Heterogeneous Gain Learning and the Dynamics of Asset Prices

Blake LeBaron *

International Business School

Brandeis University

June 2010: Preliminary

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

[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.

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

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. This includes adaptive, or momentum, strategies, two forms of mean reverting fundamental strategies, and a short range predictive liquidity strategy. These four families are designed to 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. At the moment, this simplification still seems absolutely necessary to maintain model tractability. It also gives a reasonable bound on a bubble. 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

[3] Earlier examples include Levy, Levy & Solomon (1994) and LeBaron (2001).

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) heteroscedastic. Furthermore, prices take large swings from a 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 to allowing only long memory learners who make decisions based only on long range results going into the distant past yields a very simple market which converges to a rational expectations equilibrium with returns which are independent, identically distributed (IID) Gaussians, displaying relatively stable deviations around fundamentals. These empirical summaries are presented in section 3.

The rest of the paper (section 4) examines the dynamics of the model around large falls in the asset price, or crashes. The market displays similar patterns near large market moves in terms of covariation with several internal variables. Patterns leading to a crash are documented, and compared with previous conjectures about the dynamics of bubbles and crashes. The evolutionary process over wealth that eventually leads to market instabilities is explored. Finally, section 5 will summarize, conclude and highlight future questions which can be addressed in this framework.

[4]Passive 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).

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

[5]Its 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.

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, noise traders are implemented who do not trade in a simple random fashion, but try to make money from simple short term return correlations. They base trading strategies on linear regressions run from week t to $t+1$. 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 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,

$$d_{t+1} = d_g + d_t + \epsilon_t, \tag{1}$$

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,

$$D_t = e^{d_t}. \tag{2}$$

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

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

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

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

shares is fixed, and set to unity,

$$\sum_{i=1}^{I} S_{t,i} = 1. \tag{4}$$

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 for each agent i,

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

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,

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

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

$$R^p_{t+1,i} = \alpha_{t,i} r_{t+1} + (1 - \alpha_{t,i}) r_f. \tag{8}$$

$\alpha_{t,i}$ 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,

$$\alpha_{t,i} = \frac{E_t^i(r_{t+1}) - r_f + \frac{1}{2}\sigma^2_{t,i}}{\gamma \sigma^2_{t,i}}. \tag{9}$$

All returns in this expression are in logs, and $\sigma^2_{t,i}$ is agent i's estimate of the conditional variance at time t. The derivation of this follows Campbell & Viceira (2002). It involves taking a Taylor series approximation for the log portfolio return. Details are given in Appendix B.

In the current version of the model neither leverage nor short sales are allowed. The fractional demand is restricted to $\alpha_{t,i}$ to $\alpha_L \leq \alpha_{t,i} \leq \alpha_H$. 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

$$W_{t+1,i} = (1 + R^p_{t+1})(1 - \lambda)W_{t,i}. \tag{10}$$

This also gives the current period budget constraint,

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

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.[8]

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

[8] See 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.

depending on the experiment.

All the forecasts will use long range forecasts of expected values using standard decreasing gain learning algorithms.

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

$$\bar{pd}_t = (1 - g_t)\bar{pd}_{t-1} + g_t pd_{t-1} \tag{13}$$

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

$$\bar{\sigma}_{pd,t}^2 = (1 - g_t)\bar{\sigma}_{pd,t-1}^2 + g_t(pd_t - \bar{pd}_t)^2 \tag{15}$$

$$g_t = \frac{1}{t} \tag{16}$$

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

The forecasts used will combine four linear forecasts drawn from well known forecast families.[9] The first of these is an adaptive linear forecast which corresponds to,

$$f_t^j = \bar{r}_t + m_r(f_{t-1}^j - \bar{r}_{t-1}) + g_j(r_t - f_{t-1}^j). \tag{17}$$

Forecasts of expected returns are dynamically adjusted based on the latest forecast error where $\bar{r}_t$ is the constant gain estimate of the unconditional mean return, and m_r is a common memory parameter used by all agents. This forecast format is simple and generic. It has roots connected to adaptive expectations, trend following technical trading, and also Kalman filtering.[10] In all these cases a forecast is updated given 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. The parameter m_r determines the speed with which the local forecast is

[9] This 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).

[10] A nice summary of the connections between Kalman filtering, adaptive expectations, and recursive least squares is given in Sargent (1999). The origins of this form of expectations for the stock market go back at least to Tinbergen (1939).

drawn back to the long range mean. Finally, this forecast will the 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 rules.[11]

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

$$f_t^j = \bar{r}_t + \beta_t^j (pd_t - \bar{pd}_t). \tag{18}$$

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.

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

$$f_t^j = \bar{r}_t + \sum_{k=1}^{L} \beta_{t,k}^j (r_{t-k+1} - \bar{r}_t) \tag{19}$$

This strategy plays the role of a kind of liquidity or noise trader in the market. However, this noise trader is purposeful in that he or she tries to continually monitor the current market situation, and places trades according to beliefs about local market impact and depth as represented by the short term linear return 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, $\bar{\sigma}_t^2$, 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

[11] Trimming will also take place in the regression equations below.

[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.

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.

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,

$$\tilde{r}_t = \frac{1}{M_{PD}} \sum_{j=1}^{M_{PD}} r_{t-j+1} \tag{20}$$

The fundamental regression is updated according to,

$$\begin{aligned} \beta_{t+1}^j &= \beta_t^j + \frac{g_j}{\bar{\sigma}_{pd,t}^2} pd_{t-M_{PD}} u_{t,j} \\ u_{t,j} &= (\tilde{r}_t - f_{j,t-M_{PD}}) \end{aligned} \tag{21}$$

The lagged return regression is approximated in the following updating equations. This approximation assumes that return correlations have a small impact on the estimated learning parameters.[14] For the lagged return regression this would be,

$$\begin{aligned} \beta_{t+1,k}^j &= \beta_{t,k}^j + \frac{g_j}{\bar{\sigma}_{r,t}^2} r_{t-k} u_{t,j}, \\ u_{t,j} &= (r_t - f_t^j) \end{aligned} \tag{22}$$

where g_j is again the critical gain parameter, and it varies across forecast rules. 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.

[14] This is done to keep the learning algorithm fast and simple. Further analysis will expand on this. This doesn't deviate philosophically from the other learning algorithms here, since they are all only an approximation to the unknown data generating 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 follow the adaptive expectations used for expected return forecasting. The updated conditional variance for each forecasting rule is given by,

$$\hat{\sigma}^2_{t+1,j} = \bar{\sigma^2_t} + m_\sigma(\hat{\sigma}^2_{t+1,j} - \bar{\sigma^2_t}) + g_{j,\sigma}(e^2_{t,j} - \hat{\sigma}^2_{t+1,j}) \tag{23}$$

$$e^2_{t,j} = (r_t - \bar{r_t})^2 \tag{24}$$

where $e^2_{t,j}$ is the squared forecast error at time t, for rule j.

The above conditional variance estimate is used for the first three rules.[16] There is no attempt to develop a wide range of heterogeneous variance estimators beyond differences in their gain parameters, $g_{j,\sigma}$. This reflects the fact that while there may be many ways to estimate a conditional variance, they often produce answers which are not that far from each other.[17] This forecast method has many useful characteristics as a benchmark forecast. First, it is essentially an adaptive expectations forecast on the 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. For $0 < m_\sigma < 1$ it is close to a GARCH model in terms of single period variance forecasts. The only difference is that the unconditional variance, $\bar{\sigma^2_t}$, is slowly changing over time, but if this is viewed as a constant, then this is a form of GARCH model.[18] For the case $m_\sigma = 1$, the model shifts to a simple exponential moving average of past squared forecast errors, as in the commonly used Riskmetrics framework. These two connections to common benchmark models for conditional variance forecasts continues in the spirit of the conditional return forecasts by staying close to well known and easily interpreted forecast tools.

