

Heterogeneous Gain Learning and the Dynamics of Asset Prices

Blake LeBaron *
International Business School
Brandeis University

June 2010, Revised: December 2010

Abstract

This paper presents a new agent-based financial market. It is designed to be both simple enough to gain insights into the nature and structure of what is going on at both the agent and macro levels, but remain rich enough to allow for many interesting evolutionary experiments. The model is driven by heterogeneous agents who put varying weights on past information as they design portfolio strategies. It faithfully generates many of the common stylized features of asset markets. It also yields some insights into the dynamics of agent strategies and how they yield market instabilities.

Keywords: Learning, Asset Pricing, Financial Time Series, Evolution, Memory

*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 research associate at the National Bureau of Economic Research.

1 Introduction

Models of financial markets as aggregates of dynamic heterogeneous adaptive agents faithfully replicate a large range of important stylized facts, and also offer us new insights into the underlying behavior behind asset price movements. This paper presents a new market model continuing in this tradition. It is designed with learning mechanisms that are simple enough for easier analysis and interpretation, yet rich enough to pursue many of the experiments in evolution and heterogeneity present in older, more complex setups. The goal of this balance in market design is to provide a new foundational structure for understanding financial market dynamics from this different perspective.

Heterogeneous agent-based models have been applied to financial markets for quite some time.¹ Their common theme is to consider worlds in which agents are adaptively learning over time, while they perceive and contribute to time series dynamics unfolding into the future. Endogenous price changes then feed back into the dynamic learning mechanisms. Agents are modeled as being boundedly rational, and the potential behavioral space for these systems is large. However, some distinctions in modeling strategies have emerged. One extreme of agent-based financial markets is what is known as a “few type” model where the number of potential trading strategies is limited to a small, and tractable set.² Dynamics of these markets can be determined analytically, and occasionally through computer simulations. Their simple structure often yields very easy and intuitive results. At the other extreme are what are known as “many type” models. In these cases the strategy space is large. In many cases it is infinite as agents are working to develop new and novel strategies. Obviously, the complexity of these models requires computational methods for analysis. This in itself is not a problem, but the abilities of researchers to analyze their detailed workings has been limited. The model presented here will try to seek a middle ground between these. It tries to be rich enough to generate interesting financial price and volume dynamics, but simple enough for careful analysis.

Traditionally, agent-based markets have used all kinds of internal structures to represent stock return forecasts. In this model simple linear forecasts will be used for both expected returns and conditional variances. These expectations of risk and return are the crucial inputs into a simplified, and standard portfolio choice problem. Linear forecasts will be drawn from four different forecast families which are chosen to be good general representations of what traders are doing. These include adaptive, or momentum, strategies, a mean reverting fundamental strategy, a short range predictive liquidity strategy, and a buy and hold benchmark. These four families are designed to

¹ Many examples can be found in recent surveys such as Hommes (2006), LeBaron (2006), and Chiarella, Dieci & He (2009). Another useful review is Farmer & Geanakoplos (2008) where the authors press the case for heterogeneity in modeling financial markets. Interesting theoretical results on heterogeneity and learning in multi-agent systems is in Adam & Marcet (2010b) and Frydman & Goldberg (2007).

² Early examples of these include, Day & Huang (1990), Brock & Hommes (1998) and Lux (1998).

provide a stylized representation of actual trader behavior, but should not be interpreted as literal trading strategies in actual use.

Several agent-based financial markets have highlighted the possibility for heterogeneity in the processing of past information by learning agents.³

This market will use differences in how the past is evaluated by traders to generate heterogeneous future forecasts. There are many good reasons for doing this. The most important is that the model explores the evolutionary interactions between short and long memory traders, with an interest in whether any of these types dominate. A second reason, is that this parameter is part of almost all learning algorithms. In this paper, learning will be of the constant gain variety, where a fixed gain parameter determines agents' perception of how to process past data. Setting this to a specific value, constant across all agents, would impose a very large dynamic assumption on the model.

This market monitors the evolution of wealth over traders through time. Some wealth will be removed each period as a form of consumption. Although agents are utility maximizing, the consumption decision will be taken as a fixed mechanical rule to consume a constant fraction of wealth. This could be viewed as the outcome of a fairly restrictive set of preferences, or a reasonable rule of thumb for a simple consumption dynamic. This simplification still seems absolutely necessary to maintain model tractability. It also gives a reasonable bound on bubbles. At the high end, agents are drawing consumption from dividends, and from saved cash holdings. When these holdings run out, the aggregate spending level becomes unsustainable as agents all must begin selling shares to finance consumption. As more of the population moves to this point, the bubble will have to end. The passive dynamic on wealth suggests that agents whose strategies are generating better performance over time will increase their wealth shares. One can also add to this a utility maximizing framework in which agents shift strategies over time chasing better expected utilities from observed return series. I will refer to this as active learning. This allows for some interesting contrasts and comparisons between both active and passive learning.⁴

This paper begins by demonstrating that the model generates reasonable dynamics in terms of financial time series. It will be shown that this structure is able to give returns which are: (1) uncorrelated, (2) leptokurtic, (3) heteroskedastic. Furthermore, prices take large swings from a

³Examples include Diks & van der Weide (2005), Levy, Levy & Solomon (1994), LeBaron (2001), and Thurner, Dockner & Gaunersdorfer (2002). This paper directly considers the destabilizing impact of large gain, or "short memory" traders on market dynamics. Similar questions about how much data agents should be using from the past are considered in Mitra (2005). A recent model stressing what might happen when agents overweight shorter run trends, ignoring longer length reversals, is Fuster, Laison & Mendel (2010). Hommes (2010) surveys growing experimental literature which often shows people following short term trends when reporting their expectations in controlled laboratory settings. Finally, Dacorogna, Gencay, Muller, Olsen & Pictet (2001) presents a philosophy, and some time series models, for markets populated by agents with many different time perspectives.

⁴Passive learning models have been studied extensively. Good examples are Blume & Easley (1990), Blume & Easley (2006), DeLong, Shleifer, Summers & Waldmann (1991), Evstigneev, Hens & Schenk-Hoppe (2006), Figlewski (1978), Kogan, Ross & Wang (2006), and Sandroni (2000). Some explorations of the biases present when simple passive wealth evolution is implemented are given in LeBaron (2007). Further discussion is contained in LeBaron (2011 forthcoming).

fundamental dividend process calibrated to actual dividend dynamics from the U.S. Many other agent-based markets can meet this empirical hurdle, but this market tries to do it in a more simplified fashion that clearly displays the agent dynamics leading to these results. One final check on the model is whether it can generate results which are recognizable as a rational expectations equilibrium under any set of parameters. Restricting the gain parameters by allowing only long memory learners who make decisions based only on long range results going into the distant past yields a very simple market that converges to a rational expectations equilibrium, with returns which are independent, identically distributed (IID) Gaussians. These empirical summaries are presented in section 3.

The rest of the paper (section 4) examines some details of the market. It reports on the distributions of agents surviving in the market which is critical to market dynamics. It analyzes the specific comovements, and market behavior, around crashes which is important to the overall evolutionary process of agents and learning. Finally, some simple robustness checks are performed by modifying the set of agents. Section 5 will summarize, conclude, and highlight future questions which can be addressed in this framework.

2 Model Structure

This section describes the basic structure of the model. It is designed to be tractable, streamlined, and close to well known simple financial models. The use of recognized components allows for better analysis of the impact of interactive learning mechanisms on financial dynamics. Before getting into the details, I will emphasize several key features.

First, market forecasts are drawn from two common forecasting families, adaptive and fundamental expectations. The adaptive traders base their expectations of future returns from weighted sums of recent returns. The expectation structure is related to simple adaptive expectations, but also has origins in either Kalman filter, momentum or trend following mechanisms. The fundamental traders base their expectations on deviations of the price from the level of dividends using (P_t/D_t) ratios. The impact of the price/dividend ratio on conditional expected returns is determined by running an adaptive regression using a recursive least squares learning algorithm.

Agent portfolio choices are made using preferences which correspond to standard myopic constant relative risk aversion. Portfolio decisions depend on agents' expectations of the conditional expected return and variance of future stock returns. This allows for splitting the learning task on return and risk into two different components which adds to the tractability of the model. These preferences could also be interpreted as coming from intertemporal recursive preferences subject to certain further assumptions.

The economic structure of the model is well defined, simple, and close to that for standard

simple finance models.⁵ Dividends are calibrated to the trend and volatility of real dividend movements from U.S. aggregate equity markets.⁶ The basic experiment is then to see if market mechanisms can generate the kinds of empirical features we observe in actual data from this relatively quiet, but stochastic fundamental driving process. The market can therefore be viewed as a kind of nonlinear volatility generator for actual price series. The market structure also is important in that outside resources arrive only through the dividend flows entering the economy, and are used up only through consumption. The consumption levels are set to be proportional to wealth which, though unrealistic, captures the general notion that consumption and wealth must be cointegrated in the long run. Finally, prices are set to clear the market for the fixed supply of equity shares. The market clearing procedure allows for the price to be included in expectations of future returns, so an equilibrium price level is a form of temporary equilibrium for a given state of wealth spread across the current forecasting rules.

Rule heterogeneity and expectational learning for both expected returns, and conditional variances, is concentrated in the forecast and regression gain parameters. Constant gain learning mechanisms put fixed declining exponential weights on past information. Here, the competition across rules is basically a race across different gains, or weights of the past. The market is continually asking the question whether agents weighing recent returns more heavily can be driven out of the market by more long term forecasters.

The empirical features of the market are emergent in that none of these are prewired into the individual trading algorithms. Some features from financial data that this market replicates are very interesting. This would include the simple and basic feature of low return autocorrelations. In this market traders using short range autoregressive models play the role of short run arbitragers who successfully eliminate short run autocorrelations. They continually adapt to changing correlations in the data, and their adaptation and competition with others drives return correlations to near zero. This simple mechanism of competitive near term market efficiency seems consistent with most stories we think about occurring in real markets.

Finally, learning in the market can take two different forms. First, there is a form of passive learning in which wealth which is committed to rules that perform well tends to grow over time. These strategies then play an ever bigger role in price determination. This is the basic idea that successful strategies will eventually take over the market. All simulations will be run with some form of passive learning present, since it is fundamental to the model and its wealth dynamics. Beyond this, the model can also consider a form of active learning in which agents periodically adapt their behavior by changing to forecast rules that improve their expected utility. There are

⁵Its origins are a primitive version of models such as, Samuelson (1969), Merton (1969), and Lucas (1978) which form a foundation for much of academic finance.

⁶ Dividend calibration uses the annual Shiller dividend series available at Robert Shiller's Yale website. Much of this data is used in his book Shiller (2000). Another good source of benchmark series is Campbell (1999) which gives an extensive global perspective. Early results show that the basic results are not sensitive to the exact dividend growth and volatility levels.

many ways to implement this form of adaptive learning in the model, and only a few will be explored here. Another interesting question is how precise the estimates of expected utility are that are guiding the active learning dynamics. In a world of noisy financial time series adaptations might simply generate a form of drift across the various forecasting rules. Comparing and contrasting these two different types of learning is an interesting experiment which this model is designed to explore.

2.1 Assets

The market consists of only two assets.⁷ First, there is a risky asset paying a stochastic dividend,

$$d_{t+1} = d_g + d_t + \epsilon_t, \quad (1)$$

where d_t is the log of the dividend paid at time t . Time will be incremented in units of weeks. Lower case variables will represent logs of the corresponding variables, so the actual dividend is given by,

$$D_t = e^{d_t}. \quad (2)$$

The shocks to dividends are given by ϵ_t which is independent over time, and follows a Gaussian distribution with zero mean, and variance, σ_d^2 , that will be calibrated to actual long run dividends from the U.S. The dividend growth rate would then be given by $e^{g+(1/2)\sigma_d^2}$ which is approximately $D_g = d_g + (1/2)\sigma_d^2$.

The return on the stock with dividend at date t is given by

$$R_t = \frac{P_t + D_t - P_{t-1}}{P_{t-1}}, \quad (3)$$

where P_t is the price of the stock at time t . Timing in the market is critical. Dividends are paid at the beginning of time period t . Both P_t and D_t are part of the information set used in forecasting future returns, R_{t+1} . There are I individual agents in the model indexed by i . The total supply of shares is fixed, and set to unity,

$$\sum_{i=1}^I S_{t,i} = 1. \quad (4)$$

There is also a risk free asset that is available in infinite supply, with agent i holding $B_{t,i}$ units at time t . The risk free asset pays a rate of r_f which will be assumed to be zero in most simulations. This is done for two important reasons. It limits the injection of outside resources to the dividend process only. Also, it allows for an interpretation of this as a model with a perfectly storable consumption good along with the risky asset. The standard intertemporal budget constraint holds

⁷ An interesting comparable model with learning, a similar framework, but a single agent is Adam & Marcet (2010a).

for each agent i ,

$$W_{t,i} = P_t S_{t,i} + B_{t,i} + C_{t,i} = (P_t + D_t) S_{t-1,i} + (1 + r_f) B_{t-1,i}, \quad (5)$$

where $W_{t,i}$ represents the wealth at time t for agent i .

