Financial Market Efficiency in a Coevolutionary Environment

Blake LeBaron *

Brandeis University

May 2001

Forthcoming: *Agent 2000 - The Simulation of Social Agents: Architectures and Institutions*, Argonne National Laboratory

Abstract

What does the evolutionary interaction between different types of investors look like? This paper discusses research trying to understand the evolutionary dynamics between agents using differing lengths of past data to make decisions on portfolio choice. Computer simulations of a simple agent based model show that agents taking a long perspective on past data have a difficult time dominating shorter perspective agents. The resulting dynamics replicate many features of actual markets. Furthermore, strategies become more homogeneous near sharp price declines suggesting a liquidity based explanation for market crashes and excess volatility.

*Graduate School of International Economics and Finance, 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, and an external faculty member at the Santa Fe Institute.

1 Introduction

Traditional equilibrium models for financial markets often rely on the process of evolution either explicitly or implicitly. These models generally assume that asset prices are the result of market participants holding common rational beliefs about the behavior of economic variables in the future, and acting in a well prescribed fashion on these beliefs. These economic worlds are simple, tractable, and generally at odds with the empirical evidence. Observed financial markets often appear too volatile, and too predictable to be explained as the outcome of well learned rational beliefs and strategies. Furthermore, the existence of large trading volume in financial markets adds important questions concerning heterogeneity across participants, since if all people agreed on asset valuations trade would be unnecessary in most situations. This paper uses an agent based model of a simple financial market to explore the evolutionary aspects of market dynamics with the goal of understanding the barriers to market efficiency which cannot be eliminated through evolution and learning alone.

Most modern analysis of financial markets includes two crucial assumptions. Markets are populated with rational agents, and they are in some kind of stationary equilibrium. Together these two assumptions yield tractable, testable restrictions for well crafted theories. The first assumption can be relaxed somewhat in that irrational players may appear at times, but their suboptimal strategies will be driven out of the market.[1] Weakening the second assumption causes more difficulty since it is closely linked to the rationality assumption. Out of equilibrium it becomes difficult to judge rational versus irrational strategies as the economic landscape is in a continuing state of change. The biological term for this is coevolution. Strategies evolve against this current set of strategies in the population, and not some well defined fixed fitness norm.[2] In such a world rationality can only be judged relative to the current population, and not some well defined fixed target that players should want to attain. It would be convenient to argue that these out of equilibrium dynamics can be ignored. However, this leaves open the critical question of how markets reach equilibrium in the first place.

In order to analyze out of equilibrium dynamics, and convergence properties, markets will be populated with boundedly rational, learning agents.[3] These are relatively simple agents trading and learning about

[1] Friedman (1953) provides the most often cited arguments for an evolutionary foundation for assuming rational behavior. In his case for flexible exchange rates he specifically comments on how less than rational speculators will be driven out of the market. Recent research on noise trading, (DeLong, Shleifer, Summers & Waldmann 1991), and evolution, (Blume & Easley 1990), has begun to suggest some flaws in the evolutionary argument for rationality in a financial setting.

[2] A good biological example is to think about evolution against predators versus evolution against climate. In the later case one is probably safe in assuming a fixed fitness landscape, but in the former this landscape is an ever changing target.

[3] The concept of boundedly rational agents was introduced by Simon (1969). Recent applications in macroeconomics are summarized in Sargent (1993) and Sargent (1999). Often the argument for bounded rationality rests on actual bounds on computing power in the brain. However, bounded rationality might be argued for in terms of robustness. In a complex financial

price dynamics as they go along. This facilitates the analysis of overall market dynamics and convergence properties when agents are faced with the same situations seen by ordinary people. It is important to realize that boundedly rational does not necessarily translate into stupid agents. They are often faced with situations where being completely rational may be computationally intractable, involving the beliefs of all the other market participants along with their dynamic decision making processes. The only option available is to follow simple rules of thumb which are empirically tested, and adjusted over time through learning.

This paper focuses specifically on one type of rationality, the appropriate use of past information. If financial time series were completely stationary, then more historical data would always yield better investment decisions. However, in practice it appears that many market participants chose to ignore some past information to focus on the present. Recent arguments about a "new economy" are a good example of this. This behavior might indeed be rational if markets have changed, yielding the historical data irrelevant and giving those following it the survival chances of dinosaurs.

An evolutionary struggle will be explored between short and long horizon investors in an attempt to assess when and if the long horizon types will evolutionarily dominate the market. The setting will be one with a completely stationary dividend process. In such a world it might seem obvious that the long horizon investors should dominate, but this is not necessarily the case. Prices move endogenously according to the traders' strategies, and can even move in such a way as to enhance the strategies of short term traders. It is also important to realize that in such a market it is not clear what is a rational or irrational strategy without the guidance of a market safely in equilibrium. Given a turbulent market of short run investors it may be individually rational to become one of them, as opposed to taking the more difficult path of sticking to a long run perspective.

2 Market Description

The market simulations used here are part of the class of economic models referred to as "agent based". Models of this type consist of large numbers of interacting agents each acting independently of the others often with active learning and adaptation.[4] Agent based markets share many features: many interacting

world all strategies will be incomplete in some aspect, so it can be argued that simpler strategies may do better in terms of avoiding some really big mistakes. This is a little like arguing that you don't want to get "too smart for your own good."

[4]Examples of this include Cont & Bouchaud (2000), Epstein & Axtell (1997), Palmer, Arthur, Holland, LeBaron & Tayler (1994), Arthur, Holland, LeBaron, Palmer & Tayler (1997), Kim & Markowitz (1989), Levy, Levy & Solomon (1994), Lux (1997), Tay & Linn (2001). Also, the website maintained by Leigh Tesfatsion at www.econ.iastate.edu/tesfatsi/ace.htm is an important source for agent based research in economics. Finally, the site at www.brandeis.edu/ blebaron/acf summarizes agent based research in finance, and a survey of some of the early research can be found in LeBaron (2000). Some commentary on the construction of agent based models is given in LeBaron (2001*a*).

individuals, evolutionary dynamics, learning, and bounded rationality. However, the key distinguishing feature is that heterogeneity itself is endogenous. Markets can move through periods that support a diverse population of beliefs, and others where these beliefs and strategies might collapse down to a very small set.

The market is a very simple one with a single equity like security paying a random dividend each period, and available in a fixed supply of one. This dividend follows a stochastic growth process which is calibrated to aggregate dividend series for the United States. There is a risk free asset which is available in infinite supply paying a constant real interest rate of one percent per year. Portfolios are rebalanced, and trades are made at a monthly frequency. Also, prices are determined, and dividends are paid each month, which can be thought of as the basic unit of time in the market. Therefore, this is more of an experiment concerned with longer term macroeconomic behavior as opposed to the minute by minute dynamics of day trading.

The basic actors are a set of 1000 agents. These agents adjust their portfolios and trade independently. They have well defined objectives in terms of optimal portfolio allocations.[5] However, they differ in one key respect, they have different views about how much past data is relevant in making their decisions. Some may take a long horizon perspective using an equivalent of the past 20 years of data, while others view that only the past year or two of data is important. For them previous data has become irrelevant in the investment decision making process. As trades and time go by, the agents accumulate wealth, and consume. Evolution takes place by eliminating agents with the lowest amount of wealth, and replacing them with new ones. In this way a "survival of the fittest" dynamic is imposed on the population of trading agents. As mentioned previously, it is important to understand how much pressure this puts on the market to move to a homogeneous rational outcome.

An important piece of the market is given by the trading strategies themselves. These can either be thought of as rules of thumb followed by the investors, or even as institutional managers who dynamically adjust their clients' portfolios. Since there are only two assets in the market these strategies can only be market timing strategies if they deviate at all from simple buy and hold strategies. The strategies convert a subset of current market information into a recommended portfolio allocation. The allocation gives a fraction of savings to put in the risky asset.[6] The market information includes: past returns, dividend yields, and two moving average technical indicators. The rules can be built off this information in any combination.[7] It

[5]The agents have logarithmic preferences over expected future consumption. Their time rate of discount is set to 0.95 per year. This is a well understood optimization problem in economics and finance, (Merton 1969) and (Samuelson 1969). It yields a consumption value which is a constant fraction of wealth, and an investment strategy which should maximize expected log returns.

[6]Before making asset allocation decisions agents consume a certain fraction of wealth, leaving the rest for investment. Also, they are not allowed to borrow or to sell shares short.

[7]The actual structure of the rules is given by an artificial neural network and is detailed in LeBaron (forthcoming 2001*b*).

is also important to realize that rules can ignore any piece of information too. In many ways the ability to ignore superfluous information is part of what is being tested.

Two further issues related to trading strategies remain. First, agents must decide which rules to use. At any given time there is a set of 250 active rules (investment advisors). Agents chose the rule which has performed the best in terms of their objective function. In doing so they make their decision based on past data using only what they feel is relevant. In other words, if the agent believes that the only the past 2 years of data are important (a relatively short horizon type) this is the range over which available strategies will be evaluated.[8]

The second crucial issue is how rules learn and adapt over time. If an investment advisor has at least one agent signed up for its services it will continue to exist with no change in its interpretation of market information into a dynamic strategy. If the advisor finds itself with no customers it will be eliminated, and replaced with a new advisor. The new advisor is created from the current population of active advisors using a genetic algorithm. This gives an interesting evolutionary dynamic to trading strategies. Those that are being used survive, and those that aren't are eliminated. The genetic algorithm tries to bring useful features of the current active strategies in to future ones. Success is determined purely on whether anyone is using a given strategy.[9]

Trading takes place each period. Agents all enter the market equipped with a chosen rule, and their current portfolio positions. This gives a well defined function for shares as a function of any given market price. Therefore, in principle the market could be cleared by a Walrasian auctioneer operating each period. This is essentially what is done. A numerical procedure is used to find a price which sets the demand for shares of the risky asset equal to the fixed supply in the market of one share.[10]

These provide a flexible functional form for turning the data inputs into actual portfolio weights.