[15] Several other agent-based market frameworks have explored the dynamics of risk and return forecasting. This includes Branch & Evans (2008) 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.

[16] As previously mentioned the buy and hold strategy uses the long run variance in its strategy.

[17] See Nelson (1992) for early work on this topic.

[18] See Bollerslev, Engle & Nelson (1995) or Andersen, Bollerslev, Christoffersen & Diebold (2005) for surveys of the large literature on volatility modeling.

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 agents to have a different time series perspective on returns and volatility.

2.6 Market Clearing

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

$$1 = \sum_{i=1}^{I} Z_{t,i}(P_t). \tag{25}$$

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

$$P_t Z_{t,i} = (1-\lambda)\alpha_{t,i}(P_t)W_{t,i}, \tag{26}$$

$$Z_{t,i}(P_t) = (1-\lambda)\alpha_{t,i}(P_t)\frac{(P_t + D_t)S_{t-1,i} + B_{t-1,i}}{P_t}. \tag{27}$$

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 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.

[19] A 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.

[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.

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

$$f_{t+1}^{j} = (1 - g_j) f_t^j + g_j e_{t+1}. \tag{28}$$

This easily maps to the simple exponential forecast rule,

$$f_t = \sum_{k=1}^{\infty} (1 - g_j)^k e_{t-k}. \tag{29}$$

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

$$\frac{1}{2} = (1 - g_j)^{m_h}, \tag{30}$$

or

$$g_j = 1 - 2^{-1/m_h}. \tag{31}$$

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. For most of the simulations used here these values will be chosen as 26 and 2500 weeks which corresponds to a range of 1/2 to 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.

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 will stay 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,

$$\hat{U}_{t,j} = (1 - g_u)\hat{U}_{t-1,j} + g_u U_{t,j}, \tag{32}$$

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

$$U_{t,j} = \frac{1}{1-\gamma}(1 + R^p_{t,j})^{(1-\gamma)} \tag{33}$$

with $R^p_{t,j}$ the portfolio holdings of rule j at time t. In the model agents seek to chose rules that maximize this value. They have already been simplified to have a common coefficient of relative risk aversion, and here they will be further simplified to have a common gain parameter for utility estimation, g_u.

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 polls a given number of forecast rules, L_N, and if any of these rules exceeds the estimated expected utility of the rule the agent is current using, then the agent switches to the new forecasting strategy.[21] One can therefore control the intensity of the learning process by adjusting two parameters, the fraction of agents adapting each period, and the number of rules they each search to find an improvement.

[21]This mechanism is similar to that used in LeBaron (2001). These choices correspond to the intensity of choice parameter used in many simulation models.

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 for all runs. This is a reasonable level for standard constant relative risk aversion.[22] The gain range for the learning models is set to 1/2-50 years in half-life values. This means that the 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 adaptive forecasts share a common memory which is set to $m_r = 0.995$ giving a slow convergence back to the unconditional mean. Finally, the variance memory is set to $m_\sigma = 1$. This reflects the fact that conditional variances in financial series are extremely persistent, and makes the volatility predictor a simple exponential filter of past squared errors.

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 $100,000$ weeks, or almost $2,000$ years. Statistics are drawn from the end of this simulation.

[22] Many of the results can be replicated for a range of γ from $2-4$. The value of 3 gives some of the most realistic looking series while still being a reasonable level.

3.2 Time series features

3.2.1 Weekly Series

Figure 1 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 really be said from the price figures alone. More detailed pictorial information is given in figure 2 which compares weekly return series from the CRSP value weighted index (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 3. 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. A corresponding length sample is drawn from the end of the baseline simulation for a comparable histogram. 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 out 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 though Dec 2009. All returns include dividend distributions. Mean returns are in weekly percentages. The simulation return level is slightly below the returns for the two market indices. However, this should be expected since the other returns are nominal, and are not adjusted for inflation. A quick back of the envelope adjustment for 3 percent inflation per year would increase the simulation returns to near 0.16.[23] The simulation lines up with with the VW index, and is slightly below the IBM returns.

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 2.4 which is close to the index, and slightly smaller than the individual returns of IBM. 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

[23] This comes from taking $0.03/52 * 100$.

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 $3 - 4$ are quite common for weekly asset return series, so the results here are all within reasonable ranges.[24]

Figure 4 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. It is well known that the magnitude of returns are persistent. Whether we call this volatility persistence, or correlations in squared or absolute returns, it is a common stylized fact for most asset return series. The lower panel in figure 4 shows that this positive correlation holds for the simulation. Positive correlations in returns continue out to one year which is also common for many financial series.[25] The CRSP series shows this too, but its persistence is still larger than the simulation.

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 5 presents both the S&P price/earnings (P/E) ratio, and the price/dividend ratio from the simulation. The simulation contains only one fundamental for the stock, and it can be viewed as an earnings series with a 100 percent payout. It could also be viewed as a dividend, so both series will be presented for comparison. This figure shows the simulation giving reasonable movements around the fundamental with some large swings above and below as in the actual data. There is some indication of some very large and dramatic drop offs in the simulations which appear much smoother in the data. The lower panel is repeated in figure 6 which replaces the upper panel with the S&P price dividend ratio. This appears a little smoother than the P/E ratio, but is still relatively volatile. It is dominated by the very dramatic

[24] These values give important information on higher order moment existence. They indicate that moments below 3 exists, but those above 3 may not exists. This calls into question whether reported estimates of kurtosis are reliable or meaningful.

[25] This is often conjectured to be well represented as a long memory or fractionally integrated process.

run up at the end of the century.

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 the 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 std. 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 0.76 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 7.08 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 0.18 which compares to a value of 0.17 for the S&P. The last row reports the annual Sharpe ratio for the two series.[26], yielding a smaller value than for the simulation.

4 Agents

This section will analyze some of the distributional features of agent wealth and how it moves across strategies. Figure 7 displays the wealth distributions over time for the entire $100,000$ length simulation. Several interesting features emerge from this figure. First, the market is dominated by the buy and hold strategy. It controls almost 70 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 small fraction of the noise traders. The ranking is relatively stable, but there there is a fair amount of

[26] 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, $(r_e - r_f)/\sigma_e$.

dynamics in all the fractions 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 8 presents histograms for wealth distributions across various deciles of the gain parameter for different strategies. 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 noise traders concentrate their regressions to using mostly low gain (long horizon) estimates.

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 9. The three panels again correspond to the different forecast families. The figure shows a similar pattern across all three types. Wealth has a strong bias toward high gain forecasts. This means that in terms of volatility, strategies which put a lot of weight on the recent past are dominating. Since risk is an important part of the portfolio choice problem, these distributions are a key indicator of the underlying causes of market instability.

Figure 10 shows how the strategies move with the stock price. The strategies are presented as the fraction of wealth to put 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 built to follow. As prices move 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.[27]

One possible reason for the hesitancy of the fundamental traders as the market falls, is that volatility is on the rise in exactly these periods. The risk averse fundamental traders are getting conflicting signals from rising expected returns and rising conditional variance. Some evidence in support of this is given by figure 11 that shows actual forecasts for the different strategies. This figure clearly shows that in most, but not all, major price declines, the fundamental strategy does correctly show an increased expected return. It is therefore likely that its inability to take a strong counter position at these times is due to expected risk.

