The Impact of Heterogeneous Gain Learning in an Agent-based Financial Market

Blake LeBaron *
International Business School
Brandeis University

Preliminary: June 17, 2013

Abstract

This paper explores the importance of heterogeneous gain levels in generating realistic time series mimicking well known time series features. The methodology takes a very strong stand on learning and evolution. Most forecasting rules involve dynamically adaptive parameters adjusted using a standard recursive least squares framework. Also, agents forecast both expected returns and risk as inputs to their portfolio choices. A population of agents using multiple gain levels in their learning are converged to a single "representative gain" for a small set of heterogeneous forecast strategies. Simulation results show that this dramatically alters the time series outcomes away from the realistic features generated in the heterogeneous gain model. It is therefore an early suggestion that in this class of models heterogeneity in gain levels is necessary for realistic time series replication.

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

*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

Learning involves interpreting past data, and appropriately revising current actions. In this process agents must always make assumptions about how to process this information from the past. Gains or memory lengths determine how much weight to put on the recent versus the distant past. The value of these weights is crucial in any kind of dynamic learning model, but we are still far from a definitive theory of how these values get set, and what their impact can be. Also, there is no reason to assume these gain values should be the same across market participants. This paper explores the importance of heterogeneity in gain parameters, and its impact on the dynamics in a financial market setting.

Gain parameters are present in almost all learning models, and they control the weight given to new observations as learners update their parameters. They can be thought of as controlling how much data from the past agents use to estimate new forecast parameters. They also can be interpreted as related to the signal to noise ratio in a state space framework.[1] Learning models can also entertain both decreasing and constant gain levels. In decreasing gain models, the weight on new observations approaches zero as time passes, often generating useful convergence results. Constant gain learners always are forgetting some data from the distant past, and they continue to give new data points the same weight as they come in. Committing to a fixed gain level locks agents onto a certain belief about the overall stationarity of the time series they are observing. Entertaining various beliefs about gain parameters would be suggested in a world in which priors about data stationarity, persistence, and long range dynamic features may be spread over many possibilities. Also, the data series may be short enough, and the features long enough, that the data alone does not lead to much convergence in these beliefs. In other words, the belief that we are in a "new world," and the last 5 years of data are all that is relevant will always hold some appeal in such a world, and statistical evidence firmly rejecting this idea may be very weak.

Levy, Levy & Solomon (1994) was the first agent-based financial market which explicitly constructed learning behavior with differing perspectives on past data. LeBaron (2006) continued in their spirit using a model based on learning with artificial neural networks trained with different learning memories. The model was easily able to replicate both qualitatively and quantitatively most stylized facts from financial time series.[2] [3] More explicit time series models have been constructed for financial market volatility in Corsi (2009), Ding, Granger & Engle (1993) and LeBaron (2001*b*). All of these models demonstrate a strong connection between well known long memory features in financial market volatility, and approaches using a small number of different memory perspectives.[4]

The market used in this paper is based on approaches used in LeBaron (2012) and LeBaron (2013). It

[1] In the traditional Kalman filter the gain parameter is explicitly defined by the structure of the model.

[2] One can find a lot of commentary and philosophy about this approach to markets in Dacorogna, Gencay, Muller, Olsen & Pictet (2001).

[3] Other multi gain approaches include Thurner, Dockner & Gaunersdorfer (2002), Honkapohja & Mitra (2006), and Mitra (2005). Anufriev & Dindo (2010) and Zschischang & Lux (2001) provide useful analytic results for the model of Levy et al. (1994) supporting the fact that multiple memory lengths can survive in the market. These theoretical structures follow an empirical tradition which assume markets operate with many length scales. A recent paper stressing what might happen when agents overweight shorter run trends, ignoring longer length reversals, is Fuster, Laibson & Mendel (2010). LeBaron (2002) considers the ecological battle between learners with differing perspectives, and how short term learning perspectives can survive by creating an environment in which they thrive. Hommes (2011) surveys the growing experimental literature which often shows people following short term trends when reporting their expectations in controlled laboratory settings.