2.2 Preferences

Portfolio choices in the model are determined by a simple myopic power utility function in future wealth. The agent's portfolio problem corresponds to,

$$\max_{\alpha_{t,i}} \frac{E_t^i W_{t+1,i}^{1-\gamma}}{1-\gamma}, \quad (6)$$

$$st. \quad W_{t+1,i} = (1 + R_{t+1,i}^p)(W_{t,i} - C_{t,i}), \quad (7)$$

$$R_{t+1,i}^p = \alpha_{t,i} R_{t+1} + (1 - \alpha_{t,i}) R_f. \quad (8)$$

$\alpha_{t,i}$ represents agent i 's fraction of savings ($W - C$) in the risky asset. It is well known that the solution to this problem yields an optimal portfolio weight given by,

$$\alpha_{t,i} = \frac{E_t^i(r_{t+1}) - r_f + \frac{1}{2}\sigma_{t,i}^2}{\gamma\sigma_{t,i}^2} + \epsilon_{t,i}, \quad (9)$$

with $r_t = \log(1 + R_t)$, $r_f = \log(1 + R_f)$, $\sigma_{t,i}^2$ is agent i 's estimate of the conditional variance at time t , and $\epsilon_{t,i}$ is an individual shock designed to make sure that there is some small amount of heterogeneity to keep trade operating.⁸ It is distributed normally with variance, σ_ϵ^2 .

In the current version of the model neither leverage nor short sales are allowed. The fractional demand is restricted to $\alpha_{t,i}$ with $\alpha_L \leq \alpha_{t,i} \leq \alpha_H$. The addition of both these features is important, but adds significant model complexity. One key problem is that with either one of these, one must address problems of agent bankruptcy, and borrowing constraints. Both of these are not trivial, and involve many possible implementation details.

Consumption will be assumed to be a constant fraction of wealth, λ . This is identical over agents, and constant over time. The intertemporal budget constraint is therefore given by

$$W_{t+1,i} = (1 + R_{t+1,i}^p)(1 - \lambda)W_{t,i}. \quad (10)$$

This also gives the current period budget constraint,

$$P_t S_{t,i} + B_{t,i} = (1 - \lambda)((P_t + D_t) S_{t-1,i} + (1 + r_f) B_{t-1,i}). \quad (11)$$

⁸The derivation of this follows Campbell & Viceira (2002). It involves taking a Taylor series approximation for the log portfolio return.

This simplified portfolio strategy will be used throughout the paper. It is important to note that the fixed consumption/wealth, myopic strategy approach given here would be optimal in a standard intertemporal model for consumption portfolio choice subject to two key assumptions. First, the intertemporal elasticity of substitution would have to be unity to fix the consumption wealth ratio, and second, the correlation between unexpected returns and certain state variables would have to be zero to eliminate the demand for intertemporal hedging.⁹

2.3 Expected Return Forecasts

The basic problem faced by agents is to forecast both expected returns and the conditional variance one period into the future. This section will describe the forecasting tools used for expected returns. A forecast strategy, indexed by j , is a method for generating an expected return forecast $E^j(r_{t+1})$. Agents, indexed by i , can either be fixed to a given forecasting rule, or may adjust rules over time depending on the experiment.

All the forecasts will use long range forecasts of expected values using a long range minimum gain level, g_L .

$$\bar{r}_t = (1 - g_L)\bar{r}_{t-1} + g_L r_t \quad (12)$$

$$\bar{p}d_t = (1 - g_L)\bar{p}d_{t-1} + g_L p d_{t-1} \quad (13)$$

$$\bar{\sigma}_{r,t}^2 = (1 - g_L)\bar{\sigma}_{r,t-1}^2 + g_L (r_t - \bar{r}_t)^2 \quad (14)$$

$$\bar{\sigma}_{pd,t}^2 = (1 - g_L)\bar{\sigma}_{pd,t-1}^2 + g_L (p d_t - \bar{p}d_t)^2 \quad (15)$$

The long range forecasts, \bar{r}_t , $\bar{p}d_t$, $\bar{\sigma}_{r,t}^2$, and $\bar{\sigma}_{pd,t}^2$ correspond to the mean log return, log price/dividend ratio, and variance respectively, and the gain parameter g_L is common across all agents.

The forecasts used will combine four linear forecasts drawn from well known forecast families.¹⁰ The first of these is an adaptive linear forecast which corresponds to,

$$f_t^j = f_{t-1}^j + g_j (r_t - f_{t-1}^j). \quad (16)$$

Forecasts of expected returns are dynamically adjusted based on the latest forecast and r_t . This forecast format is simple and generic. It has roots connected to adaptive expectations, trend following technical trading, and also Kalman filtering.¹¹ In all these cases a forecast is updated given

⁹See Campbell & Viceira (1999) for the basic framework. Also, see Giovannini & Weil (1989) for early work on determining conditions for myopic portfolio decisions. Hedging demands would only impose a constant shift on the optimal portfolio, so it is an interesting question how much of an impact this might have on the results.

¹⁰This division of rules is influenced by the many models in the "few type" category of agent-based financial markets. These include Brock & Hommes (1998), Day & Huang (1990), Gennotte & Leland (1990), Lux (1998). Some of the origins of this style of modeling financial markets can be traced to Zeeman (1974).

¹¹A nice summary of the connections between Kalman filtering, adaptive expectations, and recursive least squares is given in Sargent (1999).

its recent error. The critical parameter is the gain level represented by g_j . This determines the weight that agents put on recent returns and how this impacts their expectations of the future. Forecasts with a large range of gain parameters will compete against each other in the market. Finally, this forecast will be trimmed in that it is restricted to stay between the values of $[-h_j, h_j]$. These will be set to relatively large values, and are randomly distributed across the j rules.

The second forecasting rule is based on a fundamental strategy. This forecast uses log price dividend ratio regressions as a basis for forecasting future returns,

$$f_t^j = \bar{r}_t + \beta_t^j(pd_t - \bar{p}d_t). \quad (17)$$

where pd_t is $\log(P_t/D_t)$. Although agents are only interested in the one period ahead forecasts the P/D regressions will be estimated using the mean return over the next M_{PD} periods, with $M_{PD} = 52$ for all simulations.

The third forecast rule will be based on linear regressions. It is a predictor of returns at time t given by

$$f_t^j = \bar{r}_t + \sum_{i=0}^{M_{AR}-1} \beta_{t,i}^j (r_{t-i} - \bar{r}_t) \quad (18)$$

This strategy works to eliminate short range autocorrelations in returns series through its behavior, and $M_{AR} = 3$ for all runs in this paper. It will be referred to as the Short AR forecast.

The previous two rules will be estimated each period using recursive least squares. There are many examples of this for financial market learning.¹² The key difference is that this model will stress heterogeneity in the learning algorithms with wealth shifting across many different rules, each using a different gain parameter in its online updating.¹³

The final rule is a benchmark strategy. It is a form of buy and hold strategy using the long run mean, \bar{r}_t , for the expected return, and the long run variance, $\sigma_{r,t}^2$, as the variance estimate. This portfolio fraction is then determined by the demand equation used by the other forecasting rules. This gives a useful passive benchmark strategy which can be monitored for relative wealth accumulation in comparison with the other active strategies.

2.4 Regression Updates

Forecasting rules are continually updated. The adaptive forecast only involves fixed forecast parameters, so its updates are trivial, requiring only the recent return. The two regression forecasts are updated each period using recursive least squares.

¹² See Evans & Honkapohja (2001) for many examples, and also very extensive descriptions of recursive least squares learning methods.

¹³ Another recent model stressing heterogeneity in an OLS learning environment is Georges (2008) in which OLS learning rules are updated asynchronously.

All the rules assume a constant gain parameter, but each rule in the family corresponds to a different gain level. This again corresponds to varying weights for the forecasts looking at past data. The fundamental regression is run using the long range return,

$$\tilde{r}_t = \frac{1}{M_{PD}} \sum_{j=1}^{M_{PD}} r_{t-j+1} \quad (19)$$

The fundamental regression is updated according to,

$$\begin{aligned} \beta_{t+1}^j &= \beta_t^j + \frac{g_j}{\bar{\sigma}_{pd,t}^2} (pd_{t-M_{PD}} u_{t,j}) \\ u_{t,j} &= (\tilde{r}_t - f_{j,t-M_{PD}}) \end{aligned} \quad (20)$$

For the lagged return regression this would be,

$$\begin{aligned} \beta_{t+1,i}^j &= \beta_{t,i}^j + \frac{g_j}{\bar{\sigma}_{r,t}^2} (r_{t-i} u_{t,j}), \\ u_{t,j} &= (r_t - f_t^j) \end{aligned} \quad (21)$$

where g_j is again the critical gain parameter, and it varies across forecast rules.¹⁴ In both forecast regressions the forecast error, $u_{t,j}$, is trimmed. If $u_{t,j} > h_j$ it is set to h_j , and if $u_{t,j} < -h_j$ it is set to $-h_j$. This dampens the impact of large price moves on the forecast estimation process.

2.5 Variance Forecasts

The optimal portfolio choice demands a forecast of the conditional variance as well as the conditional mean.¹⁵ The variance forecasts will be generated from adaptive expectations as in,

$$\hat{\sigma}_{t,j}^2 = \hat{\sigma}_{t-1,j}^2 + g_{j,\sigma} (e_{t,j}^2 - \hat{\sigma}_{t-1,j}^2) \quad (22)$$

$$e_{t,j}^2 = (r_t - f_{t-1}^j)^2, \quad (23)$$

where $e_{t,j}^2$ is the squared forecast error at time t , for rule j . The above conditional variance estimate is used for all the rules. There is no attempt to develop a wide range of variance forecasting

¹⁴This format for multivariate updating is only an approximation to the true recursive estimation procedure. It is assuming that the variance/covariance matrix of returns is diagonal. Generated returns in the model are close to uncorrelated, so this approximation is probably reasonable. This is done to avoid performing many costly matrix inversions.

¹⁵ Several other papers have explored the dynamics of risk and return forecasting. This includes Branch & Evans (2011 forthcoming) and Gaunersdorfer (2000). In LeBaron (2001) risk is implicitly considered through the utility function and portfolio returns. Obviously, methods that parameterize risk in the variance may miss other components of the return distribution that agents care about, but the gain in tractability is important.

rules, reflecting the fact that while there may be many ways to estimate a conditional variance, they often produce similar results.¹⁶ This forecast method has many useful characteristics as a benchmark forecast. First, it is essentially an adaptive expectations forecast on second moments, and therefore shares a functional form similar to that for the adaptive expectations family of return forecasts. Second, it is closely related to other familiar conditional variance estimates.¹⁷ Finally, the gain level for the variance in a forecast rule, $g_{j,\sigma}$, is allowed to be different from that used in the mean expectations, g_j . This allows for rules to have a different time series perspective on returns and volatility.

There is one further aspect of heterogeneity that is important to the market dynamics. Agents do not update their variance estimates immediately. They do it with a lag using a stochastic updating processes. Agent i will update to the current variance estimate for rule j , $\hat{\sigma}_{t+1,j}^2$, with probability p_σ .¹⁸ This allows for a greater amount of heterogeneity in variance forecasts, and mitigates some extreme moves in price which can be caused by a simultaneous readjustment in market risk forecasts. This is a form of simulating more heterogeneity in the variance forecasting process, but in a stochastic fashion.

2.6 Market Clearing

The market is cleared by setting the individual share demands equal to the aggregate share supply of unity,

$$1 = \sum_{i=1}^I Z_{t,i}(P_t). \quad (24)$$

Writing the demand for shares as its fraction of current wealth, remembering that $\alpha_{t,i}$ is a function of the current price gives

$$P_t Z_{t,i} = (1 - \lambda) \alpha_{t,i}(P_t) W_{t,i}, \quad (25)$$

$$Z_{t,i}(P_t) = (1 - \lambda) \alpha_{t,i}(P_t) \frac{(P_t + D_t) S_{t-1,i} + B_{t-1,i}}{P_t}. \quad (26)$$

This market is cleared for the current price level P_t . This needs to be done numerically given the complexities of the various demand functions and forecasts, and also the boundary conditions on $\alpha_{t,i}$.¹⁹ It is important to note again, that forecasts are conditional on the price at time t , so the

¹⁶See Nelson (1992) for early work on this topic.

¹⁷ See Bollerslev, Engle & Nelson (1995) or Andersen, Bollerslev, Christoffersen & Diebold (2006) for surveys of the large literature on volatility modeling.

¹⁸Also, the agents do not use p_t information in their forecasts of the conditional variance at time t . This differs from the return forecasts which do use time t information. Incorporating time t information into variance forecasts will cause the market not to converge as prices can spiral far from their current levels, causing market demand for shares to crash to zero.

¹⁹A binary search is used to find the market clearing price using starting information from P_{t-1} . The details of this algorithm are given in Appendix A.

market clearing involves finding a price which clears the market for all agent demands, allowing these demands to be conditioned on their forecasts of R_{t+1} given the current price and dividend.²⁰

2.7 Gain Levels

An important design question for the simulation is how to set the range of gain levels for the various forecast rules. These will determine the dynamics of forecasts. Given that this is an entire distribution of values it will be impossible to accomplish much in terms of sensitivity analysis on this. Therefore, a reasonable mechanism will be used to generate these, and this will be used in all the simulations.

Gain levels will be thought of using their half-life equivalents, since the gain numbers themselves do not offer much in the way of economic or forecasting intuition. For this think of the simple exponential forecast mechanism with

$$f_{t+1}^j = (1 - g_j)f_t^j + g_j e_{t+1}. \quad (27)$$

This easily maps to the simple exponential forecast rule,

$$f_t = \sum_{k=1}^{\infty} (1 - g_j)^k e_{t-k}. \quad (28)$$

The half-life of this forecast corresponds to the number of periods, m_h , which drops the weight to 1/2,

$$\frac{1}{2} = (1 - g_j)^{m_h}, \quad (29)$$

or

$$g_j = 1 - 2^{-1/m_h}. \quad (30)$$

The distribution of m_h then is the key object of choice here. It is chosen so that $\log_2(m_h)$ is distributed uniformly between a given minimum and maximum value. The gain levels are further simplified to use only 5 discrete values. These are given in table 1, and are [1, 2.5, 7, 18, 50] years respectively. In the long memory (low gain) experiments these five values will be distributed between 45 and 50 years.