5 Robustness checks

This section explores several modifications to the baseline simulation for comparison. In the first case the distribution of trader gain levels is modified. The high gain, or short memory, strategies are eliminated. Only strategies with half lives between 2000 and 2500 weeks are used. Also, the level of risk aversion is increased from 3 to 5.[28] Figure 12 shows the returns, histogram of returns, and autocorrelations for the returns and absolute returns in its three panels. The simulation generates a pattern of nearly independent, Gaussian returns. There is little evidence for correlation in either the returns or absolute returns series. Eliminating the short memory traders, has had a dramatic impact on the market in terms of not generating any of the interesting features contained in real financial series.

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 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 13 which shows all the usual features of a standard financial returns series.

[27] Some analysis suggests that this is due to the estimated regression coefficient eventually going to zero as a bubble persists. For larger gain regressions, the agents begin to lose empirical faith in their dividend/price regression model.

[28] At 3 this model converges to the maximum equity level for all the strategies. This outcome turns out to be unstable since the strategies are at their boundary. To compensate for this the level of risk aversion in increased.

The similarities continue in figure 14 which is comparable with the previous results from figure 7. The market is again dominated by the buy and hold strategy. The other strategies now take larger shares, the adaptive strategy is now close to 20 percent of total wealth. One difference is that the variability over time appears to be reduced as it should be since agents are not moving around across the different strategies.

The final roubustness check examines the importance of the noise trading strategy. This trial strategy forms short term forecasts using a simple linear regression on lagged returns. It appears to be somewhat inconsequential in terms of wealth accumulation. Figure 15 reports the results from eliminating this single strategy. The results 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 noise 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 is a kind of hard working simple momentum trader, who most of the time works to make itself almost disappear, but its presence is necessary for realism in the generated series.

6 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 a 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 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 A: 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 27,

$$Z_{t,i}(P_t) = (1 - \lambda)\alpha_{t,i}(P_t)\frac{(P_t + D_t)S_{t-1,i} + B_{t-1,i}}{P_t}. \qquad (34)$$

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}.

Appendix B: Portfolio choice

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

$$\max_{\alpha_{t,i}} \tfrac{E_t^i W_{t+1}^{1-\gamma}}{1-\gamma}, \tag{35}$$

$$st. \quad W_{t+1} = (1 + R_{t+1}^p)W_t \tag{36}$$

Dropping out constant values known at time t, this becomes,

$$\max \frac{1}{1-\gamma} E_t^i (1 + R_{t,i}^p)^{1-\gamma}. \tag{37}$$

If portfolio returns were log normal, this could be transformed. Unfortunately, portfolio returns are not log normal. Campbell & Viceira (2002) show using a Taylor series approximation that the log portfolio return is approximated by,

$$r_{p,t+1} = r_f + \alpha_t(r_{t+1} - r_f) + (1/2)\alpha_t(1-\alpha_t)\sigma_t^2 \tag{38}$$

where $r_{p,t} = \log(1 + R_t^p)$, $r_t = \log(1 + R_t)$, and $\sigma_t^2 = var(r_t)$. Assuming the return on the risky asset is log normal, then the approximate portfolio return is also log normal. This allows the use of the well known fact for log normal random variables that

$$\log(E(Y)) = E\log(Y) + (1/2)\sigma_y^2. \tag{39}$$

Returning to the maximization problem in equation 37, taking logs of the expectation, and using 39 we get,

$$\max \frac{1}{1-\gamma} \log(E_t^i (1 + R_{t,i}^p)^{1-\gamma}). \tag{40}$$

$$\max \frac{1}{1-\gamma} E_t^i (1-\gamma)\log(1 + R_{t+1}^p) + (1/2)(1-\gamma)^2\sigma_{r_p}^2 \tag{41}$$

$$\max E_t^i r_p + (1/2)(1-\gamma)\sigma_{r_p} \tag{42}$$

Using the approximation in 38 gives,

$$\max_{\alpha_t} r_f + \alpha_t(E^i_t r_{t+1} - r_f) + (1/2)\alpha_t(1-\alpha_t)\sigma_t^2 + (1/2)(1-\gamma)\alpha_t^2\sigma_t^2 \tag{43}$$

Dropping the constant, r_f, and solving for the optimal holding, α, gives

$$\alpha_{t,i} = \frac{E^i_t(r_{t+1}) - r_f + \sigma_t^2/2}{\gamma\sigma_t^2} \tag{44}$$

References

Adam, K. & Marcet, A. (2010), Booms and busts in asset prices, Technical Report 2010-E-2, Bank of Japan, Tokyo, Japan.

Andersen, T. G., Bollerslev, T., Christoffersen, P. F. & Diebold, F. X. (2005), Volatility forecasting, Technical report, NBER.

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. (2008), Learning about risk and return: A simple model of bubbles and crashes, Technical report, University of California - Irvine, Irvine, CA.

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.

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.

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.

Gaunersdorfer, A. (2000), 'Endogenous fluctuations in a simple asset pricing model with heterogeneous agents', *Journal of Economic Dynamics and Control* **24**, 799–831.

Gennotte, 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.

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 survial 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.

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.

Nelson, D. B. (1992), 'Filtering and forecasting with misspecified ARCH modelsI: 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.

Tinbergen, J. (1939), 'The dynamics of share-price formation', *Review of Economics and Statistics* **21**, 153–160.

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	35
λ	0.0005
I	16000
J	4000
g_j	$[1/2, 50]$ years
g_u	5 years
L	5 percent/year
L_N	1
$[\alpha_L, \alpha_H]$	$[0.05, 0.95]$
σ_ϵ	0.015
M_{PD}	52 weeks
m_r	0.995
m_σ	0.995
h_j	$[0.025, 0.15]$

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	Dow Weekly	CRSP VW Weekly
Mean (percentage)	0.11	0.17	0.25
Std	2.40	2.44	3.41
Skewness	-0.96	-0.70	-0.59
Kurtosis	17.67	8.97	12.76
Tail exponent (left)	2.47	3.14	2.98
Tail exponent (right)	2.38	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)	24.64	15.32	26.62
Std (P/D)	5.95	5.97	13.81
Autocorrelation(1)	0.76	0.68	0.93
Mean(Return)	7.08		7.95
Mean(Log(Return))	5.53		6.22
Std(Return)	0.18		0.17
Annual Sharpe	0.40		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.