[8]To introduce further heterogeneity in trader behavior it is assumed that traders only evaluate a small number of rules each period. They do not do a complete search over all possible advisors. Specifically, their current advisor is compared to one chosen randomly from the upper half of the advisor distribution measured using the agent's own view of how much history to use.

[9]It would be difficult to use any other fitness measure such as expected returns since agents don't view these from a common perspective. They are estimating returns over differing time horizons.

[10]A more realistic market might consider the actual market microstructure. However, since this market will be viewed as a fairly long range (monthly) pricing series, the temporary equilibrium assumption does not seem unreasonable.

3 Computational Experiments

3.1 Benchmark Runs

The following sections provide example runs of the market. In all cases the dividend series follows a random walk which is roughly calibrated to postwar U.S. aggregate dividends with an annual growth rate of 2 percent, and a annual standard deviation of 6 percent.[11] The risk free rate of interest is fixed to a constant 1 percent per year. All these rates are adjusted to monthly frequency which is the benchmark time horizon for trading and dividend/interest payments.

The key variable of interest in these experiments is the horizon length of the agents. This represents the distance they look into the past. Two different experiments will be considered. The first, referred to as *all horizon* uses agents drawn randomly from 5 to 250 months in length. This allows for a diverse population with many different investment horizons competing against eachother. The second experiment, referred to as *long horizon* loads the market with a set of agents using a relatively long time horizon. In this case agents are drawn from a distribution between 220 and 250 months. This loads the market with only long horizon investors. The objective is to see if this group performs differently from the diverse investor horizon case.

The market is run for 10000 periods which in calibrated time is actually over 800 years. Figure 1 shows a run for the first 3000 periods for a set of all horizon agents. The figure displays both price and volume. The figure shows three different phases for the market. In the first, the market is slowly adjusting, and the price is catching up to get on the correct exponential growth path. In the second phase the market appears almost cyclical in its fluctuations about the constant growth trend. There appear to be smooth cycles which don't look very reminiscant of actual markets. In the final phase, after period 1500, the market begins to look more normal with some run ups in the price followed by some sudden crashes. Trading volume also increases during this later period too.

Figure 2 presents the same information for the long horizon investors. Here we see a very different picture. After an initial adjustment phase the price maintains a relatively steady growth rate, and trading volume drops to near zero.[12] This picture shows a glimpse of a market more closely resembling a traditional efficient market equilibrium. Further results will reinforce these early pictures.

In figure 3 a detailed picture of the two price series is compared. This figure displays prices taken from the final 1000 periods (roughly 80 years) of a 10000 period run. This is done to capture behavior after

[11] See Campbell (1999) for a summary of values for several different countries.

[12] Investors enter the market with no innate knowledge of how it works, or how prices move with the fundamental, so it is sensible that some learning must take place for a short time.

all agents have settled down in their learning periods. The first two figures demonstrated that during the early periods behavior might not reflect the final outcomes in the different market experiments. This picture very clearly repeats the message of the early figures. The top panel, which corresponds to the case where traders are using diverse horizon lengths, shows a market price that occasionally takes some large swings often ending in dramatic crashes. In the lower panel, where traders are only long horizon, the picture shows a much smoother price dynamic with fewer large moves. The price still does indeed move, which reflects the fact that the fundamental dividend series is a somewhat volatile random walk. However, the amplification of volatility from the lower to the upper panel is clear.

3.2 Returns, Volume, and Volatility

Figure 4 compares returns across the two different computer experiments. The plot shows monthly returns inclusive of dividends. The return comparisons in figure 4 confirm the earlier plots showing a much more variable return in the all horizon case. Visually, returns also appear to exhibit some very large moves with a few months yielding returns over 50 percent. Also, there appears to be some clumping in that periods of large movements are grouped together. None of these features are present in the lower panel which corresponds to the long horizon case. The computer generated market yields a fairly homogenous set of returns with few if any large moves.

Figure 5 presents more evidence on the return distributions. Here, return histograms are compared with normal distributions. In the bottom panel it is clear that the normal distribution provides a relatively good fit to the computer generated returns. In the upper panel the unusual aspects of the large moves in the all horizon case become clear. The histogram is more peaked, and contains several very unusual observations which are well outside the normal range. This replicates the sorts of distributions which are often observed in actual markets.

Trading volume for the final 1000 periods is shown in figure 6. This plot shows a very strong distinction between the two cases with very large and fluctuating volume in the all horizon case, and nearly zero volume in the long horizon case. Obviously, trading volume is an important part of actual financial markets. The lower panel is merely reflecting the fact that the set of agents is in close agreement for asset valuations. In this case they have no interest in trading with each other while in the all horizon case differences of opinion and the desire to trade do not disappear.

Table 1 summarizes some of the results on the equity returns and trading volume. The mean return in

Table 1: *Excess Return Summary Statistics*

	Mean	Std	Skewness	Kurtosis	Volume
All Horizon	10.6	27.1	0.3	32.6	8.7
Long Horizon	6.2	10.4	0.1	3.2	1.3
S&P	8.9	19.6	0.5	12.9	[15,78]

Summary statistics: Mean and Std. are the annualized mean and standard deviation of the returns series inclusive of dividends. Skewness and kurtosis are estimated at the monthly horizon. Values for the S&P are the total return less the 30 day T-bill rate monthly from Jan 1926 through June 1998. Trading volume reflects the percentage turnover at an annual rate. The value corresponding to the line S&P is the range of NYSE reported values from 1958 through 1999. It is taken from *NYSE Fact Book 1999* (2000).

the all horizon case is larger than the long horizon case. The standard deviation shows a large increase in moving from the long horizon to the all horizon case. Neither case demonstrates any significant skewness in returns. The most dramatic difference appears in the kurtosis estimates. These measure the thick tailed aspects of the return distribution which have been displayed in figure 5. For a normally distributed random variable kurtosis is 3. For the long horizon case it is very close to 3 with an estimate of 3.2. The all horizon case shows clear evidence for leptokurtosis or "fat tails" with an estimated kurtosis of 32. The volume numbers report turnover at an annual rate. The all horizon case displays trading volume that is nearly 7 times the value for the long horizon case. Comparison numbers for the S&P are also presented. On means and standard deviations the S&P values are between the all horizon and long horizon cases. The all horizon market is actually generating more volatility, and large moves than in actual monthly data. In terms of trading volume the NYSE clearly generates more turnover per year than the market simulations.[13]

Returns and trading volume in actual markets display several interesting dynamic features that have been hinted at in some of the earlier figures. First, returns are close to uncorrelated. Second, volatility, or the absolute value of returns, is positively correlated. In other words, large moves in either direction tend to follow large moves. Finally, trading volume is also positively correlated. Figure 7 displays the autocorrelations for returns, absolute returns, and trading volume for the long horizon case. As in real financial data, returns show little or no correlations, positive or negative. However, both volume and volatility display strong positive correlations going out several periods as would be the case with actual financial time series. Figure 8 displays the autocorrelations for the market populated with long horizon investors only. This displays a picture quite different from actual markets. Returns show a small amount of negative correlation at one lag, and then zero after that. The volatility and volume series show only negligible correlation compared

[13]This number is presented more for information purposes. This two asset market is difficult to compare with the NYSE in terms of trading volume since the latter obviously has many more opportunities for trade.

with those from the all horizon experiment. These figures clearly show a very different dynamic in the two different cases with the all horizon case showing a picture that more accurately reflects real markets.

3.3 Dividend Yields

Another feature of real markets is that they appear to deviate from accepted fundamental models of valuation yielding some predictability from classical ratios such as price/earnings and dividend yields.[14] In these simple simulated markets this a very easy and direct experiment as to how well these markets are doing in terms of reflecting fundamental valuation. Since the equity asset only reflects a stochastic dividend stream the dividend yield is the ratio of choice for valuation. In a constant growth situation, both the price and dividend will be growing at the same rate, and the ratio should be constant. Figure 9 compares the dividend yields in the two cases. When investors are long horizon in nature, the dividend yield is nearly flat with only small variation around a value near 5 percent.[15] The situation is very different from the all horizon case. It is clear that the dividend yield takes some very wide swings, going as high as 12 and as low as 2 percent. It is far from a stable series. This compares much more favorably to dividend yields and price/earnings ratios in actual financial series.

3.4 Crash Dynamics

The last two figures present an initial picture of what may be behind the large price changes in this model.[16] They take a snapshot of a short time series of prices including some large crash periods and compare these with two other related series.

Two measures will be used which try to assess some aspect of trader heterogeneity. As mentioned earlier a key aspect of agent based models is that the actual level of heterogeneity in the market is endogenous. It is possible this may be a precursor to market instability. First, the most obvious magnitude to check around crashes is trading volume. In figure 10 several crashes are plotted along with the trading volume series. There is a weak indication that trading volume increases greatly after crashes, but it doesn't appear to show a very strong pattern before any of the large price drops in the figure so it would be difficult to blame crashes on trading volume.