These distributions are used for all forecasting rules. All forecast rules need a gain both for the expected return forecast, and the variance forecast. These will be chosen independently from each other. This allows for agents to have differing perspectives on the importance of past data for the expected return and variance processes.

²⁰ The current price determines R_t which is an input into both the adaptive, and noise trader forecasts. Also, the price level P_t enters into the P_t/D_t ratio which is required for the fundamental forecasts. All forecasts are updated with this time t information in the market clearing process.

2.8 Adaptive rule selection

The design of the models used here allows for both passive and active learning. Passive learning corresponds to the long term evolution of wealth across strategies. Beyond passive learning, the model allows for active learning, or more adaptive rule selection. This mechanism addresses the fact that agents will seek out strategies which best optimize their estimated objective functions. In this sense it is a form of adaptive utility maximization.

Implementing such a learning process opens a large number of design questions. This paper stays with a relatively simple implementation. The first question is how to deal with estimating expected utility. Expected utility will be estimated using an exponentially weighted average over the recent past,

$$\hat{U}_{t,j} = \hat{U}_{t-1,j} + g_u^i (U_{t,j} - \hat{U}_{t-1,j}), \quad (31)$$

where $U_{t,j}$ is the realized utility for rule j received at time t . This corresponds to,

$$U_{t,j} = \frac{1}{1-\gamma} (1 + R_{t,j}^p)^{(1-\gamma)} \quad (32)$$

with $R_{t,j}^p$ the portfolio holdings of rule j at time t . Each rule reports this value for the 5 discrete agent gain parameters, g_u^i . Agents choose rules optimally using the objective that corresponds to their specific perspective on the past, g_u^i , which is a fixed characteristic. The gain parameter g_u^i follows the same discrete distribution as that for the expected return and variance forecasts.

The final component to the learning dynamic is how the agents make the decision to change rules. The mechanism is simple, but designed to capture a kind of heterogeneous updating that seems plausible. Each period a certain fraction, L , of agents is chosen at random. Each one randomly chooses a new rule out of the set of all rules. If this rule exceeds the current one in terms of estimated expected utility, then the agent switches forecasting rules.

3 Results and Experiments

3.1 Calibration and parameter settings

Table 1 presents the key parameters used in the simulation. As mentioned the dividend series is set to a geometric random walk with drift. The drift level, and annual standard deviation are set to match those from the real dividend series in Shiller's annual data set. This gives a recognizable real growth rate for dividends of 2 percent per year. The level of risk aversion, γ will be fixed at 3.5 for all runs. This is a reasonable level for standard constant relative risk aversion.²¹ The gain range for the learning models is set to 1-50 years in half-life values. This means that the

²¹Many of the results can be replicated for a range of γ from 2 – 4. The value of 3.5 gives some of the most realistic looking series while still being a reasonable level.

largest gain values one year in the past, at one half the weight given to today, and the smallest gain weights data 50 years back at 1/2 today's weight. For all runs there will be $I = 16000$ agents, and $J = 4000$ forecast rules. The value of λ , the consumption wealth ratio, was chosen to give both a reasonable P/D ratio, and also reasonable dynamics in the P/D time series.

The basic simulations using these parameters with agent adaptation will be referred to as the *baseline* model. It will be shown that this model replicates most of the common features in financial series, and yields a large amount of intuition into price dynamics. Extensions and robustness checks will build and add to this *baseline* case. All simulations will be run for 200,000 weeks, or almost 4,000 years. Statistics are drawn from the end of this simulation.

Before beginning with this baseline experiment, the model is tested to see if it can converge to a reasonable, recognizable equilibrium for some parameter values. The model is simulated with learning half lives ranging from 40 to 50 years.²² The objective is to look at a population of agents restricted to only using long time series in their decision making and learning dynamics. Figure 1 presents a basic summary of the results for this case using data from a run length of 200,000 weeks. The top panel displays a subset of weekly returns which looks relatively uniform. The second panel shows that the returns are close to a Gaussian in terms of distribution. Finally, the bottom panel shows the autocorrelations for both the returns and absolute returns. Both are near zero at all lags. This shows the model generating time series which appear independent and close to Gaussian. Obviously, neither of these patterns is representative of actual return series. However, this is an important test of the learning algorithms in the model. Forcing the populations to only low gain types alone, gives the learning algorithms enough structure to converge to a reasonable equilibrium.

3.2 Time series features

3.2.1 Weekly Series

The simulations now turn to baseline runs using the parameter values from table 1. This performs the main test of the paper which is to see how learning algorithms of different gain levels interact with each other. Figure 1 is now repeated for this case in figure 2. The features are dramatically different from those in the first figure. The returns now show pockets of clustered volatility, and they are not close to a Gaussian, exhibiting fat tailed behavior. The bottom panel reports the autocorrelations, and shows that the returns close to uncorrelated, but absolute returns show strong positive correlations.

²² For these runs only the coefficient of relative risk aversion is increased to 8. This is done since the model drives returns to a very low volatility. At this level, for the baseline risk aversion, agents will be up against the portfolio constraints at the maximum holding level. In this case, market dynamics can occasionally become unstable. It is necessary to move the agents into the interior of their choice space to maintain stability. In some ways this is a form of the equity premium puzzle in a dynamic learning setting.

Figure 3 presents a simple price series from the last 10 years of the baseline simulation along with the most recent 10 years for the S&P 500 index. The two figures look similar, but not much can be said from the price figures alone. More detailed pictorial information is given in figure 4 which compares weekly continuously compounded return series from the CRSP value weighted index with dividends (1926-2009), and the baseline simulation. Both display some extreme movements, and some pockets of increased volatility which are common features of most financial series. These returns are further compared in two histograms in figure 5. For these the full sample is used again for the CRSP weekly returns and a similar length period from the end of the simulation run. Both show visually comparable levels of leptokurtosis relative to a standard Gaussian which is drawn for comparison. They display a large peak near zero, and too many observations in the tails.

Table 2 presents weekly summary statistics which reflect most of these early graphical features. They use the full 1926-2009 series for the CRSP index, and an even longer series, corresponding to the final 25,000 weeks in the baseline simulation. The table also reports results for an individual stock series using IBM returns from 1926 through Dec 2009. All returns include dividend distributions. Mean returns are in weekly percentages. The simulation return level is above the return for the market index, but below that for the IBM return. These mean return comparisons are casual since the data returns are nominal, and the model returns are real since there is no inflation. The model displays one of its important characteristics in the second line which reports the standard deviation. The weekly standard deviation for the model is 3.54 which is higher than the index, but close to the level of volatility in the IBM return series. In row three all series show evidence for some negative skewness. Row four shows the usual large amount of kurtosis for all 3 series. This is consistent with the visual evidence already presented.

The last two rows in the table present the tail exponent which is another measure of the shape of the tail in a distribution. This estimate uses a modified version of the Hill estimator as developed in Huisman, Koedijk, Kool & Palm (2001), and further explored in LeBaron (2008) who shows it gives a very reliable estimate of this tail shape parameter. Values in the neighborhood of 2 – 4 are common for weekly asset return series, so the results here are all within reasonable ranges. There is some indication that the simulations are producing more extreme tails than the actual data which is consistent with the graphical evidence in figure 5.

Figure 6 displays the return autocorrelations for the two series. The top panel displays autocorrelations for returns on the baseline simulation and the weekly CRSP index. They reveal the common result of very low autocorrelations in returns. The lower panel in figure 6 reports autocorrelations for absolute returns. Positive correlations in absolute returns continue out to one year for both the simulation and actual return series. Persistence in the CRSP series is slightly smaller at the lower lag lengths.

3.2.2 Annual Series

This section turns to the longer run properties of the simulation generated time series. For comparisons, the annual data collected by Shiller are used. Figure 7 presents both the S&P price/dividend (P/D) ratio, and the price/dividend ratio from the simulation. The simulation contains only one fundamental for the stock, and it can also be viewed as an earnings series with a 100 percent payout. This figure shows the simulation giving reasonable movements around the fundamental with some large swings above and below as in the actual data. The actual series is truncated to yield a better scale on the two plots. Its maximum value at the top of the dot com bubble is near 90.

Quantitative levels for these long range features are presented in table 3. The first two rows give the mean and standard deviation for the P/E and P/D ratios at annual frequency. The first two columns show a generally good alignment between the simulation and the annual P/E ratios. The P/D ratio from the actual series is slightly more volatile with an annual standard deviation of over 12. Deviations from fundamentals are very persistent, and these are displayed in all three series by the large first order autocorrelation. Again, the simulation and the P/E ratio are comparable with values of 0.72 and 0.68 respectively. The P/D ratio is slightly larger with an autocorrelation of 0.93. The last three rows present the annual mean and standard deviations for the total real returns (inclusive of dividends) for the simulation and annual S&P data. The returns generate a real return of 12.4 percent as compared to 7.95 percent for the S&P. The mean log returns are given in the next row, and are also comparable between the simulation and data. The simulation gives an annual standard deviation of 0.26 which is large relative to the value 0.17 for the S&P. The last row reports the annual Sharpe ratio for the two series with the simulation showing a value of 0.44 which is larger than the 0.30 from the actual data.²³,

4 Market internals and robustness

4.1 Agents

This section will analyze some of the distributional features of agent wealth and how it moves across strategies. Figure 8 displays the wealth distributions over time for the entire 200,000 length simulation. Several interesting features emerge from this figure. First, the market is dominated by the buy and hold strategy. It controls almost 50 percent of the wealth. It is very interesting that there is still enough wealth controlled by the dynamic strategies to have an impact on pricing, even though they are only about 30 percent of the market. This emphasizes the importance of certain marginal types in price determination in a heterogeneous world. The adaptive strategies are generally ranked second, in terms of wealth, followed by the fundamental, and then a very

²³ For the simulation this is simply the annual return divided by the standard deviation. For the S&P the annual interest rate from the Shiller series is used in the standard estimate, $(r_e - r_f)/\sigma_e$.

small fraction of the noise traders. The ranking is relatively stable, but the fractions do exhibit some interesting dynamics over time. Fundamental strategy wealth is particularly volatile, and is generally counter cyclical to the adaptive strategies.

Wealth distributions across gain levels in forecasts, and volatility forecasts are as important as the actual strategy types. High gain forecasts are sensitive to recent moves in prices and convert small price changes into relatively large changes in their forecasts. Figure 9 presents histograms for wealth distributions across the five different gain levels for each of the forecast strategy types. Moving left to right goes from smallest (largest half-life) to largest (smallest half-life). The distributions are constructed from 100 snapshots taken off the market at different times. They represent the means across these 100 snapshots. The purpose of this is to get a better picture of the unconditional time averages on these densities, as opposed to the one time densities which may vary a lot over time. The patterns for the strategies are very interesting. The adaptive and fundamental forecasts support a wide range of gain parameters. Wealth is not drawn to any particular value, and the market is composed of both long a short horizon traders. Interestingly, the short AR strategies concentrate their regressions to using mostly low gain (long horizon) estimates. These simple linear forecasts appear to be functioning well in terms of selecting for low gain levels in their learning models.

Gain parameters are also part of volatility forecasting too. They control the impact of recent squared returns on forecast volatility estimates. Unlike the previous plot, these gain parameters are used in the same fashion by all three forecast families. Density plots are given for these in figure 10. The three panels again correspond to the different forecast families. The adaptive strategies support all gain levels, but there is some indication of a bias toward larger gain volatility models for both the fundamental and Short AR forecasts types. Since risk is an important part of the portfolio choice problem, these distributions are a key indicator of the underlying causes of market instability. This evidence suggests that a large amount of wealth is concentrated on strategies which put a large amount of weight on recent volatility when estimating risk.

Figure 11 shows how the strategies move with the stock price. The strategies are presented as the fraction of wealth invested in the risky asset. The top panel is a price snapshot from the baseline simulation run. The second panel displays the strategies for the adaptive and fundamental strategies. These are wealth weighted averages across the entire forecast family. The adaptive strategy moves with the price trends that it is designed to follow. As price moves up, it locks in on the trend, and often maxes out the portfolio to the risky asset. As a market crashes, these strategies quickly withdraw from the risky asset. The fundamental strategy is less precise in its behavior. It generally takes a strong position after a market fall, but not all the time. It can also take a strong position just before a fall.

There are two possible reasons for this hesitancy on the part of fundamental traders. First, while market crashes should drive up their conditional returns, they will also increase their risk

estimates. Confirmation of this is given in figure, 12. The top panel is again the price time series. The middle series are the portfolio fractions which correspond to fundamental strategies using the highest and lowest gain variance forecasts. The low gain forecasts are not sensitive to recent changes in volatility, and show expected strong portfolio swings on market declines. The high gain forecasts are very sensitive to recent changes in volatility, which diminishes the impact of their changes in conditional mean returns. In several cases their response to increases in volatility dominates, in that they reduce their exposure to the risky asset.

A second possible explanation for the loose connections between market crashes and the behavior of fundamental traders is the fact that their estimated models will weaken as the market goes into a bubble period. The lower panel in 12 shows the estimate of the coefficient from the regression of returns on the logged price dividend ratio. The figure displays a time series for this recursively estimated parameter both for the lowest and highest gain learners. The low gain learners display a reliably negative, and stable value for this parameter. This would yield a reversal strategy for fundamental traders. The high gain counterpart for this forecast family shows a fluctuating value which can occasionally moves toward zero as a bubble proceeds. This indicates that agents using relatively short series in these price/dividend regressions begin looking at sets of data which no longer contain evidence of a reversion of prices toward fundamentals. They have lost faith in the basic fundamental forecast weakening their stabilizing trading strategies.