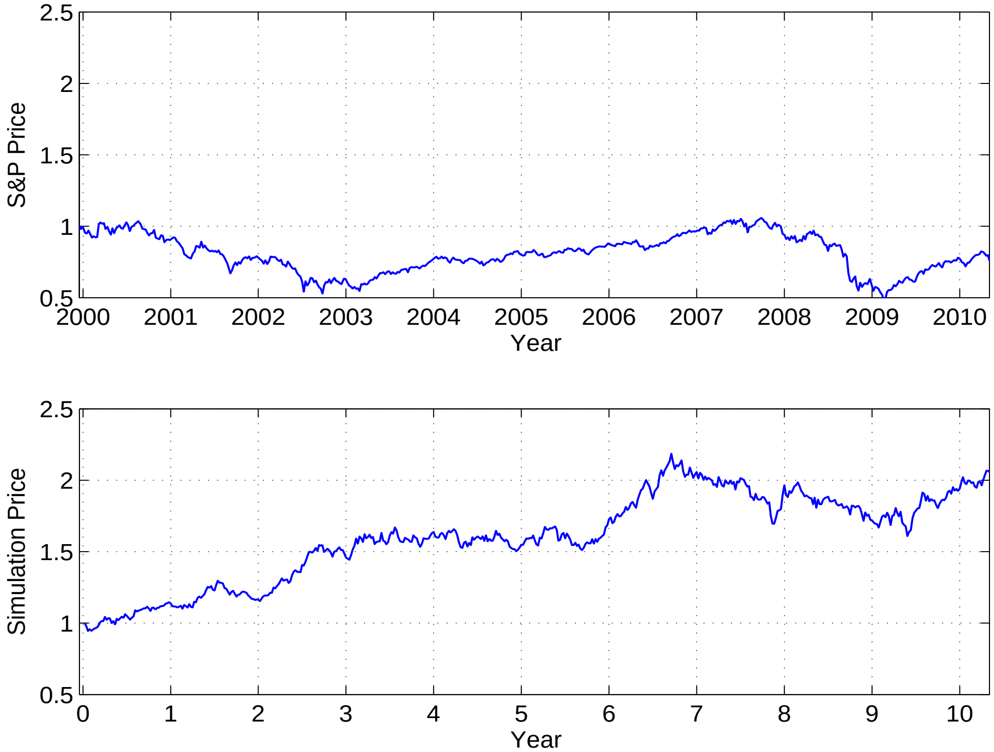

Figure 1: **Price level comparison**

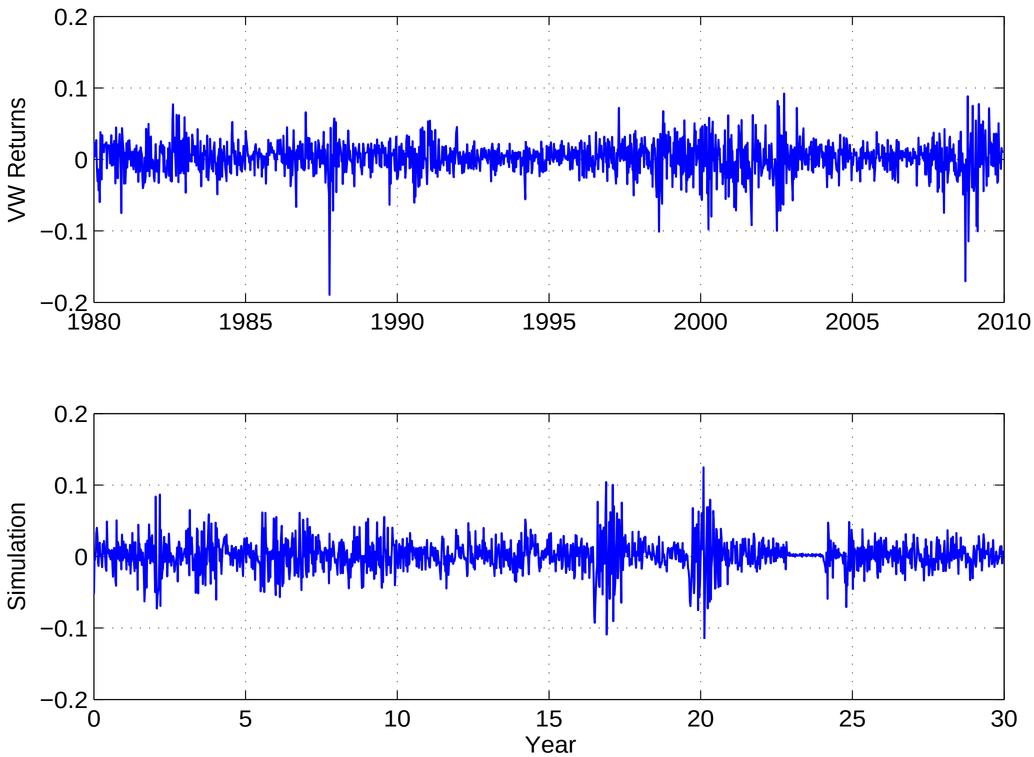

Figure 2: **Return comparison**

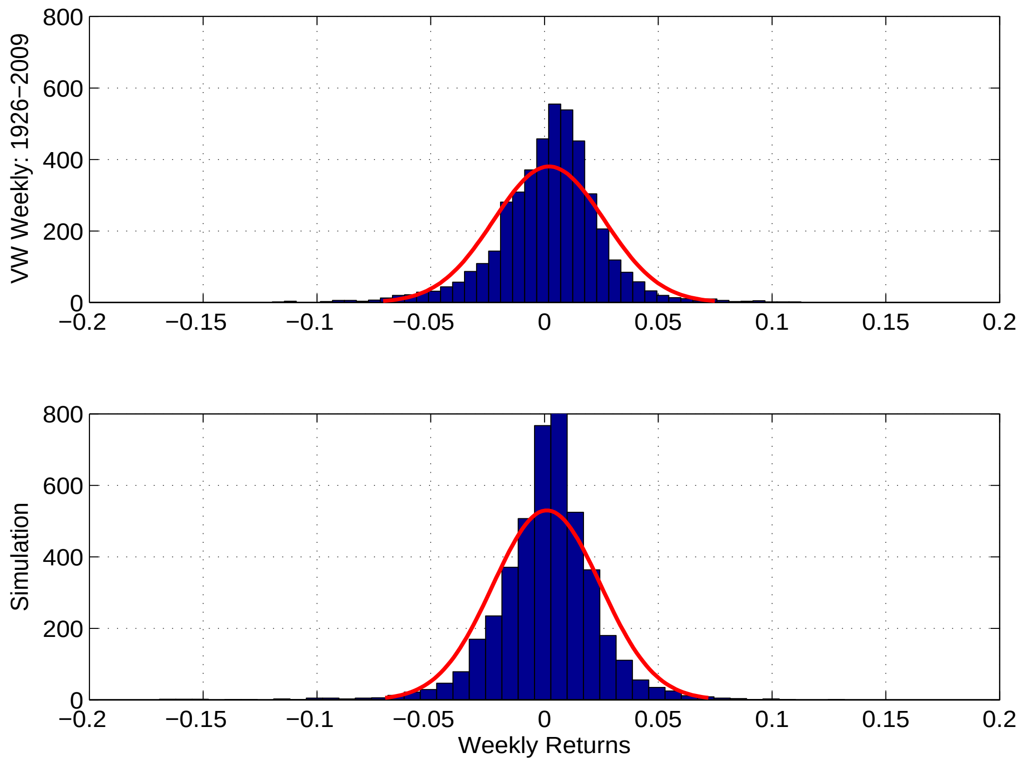

Figure 3: **Weekly Return densities and Gaussian: CRSP VW Index 1926-2009**