In figure 11 the same price series is plotted against another measure of agent activity, rule dispersion.

[14] See for example Campbell & Shiller (1988), and for a recent commentary on valuation ratios and today's markets see Campbell & Shiller (2001).

[15] The monthly dividend yield is annualized by multiplying by 12. The historical average for the S&P is close to 5 percent.

[16] For a much more detailed statistical analysis see LeBaron (2001*c*).

This variable measures the fraction of the rules which are currently being used by agents in their active trading strategies. This comparison shows a very dramatic pattern occurring before both large crashes. Rule dispersion begins to fall long before the crash, and reaches historically low levels at or near the crash date.

Armed with only this small picture of the overall dynamics it is difficult to confirm an exact cause for market crashes. However, a simple story is starting to emerge. During the run up to a crash population diversity falls. Agents begin to use very similar trading strategies as their common good performance begins to self-reinforce. This makes the population very brittle in that a small reduction in the demand for shares could have a strong destabilizing impact on the market. The economic mechanism here is clear. Traders have a hard time finding anyone to sell to in a falling market since everyone else is following very similar strategies. In the Walrasian setup used here this forces the price to drop by a large magnitude to clear the market. The population homogeneity translates into a reduction in market liquidity.[17]

4 Summary and Conclusions

The results in this paper can be summarized along two dimensions. Its most important result is to provide a counter example to the the argument that evolutionarily less rational strategies should be driven out of the market. Its second result is in generating time series that appear reasonable for fitting certain difficult to replicate features from actual markets.

As a counter example to evolutionary arguments about market efficiency this model calls into question the basic structure of this argument. Who exactly is "less rational" in a world of heterogeneous agent investors? Will it be clearly obvious to investors to take a long run perspective in a market dominated by short run investors? In many ways these computer experiments may simply be demonstrating that it is very difficult to go against the flow of the current market even if you feel you must be right. Eventually, your performance relative to others will induce you to take a different (shorter) run perspective, and possibly further add to market instability and deviations from fundamental values.

It is important to realize that the evolutionary experiments implied in Friedman (1953) are quite a bit simpler than the actual market population dynamics. These arguments suggest a well defined population of super rational traders that is getting invaded by less rational types. In such a situation the rational types

[17]This overall dynamic has some interesting parallels to the problems encountered by Long Term Capital Management. This hedge found it difficult to reduce it positions since many other traders had similar trades in place. See Lowenstein (2000) for a detailed description.

may have already defined the environment since their numbers allow them to dominate the pricing of traded assets. In other words they have defined the rational world. In such a situation the evolutionary argument is correct, and it should be difficult for the short run types to stage a successful invasion of the market. The problem is that the more rational types often do not get to start with the luxury of having dominated the market. They must take over in a sea of noise from a heterogeneous population of less than rational types. To further complicate matters, this population may be changing its character over time, so the question of optimality gets quite murky.

Several time series features of financial data have proved to be difficult to replicate in standard models of asset pricing. This agent based simulation easily handles many of these empirical regularities. These include "fat tailed" return distributions, increased volatility and trading volume, volatility and volume persistence, and highly persistent and variable dividend yields. Although the model may not match all of these entirely it is still impressive that it can attack such a wide range of financial puzzles.

The model also takes a stand on the causes for large moves or market crashes. In the agent based laboratory the evidence suggests that crashes are caused by a move toward generally homogeneous markets. As agent strategies become aligned market liquidity falls, and the market dynamics become brittle as traders cannot find counter parties for their trades. Hopefully, this will lead to further testable hypothesis concerning large price moves, and to more general models of market liquidity.

The technology of agent based financial markets is still in its infancy, but it appears to be a promising route for increasing our understanding of the dynamics of financial markets. It provides new incites for trying to understand financial theories in a world which may be far from the equilibrium that the theories were designed to describe. Features of financial series such as large moves and excess volatility may become easier to understand when viewed from a multiagent evolutionary perspective. There is a long way to go in terms of model fitting and assessment, but for the moment agent based theories should take their place as a viable alternative to more traditional financial theories.

References

Arthur, W. B., Holland, J., LeBaron, B., Palmer, R. & Tayler, P. (1997), Asset pricing under endogenous expectations in an artificial stock market, *in* W. B. Arthur, S. Durlauf & D. Lane, eds, 'The Economy as an Evolving Complex System II', Addison-Wesley, Reading, MA, pp. 15–44.

Blume, L. & Easley, D. (1990), 'Evolution and market behavior', *Journal of Economic Theory* **58**, 9–40.

Campbell, J. Y. (1999), Asset prices, consumption, and the business cycle, *in* J. B. Taylor & M. Woodford, eds, 'Handbook of Macroeconomics', North Holland, pp. 1231–1303.

Campbell, J. Y. & Shiller, R. (1988), 'The dividend-price ratio and expectations of future dividends and discount factors', *Review of Financial Studies* **1**, 195–227.

Campbell, J. Y. & Shiller, R. J. (2001), Valuation ratios and the long-run stock market outlook: An update, Technical Report 8221, National Bureau of Economic Research.

Cont, R. & Bouchaud, J. P. (2000), 'Herd behavior and aggregate fluctuations in financial markets', *Macroeconomic Dynamics* **4**, 170–196.

DeLong, J. B., Shleifer, A., Summers, L. H. & Waldmann, R. (1991), 'The survival of noise traders in financial markets', *Journal of Business* **64**, 1–19.

Epstein, J. M. & Axtell, R. (1997), *Growing Artificial Societies*, MIT Press, Cambridge, MA.

Friedman, M. (1953), The case for flexible exchange rates, *in* 'Essays in positive economics', University of Chicago Press, Chicago, IL.

Kim, G. & Markowitz, H. (1989), 'Investment rules, margin, and market volatility', *Journal of Portfolio Management* **16**(1), 45–52.

LeBaron, B. (2000), 'Agent based computational finance: Suggested readings and early research', *Journal of Economic Dynamics and Control* **24**(5-7), 679–702.

LeBaron, B. (2001*a*), 'A builder's guide to agent-based financial markets', *Quantitative Finance* **1**(2), 254–261.

LeBaron, B. (2001*c*), Volatility magnification and persistence in an agent based financial market, Technical report, Graduate School of International Economics and Finance, Brandeis University, Waltham, MA.

LeBaron, B. (forthcoming 2001*b*), 'Empirical regularities from interacting long and short memory investors in an agent based stock market', *IEEE Transactions on Evolutionary Computation* .

Levy, M., Levy, H. & Solomon, S. (1994), 'A microscopic model of the stock market: cycles, booms, and crashes', *Economics Letters* **45**, 103–111.

Lowenstein, R. (2000), *When Genius Failed: The Rise and Fall of Long-Term Capital Management*, Random House, New York, NY.

Lux, T. (1997), 'Time variation of second moments from a noise trader/infection model', *Journal of Economic Dynamics and Control* **22**, 1–38.

Merton, R. (1969), 'Lifetime portfolio selection under uncertainty:The continuous time case', *Review of Economics and Statistics* **51**, 247–257.

NYSE Fact Book 1999 (2000), New York Stock Exchange, New York, NY.

Palmer, R., Arthur, W. B., Holland, J. H., LeBaron, B. & Tayler, P. (1994), 'Artificial economic life: A simple model of a stock market', *Physica D* **75**, 264–274.

Samuelson, P. (1969), 'Lifetime portfolio selection by dynamic stochastic programming', *Review of Economics and Statistics* **51**, 239–246.

Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.

Sargent, T. (1999), *The Conquest of American Inflation*, Princeton University Press, Princeton, NJ.

Simon, H. (1969), *The Sciences of the Artificial*, MIT Press, Cambridge, MA.

Tay, N. S. P. & Linn, S. C. (2001), 'Fuzzy inductive reasoning, expectation formation and the behavior of security prices', *Journal of Economic Dynamics and Control* **25**, 321–362.