4.2 Crashes

This section examines the dynamics of the market around extreme price declines or crashes. Much of the market behavior can be summarized as moving through slow expansions increasing well above fundamentals, followed by large and sudden price declines which move the market well below its fundamental value. Figure 13 displays the time dynamics of a short snapshot of the market. The top panel repeats the price time series for the market. The 3rd panel displays information on the total market trading volume each period. Volume comoves with prices in interesting ways. Market crashes are usually followed by large increases in trading volume. Also, as bubbles increase, trading volume slowly drops off, often reaching a local minimum just before a market crash. The 4th panel displays the wealth weighted conditional variance estimate across strategies. This moves as expected with sharp, and persistent increases after large market declines. Similar to trading volume, volatility, or market risk perception is often at a local minimum near the top of a bubble.

The 2nd panel is the most interesting in its connections to market dynamics. Since the market demand curve is well defined numerically in the simulation, one can estimate the demand elasticity. Magnitudes will not be a major concern here, but the sign will be. Negative values indicate well behaved downward sloping demand curves, but positive values indicate that the demand curve has, locally, a positive slope at the current equilibrium price. Obviously, this will contribute

to an unstable dynamic. Though the market often stays comfortably in the negative region, it can swing positive occasionally. This is often at or near periods of extreme market instability.

An intuitive picture of what happens to demand curves is presented in figure 14. Two excess demand curves are displayed. One is during normal times, and one is just after a large price decline. The crash period demand is shows a kink which could easily be related to market instabilities. For example, imagine the curve shifting slightly left with the cusp just moving off the point where excess demand is zero. This would cause a sudden downward shift in the market clearing price. Given the set of strategies in the market, it is not surprising that the market might be unstable in some periods. The simple adaptive forecasts are a major force in market instability since they will increase demands on price rises, and decrease on price falls. The behavior of aggregate demand will depend critically on their fraction in the population.

One part of the dynamics which probably contributes to some of the instability is the interaction between agents' risk perception and crash dynamics. As a bubble continues, risk perception, measured as estimated variance, falls (figure 13). In the aggregate this will push agents to more aggressive investments which can contribute to instability through the following channel. Overall market demand for shares is given by,

$$S_t = \sum_{i=1}^I \frac{\alpha_i(P_t)(S_{t-1,i}(P_t + D_t) + B_{t-1,i})}{P_t}. \quad (33)$$

Which can be split into two parts,

$$S_t = \sum_{i=1}^I \alpha_i(P_t)S_{t-1,i} + \sum_{i=1}^I \alpha_i(P_t) \frac{(S_{t-1,i}D_t + B_{t-1,i})}{P_t}. \quad (34)$$

The actual shape of the demand curve is complicated and depends on the distribution of wealth across various strategies represented by $\alpha_i(P_t)$. However, the two parts of the demand give an intuitive picture of why, as strategies move closer to $\alpha_i(P_t) = 1$, demand will more likely be in the unstable region. The second part of the demand function includes both a component based on $\alpha_i(P_t)$ which will vary according to the specific strategy, and a kind of "rebalancing" component, depending on $B_{t-1,i}/P_t$ which will generate a well behaved downward component in market demand. This represents the natural adjustment investors make with a fall in price for a given α strategy. To maintain the same portfolio composition, a price fall necessitates new share purchases, and buying pressure comes into the falling market. This is a kind of mechanical stabilizing force. As α gets closer to zero this force is a smaller part of the market. It is also important to note that if one allowed leverage in the model this is where it would enter. Leverage would allow for α to be greater than 1. In this case the mechanical "rebalancing" would shift from a stabilizing to a destabilizing force since $B_{t-1,i}$ would be negative. The nice feature is that the unstable selling

pressure driven by leverage can be viewed as part of a continuum in the general portfolio decision making process. This is why leverage per se is not necessary in this market to generate instability. However, it is likely that it would magnify the extreme moves in the market.²⁴

The analysis of figure 13 is continued more formally in figure 15. This picture analyzes the dynamics of the market near large price declines. A large decline is defined as a return strictly less than the 0.005 quantile of the return distribution using the last 50,000 weeks in the simulation. This point is dated as a crash, and the next 10 weeks after this are skipped. This is obviously an imprecise measure, but it is relatively simple and effective. The upper left panel displays the dynamics of volume around a crash. The dashed line indicates the unconditional mean. Trading volume is low, but rising before the crash, and then hits a large upward spike which leads to a period of persistent and high volume for up to 100 weeks after the crash.

The upper right panel performs the same experiment with the aggregate, wealth weighted, trading strategy, recorded as the fraction of wealth in the equity market. As expected it shows a sudden and large drop when the crash hits, and slowly starts to increase, though it is still well below its mean value even after 100 weeks have gone by. There is also some indication these portfolio strategies begin dropping well ahead of the crash.

The lower left panel displays the coefficient for the price/dividend regression. It moves in a generally expected way, in that it is above its mean before the crash, and the drops dramatically as the recent high P/D ratio is confirmed in terms of low conditional expected returns by the sudden drop in price. Its actual dynamics is a little unusual, because it starts falling before the crash, and reaches a minimum about 52 weeks after the crash. This latter feature makes some sense since the crash P/D level doesn't enter into the P/D regression until 52 weeks after a crash has gone by. It's unusual dynamics may be related to the fact that agents are running long range (52 week) regressions in their models.

The lower right panel shows the dynamics of the elasticity near a crash. It sweeps through zero, and moves into the positive, unstable range in a period near the crash. It is interesting that it is well above its unconditional mean in both the 100 weeks before and after the crash. An important question is whether the elasticity sign change actually leads a crash. The figure draws a line at the point where the sign changes. This turns out to be 3 weeks prior to the crash. It is not clear whether in reality this would be enough information to forecast an eminent period of market instability, and further analysis will be necessary. One tricky aspect in all four of these panels is the dating of crises. Imprecise dating may lead to some spurious indications of crash predictability for some indicators. It is also important to realize that aside from trading volume, most of the indicators are not observable in actual market data.

²⁴See Thurner, Farmer & Geanakoplos (2010) for a model which requires leverage for market instability.

4.3 Robustness checks

Analysis of the wealth and strategy compositions in the market suggests that high gain learning algorithms are difficult to eliminate, and their presence is important for market instability. These learners seem like they could potentially be irrational, because they rely on relatively short time series in their forecasts. Figure 16 tests this conjecture by presenting some comparisons of the time series forecast performance for conditional variances for all rules split across the 5 different gain levels used in the variance forecast. If there were no predictability in variances, then the minimum gain forecasts should do the best. The two upper panels display the mean squared error (MSE) and mean absolute error (MAE) for the 5 variance forecast groups. Forecasts are normalized by the unconditional value, so a value of 1 corresponds to using the unconditional variance. The high gain, low half-life, forecast shows forecast improvements near 15 percent. The appeal of high gain variance forecasts in our populations is therefore driven by the empirical features of the time series. The lower left figure looks at an estimate of expected utility levels for the different variance forecasts. It is reported as an annual certainty equivalent return. Again, the indication of the usefulness of high gain strategies is evident in the figure since there is no dramatic pull toward small gain variance estimates. There are relatively balanced utility levels across all the forecasts.²⁵

Since high gain learners appear to be contributing to market instability it would be interesting to see if they alone could generate reasonable market dynamics. In the first experiment of section three, low gain only learners were shown to converge to a reasonable market equilibrium. What would happen if only high gain learners were present? To explore this, a reverse of the first benchmark, low gain, experiment is performed. Agents with only high gain, or short half-life rules are used. The set of rules is again concentrated on 5 different gain levels, but instead of being distributed between 1 and 50 years, they are reduced to a range of 1-5 years only. Figure 17 displays a 100 year period of a run with these gain parameters. The market still displays significant instability. However, the dynamics do not appear reasonable for lining up with real data. There are quick bursts in the price level which suddenly take off, and crash almost as quickly. After the sharp drop in prices, prices return relatively quickly to a central P/D level. Returns are punctuated by large tail events, but prolonged periods of high volatility are not evident. The competing dynamics of learning agents of all gain levels would appear essential to spread out some of the market instability, making it less dramatic, and more persistent in all dimensions. This is a critical requirement if one is interested in modeling deviations from fundamentals which are not just large, but are persistent.

The second experiment looks at the impact of the adaptive learning component of the model. This will be turned off so that agents stay with the strategy that they start out with. Wealth moves

²⁵The reader should be cautioned that this plot must be subject to a large amount of noise. The key point here is that there will be no dramatic selection drawing wealth to any forecast, and leaving variance forecasts at all gain levels active in terms of wealth.

only due to the relative performance of strategies over time. In this case the simulation shows a pattern similar to that from the original runs in terms of qualitative performance. The results are summarized in figure 18 which shows all the usual features of a standard financial returns series.²⁶

The final robustness check examines the importance of the short AR trading strategy. This strategy forms short term forecasts using a simple linear regression on lagged returns. It appears to be somewhat inconsequential in terms of wealth accumulation. Also, it does not appear to be a foundational strategy in terms of contributing to overall price dynamics. Figure 19 reports the results of eliminating this single strategy, and they are quite dramatic. The returns exhibit somewhat strange and spiky behavior. The density plot looks similar to the previous cases, but the autocorrelation patterns are different in a very important way. Now the returns series shows strong positive autocorrelation out to almost 5 lags. The values are on the order of 0.2 at lag 1 which is very unusual for a financial series. This shows that while the short AR strategy is only a small part of wealth, it is still working hard to make sure no obvious trading patterns appear in the returns series. It works hard to drive short range linear predictability to near zero, and in doing so it almost puts itself out of business. However, it always remains available, and ready to arbitrage away any kind of predictability that might appear. It is interesting that this strategy actually fits well into our traditional assumptions about the dynamics of strategies in an efficient market. This does not appear to be working as well for longer term strategies in this market.

5 Conclusions

This paper has presented results from a new agent-based financial market. It is argued that this model can play the role of a useful benchmark for experiments in agent-based finance. It is designed to bridge the gap between complex “many type” models, with many pieces and parameters, and the simpler “few type” models, with relatively few strategies.

The model is shown to be rich enough to meet the hurdle of generating most of the basic stylized facts of asset returns and trading volume. It does this in a way that is much more amenable to detailed analysis than for some of the larger more complex models used in the past. However, it maintains a rich evolutionary flavor which stays close to the spirit of evolutionary finance and economics. Further, it connects to the important time series dimension of learning models, and their perception of the past. Much of the observed dynamics comes from the fact that traders putting relatively large weights on the recent past are not easy to remove from the population in an evolutionary struggle for survival.

The results show several interesting features about the set of agents surviving in the market. First, the buy and hold strategy controls a large fraction of wealth, though it is not crucial in actual

²⁶More detailed comparisons of these runs to those with active learning activated is an interesting question left for future experiments.

price setting. Second, the adaptive and fundamental strategies maintain a large fraction of high gain learners who are using only recent data in their forecasting updates. For the adaptive strategies, this would correspond to short range momentum strategies who follow recent trends. For the fundamental types, it means they are weighing recent events heavily in their dividend/price ratio forecast regressions. In terms of volatility estimates, all the strategies, put heavy weight on the recent past. This seems unusual, and may drive the intensity of the market's sudden price drops. Further work will try to examine the contribution of this specific form of short-memory in the model, and where it is coming from.

Agent-based markets offer an important technology for exploring conjectures about evolution and rationality in finance. They allow for computational experiments which can reveal the underlying dynamics in a world of heterogeneous and learning agents. Understanding the dynamics of these markets as thought experiments is necessary for building up our intuition for what is going on in real markets, and influencing better policy choices.

Appendix: Price search mechanism

The artificial market is cleared each period by matching the demand for shares with the supply. Given that agents have chosen their forecasts, and these forecast rules are fixed in period t , the demand for shares is given as in equation 26,

$$Z_{t,i}(P_t) = (1 - \lambda)\alpha_{t,i}(P_t) \frac{(P_t + D_t)S_{t-1,i} + B_{t-1,i}}{P_t}. \quad (35)$$

It is important to remember that $Z_t(P_t)$ includes changes in demand that recognize both new portfolio decisions that come from maintaining optimal portfolio fractions at the new price level, and also changes that come from modifying the optimal portfolio fractions given the new forecasts consistent with P_t . This will require a numerical solution for the market clearing price.

The algorithm used is a simple binary search procedure which corresponds to the standard search method from computer science. The search is started in a range of prices around P_{t-1} .

References

- Adam, K. & Marcet, A. (2010a), Booms and busts in asset prices, Technical Report 2010-E-2, Bank of Japan, Tokyo, Japan.
- Adam, K. & Marcet, A. (2010b), Internal rationality, imperfect market knowledge and asset prices, Technical report, Mannheim University.
- Andersen, T. G., Bollerslev, T., Christoffersen, P. F. & Diebold, F. X. (2006), Volatility and correlation forecasting, in G. Elliott, C. W. J. Granger & A. Timmermann, eds, 'Handbook of Economic Forecasting', North Holland.
- Blume, L. & Easley, D. (1990), 'Evolution and market behavior', *Journal of Economic Theory* **58**, 9–40.
- Blume, L. & Easley, D. (2006), 'If you're so smart, Why aren't you rich? Belief selection in complete and incomplete markets', *Econometrica* **74**, 929–966.
- Bollerslev, T., Engle, R. F. & Nelson, D. B. (1995), ARCH models, in 'Handbook of Econometrics', Vol. 4, North-Holland, New York, NY.
- Branch, W. A. & Evans, G. W. (2011 forthcoming), 'Learning about risk and return: A simple model of bubbles and crashes', *American Economic Journal: Macroeconomics* .
- 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. (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. & Viceira, L. M. (1999), 'Consumption and portfolio decisions when expected returns are time varying', *Quarterly Journal of Economics* **114**, 433–495.
- Campbell, J. Y. & Viceira, L. M. (2002), *Strategic Asset Allocation*, Oxford University Press, Oxford, UK.
- Chiarella, C., Dieci, R. & He, X.-Z. (2009), Heterogeneity, market mechanisms, and asset price dynamics, in T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', Elsevier, USA, pp. 277–344.
- Dacorogna, M. M., Gencay, R., Muller, U. A., Olsen, R. B. & Pictet, O. V. (2001), *An Introduction to High-Frequency Finance*, Academic Press, San Diego, CA.
- Day, R. H. & Huang, W. H. (1990), 'Bulls, bears, and market sheep', *Journal of Economic Behavior and Organization* **14**, 299–330.