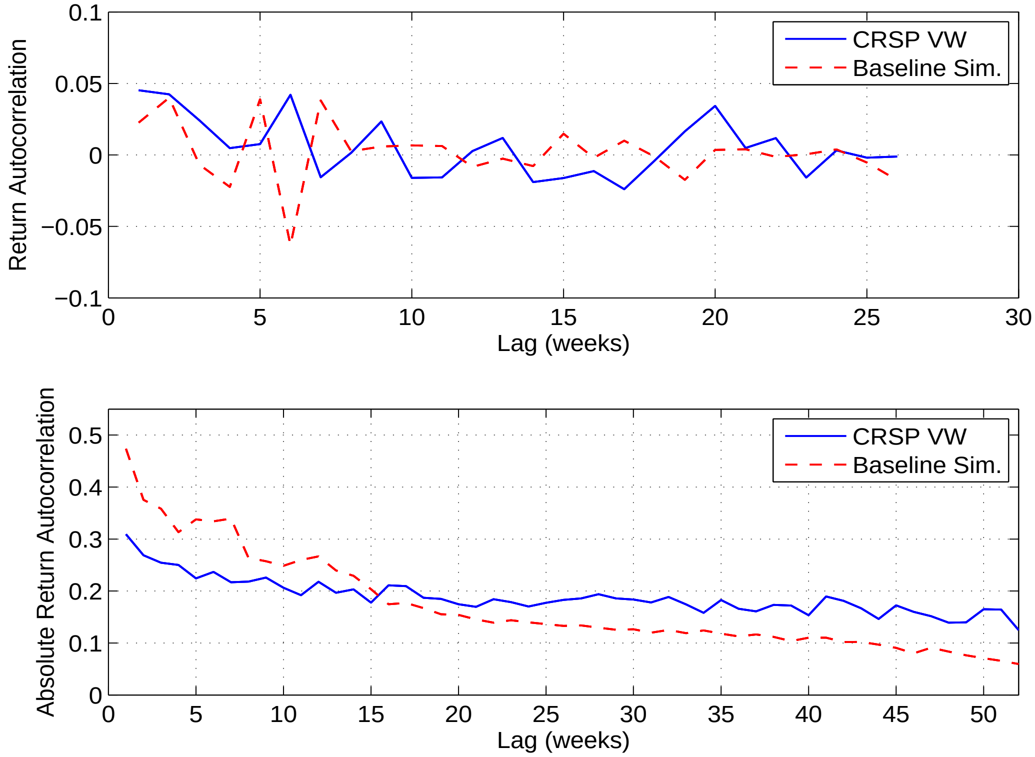

Figure 4: **Return Autocorrelations: Returns and absolute value of returns**

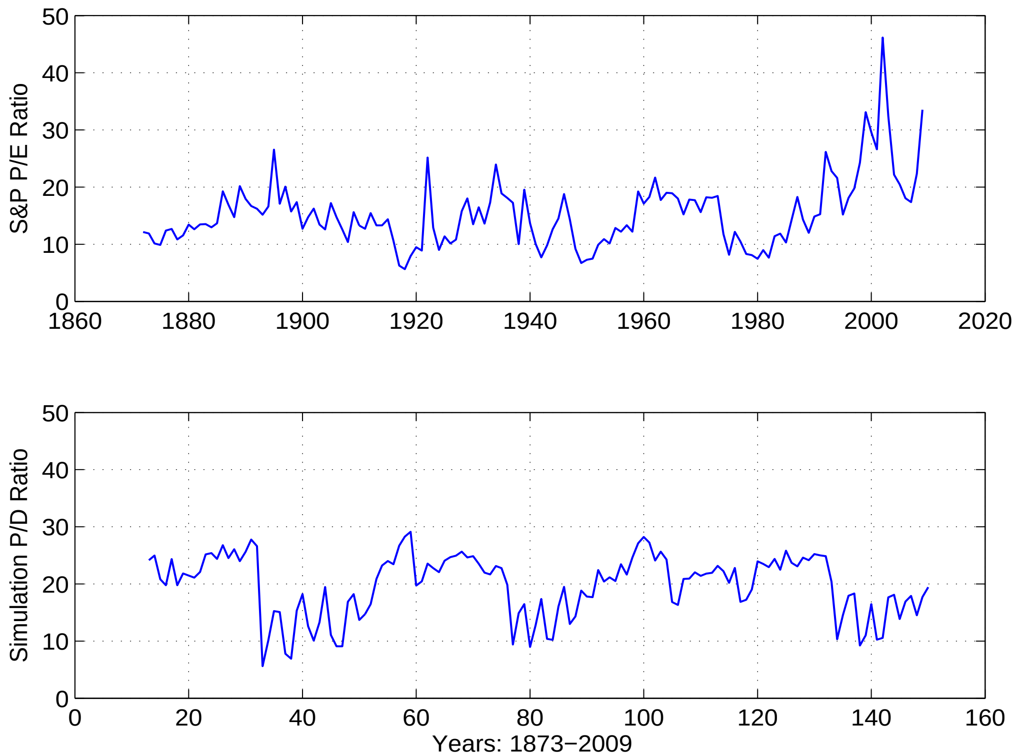

Figure 5: **P/E ratios: Annual 1871-2009**

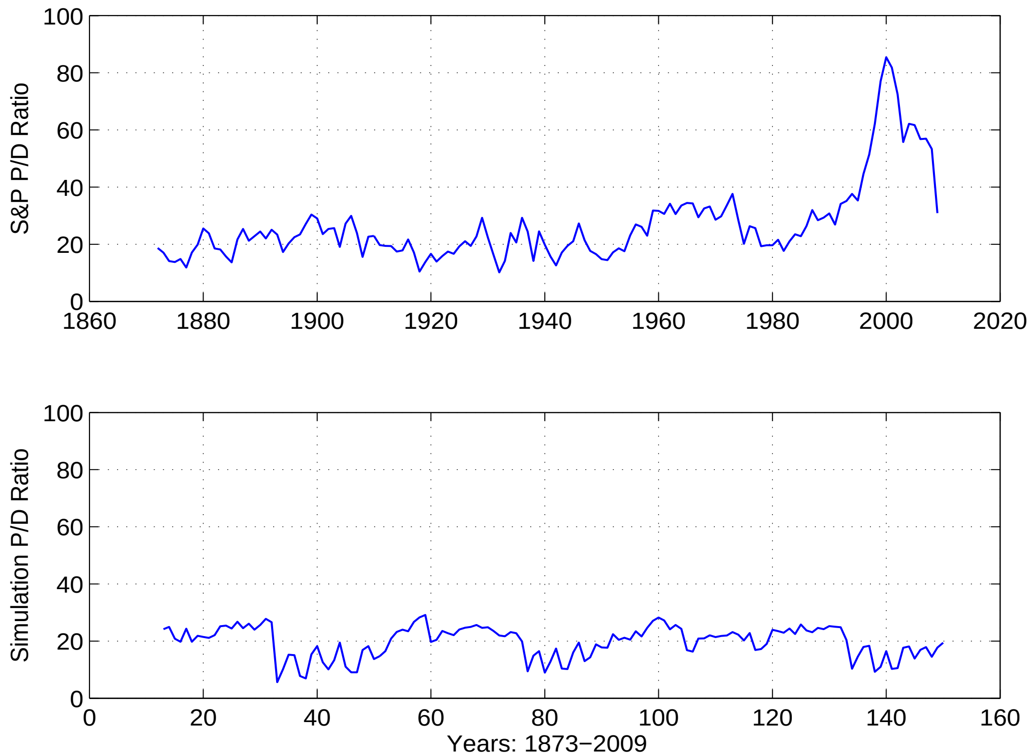

Figure 6: **P/D ratios: Annual 1871-2009**