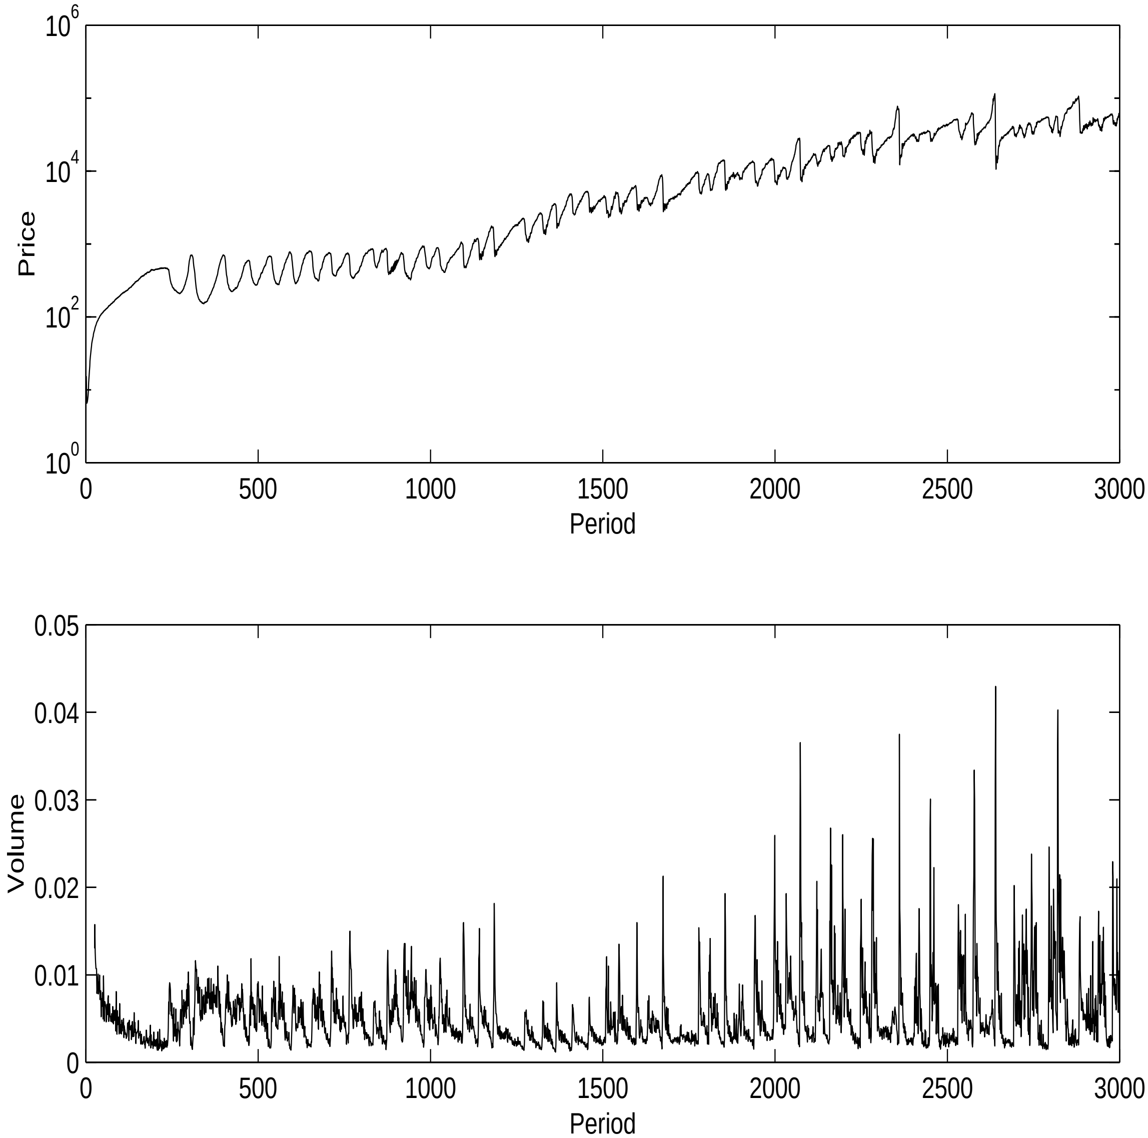

Figure 1: **Price/Volume: All Horizons**

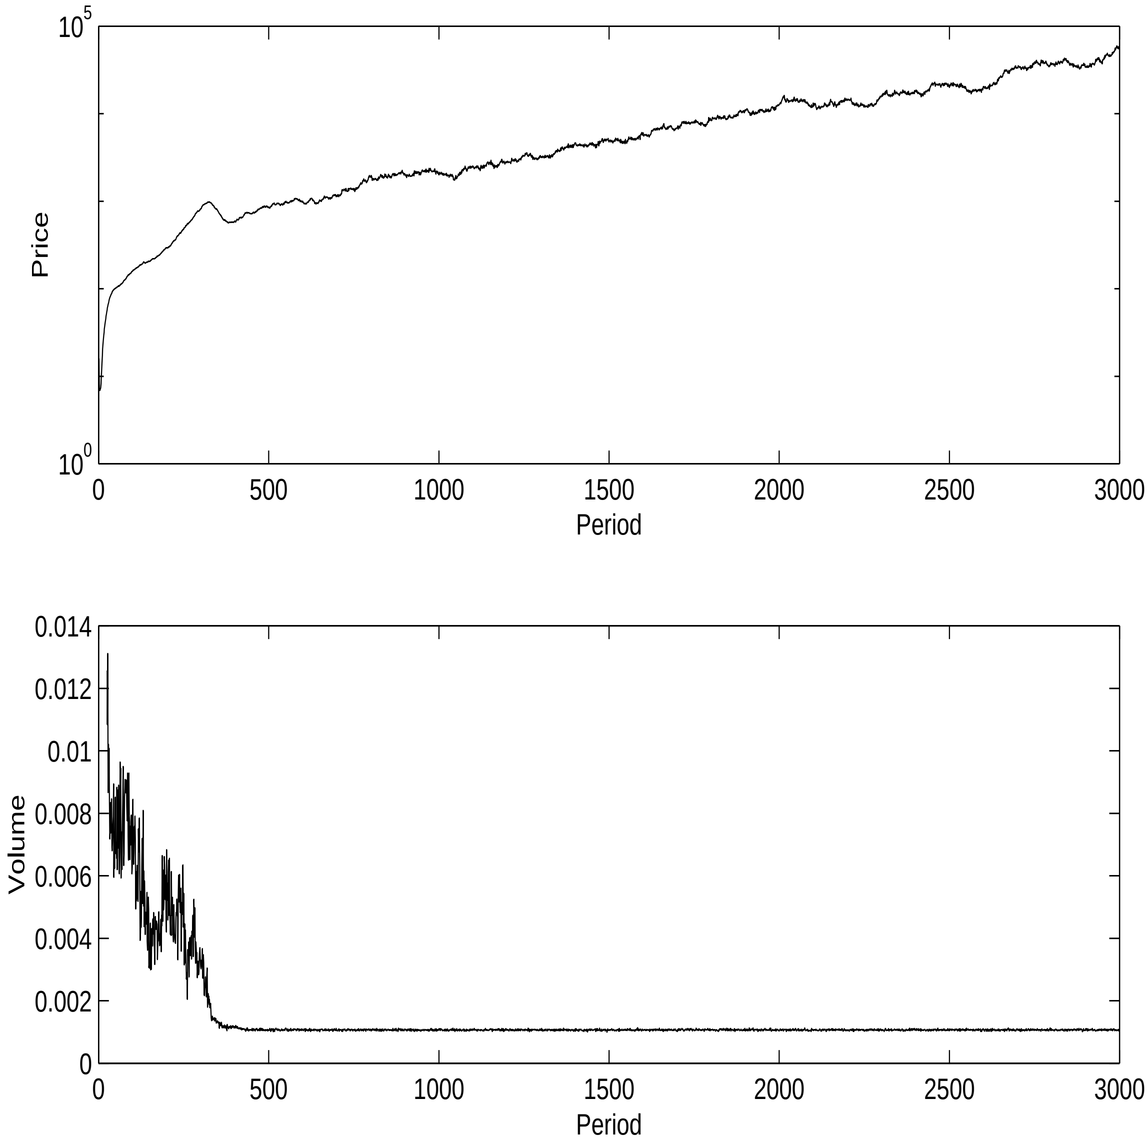

Figure 2: **Price/Volume: Long Horizons**

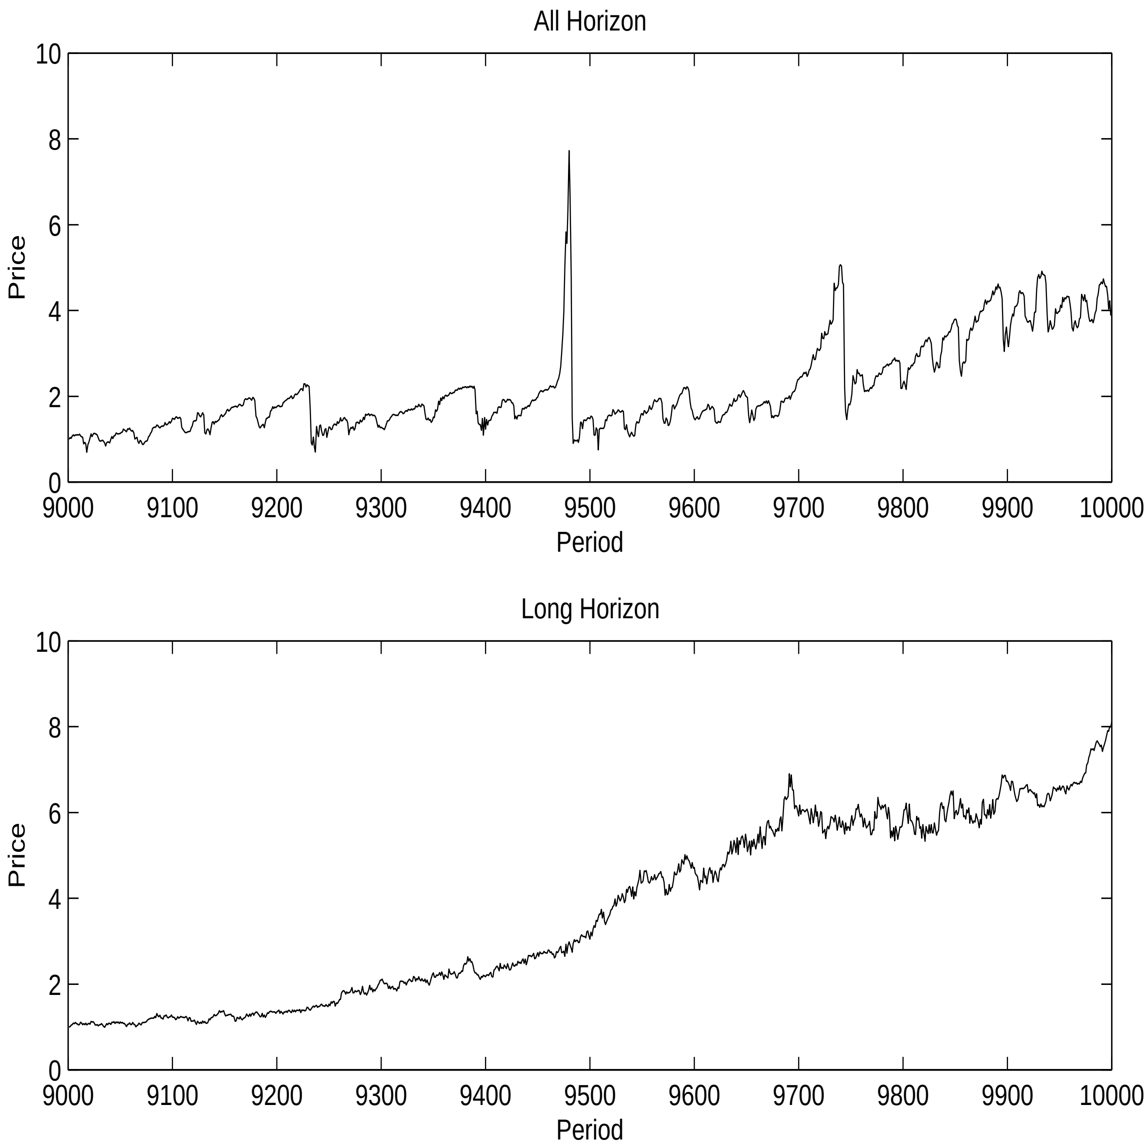