- 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.
- Diks, C. & van der Weide, R. (2005), 'Herding, a-synchronous updating and heterogeneity in memory in a cbs', *Journal of Economic Dynamics and Control* **29**, 741–763.
- Evans, G. W. & Honkapohja, S. (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.
- Evstigneev, I. V., Hens, T. & Schenk-Hoppe, K. R. (2006), 'Evolutionary stable stock markets', *Economic Theory* **27**, 449–468.
- Farmer, J. D. & Geanakoplos, J. (2008), 'The virtues and vices of equilibrium and the future of financial economics', *Complexity* **14**, 11–38.
- Figlewski, S. (1978), 'Market "efficiency" in a market with heterogeneous information', *Journal of Political Economy* **86**(4), 581–597.
- Frydman, R. & Goldberg, M. D. (2007), *Imperfect Knowledge Economics*, Princeton University Press, Princeton, New Jersey.
- Fuster, A., Laison, D. & Mendel, B. (2010), 'Natural expectations and macroeconomic fluctuations', *Journal of Economic Perspectives* **24**(4), 67–84.
- Gaunersdorfer, A. (2000), 'Endogenous fluctuations in a simple asset pricing model with heterogeneous agents', *Journal of Economic Dynamics and Control* **24**, 799–831.
- Genotte, G. & Leland, H. (1990), 'Market liquidity, hedging, and crashes', *American Economic Review* **80**, 999–1021.
- Georges, C. (2008), 'Staggered updating in an artificial market', *Journal of Economic Dynamics and Control* **32**, 2809–2825.
- Giovannini, A. & Weil, P. (1989), Risk aversion and intertemporal substitution in the capital asset pricing model, Technical Report 2824, National Bureau of Economic Research, Cambridge, MA.
- Hommes, C. H. (2006), Heterogeneous agent models in economics and finance, in K. L. Judd & L. Tesfatsion, eds, 'Handbook of Computational Economics', Elsevier.
- Hommes, C. H. (2010), The heterogeneous expectations hypothesis: Some evidence from the lab, Technical report, CeNDEF, University of Amsterdam.

- Huisman, R., Koedijk, K. G., Kool, C. J. M. & Palm, F. (2001), 'Tail-index estimates in small samples', *Journal of Business and Economic Statistics* **19**, 208–216.
- Kogan, L., Ross, S. A. & Wang, J. (2006), 'The price impact and survival of irrational traders', *Journal of Finance* **61**, 195–229.
- LeBaron, B. (2001), 'Evolution and time horizons in an agent based stock market', *Macroeconomic Dynamics* **5**(2), 225–254.
- LeBaron, B. (2006), Agent-based computational finance, in K. L. Judd & L. Tesfatsion, eds, 'Handbook of Computational Economics', Elsevier, pp. 1187–1233.
- LeBaron, B. (2007), Wealth evolution and distorted financial forecasts, Technical report, International Business School, Brandeis University.
- LeBaron, B. (2008), Robust properties of stock return tails, Technical report, International Business School, Brandeis University.
- LeBaron, B. (2011 forthcoming), 'Active and passive learning in agent-based financial markets', *Eastern Economic Journal* .
- Levy, M., Levy, H. & Solomon, S. (1994), 'A microscopic model of the stock market: cycles, booms, and crashes', *Economics Letters* **45**, 103–111.
- Lucas, R. E. J. (1978), 'Asset prices in an exchange economy', *Econometrica* **46**, 1429–45.
- 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.
- Merton, R. (1969), 'Lifetime portfolio selection under uncertainty: The continuous time case', *Review of Economics and Statistics* **51**, 247–257.
- Mitra, K. (2005), 'Is more data better?', *Journal of Economic Behavior and Organization* **56**, 263–272.
- Nelson, D. B. (1992), 'Filtering and forecasting with misspecified ARCH models I: Getting the right variance with the wrong model', *Journal of Econometrics* **52**, 61–90.
- Samuelson, P. (1969), 'Lifetime portfolio selection by dynamic stochastic programming', *Review of Economics and Statistics* **51**, 239–246.
- Sandroni, A. (2000), 'Do markets favor agents able to make accurate predictions?', *Econometrica* **68**, 1303–1341.

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

Shiller, R. J. (2000), *Irrational Exuberance*, Princeton University Press, Princeton, NJ.

Thurner, S., Dockner, E. J. & Gaunersdorfer, A. (2002), Asset price dynamics in a model of investors operating on different time horizons, Technical report, University of Vienna.

Thurner, S., Farmer, J. D. & Geanakoplos, J. (2010), Leverage causes fat tails and clustered volatility, Technical Report Cowles discussion paper 1745, Yale University, New Haven, CT.

Zeeman, E. (1974), 'On the unstable behavior of stock exchanges', *Journal of Mathematical Economics* **1**, 39–49.

Table 1: *Parameter Definitions*

Parameter	Value
d_g	0.0126
σ_d	0.12
r_f	0
γ	3.5
λ	0.0007
I	16000
J	4000
g_j	[1, 2.5, 7, 18, 50] years
g_L	50 years
g_u	[1, 2.5, 7, 18, 50] years
p_σ	0.5
L	5 percent/year
$[\alpha_L, \alpha_H]$	[0.05, 0.95]
σ_ϵ	0.02
M_{PD}	52 weeks
M_{AR}	3
h_j	[0.05, 0.15]

The annual standard deviation of dividend growth is set to the level from real dividends in Shiller's annual long range data set. The growth rate of log dividends, 0.0126 corresponds to an expected percentage change of $D_g = 0.02 = d_g + (1/2)\sigma_d^2$ in annual dividends. This corresponds to the long range value of 0.021 in the Shiller data.

Table 2: *Weekly Statistics*

	Baseline	CRSP VW Weekly	IBM Weekly
Mean (percentage)	0.22	0.17	0.26
Std	3.54	2.49	3.40
Skewness	-0.77	-0.70	-0.59
Kurtosis	19.63	8.97	12.76
Tail exponent (left)	2.12	3.14	2.98
Tail exponent (right)	2.25	3.43	3.93

Basic summary statistics. CRSP corresponds to the CRSP value weighted index nominal log returns with dividends, Jan 1926 - Dec 2009 with a sample length of 4455. IBM corresponds to nominal weekly log returns with dividends over the same period. The simulations use a sample length of 25,000, drawn after 75,000 weeks have gone by.

Table 3: *Annual Statistics*

	Baseline	Shiller Earnings	Shiller Dividends
Mean(P/D)	16.75	15.32	26.62
Std (P/D)	4.76	5.97	13.81
Autocorrelation(1)	0.72	0.68	0.93
Mean(Return)	12.4		7.95
Mean(Log(Return))	8.33		6.22
Std(Return)	0.26		0.17
Annual Sharpe	0.44		0.30

Baseline model uses a sample of 100,000 weeks or about 1900 years. The Shiller series are annual from 1872-2009. P/D refers to the price dividend ratio for the model and the Shiller dividends column. The earnings column uses the annual P/E ratio instead. All returns are real including dividends. The Sharpe ratio estimated from the Shiller annual data uses the 1 year interest rates from that series.