[4] In many ways this literature follows in the spirit of Granger (1980) which formally shows the connection between aggregation of memory perspectives and long memory.

is a form of hybrid model that attempts to incorporate the richness and adaptive parameter estimation of more complex computational learning models, while remaining simple and tractable enough for detailed analysis. At its core is the usual interaction between stabilizing and destabilizing forces that determine market price dynamics.[5] Following in a long tradition the model contains adaptive, or tend following strategies, along with fundamental, or mean reverting strategies. These are augmented with two additional forecasting types. They are a short term arbitrage strategy that effectively eliminates short run predictability in return series, and a passive buy and hold strategy that invests on long term market averages. Each strategy type forms the core of a forecasting family that implements the predictor using different gain levels. While simple enough for detailed analysis the model is still requires a computational approach in its implementation.

A second key feature in this model is that learning occurs both about future returns and risk. This was part of the original SFI artificial stock market, LeBaron, Arthur & Palmer (1999), and has been used explicitly in a simple learning model by Branch & Evans (2011). This aspect of learning is critical to the dynamics and wealth evolution in many of these markets.[6] One feature of this is that while agents are modeled across several forecast families for the conditional mean, when it comes to modeling the variance they all resort to relatively similar approaches. A large move in the market today contributes to increased conditional variance for all learning agents. This adds an important learning based comovement which is difficult to assume away in both real and artificial markets.

The market also follows two basic design principles. First, aggregate shocks are kept to a minimum. There is only one aggregate shock, the dividend, which is calibrated to the actual dividend in aggregate U.S. data. Adding other aggregate shocks is dangerous from the standpoint of using up degrees of freedom, and questions about just what these shocks are.[7] It also seems to go against the spirit of agent-based modeling that aims to make most of the aggregate variability an endogenous part of the model. A second principle that is followed is to demand a convergence to a relatively recognizable equilibrium for some subset of parameters. In this model, a homogeneous gain world, where agents use only long memory perspectives in their data processing, converges to a situation close to rational expectations.

The basic experiment performed in this paper will be to allow the market to run for some time, settling to a steady state distribution in terms of learning gains. This is done in an environment where gains are allowed to vary across forecasts and agents. Eventually these diverse gains are slowly converged to a common gain represented by their wealth weighted average. The market continues, but now forecasts are generated by a single "representative" gain level, used by all the different forecasts. It is still a multi-agent market since forecasts from the four families are still present, but each uses only a single, and common, gain level. The dynamics of the market are shown to change dramatically in this situation. Large swings in prices are now more regular, and crashes are more dramatic. The time series reveals structure which differs significantly from actual financial markets.

Section 2 describes the model. Section 3 will present simulation results and comparisons with time series features from a long range U.S. data set. Section 3.3 performs the key gain comparisons in the paper,

[5] The model shares much basic intuition in this area with the early work in Brock & Hommes (1998), Day & Huang (1990), Frankel & Froot (1988), Kirman (1991), and Lux (1998). Recent surveys to this literature can be found in Chiarella, Dieci & He (2009) and Hommes & Wagener (2009), Lux (2009), and Tesfatsion & Judd (2006).

[6] See LeBaron (2012) for commentary on this.

[7] See Kocherlakota (2010) for commentary on this issue in relation to modern DSGE models.

and section 4 concludes.

2 Model description

This section will briefly describe the structure of the model. It combines components from many parts of the learning literature. The goal is to build a heterogeneous agent asset market which is as parsimonious as possible, but can still do a good job of replicating financial market features. Also, its inner workings should be simple enough for detailed analysis, meet a general plausibility test, and yet be rich enough to understand several aspects of how wealth moves around in a learning investment environment.

2.1 Assets

The market consists of only two assets. 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 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. 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,

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

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

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

$\alpha_{t,i}$ 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,

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

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 $\alpha_{t,i}$ 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. One further modification is that portfolio choices are tapered in their extremes.

$$\tilde{\alpha}_{t,i} = \begin{cases} \alpha_{t,i} & 0.1 \leq \alpha_{t,i} \leq 0.9 \\ 0.9 + 0.02(\alpha_{t,i} - 0.9) & \alpha_{t,i} > 0.9 \\ 0.1 + 0.02(\alpha_{t,i} - 0.1) & \alpha_{t,i} < 0.1 \end{cases} \tag{10}$$