Figure 3: **Price Comparison**

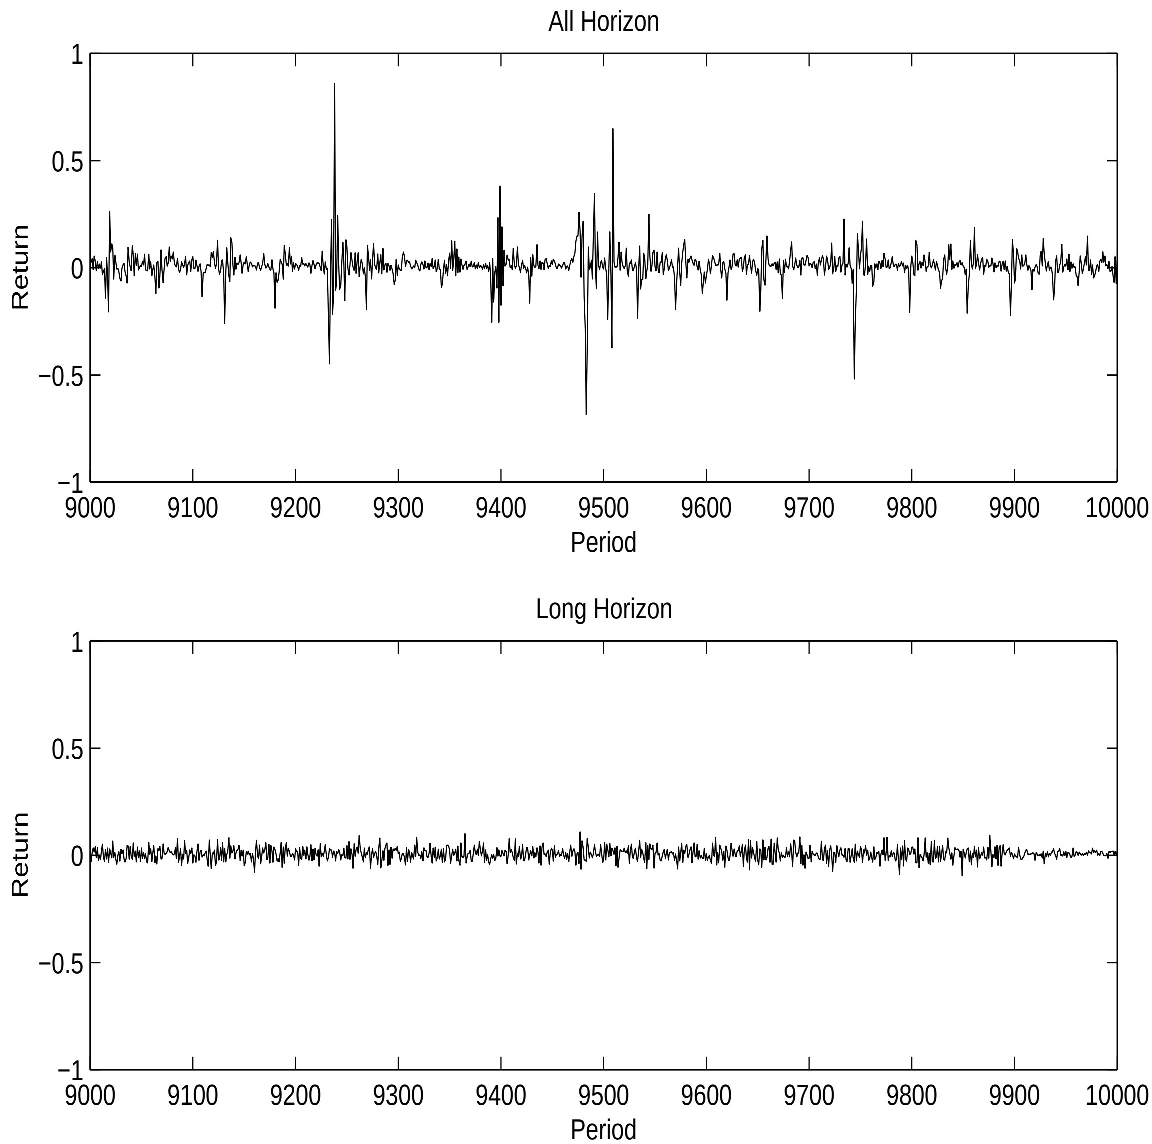

Figure 4: **Return Comparison**

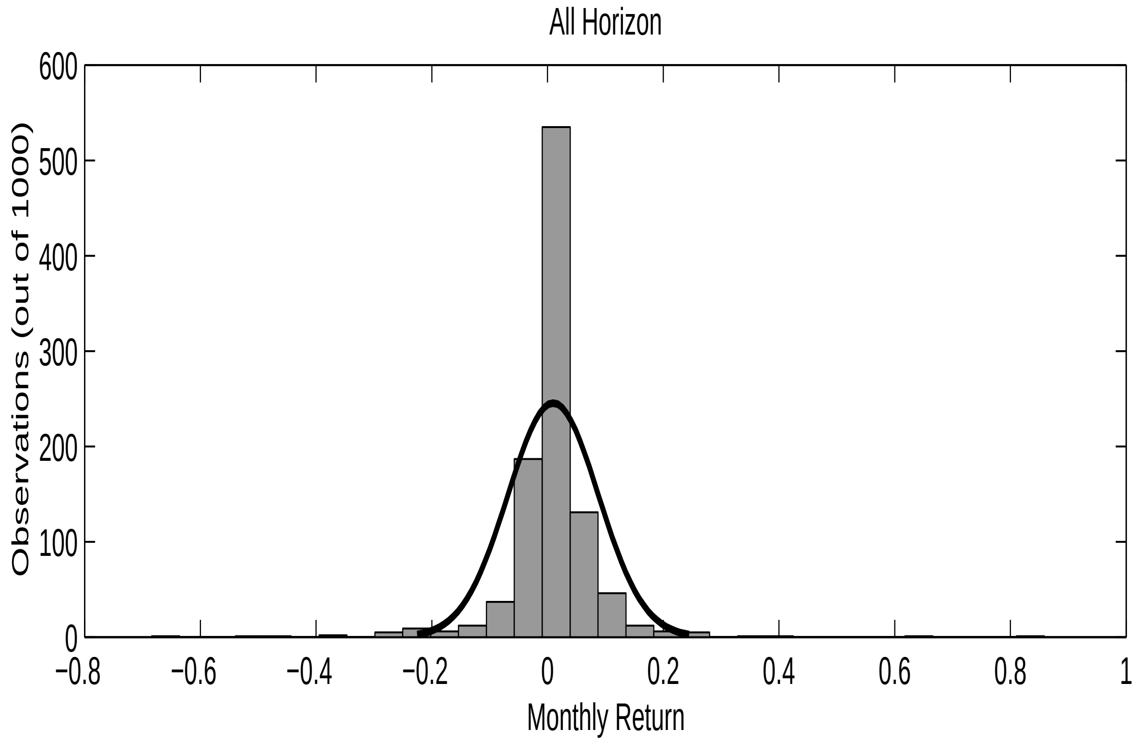

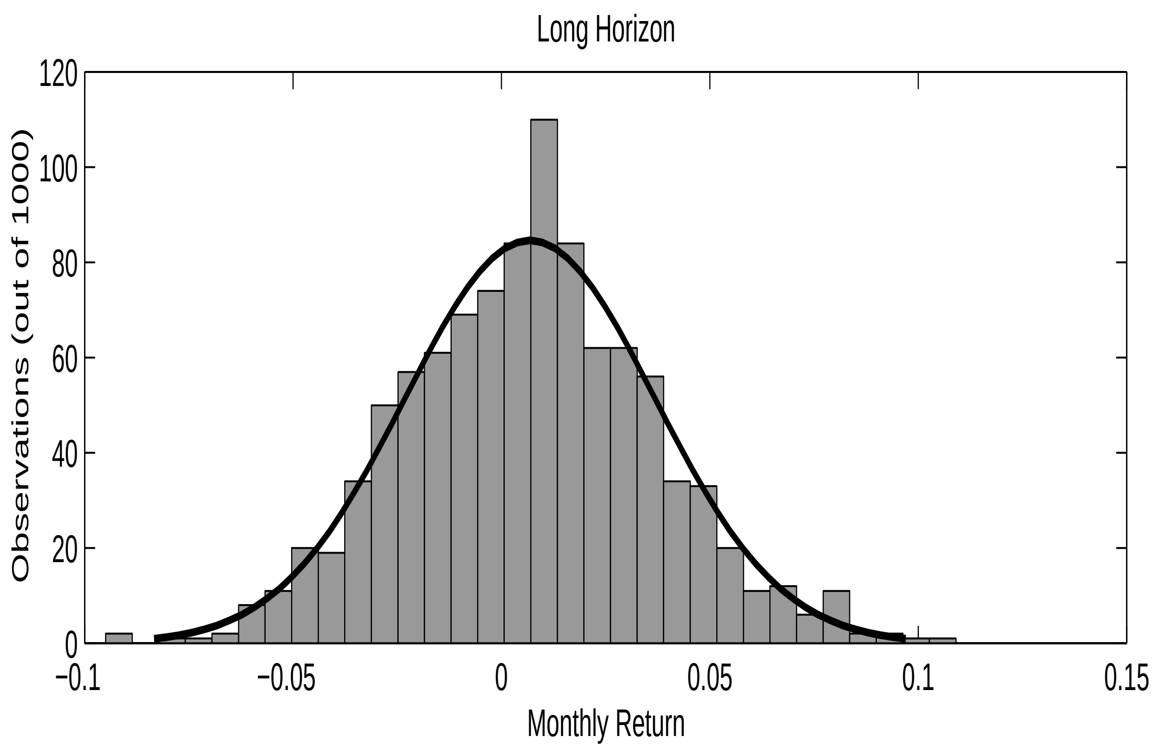