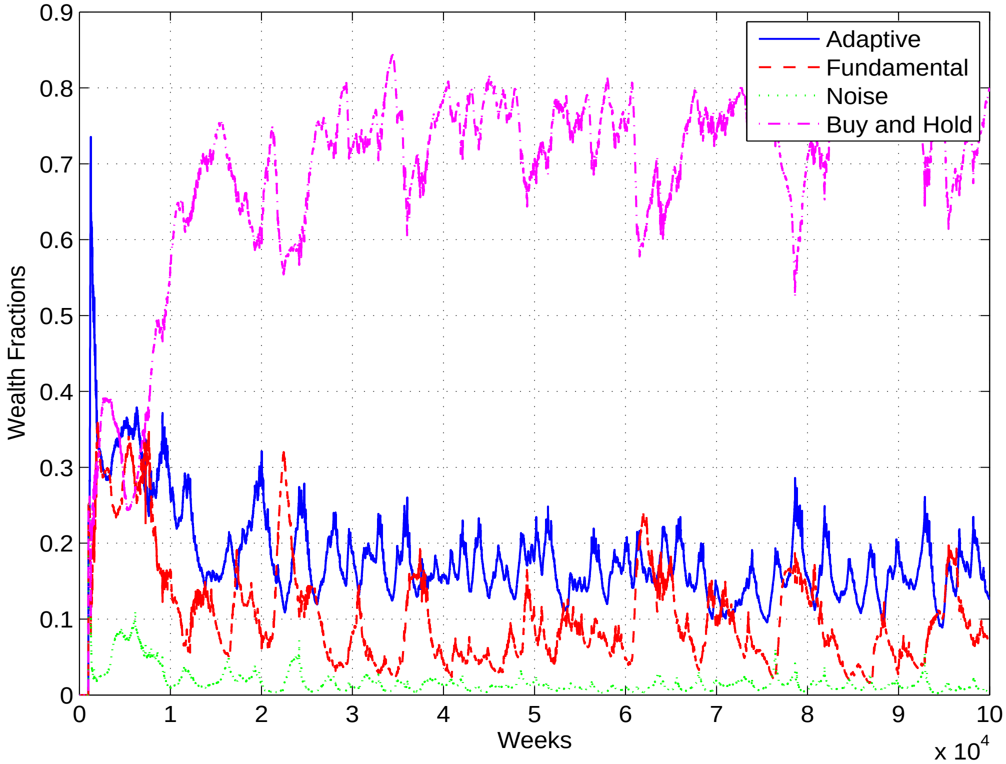

Figure 7: **Wealth Fraction Time Series**

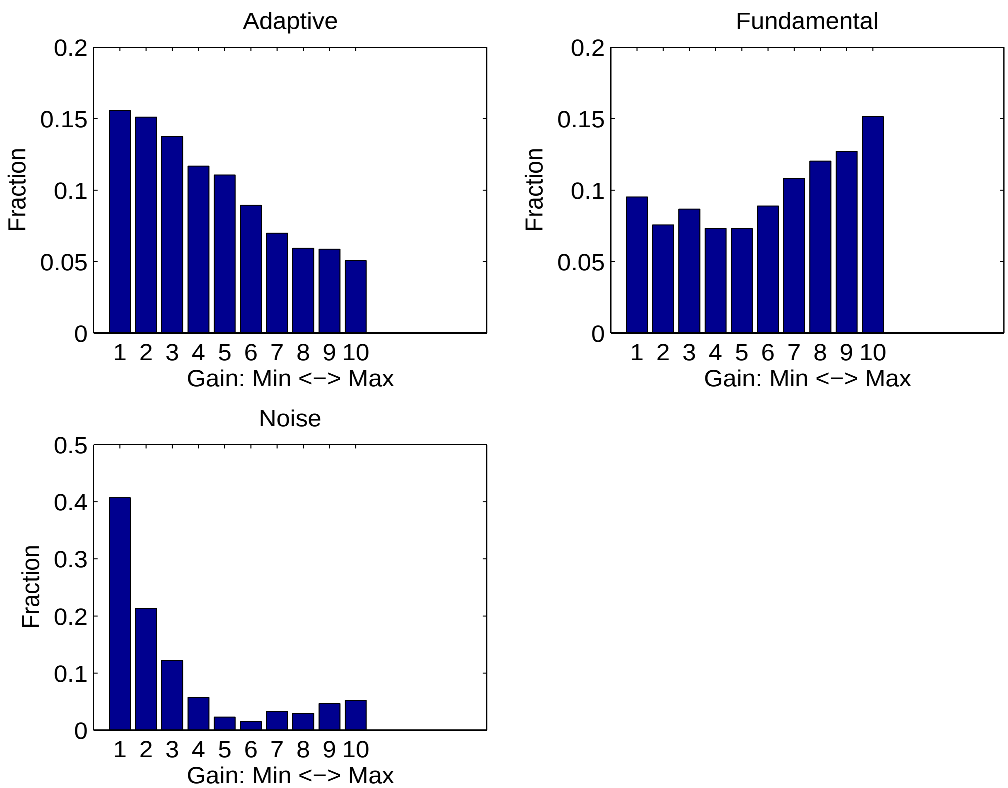

Figure 8: **Gain Wealth Distributions**

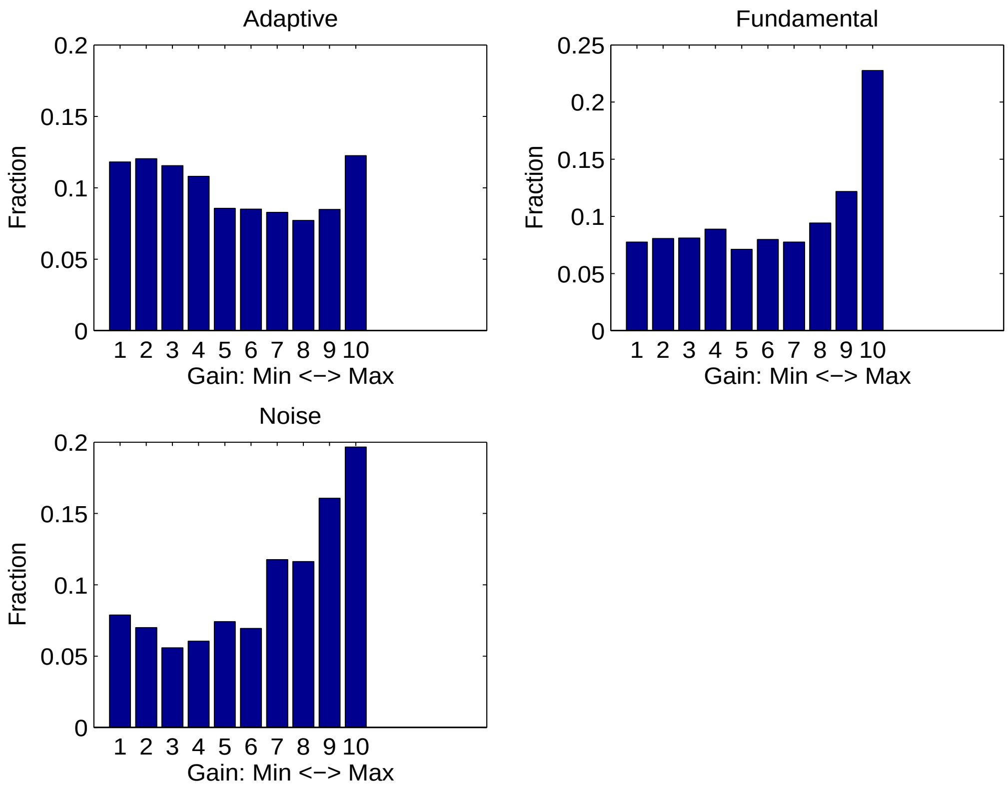

Figure 9: **Volatility Gain Wealth Distributions**