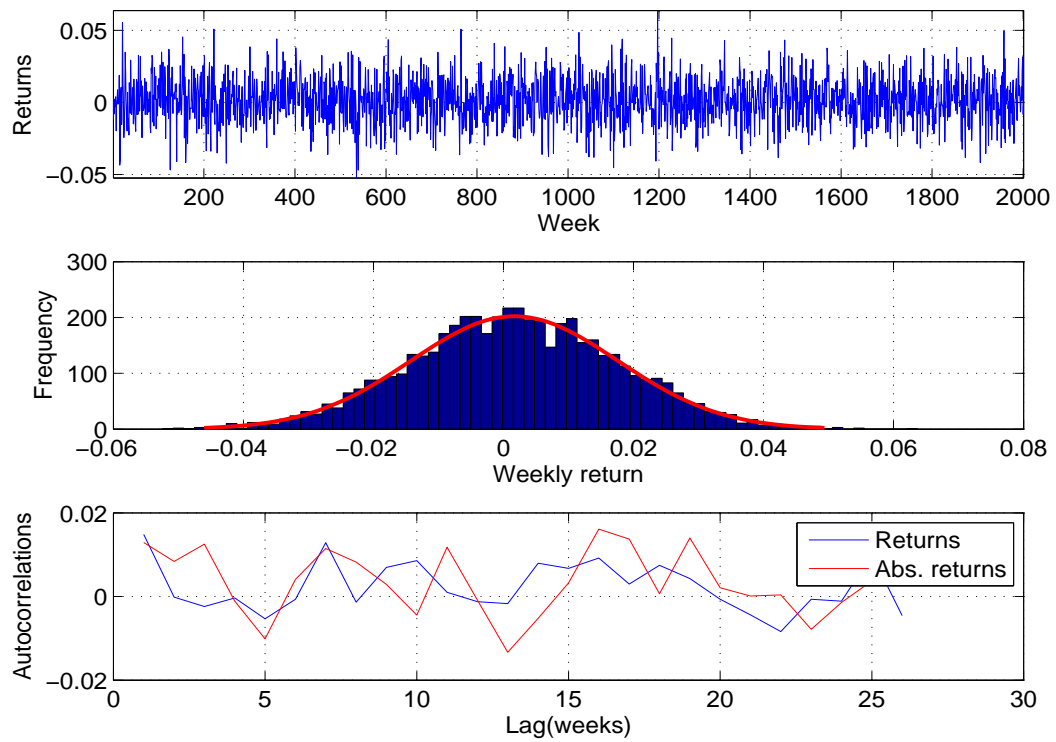


Figure 1: Summary stats: Low gain only

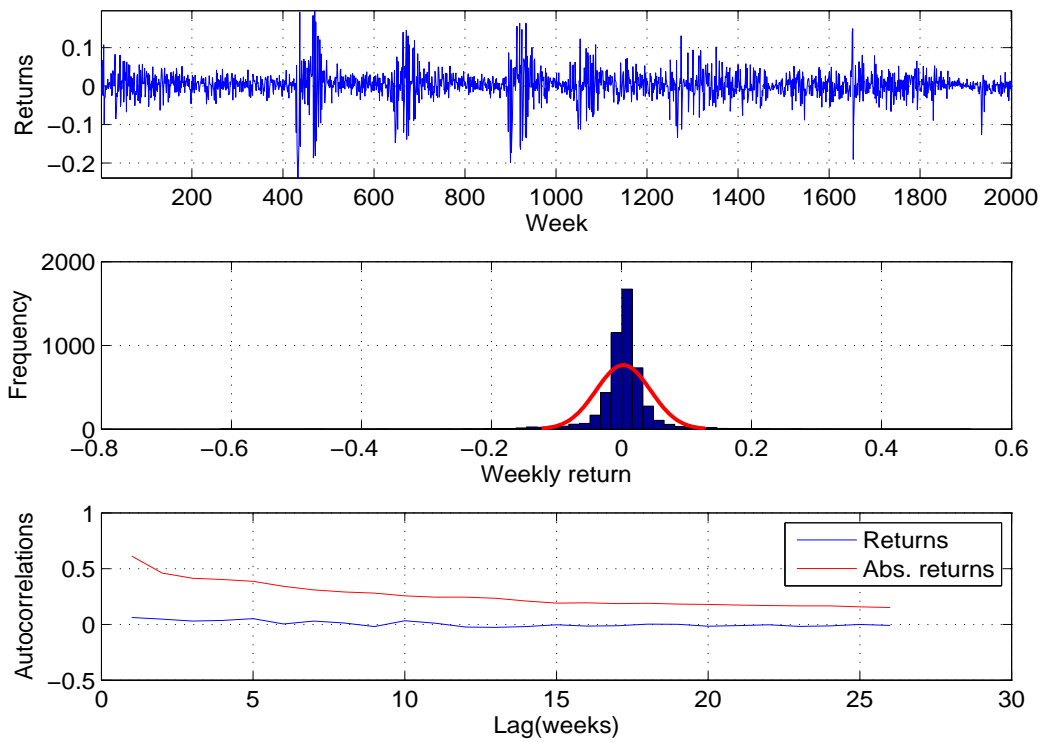


Figure 2: Summary stats: All gain baseline

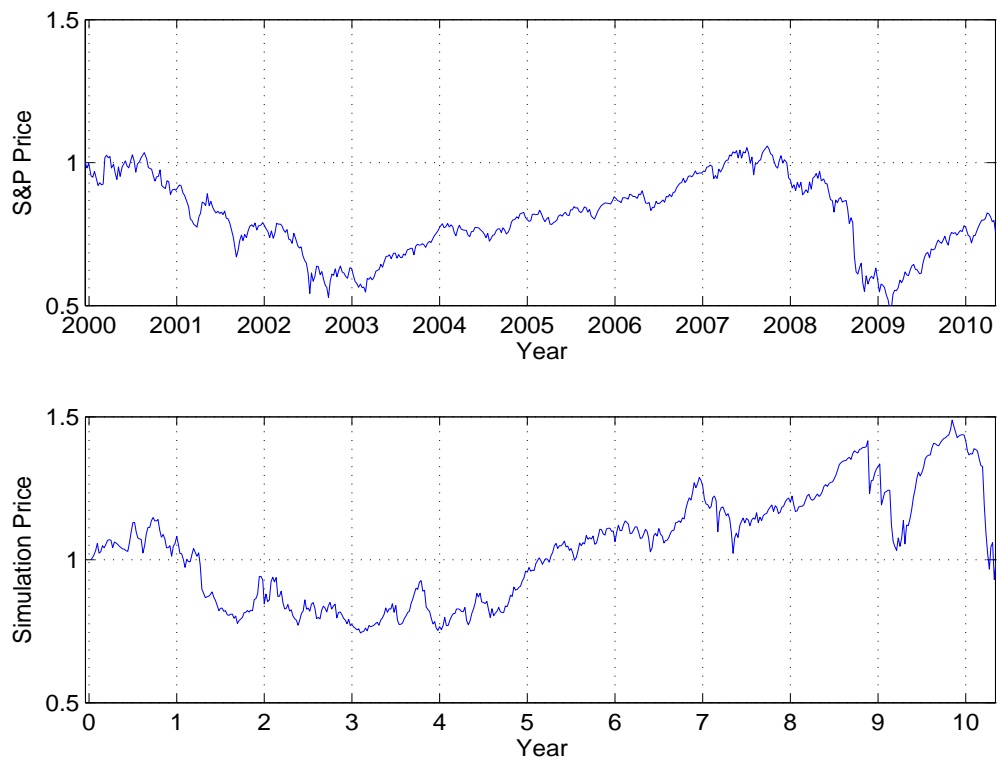


Figure 3: Price level comparison: S&P index versus simulation

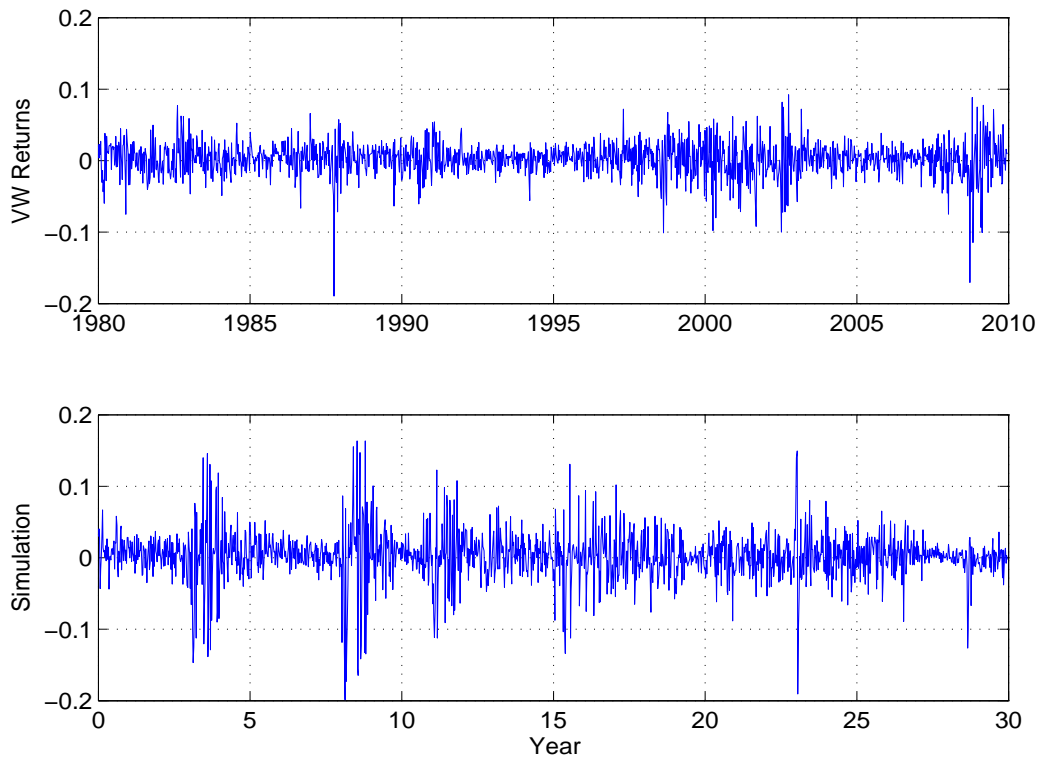


Figure 4: Return comparison: CRSP valued weith

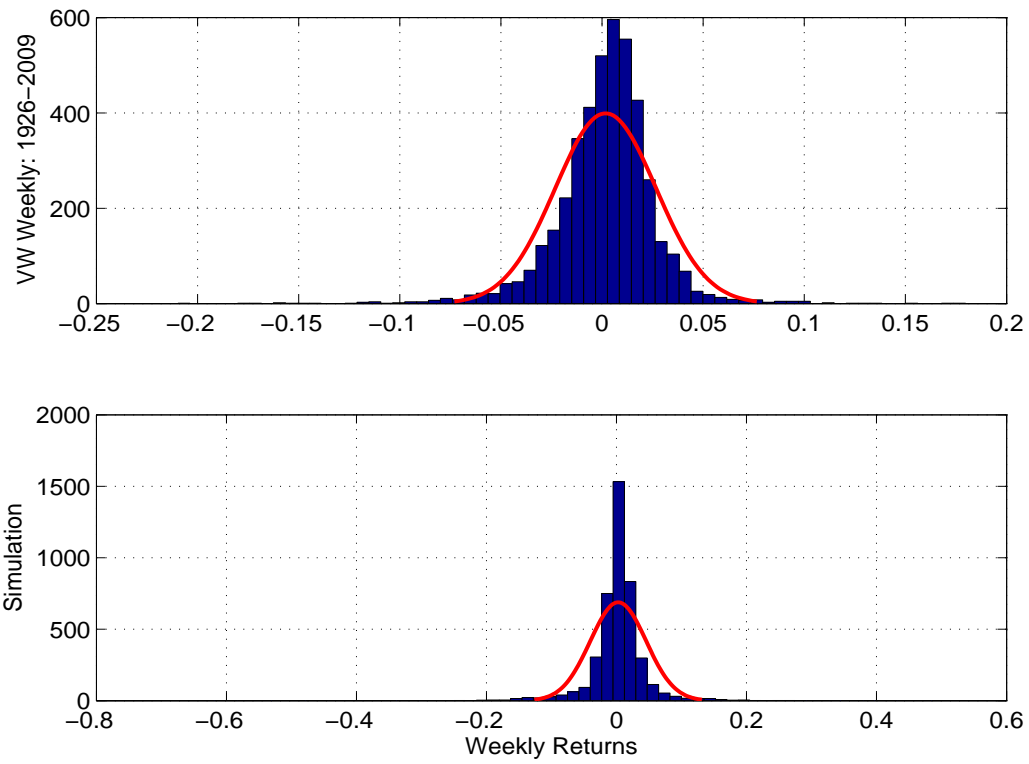


Figure 5: Weekly Return densities and Gaussian: CRSP VW Index 1926-2009

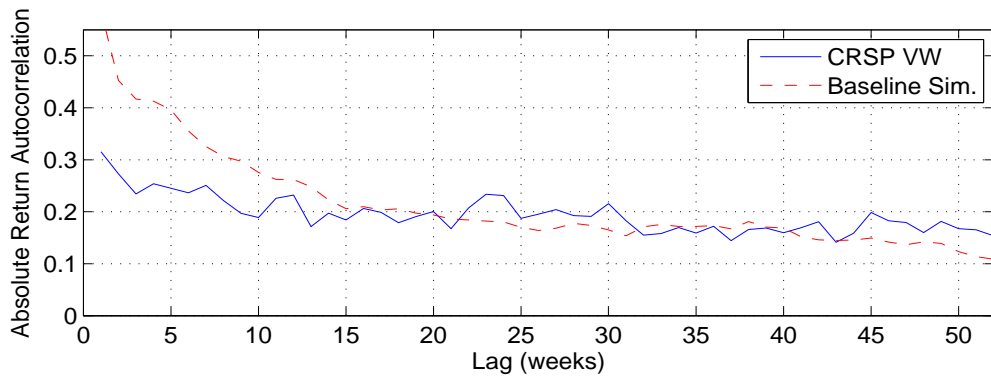
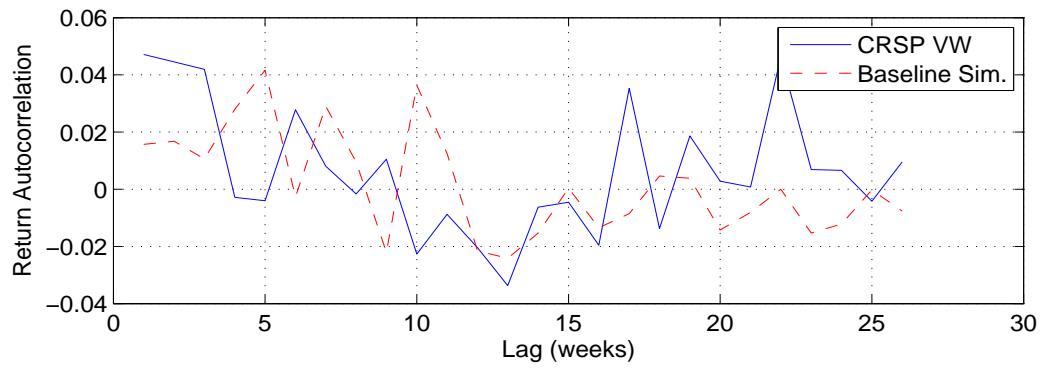


