying the cross correlations between trading volume and absolute values of returns.¹⁰ It displays the strong positive contemporaneous correlation which is clearly present in the IBM series as well. Also, there is an interesting asymmetry indicating a stronger impact from today's volatility into future trading volume than there is from today's trading volume into future volatility. This is another important feature replicated in both series.

Beyond short term dynamics, an important hurdle for agent-based markets is to replicate return properties at the long run. Table 2 displays simple summary statistics for the all memory market runs, and compares these with several examples drawn from a long horizon return series constructed by Robert Shiller.¹¹ The simulated markets are in a reasonable range for all of the annual return numbers. This includes the reduction in leptokurtosis that is observed in all long horizon return series. Another important measure is the Sharpe ratio, which looks at the excess return over the risk free rate divided by the return standard deviation. This level of the risk reward ratio is another important fact from the actual data that has been traditionally difficult to match.¹²

Traditional models have had difficulty in matching many of these facts. Also, most have only concentrated

⁸Another model that is designed specifically to generate long memory is Kirman & Teyssiere (2001).

⁹Evidence for long memory in volatility can be found in papers such as Ding, Granger & Engle (1993) and Baillie, Bollerslev & Mikkelsen (1996). It is not clear why volatility shows such persistence, or persistence at all for that matter.

¹⁰Connections between volatility and trading volume have been extensively documented. See papers such as ? and also Liesenfeld (2001).

¹¹These are used in Shiller (2000) and are available on his website.

¹²It is directly related to Hansen/Jagannathan bounds as shown in Hansen & Jagannathan (1991) and Cochrane & Hansen (1992). The generally large magnitude of estimated Sharpe ratios puts a lower bound on the volatility of risk measures in a standard representative agent asset pricing world.

on a subset of this set of features. Financial economists often prefer the realm of pricing anomalies such as the equity premium, while physicists prefer to look at distribution tails and power law related features. A recent paper that tries to look at returns at many horizons and performs detailed calibration is Campbell & Cochrane (1999). This is a representative agent framework that faithfully matches many of these empirical features. It does this by implementing a representative agent with habit persistent preferences. The form of the habit persistence is carefully designed to make the model work. It is not clear how robust the model might be to other forms of habit, and our knowledge of the habit structure is probably small. It is not clear how a diverse set of agents might aggregate up into a representative agent with habit preferences. Also, many of the features listed here are not touched in their framework. They do not deal with the extreme persistence of volatility, and its possible long memory behavior. Finally, since it is a representative agent model it cannot deal with any issues related to trading volume.

Incomplete markets and idiosyncratic risks are another major explanation for asset pricing anomalies. Mankiw (1986) is an early example of how an incomplete market with uninsurable risks hitting individuals might impact the aggregate risk level, and aggregate asset pricing. This has been explored in detail in Constantinides & Duffie (1996). Similar to some of the other explanations, this one also may not be all that robust. It requires that shocks be persistent, and they must be heteroskedastic with countercyclical conditional variance. This explanation still has not attempted to explain heteroskedasticity in overall price levels, or excess kurtosis, and it also does not attempt to explain trading volume.

This paper's main point is that most of these rational/equilibrium based models will have a difficult time replicating these features.¹³ In the next section attention turns to more behavioral approaches. These need to be treated with some care since many of their foundations have overlaps with agent-based approaches. The relationship between the two is interesting, and it is likely to be complementary since the tools of heterogeneous agent modeling will be necessary to help sort out some of the aspects of behavioral models.

4 Aggregation and Behavior

Agent-based models and Post Walrasian approaches both deviate from more traditional approaches based on complete rationality in a well understood equilibrium environment. Recently, the field of behavioral finance has weakened strict restrictions on rationality in economics.¹⁴ Dynamic models of interacting boundedly

¹³A recent survey that explores recent rational explanations is Constantinides (2002).

¹⁴These restrictions often include the entire mechanism of intertemporal expected utility maximization. Many surveys on behavioral finance have appeared including Mullainathan & Thaler (2000). Also, Thaler (1993) is a good early collection.

rational agents are behavioral in that the agents do not optimize their behavior and do follow rules of thumb. However, they often do not take on all the behavioral quirks suggested in the psychology literature. They form an important test bed for exploring behavior outside of the standard rationality paradigm.