This shrinks the extreme levels of α to the central range. This is can be viewed as a slight behavioral modification that agents require extra empirical evidence to push their portfolios weights to the most extreme positions. This has little impact on most of the simulation runs, and it could be easily removed without affecting most runs. However, it is important in near rational expectations benchmarks which cause the agents to hold almost all equity since volatility is driven to very low levels.[9] The above restriction keeps the agents in the interior of their portfolio decision and helps the market to converge to a well defined equilibrium in this case.[10]

Consumption will be assumed to be a constant fraction of wealth, λ. This is identical over agents, and

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

[9]This is an appearance of the equity premium puzzle when this model is near the rational expectations equilibrium.

[10]It also corresponds to the neural network based strategies used in LeBaron (2001*a*), but in a crude piecewise linear fashion.

constant over time. The intertemporal budget constraint is therefore given by

$$W_{t+1,i} = (1 + R^p_{t+1})(1 - \lambda)W_{t,i}. \tag{11}$$

This also gives the current period budget constraint,

$$P_t S_{t,i} + B_{t,i} = (1 - \lambda)((P_t + D_t)S_{t-1,i} + (1 + R_f)B_{t-1,i}). \tag{12}$$

This simplified portfolio strategy will be used throughout the chapter. 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.[11]

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, will choose forecast rules, indexed by j, according to expected utility objectives.

All the forecasts will use long range forecasts of expected values using constant gain learning algorithms equipped with the minimum gain level, denoted g_L.

$$\bar{r}_t = (1 - g_L)\bar{r}_{t-1} + g_L r_t \tag{13}$$

$$\overline{pd}_t = (1 - g_L)\overline{pd}_{t-1} + g_L pd_{t-1} \tag{14}$$

$$\bar{\sigma}^2_t = (1 - g_L)\bar{\sigma}^2_{t-1} + g_L(r_t - \bar{r}_t)^2 \tag{15}$$

$$\bar{\sigma}^2_{pd,t} = (1 - g_L)\bar{\sigma}^2_{pd,t-1} + g_L(pd_t - \bar{pd}_t)^2 \tag{16}$$

$$\overline{\Delta p_t} = (1 - g_L)\overline{\Delta p_{t-1}} + g_L \Delta p_t \tag{17}$$

$$\Delta p_t = p_t - p_{t-1} \tag{18}$$

The log price/dividend ratio is given by $pd_t = \log(P_t/D_t)$. The forecasts used will combine four linear

[11]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.

forecasts drawn from well known forecast families. The first of these is a form of adaptive expectations which corresponds to,

$$f_t^j = f_{t-1}^j + g_j(r_t - f_{t-1}^j). \tag{19}$$

Forecasts of expected returns are dynamically adjusted based on the latest forecast error. This forecast format is simple and generic. It has roots connected to adaptive expectations, trend following technical trading, and also Kalman filtering.[12] 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 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.

The second forecasting rule is based on a classic fundamental strategy. It is based on the log linear approximation for returns,[13]

$$r_{t+1} \approx k + \rho p_{t+1} + (1-\rho)d_{t+1} - p_t, \tag{20}$$

$$\rho = \frac{1}{1 + e^{-(\overline{p-d})}} \qquad k = -\log(\rho) - (1-\rho)\log(\frac{1}{\rho} - 1). \tag{21}$$

This approximation allows for a separation of forecasts into a dividend and price component. Agents are assumed to have observed the random walk dividend process, and d_{t+1} is easy to forecast. The mapping of dividends into price movements is much more complicated, and here they resort to a linear regression of $\Delta p_{t+1} = p_{t+1} - p_t$ on current price dividend ratios, pd_t.

$$f_t^j = E_{t,j}(\Delta p_{t+1}) = \overline{\Delta p_t} + \beta_t^j(pd_t - \overline{pd}_t). \tag{22}$$

where pd_t is $\log(P_t/D_t)$, and

$$E_{t,j}(p_{t+1}) = f_t^j + p_t \tag{23}$$

$$E_{t,j}(d_{t+1}) = d_t \tag{24}$$

$$E_{t,j}r_{t+1} \approx k + \rho E_{t,j}p_{t+1} + (1-\rho)E_{t,j}d_{t+1} - p_t. \tag{25}$$