Figure 5: **Return Distributions**

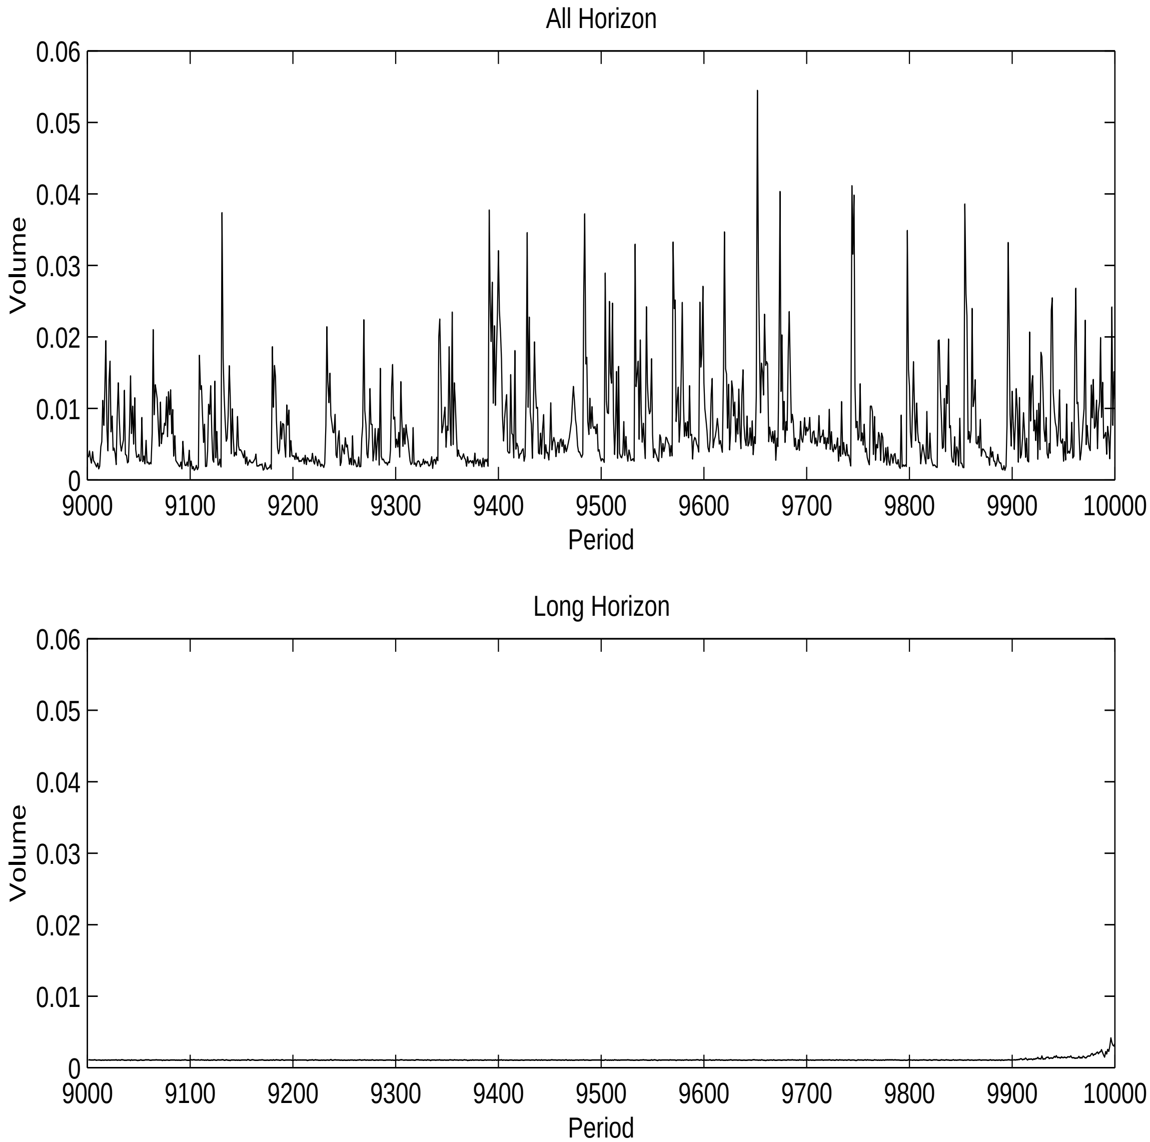

Figure 6: **Trading Volume**

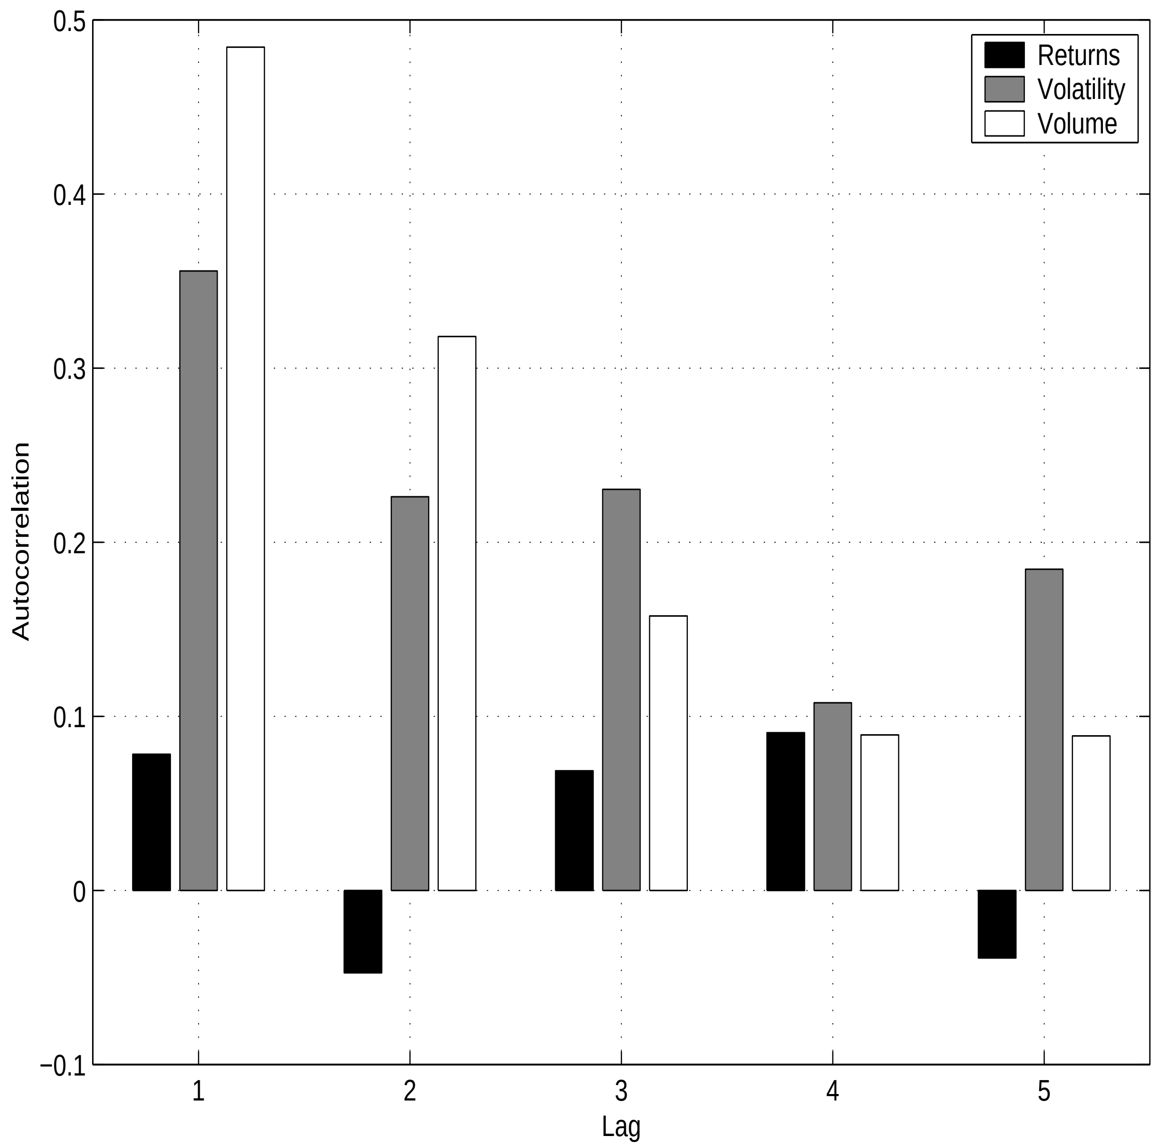

Figure 7: **All Horizon: Return, Volatility, and Volume Autocorrelations**

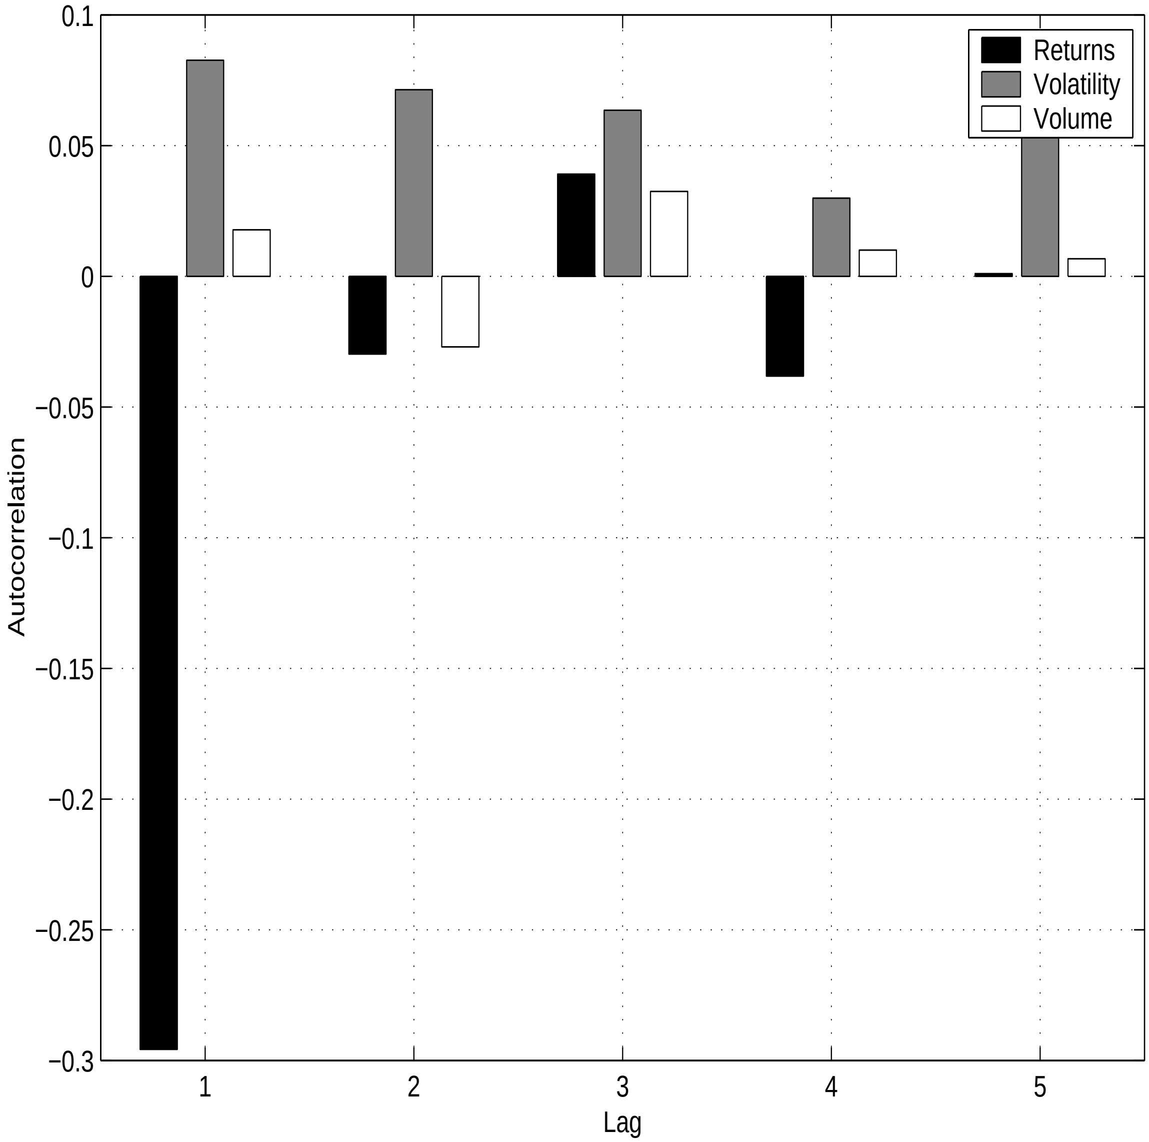

Figure 8: **Long Horizon: Return, Volatility, and Volume Autocorrelations**

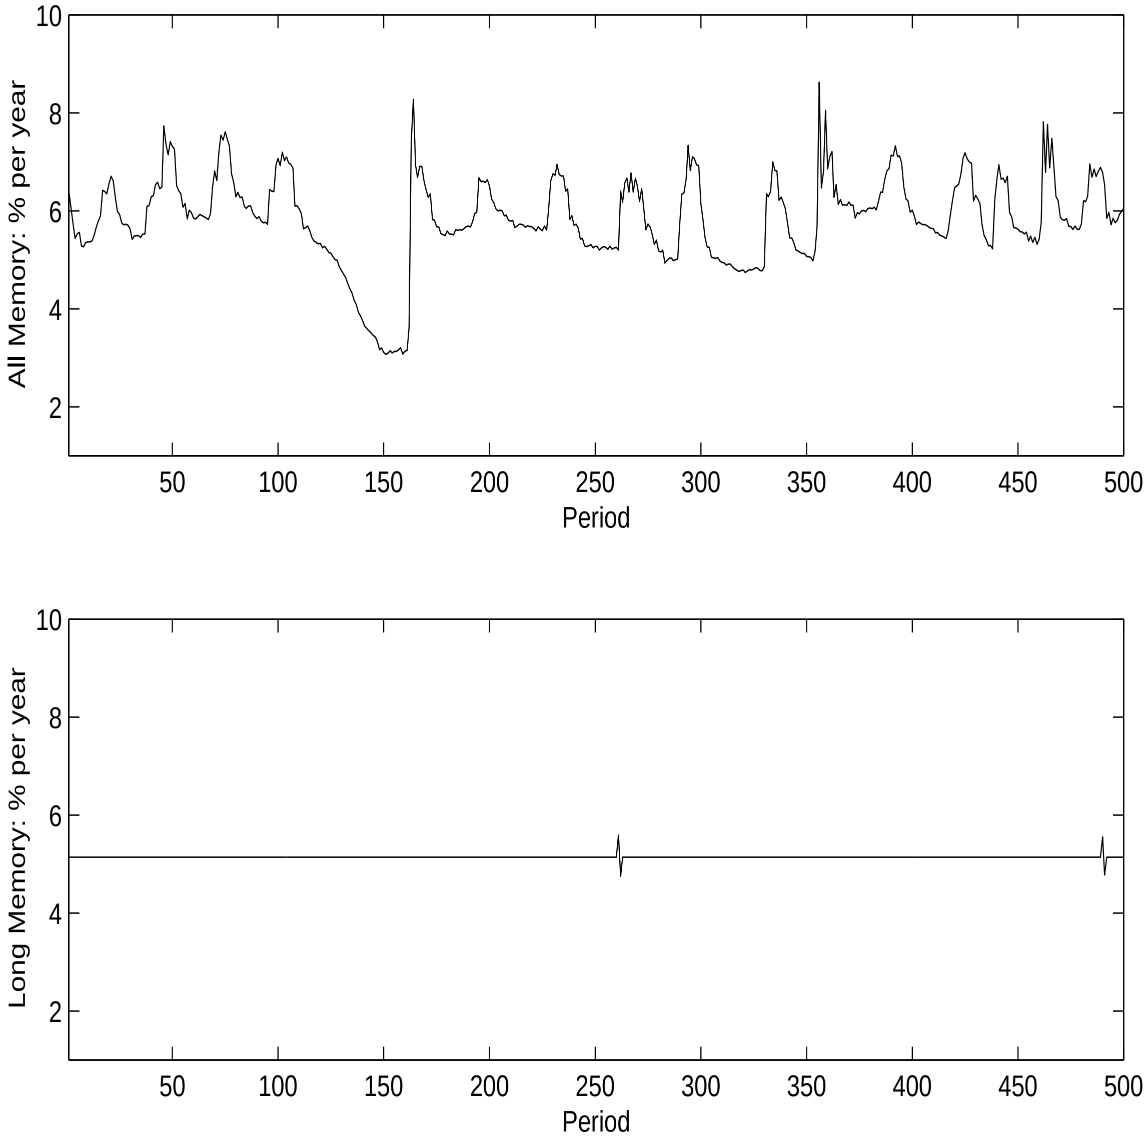