There are several reasons individual micro irrationality might not appear in aggregate financial prices.¹⁵ The first is related to aggregation. Will features beyond rationality and optimality appear at the macro level? This is a difficult question, and much of the behavioral finance literature ignores this. Should one expect to see loss aversion, or regret at the aggregate level? ¹⁶ In many cases these modified preferences depend critically on a reference point where a stock is purchased defining gains and losses. If a heterogeneous population has loss aversion preferences, but also makes equity purchases at different prices, then what will be the impact on aggregate returns? This would need to be answered in an explicitly heterogeneous agent framework. It is possible that certain types of behavior will not appear in the aggregate, and explicit micro level modeling is necessary to make this determination. Ideas from complex systems, which often deal with how and when individual components in a system become correlated, suggest that this will not be an easy problem.

A second reason for aggregate behavior to appear more rational than individual psychology would suggest is that less than rational strategies may not survive in an evolutionary race.¹⁷ If certain strategies are less than rational they will eventually be eliminated by market selection forces. Evolution is a very powerful and important idea in both economics and finance, but whether it works as this conjecture states is a difficult problem. There are several dimensions to this. First, in markets that are not in an equilibrium, the determination of which strategy is “rational” is very difficult if not impossible since rationality must be defined relative to a current population of strategies. Second, the evolutionary dynamic in a multi-agent setting is very complex in that strategies are all evolving against each other, and there is no guarantee of a simple dynamic converging to a homogeneous equilibrium. The model used in the previous section is a good example of this. There are agents who, ex ante, appear rational and less than rational. The market allows wealth to shift between these different types, and therefore select for the “best” strategy. This dynamic doesn’t settle down, and the short memory agents are not driven out of the market. Therefore, they demonstrate a type of behavior that withstands two restrictions. First they are observable at the aggregate level, and also they are not eliminated by evolution.

¹⁵Some of these are mentioned in recent critiques of the behavioral finance literature as in Rubinstein (2001).

¹⁶A example of a representative agent model of this is in Benartzi & Thaler (1996) or more recently, Barberis, Huang & Santos (2001).

¹⁷This idea goes back at least as far as Alchian (1950) and Friedman (1953).

A final restriction on behavior concerns institutions that can often play a key roll in coordinating behavior.¹⁸ Financial markets are constrained by their institutional structure, and are not free from this argument. A good example of this is Gode & Sunder (1993), which shows that markets with nearly random behavior can still be very efficient in the sense of making trades. Once again the use of an agent-based approach can be important in analyzing behavior when the outcome of institutions in a heterogeneous world is not completely obvious.¹⁹

5 Conclusions

Agent-based financial models can be fit to a large range of empirical features from financial markets. Some of these facts have been tackled by more traditional models. However, these models are fine tuned to carefully fit the facts, and it is not clear how reasonable their preference structures are for a representative consumer. Beyond standard setups, behavioral finance may offer some promise for understanding some of the empirical puzzles of modern finance. However, these approaches will need the addition of an agent-based, or Post Walrasian, approach to markets to make sense of which aspects of individual psychology remain relevant at the macro level.

At the moment agent-based markets can match many features of finance in a fairly “stylized manner.” They do this by setting forth agents who in some ways are quite simple. Trying to do the best they can in a complex world, and working hard to adjust and adapt and maximize relatively simple objectives. These approaches are still new, often involving computational tools that are still not well understood by agent-based modelers and the economics profession. Also, the construction of heterogeneous economics models still remains uncharted territory. This should not be a major concern since the field and its tools are still very young.²⁰

There are many policy questions that can be answered by well constructed agent-based models. As we move toward models taking greater license with rationality, the importance of many policy choices, and quantitative modeling will increase. In financial markets questions about institutions and trading become more important in a Post Walrasian setting where institutions can have a profound impact on outcomes, and restrictions that might appear inefficient in a market clearing setting can be beneficial in these out

¹⁸Colander (1996) comments on the importance of institutions in a Post Walrasian setting. Also, cognitive psychologists such as Clark (1997) stress the importance of institutions in guiding behavior.

¹⁹Gode & Sunder (2004) is an example of this, using nonbinding price controls which should have no impact in an equilibrium price setting world.

²⁰It might be compared to the early stages of experimental economics.

of equilibrium settings.²¹ Applied economists and consultants have already entered this area in advising several institutions on how to deal with financial trading. An example is NASDAQ's questions about decimal trading.²² More questions like this will certainly follow.