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

$$f_t^j = \bar{r}_t + \sum_{i=0}^{M_{AR}-1} \beta_{t,i}^j (r_{t-i} - \bar{r}_t) \tag{26}$$

[12]See Cagan (1956), Friedman (1956), Muth (1960), and Phelps (1967) for original applications. A nice summary of the connections between Kalman filtering, adaptive expectations, and recursive least squares is given in Sargent (1999). A recent paper demonstrating a more general connection to state space models and expected returns is Pastor & Stambaugh (2009). Empirically, Frankel & Froot (1987) and Taylor & Allen (1992) provide evidence that at least some forecasters do use these rules. Finally, Hommes (2011) surveys some of the laboratory evidence that experimental subjects also use extrapolative methods for forecasting time series.

[13] See Campbell, Lo & MacKinlay (1996) for a textbook exposition.

This strategy works to eliminate short range autocorrelations in returns series through its behavior, and $M_{AR} = 3$ for all runs in this chapter. It will be referred to as the Short AR forecast.[14]

The previous two rules will be estimated each period using recursive least squares. There are many examples of this for financial market learning.[15] 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.[16]

The final rule is a 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 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 updated according to,

$$\begin{aligned} \beta_{t+1}^j &= \beta_t^j + \frac{g_j}{\bar{\sigma}_{pd,t}^2}(pd_{t-1} - \bar{pd}_{t-1})u_{t,j} \\ u_{t,j} &= ((p_t - p_{t-1}) - f_{j,t-1}) \end{aligned} \tag{27}$$

Also, β_t^j is restricted to be between 0 and -0.05. The zero upper bound on β makes sure this strategy is mean reverting, with an overall stabilizing impact on the market.

For the lagged return regression the update follows,

$$\begin{aligned} \beta_{t+1,i}^j &= \beta_{t,i}^j + \frac{g_j}{\bar{\sigma}_{r,t}^2}(r_{t-i} - \bar{r}_{t-i})u_{t,j}, \\ u_{t,j} &= (r_t - f_t^j) \end{aligned} \tag{28}$$

where g_j is again the critical gain parameter, and it varies across forecast rules.[17]

[14] These simple forecasting agents who use only recent returns in their models, fighting against return correlations, share some features with the momentum traders of Hong & Stein (1999).

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

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

[17] This 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 reasonable. This is done to avoid performing many costly matrix inversions. Also, the standard recursive least squares is simplified by using the long run estimates for the mean in both regressions. Only the linear coefficient is estimated with a heterogeneous learning model. This is done to simplify the learning model, and concentrate heterogeneity on the linear parameters, β_t^j.

2.5 Variance Forecasts

The optimal portfolio choice demands a forecast of the conditional variance as well as the conditional mean.[18] The variance forecasts will be generated from adaptive expectations as in,

$$\hat{\sigma}^2_{t+1,j} = \hat{\sigma}^2_{t,j} + g_{j,\sigma}(e^2_{t,j} - \hat{\sigma}^2_{t,j}) \tag{29}$$

$$e^2_{t,j} = (r_t - f^j_{t-1})^2, \tag{30}$$

where $e^2_{t,j}$ 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 rules, reflecting the fact that while there may be many ways to estimate a conditional variance, they often produce similar results.[19] 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.[20] 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.

Finally, agents do not update their estimates of the variance each period. They do this stochastically by updating their variance estimate each period with probability 0.25. This is done for several reasons. First it introduces more heterogeneity into the variance estimation part of the model since its construction yields a lot of similarity in variance forecasts. Also, if variance updating occurred simultaneous to return forecasts, the market would be unstable. Spirals of ever falling prices, and increasing variance estimates would be impossible to avoid in this case.

2.6 Market Clearing

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

$$1 = \sum_{i=1}^{I} Z_{t,i}(P_t). \tag{31}$$

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

$$P_t Z_{t,i} = (1-\lambda)\alpha_{t,i}(P_t)W_{t,i}, \tag{32}$$

$$Z_{t,i}(P_t) = (1-\lambda)\alpha_{t,i}(P_t)\frac{(P_t + D_t)S_{t-1,i} + B_{t-1,i}}{P_t}. \tag{33}$$

[18] Several other papers have explored the dynamics of risk and return forecasting. This includes Branch & Evans (2011) and Gaunersdorfer (2000). In LeBaron (2001*a*) 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.