Figure 9: **Dividend Yields**

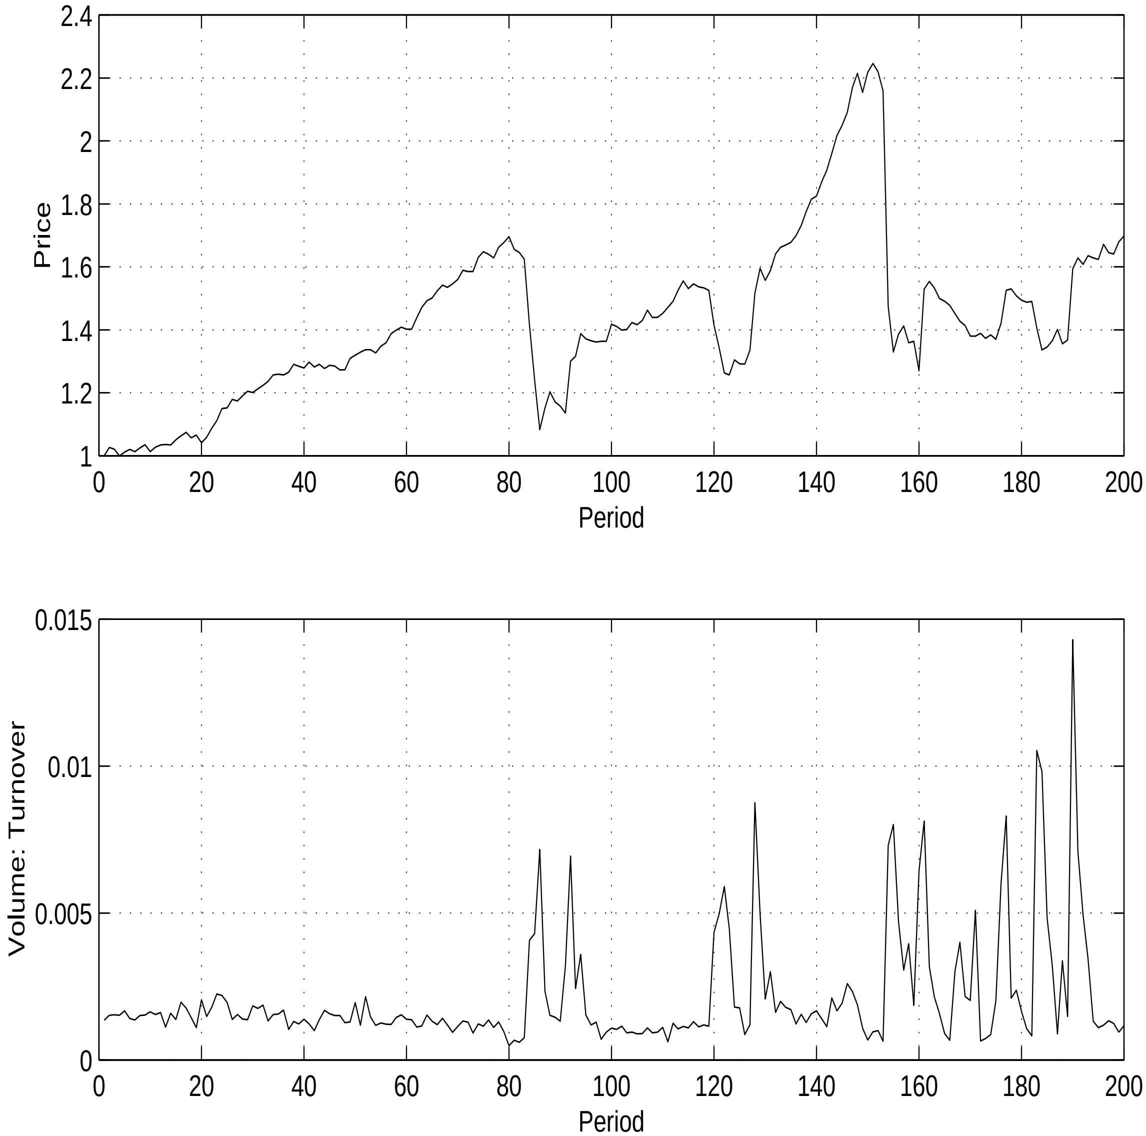

Figure 10: **Crashes and Volume**

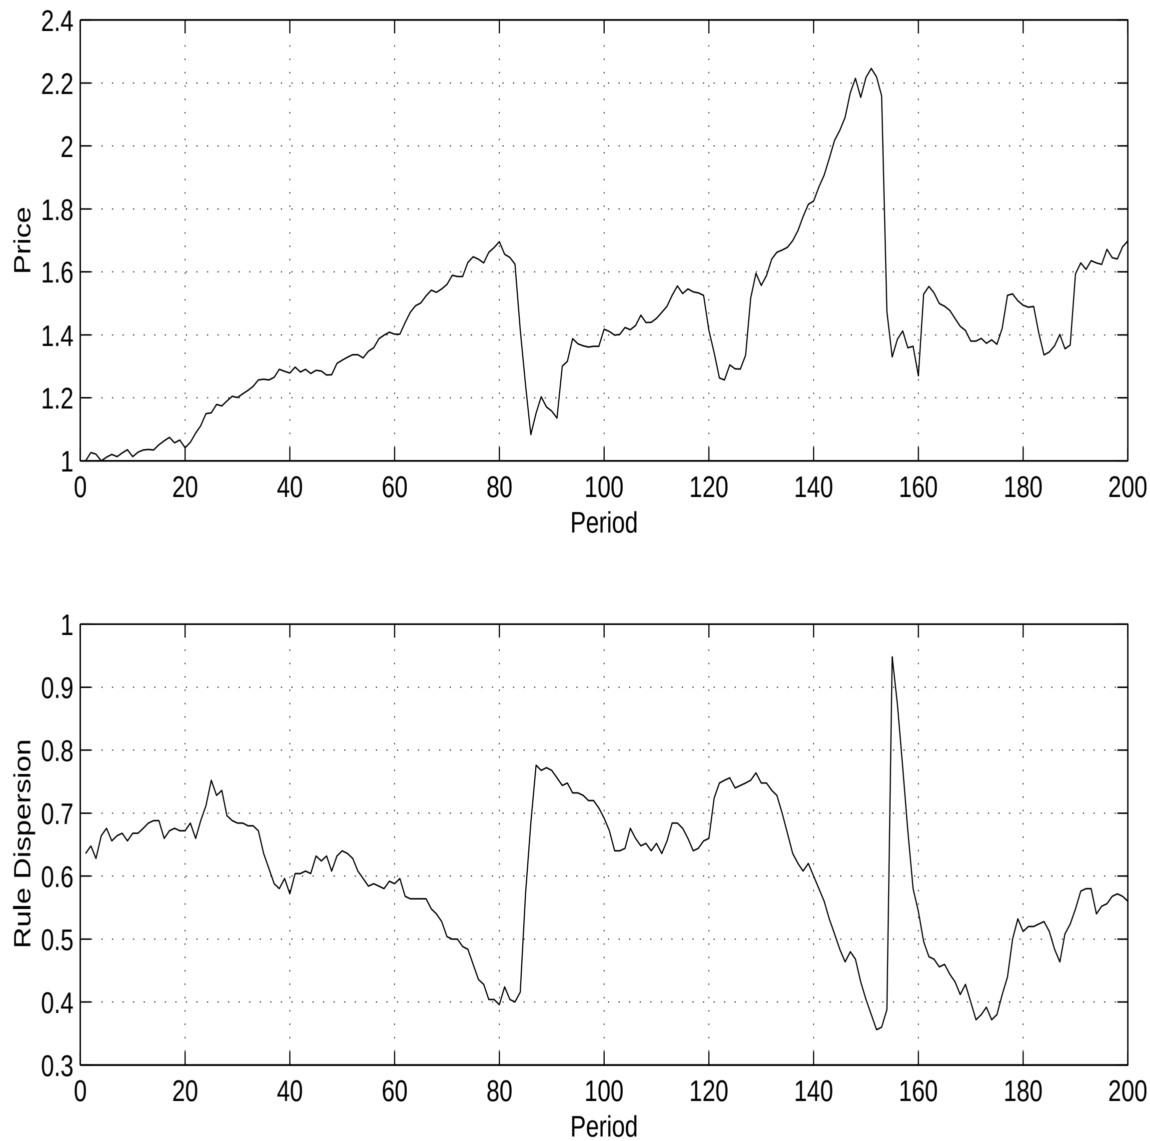

Figure 11: **Crashes and Dispersion**