Figure 6: Return Autocorrelations: Returns and absolute value of returns

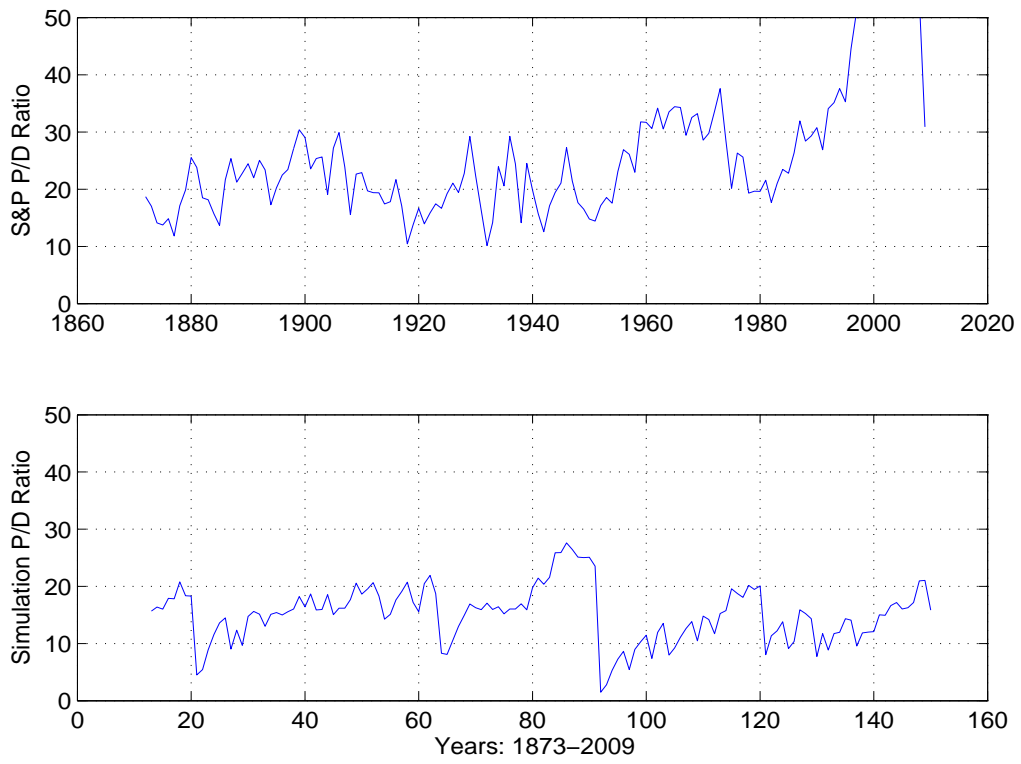


Figure 7: P/D ratios: Annual 1871-2009

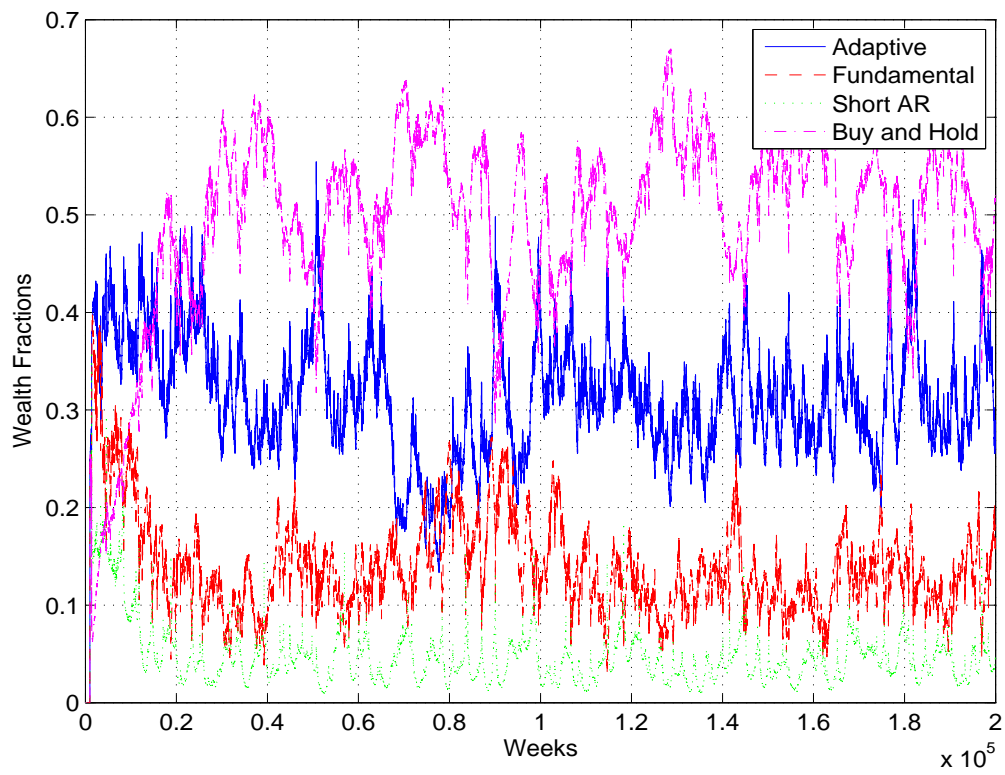


Figure 8: Wealth Fraction Time Series

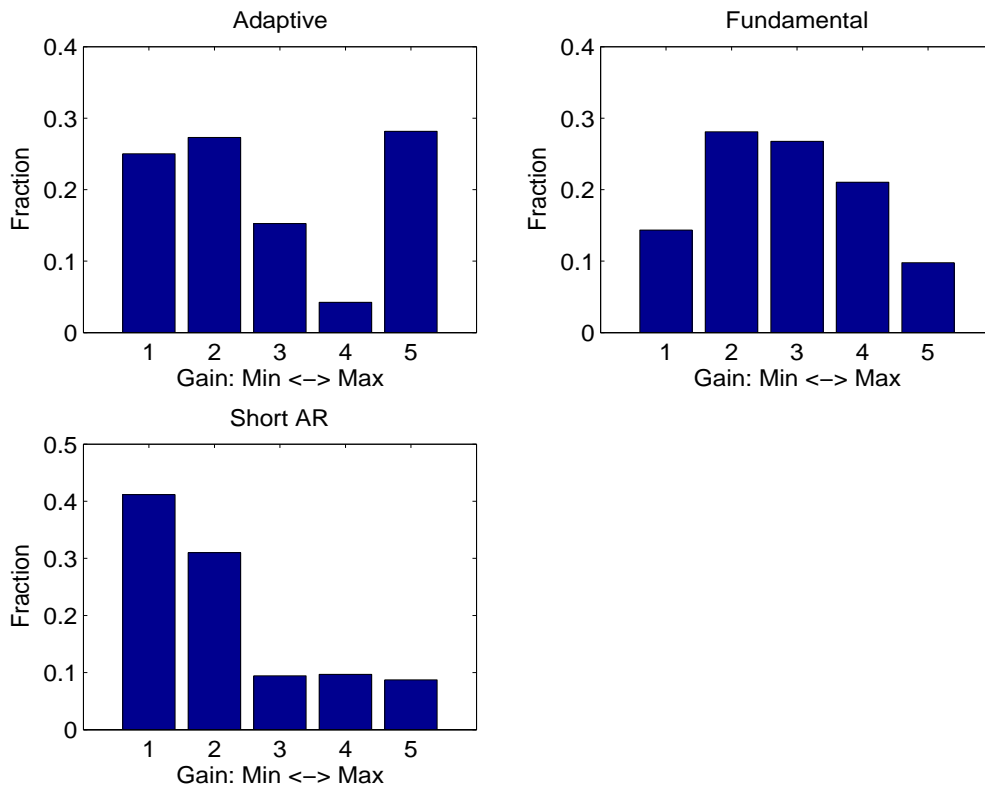


Figure 9: Gain Wealth Distributions

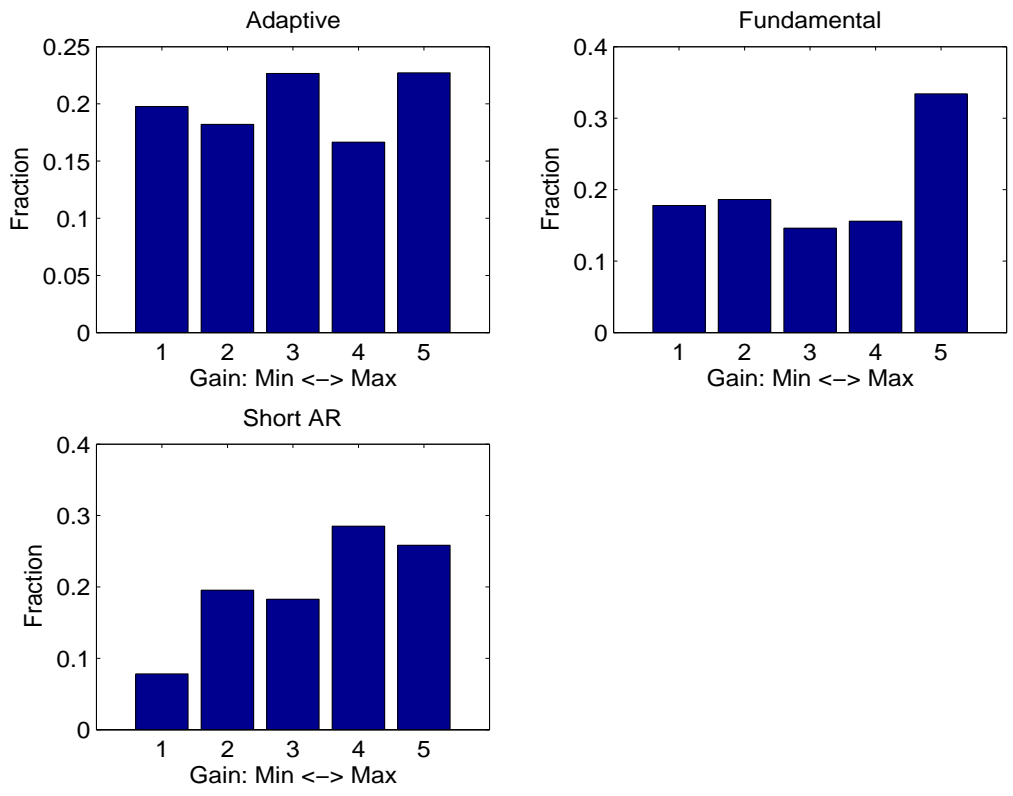


Figure 10: Volatility Gain Wealth Distributions

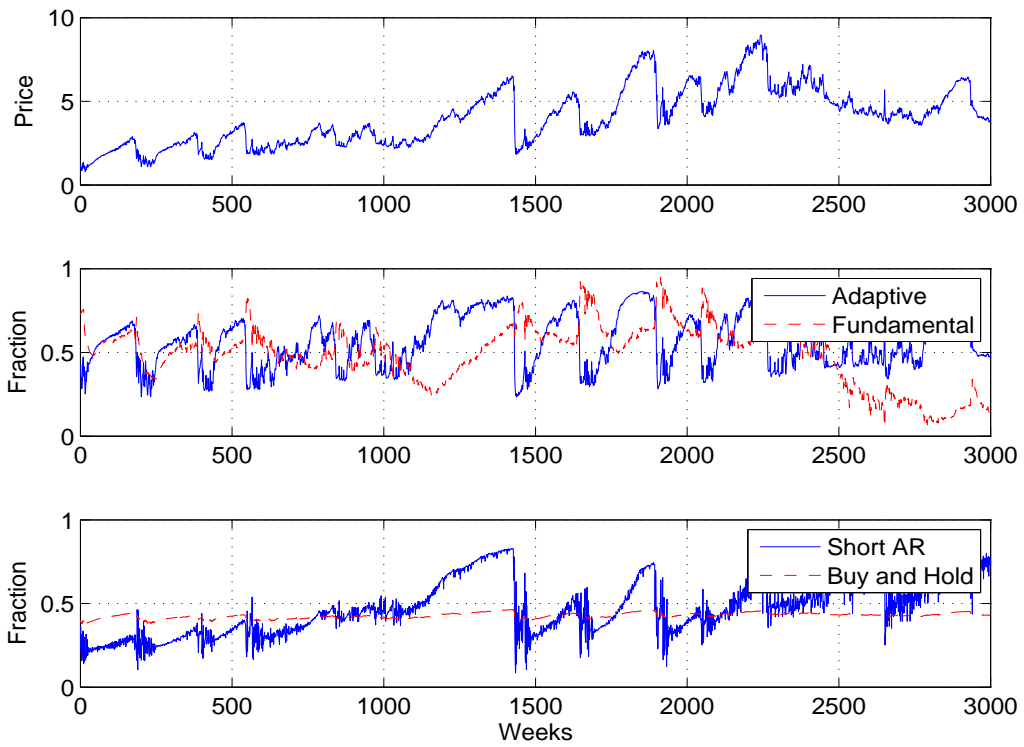


Figure 11: Strategy Fractions

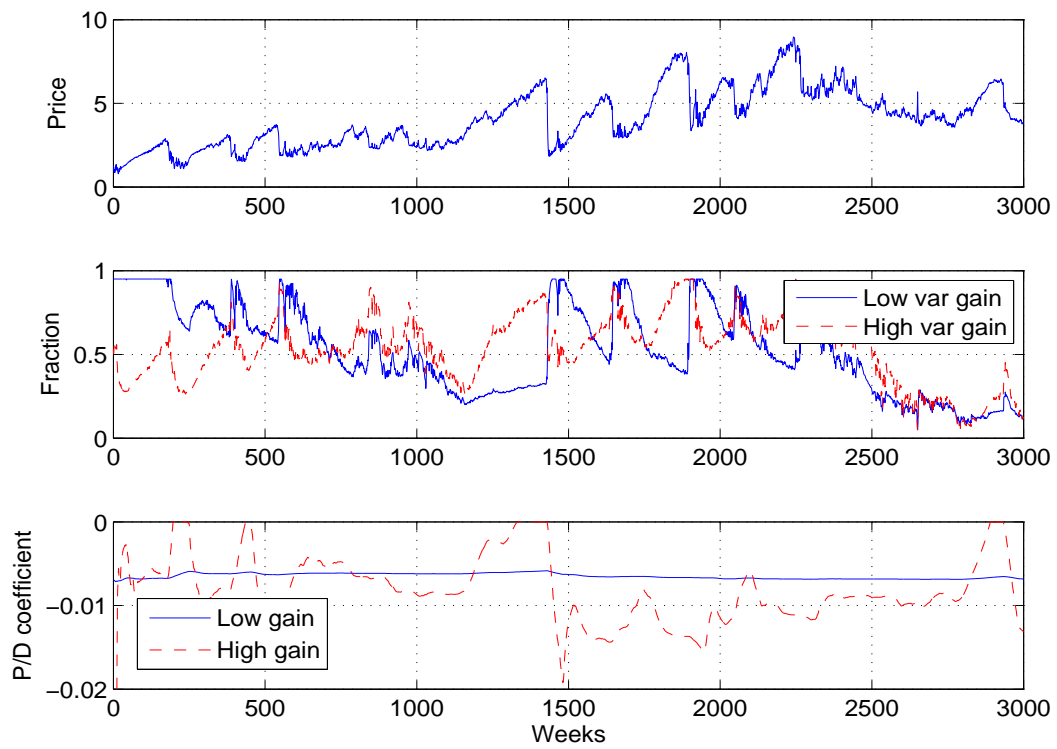


Figure 12: Strategy Fractions by Gain Levels

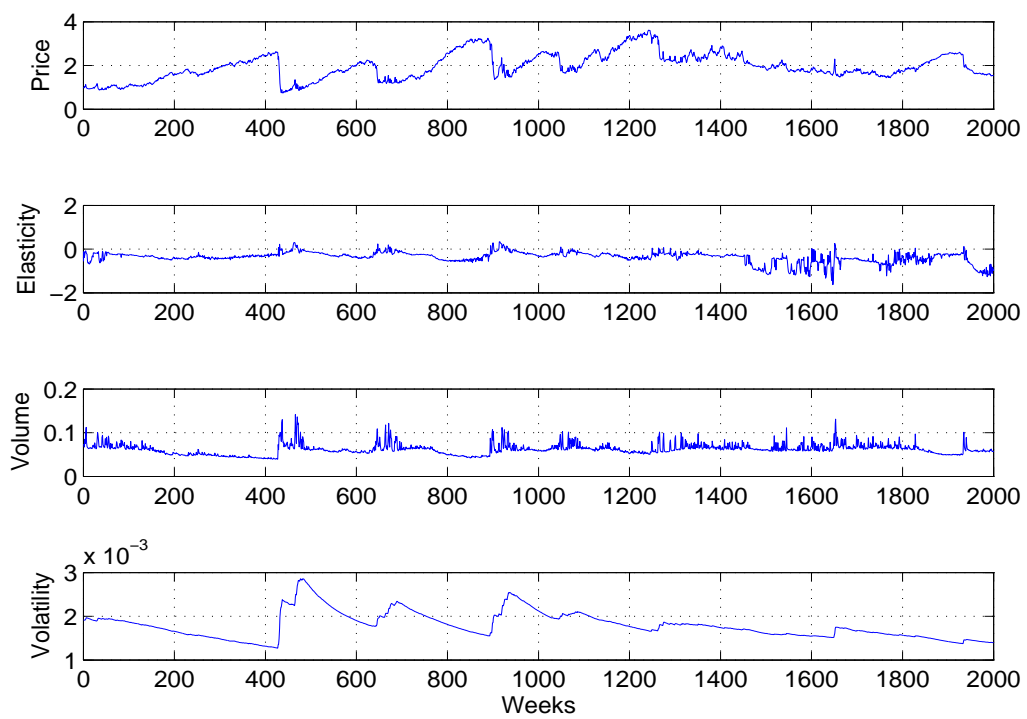