While there are problems that need Post Walrasian approaches in all of economics, the problems in finance have some of the most pressing empirical questions. The data in these markets appear extremely far from the usual economic worlds of stability and equilibrium. Agent-based models make more progress than other frameworks in explaining these features due to that fact that at their core is a world of people who process information differently, and try hard to continually adjust and adapt their behavior over time. This market may never reach anything that looks like an equilibrium efficient market, but it is in a continual struggle toward this. The range of facts these models explain, and the robustness of their explanations to different structures and parameters, is impressive. At the moment, no other models can capture this many facts with this kind of simplicity and style.

²¹An example might be trading halts which would seem to be a kind of constraint on agents' trading abilities.

²²See Meyer & Davis (2003) for examples.

References

- Alchian, A. (1950), 'Uncertainty, evolution, and economic theory', *Journal of Political Economy* **58**, 211–221.
- Arifovic, J. & Gencay, R. (2000), 'Statistical properties of genetic learning in a model of exchange rate', *Journal of Economic Dynamics and Control* **24**, 981–1005.
- Baillie, R. T., Bollerslev, T. & Mikkelsen, H.-O. (1996), 'Fractionally integrated generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **74**, 3–30.
- Barberis, N., Huang, M. & Santos, T. (2001), 'Prospect theory and asset prices', *Quarterly Journal of Economics* **116**, 1–53.
- Benartzi, S. & Thaler, R. (1996), 'Myopic loss aversion and the equity premium puzzle', *Quarterly Journal of Economics* **110**, 73–92.
- Brock, W. A. & Hommes, C. H. (1998), 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model', *Journal of Economic Dynamics and Control* **22**(8-9), 1235–1274.
- Campbell, J. Y. & Cochrane, J. H. (1999), 'By force of habit: A consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy* **107**, 205–251.
- Clark, A. (1997), *Being There: Putting Brain, Body and World Together Again*, MIT Press, Cambridge, MA.
- Cochrane, J. H. & Hansen, L. P. (1992), 'Asset pricing explorations for macroeconomists', *NBER Macroeconomics Annual 1992* **7**, 115–164.
- Colander, D. (1996), Overview, in 'Beyond Microfoundations: Post Walrasian Macroeconomics', Cambridge University Press, Cambridge, UK, pp. 1–20.
- Constantinides, G. M. & Duffie, D. (1996), 'Asset pricing with heterogeneous consumers', *Journal of Political Economy* **104**, 219–240.
- Constantinides, G. M. (2002), 'Rational asset prices', *Journal of Finance* **57**, 1567–1591.
- Ding, Z., Granger, C. & Engle, R. F. (1993), 'A long memory property of stock market returns and a new model', *Journal of Empirical Finance* **1**, 83–106.

- Friedman, M. (1953), The case for flexible exchange rates, in 'Essays in positive economics', University of Chicago Press, Chicago, IL.
- Gallant, A. R., Rossi, P. E. & Tauchen, G. (1992), 'Stock prices and volume', *The Review of Financial Studies* **5**, 199–242.
- Gallant, A. R., Rossi, P. E. & Tauchen, G. (1993), 'Nonlinear dynamic structures', *Econometrica* **61**, 871–908.
- Gode, D. K. & Sunder, S. (1993), 'Allocative efficiency of markets with zero intelligence traders', *Journal of Political Economy* **101**, 119–37.
- Gode, D. K. & Sunder, S. (2004), 'Double auction dynamics: structural effects of non-binding price controls', *Journal of economic dynamics and control* **28**, 1707–1731.
- Hansen, L. & Jagannathan, R. (1991), 'Implications of security market data for models of dynamic economies', *Journal of Political Economy* **99**, 225–262.
- Hansen, L. & Singleton, K. (1982), 'Generalized instrumental variables estimation of nonlinear rational expectations models', *Econometrica* **50**, 267–268.
- Kirman, A. P. & Teyssiere, G. (2001), 'Microeconomic models for long-memory in the volatility of financial time series', *Studies in Nonlinear Dynamics and Econometrics* **5**, 281–302.
- Kocherlakota, N. (1996), 'The equity premium: It's still a puzzle', *Journal of Economic Literature* **34**(1), 42–71.
- LeBaron, B. (2001), 'Stochastic volatility as a simple generator of apparent financial power laws and long memory', *Quantitative Finance* **1**, 621–631.
- LeBaron, B. (2002a), Calibrating an agent-based financial market, Technical report, International Business School, Brandeis University, Waltham, MA.
- LeBaron, B. (2002b), 'Short-memory traders and their impact on group learning in financial markets', *Proceedings of the National Academy of Science: Colloquium* **99**(Supplement 3), 7201–7206.
- Leland, W. E., Taqqu, M. S., Willinger, W. & Wilson, D. V. (1994), 'On the self-similar nature of Ethernet traffic', *IEEE/ACM Transactions on Networking* **2**, 1–14.