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

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

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 $\alpha_{t,i}$.[21] 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. [22]

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

$$f_{t+1}^j = (1-g_j)f_t^j + g_j e_{t+1}. \tag{34}$$

This easily maps to the simple exponential forecast rule,

$$f_t = \sum_{k=1}^{\infty}(1-g_j)^k e_{t-k}. \tag{35}$$

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

$$\frac{1}{2} = (1-g_j)^{m_h}, \tag{36}$$

or

$$g_j = 1 - 2^{-1/m_h}. \tag{37}$$

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. 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

[21] A binary search is used to find the market clearing price using starting information from P_{t-1}.

[22] The current price determines R_t which is an input into both the adaptive, and short AR 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.

for active learning, or 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 chapter 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,

$$\hat{U}_{t,j} = \hat{U}_{t-1,j} + g_u^i (U_{t,j} - \hat{U}_{t-1,j}), \tag{38}$$

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

$$U_{t,j} = \frac{1}{1-\gamma}(1 + R_{t,j}^p)^{(1-\gamma)} \tag{39}$$

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

3.1 Benchmarks

This paper uses a data set for U.S. stock returns and dividends as a comparison. It is is designed to both to cover as long a horizon as is reasonably possible for dividends, prices, and return volatility. It merges three different data sets. For the period from 1926 through 2012, the benchmark CRSP data is used with the value weighted index as the proxy for market returns. The difference between the returns on the index with and without dividends is used to construct price dividend ratios as is standard in the finance literature. Also, the daily returns data are used to build a monthly volatility series with a simple monthly realized volatility.[23] For the period from 1871 to 1925 the data set uses the long data set constructed by Shiller, and available at his website.[24] The Shiller data is only at monthly frequency, so it is also augmented with an early daily series from Schwert which is used to generate early monthly realized volatilities back to 1872.[25]

Results will be presented for three different simulation specifications. First, the multiple gain runs correspond to a full run using all the different gain levels across forecast rules and agents. These runs serve

[23] This method is very commonly used with intraday data, but was originally used in French, Schwert & Stambaugh (1987) at the daily frequency.

[24] This is the main input for Shiller (2000), and many papers.

[25] Schwert (1990) gives extensive details on the construction of this series which is available at his website.

as a benchmark for replicating actual returns in the U.S. data set. The simulation is allowed to run for 200,000 weeks, and the last 50,000 weeks are used in most of the time series analysis, giving approximately 950 years of data.

A figure showing the final 100 years of a run is presented in 1. The upper panel shows the price dividend ratio which is moving through time, displaying large and persistent swings. The pattern is somewhat irregular with no clear periodicity or maximum or minimum level which would allow investors to call the top or the bottom. The middle panel presents the weekly returns corresponding to the price/dividend ratios. These give temporal pockets of high and low volatility common in most financial time series. The lower panel displays trading volume which also moves counter cyclically to the P/D ratio, and is strongly correlated with the level of volatility in the return series.

The second set of runs is a counter factual run with the gain levels set to very long half lives of about 100 years. This market converges to something very different from the multiple gain runs. It is again run for 200,000 weeks, and the final 100 years are summarized in figure 2. The three panels display P/D ratios, returns, and trading volume. It is obvious this run is very different from the previous one. The price dividend ratio stays relatively constant around a level of 35. Returns are uniform with no dramatic changes in volatility. Trading volume is also flat, and more than an order of magnitude lower than that in the multiple gain benchmark. This casual picture appears to show a market which has converged to something close to a simple rational expectations equilibrium. Further, evidence in this paper will give more information supporting this.

Finally, to get a sense about actual price movements relative to dividends, the price dividend ratio is presented from the merged data set in figure 3. This shows the large and very erratic swings in prices (relative to fundamentals) which are common in most financial markets. There is an obvious unusual period in the 1990's and beyond as the ratio peeks near 90 in 2000. Occasionally some statistics will be calculated by truncating this series in 1990.