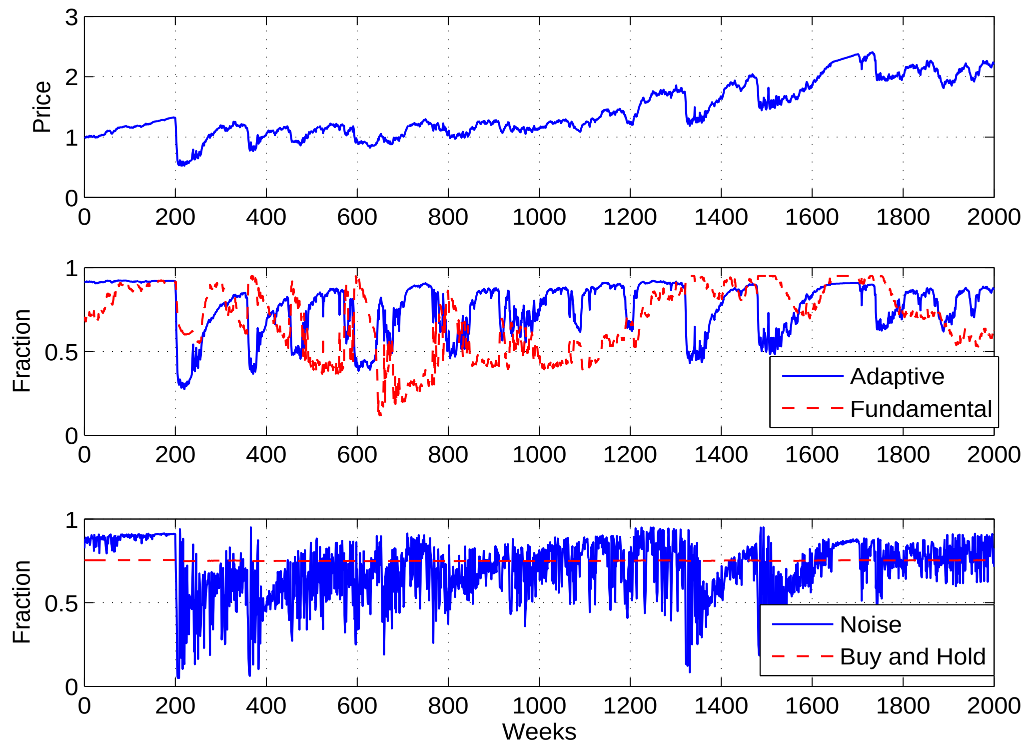

Figure 10: **Strategy Fractions**

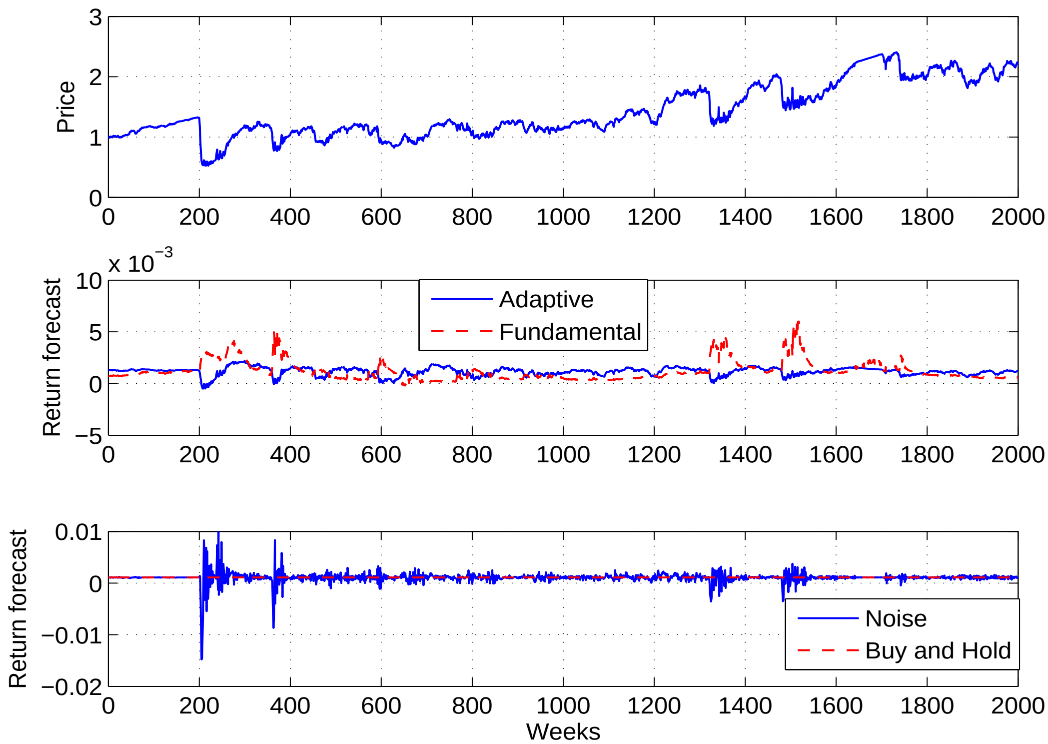

Figure 11: **Strategy Forecasts**

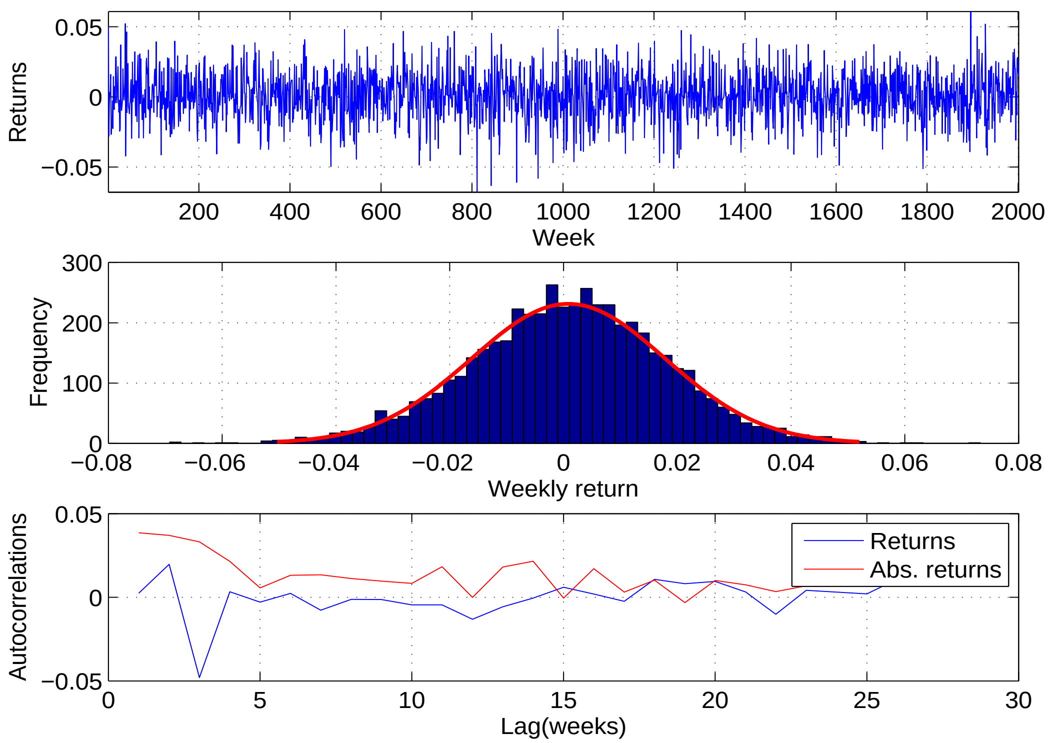

Figure 12: **Summary stats: Low gain only**

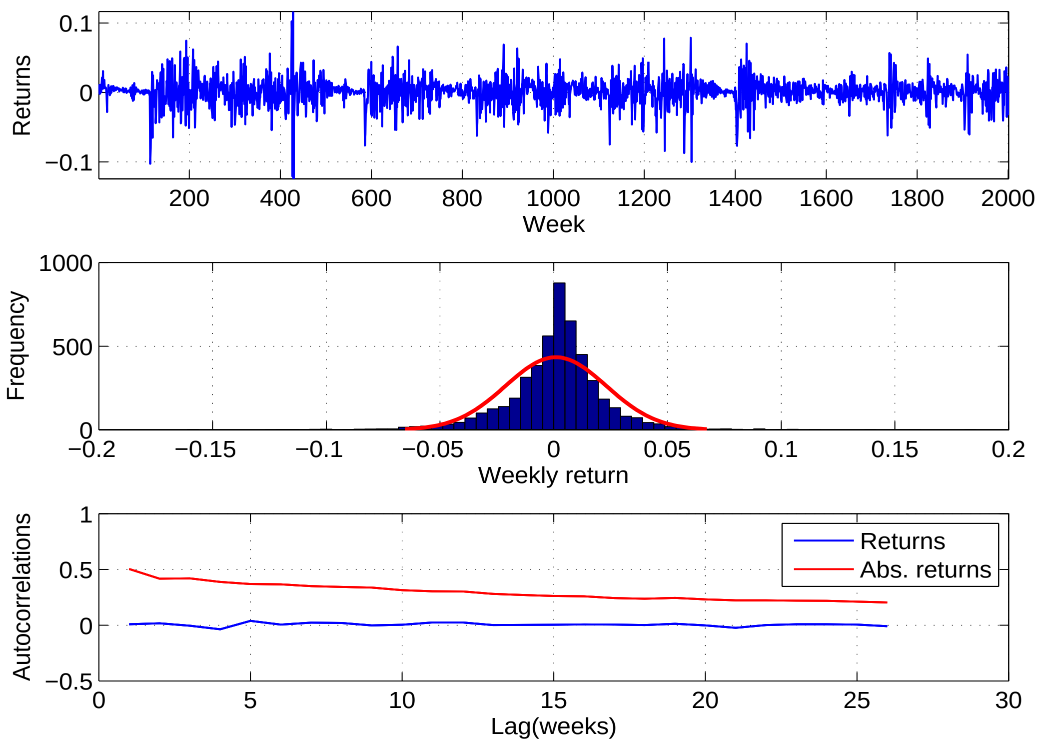