- Levy, M., Levy, H. & Solomon, S. (2000), *Microscopic Simulation of Financial Markets*, Academic Press, New York, NY.
- Liesenfeld, R. (2001), 'A generalized bivariate mixture model for stock price volatility and trading volume', *Journal of Econometrics* **104**, 141–178.
- Lobato, I. & Velasco, C. (2000), 'Long memory in stock-market trading volume', *Journal of Business and Economic Statistics* **18**, 410–426.
- Lux, T. (1998), 'The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions', *Journal of Economic Behavior and Organization* **33**, 143–165.
- Mankiw, N. G. (1986), 'The equity premium and the concentration of aggregate shocks', *Journal of Financial Economics* **17**, 211–219.
- Mantegna, R. N. & Stanley, H. E. (1999), *An introduction to econophysics: Correlations and complexity in Finance*, Cambridge University Press, Cambridge, UK.
- Mehra, R. & Prescott, E. C. (1988), 'The equity risk premium: A solution?', *Journal of Monetary Economics* **22**, 133–136.
- Meyer, C. & Davis, S. (2003), *It's Alive: The coming convergence of Information, Biology, and Business*, Crown Business.
- Mullainathan, S. & Thaler, R. H. (2000), Behavioral economics, Technical Report 7948, National Bureau of Economic Research.
- Rubinstein, M. (2001), 'Rational markets: Yes or no? The affirmative case', *Financial Analyst Journal* **17**, 15–29.
- Sargent, T. (1999), *The Conquest of American Inflation*, Princeton University Press, Princeton, NJ.
- Shiller, R. (1981), 'Do stock prices move too much to be justified by subsequent changes in dividends?', *American Economic Review* **71**, 421–36.
- Shiller, R. J. (2000), *Irrational Exuberance*, Princeton University Press, Princeton, NJ.
- Shiller, R. J. (2003), 'From efficient market theory to behavioral finance', *Journal of Economic Perspectives* **17**, 83–104.

Thaler, R. H. (1993), *Advances in Behavioral Finance*, Russel Sage Foundation, New York, NY.

Table 1: *Weekly Return Summary Statistics*

	Mean (%)	Std (%)	Skew	Kurtosis	1% VaR (%)	ARCH(5)	Volume
All Memory	0.114	2.51	-0.567	10.23	-7.47	315.3	0.065
Long Memory	0.083	0.75	-0.012	3.01	-1.69	2.7	$0.053x10^{-8}$
S&P (47-2000)	0.175	1.88	-0.380	5.92	-4.69	211.1	
S&P (28-2000)	0.140	2.56	-0.214	11.68	-7.38	738.6	

Summary statistics for weekly returns. Simulation returns include dividends. S&P returns are nominal without dividends. 1% VaR is the Value-at-Risk at the one percent level, or the 1% quantile of the return distribution. ARCH(5) is the Engle test for ARCH using an autoregression of squared returns on 5 lags. It is distributed asymptotically χ^2_5 with a 1% critical value of 15.1

Table 2: *Annual Return Summary Statistics*

	Excess Return (%)	Return Std (%)	Sharpe Ratio	Kurtosis
All Memory	6.8	20.5	0.33	3.49
S&P 47-01	7.1	15.3	0.47	2.86
S&P 28-01	6.9	19.4	0.36	3.25
S&P 1871-2001	5.8	17.9	0.32	3.21

Annual summary statistics. All returns include dividends and are excess of a T-bill. Simulation returns are compounded over 52 weeks, and are non-overlapping.




Figure 1: Price, equilibrium price, and d/p ratio for all memory case




Figure 2: Price, Shiller P*, and d/p ratio for S&P




Figure 3: Price, equilibrium price, and d/p ratio for long memory case




Figure 4: Absolute Return Autocorrelations: All Memory, Long Memory, S&P




Figure 5: Trading Volume Autocorrelations




Figure 6: Absolute return/volume cross correlations