3.2 Wealth distributions

The model depends on where agent wealth is distributed across strategies and gain levels. Figure 4 shows the long run dynamics of wealth through the entire multiple gain simulation run. It is an interesting picture with the distribution of wealth stabilizing to a relatively smooth pattern across strategies. Buy and hold dominates with nearly 50 percent of the wealth. It is followed by the adaptive forecast rule which controls about 35 percent of wealth. Beyond these two strategies the fraction drops off dramatically with the other two strategies holding less than 10 percent of wealth in the market.

There are two important facts to note on this figure. First, the adaptive forecast rule is generally a destabilizing or positive feedback rule buying on price increases, and selling when prices fall. The fundamental rule is a stabilizing strategy following the reverse pattern. The relative concentration of wealth suggests that the fundamental types will have a difficult time smoothing out market instabilities. Also, the relative strength of the buy and hold strategies reflect the fact that while the market has some predictability to it, the strategies are still somewhat difficult to implement. In other words there is no strategy that can strongly dominate simply buying the market. This feature is consistent with actual market data.

Inside each strategy wealth is also distributed across the forecast and volatility forecast gain levels.

Figure 5 shows how wealth is distributed across the different discrete gain levels used in the market. This is a snapshot taken at the end of the $200,000$ week run. For the adaptive and fundamental strategies wealth is neither uniform, nor concentrated on one gain level alone. The arbitrage strategy which runs short range linear regressions is an interesting exception. Here wealth concentrates on the lowest gain (longest memory) level which shows agents using long histories of data to make their forecasts.[26]

Figure 6 repeats the analysis for the variance forecast gain levels. Here, not only is the distribution dispersed across all levels, it is not monotonic. There appear to be local peaks for both the lowest and highest gain levels. This is begins to make it seem unlikely that a single representative gain level will suffice to replicate these diverse types.

3.3 Converging gain levels

This section presents the experiments for converging gain levels. The multiple gain simulation is run as before for $100,000$ weeks. At this point the various diverse gain levels are slowly converged to the wealth weighted average gain at that point of the run.[27] Figure 7 shows the time series of the wealth weighted gain level. It is converted to the half life of the memory $(1 - gain)$ level.

It is interesting to note in the figure that both rule average gain levels continue to show some variability before the convergence experiment is started. After week $100,000$ the gain levels are all forced to converge to the current average level. In the case of the fundamental strategy this is about 400, and for the adaptive it is only about 250. This is consistent with the distributions displayed in figure 5. Figure 8 repeats this for the volatility gains. These also display some variability before convergence starts along with an interesting runup in the early part of the simulation. The ordering in level is reversed with the adaptive types using lower gain (longer memory) variance forecasts. Again, this is consistent with the distributions in figure 6.

The question is now whether this run changes at all in terms of its price, return and volume dynamics. This is dramatically answered in figure 9. This should be compared to the same components plotted in figure 1. The dynamics have clearly changed. The bubble like price dynamics now are longer (on the order of 30 years), more regular, and more intense. The P/D ratio runs up to above 40, and then quickly drop to nearly zero. A similar increase in intensity and regularity is present for the return series and volume series. Volume dramatically drops to near zero as the bubble reaches its peak, and then spikes as the market crashes, remaining high for almost fifteen years after the crash.

The move to a single representative gain level clearly changes price/return/volume dynamics, but has it changed anything else? Figure 10 repeats the time series plot for wealth fractions previously displayed in figure 4. The convergence begins in week $100,000$, and this is evident in the plot. The previous fraction rankings hold well until period $100,000$. At that point they change dramatically with the adaptive (or trend following) strategy taking over almost 60 percent of the wealth, with most of this coming at the expense of the buy and hold strategy which drops to only 20 percent. It also appears that the variability of these strategies over time has increased as well. Further analysis of this effect is underway, but it would appear that the single adaptive strategy which follows well defined, and homogeneous, trend patterns is better able to self reinforce its behavior in this market. Trends are longer and stronger, and it is able to take advantage

[26]This is consistent with agents learning that their time series are stationary in terms of this strategy. Also, this strategy can be replaced with a single low gain type, and none of the results presented here will change.