Figure 13: **Summary stats: No adaptive learning**

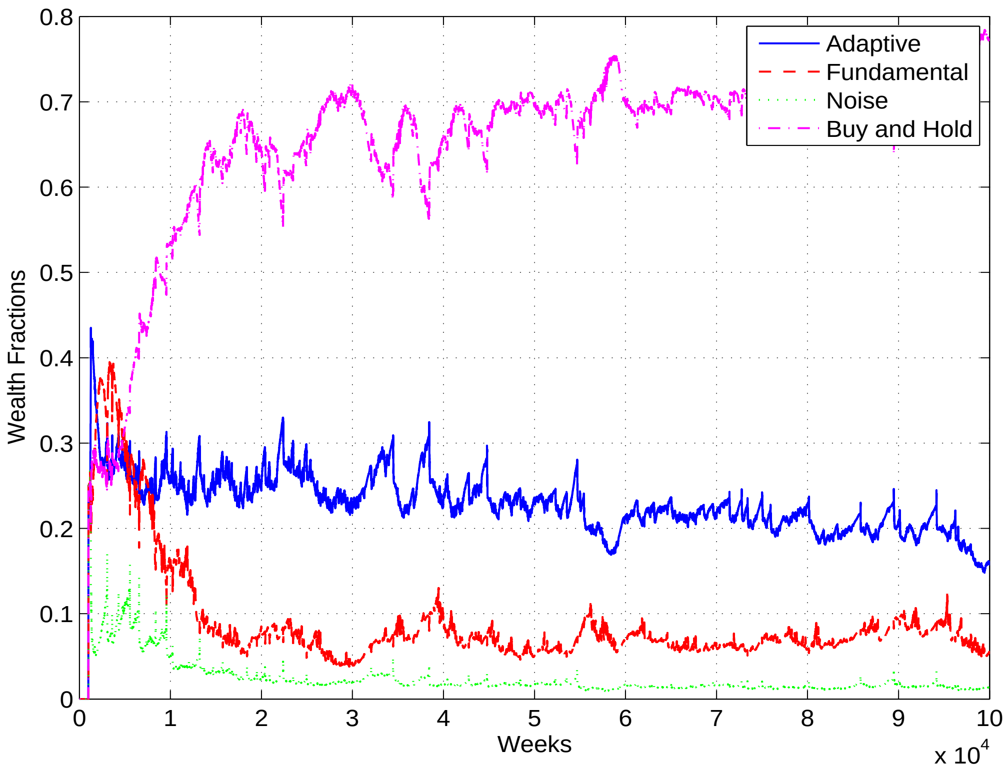

Figure 14: **Wealth Fraction Time Series: No active learning**

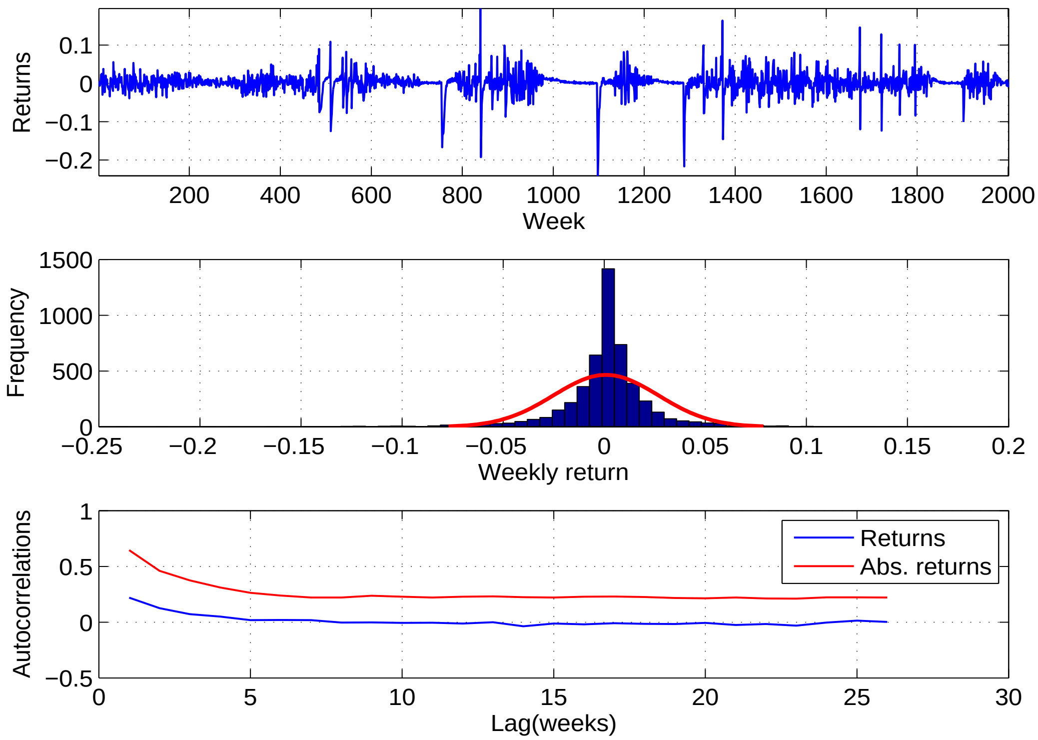

Figure 15: **Wealth Fraction Time Series: No noise trading**