Figure 13: Elasticity, Volume, and Volatility

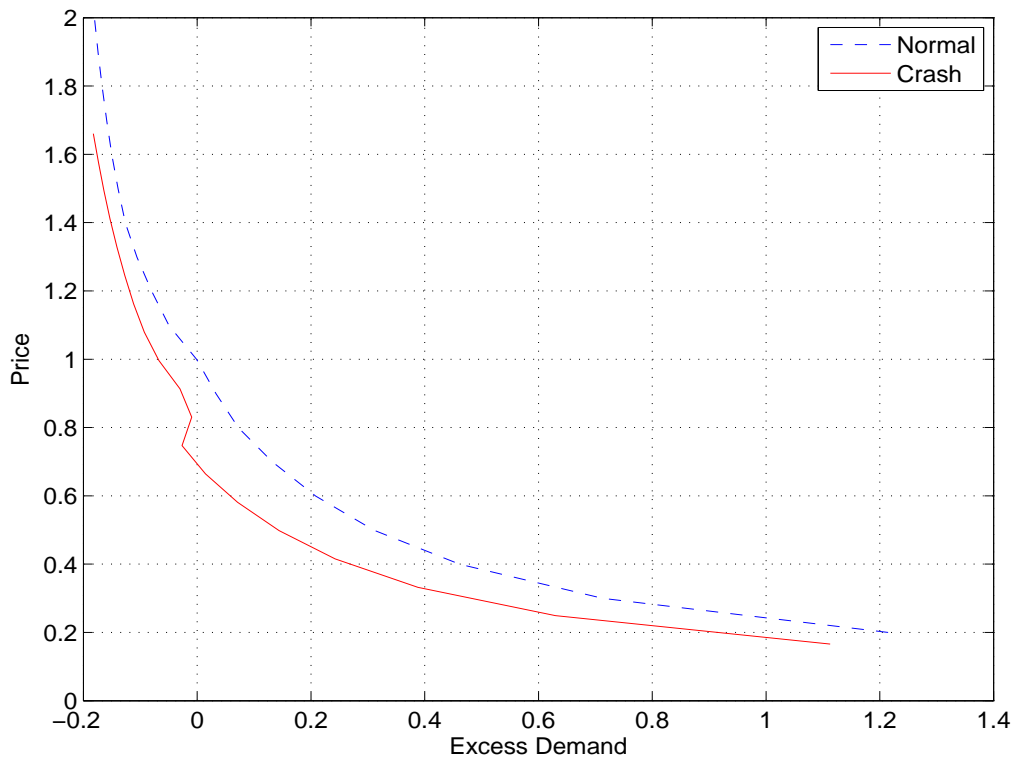


Figure 14: Excess Demand Curves: Normal and crash periods

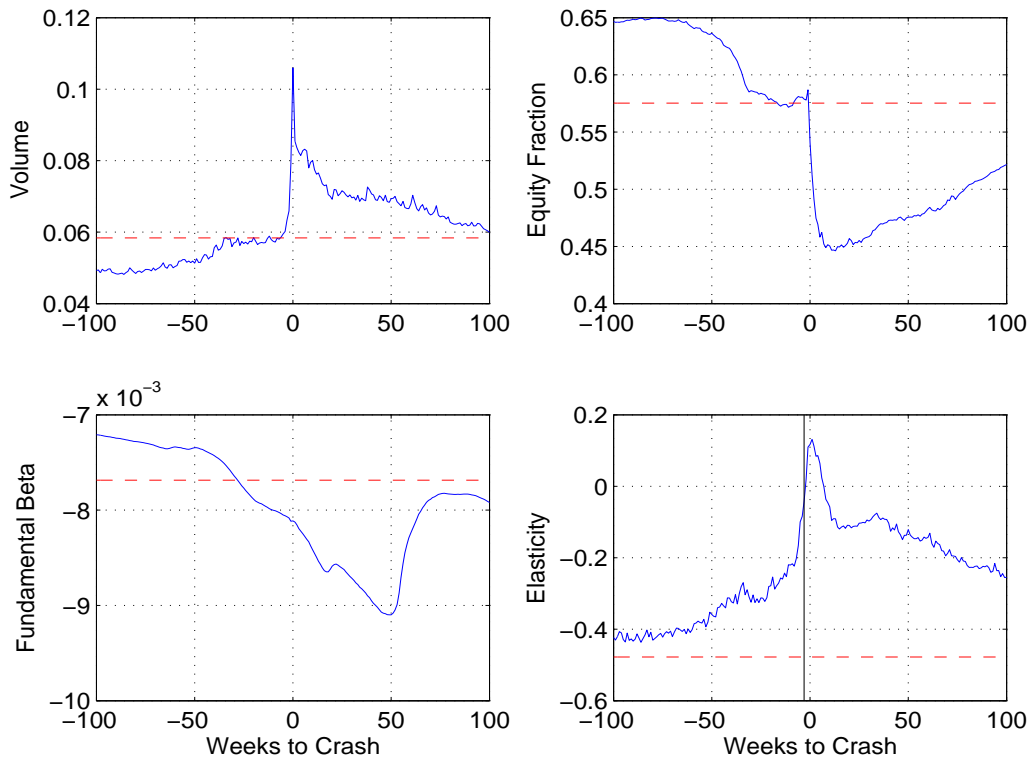


Figure 15: Dynamics Around Crashes

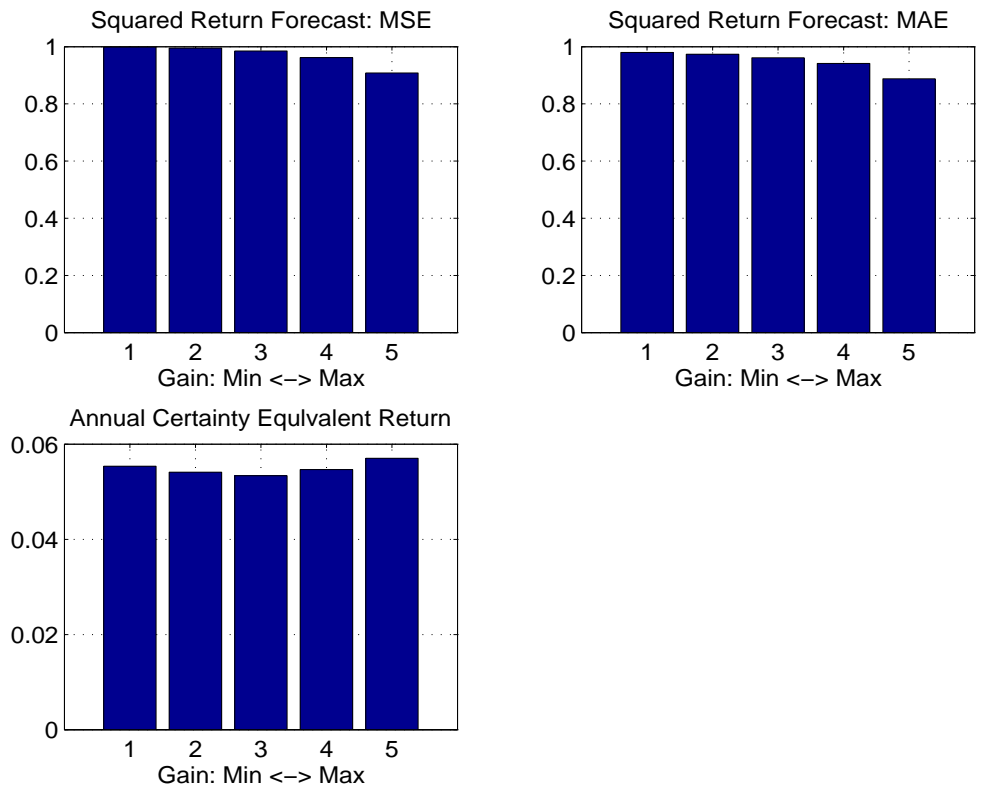


Figure 16: Strategy Forecasts




Figure 17: **High Gain Only: Short memory learning:** Simulation results for gain parameters in the range [1-5] years only.

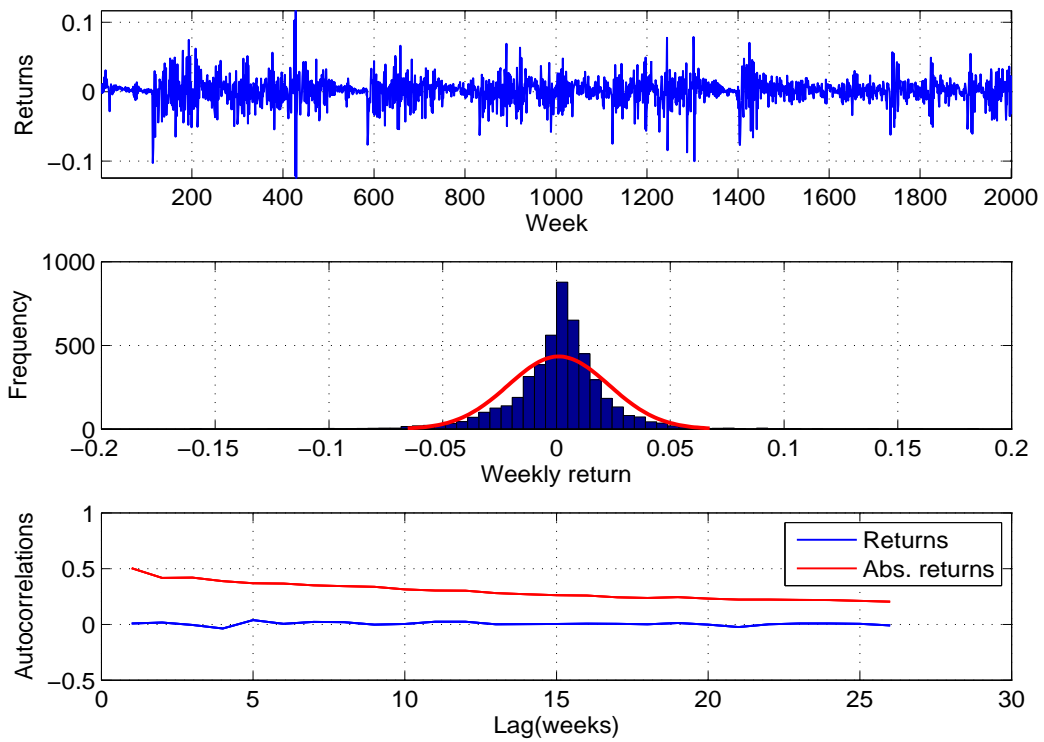


Figure 18: Summary stats: No adaptive learning

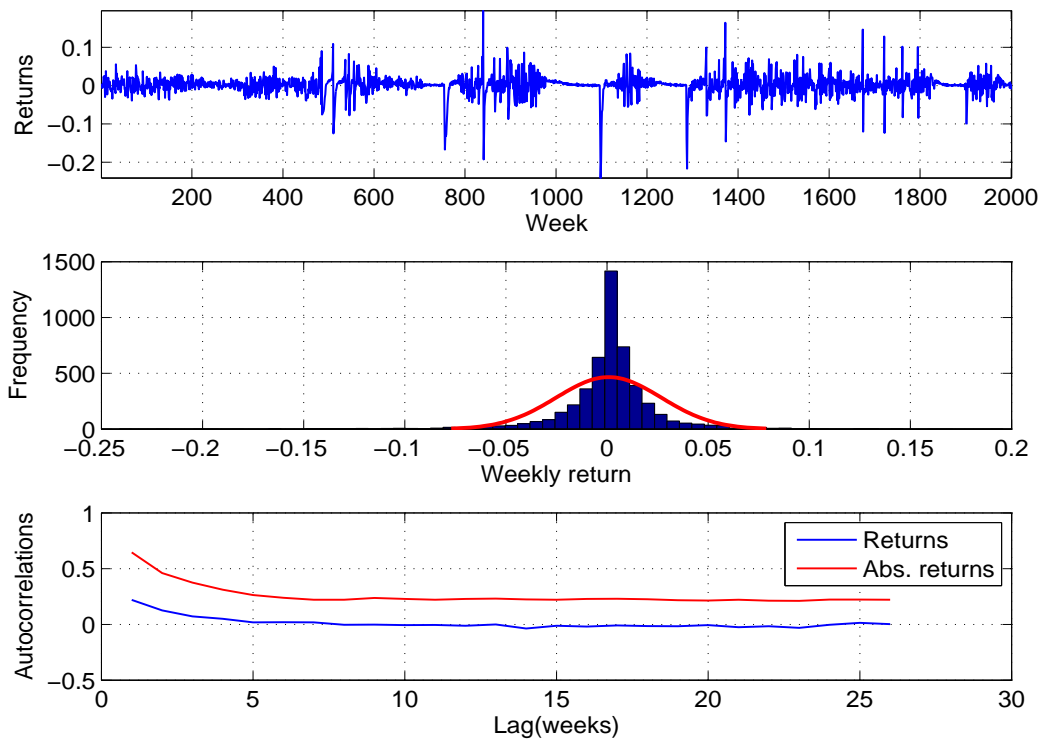


Figure 19: Wealth Fraction Time Series: No Short AR trading