[27]Actually, a mean from the last 5,000 weeks is taken as the target gain. This smooths out any short term movements in gains.

of this. In the heterogeneous world, a diverse set of trend followers does not lay out the pattern of a single well defined trend.

Table 2 presents some summary statistics from the merged U.S. data series, and the three simulation runs described so far. Columns one and two compare the actual data to the multiple gain simulation runs. There is reasonably good alignment here with means and standard deviations of the same order of magnitude. The simulation shows a slight increase in both, but it should be remembered that there is considerable noise around the point estimates from the U.S. data. Also, the simulated market represents trading in a single security while the actual data is a value weighted index that smooths over some variability in individual stocks. The single gain run shows a large increase in the excess return to 16.4 percent which is out of line with any estimate of U.S. excess stock returns. Its volatility is also almost twice that of the actual data. The extreme variability shown in figure 9 carries over into the simple summary statistics. The low gain run displays a slightly lower excess return, and a dramatically lower volatility at 5 percent which is well below any estimate of aggregate stock volatility. This shows that it is no longer amplifying its true fundamental variability driven by the dividend process.

The final two rows in table 2 report statistics for the price/dividend ratios. All are relatively similar in terms of the mean level except for the low gain series which displays an unusually high P/D ratio at nearly 36. This is driven by the very low volatility levels in the market which push the prices higher relative to dividends on average. The overall variability of the P/D ratio is similar for the data and the multiple gain simulations. The single gain runs generate much more variability, as was clear in the figure, and similarly the low gain runs generate much less P/D variability.

Further time series differences in the dynamics of the P/D ratio are shown in figure 11 which reports the autocorrelations of the annual price dividend ratios. All show a very persistent process with positive autocorrelations out to nearly 100 months. At this point they diverge. The single gain runs generate large negative autocorrelations in a regular pattern with a minimum at about 200 months or roughly 17 years. The pattern seems to line up well with the frequencies seen in figure **??**, and reflects the much more regular behavior exhibited by the single gain simulations. Both the data and the multiple gain drop to near zero, but both do this in a somewhat erratic fashion.[28]

Shorter horizons returns are analyzed in the next figures and tables. Figure **??** presents the histograms for the daily returns in the 3 different simulations and U.S. weekly returns built from the CRSP value weighted (VW) daily return series from 1926 through 2012. The multiple gain distribution is most similar to the actual return series, although visually it is slightly more leptokurtic. The low gain simulation run generates weekly returns which are nearly normal. The single gain run generates returns which are much more leptokurtic with a sharp peak, and some very extreme returns.

These graphical features are repeated in table 4 which again compares the different runs. Once again the actual U.S return series and the multiple gain runs are similar displaying moderate kurtosis, and a slight negative skew. The single gain series is well off the actual data, with an extreme kurtosis over 160.[29] The low gain series again appears close to a Gaussian with zero skew, and a kurtosis of 3.1. The last two lines

[28]Two things should be noted on the U.S. data. First, the series is stopped in 1992 respecting the fact that it is arguably nonstationary after this point. Also, taking long range autocorrelations with such a short time series should be viewed with caution. The standard errors on these estimates must be huge.

[29]This kurtosis level and the tail exponents suggest that fourth moments don't exist in this series, so the actual estimated kurtosis may be a meaningless number.

present tail exponent estimates which give further information on the fatness of the extreme tails of the distributions.[30] The multiple gain runs are reasonably close to the actual data for the left fail, but display a larger index (thinner tail) on the right tail. The single gain case displays a lot of action out in the extreme tails with exponent estimates of 1.5 and 2.5 for the left and right tails respectively. 1.5 is a very unusual number, since it indicates that all moments above 1.5 don't exist for this distribution which would mean that variances don't exist. There is no question that this series is generating tail activity on a level much greater than the actual data, but how much these extreme tail estimates can be relied on will require some further analysis. In the final case of low gain forecasts the tail exponents are large as should be expected. Given that the series appears near Gaussian these values should be infinite.

The next two figures look at the autocorrelations in volatility. For the actual data, monthly realized volatility measures are constructed using daily returns data. This gives a volatility estimate corresponding to each month. For the simulations a similar measure is built from the four weekly observations to construct a realized monthly volatility value. The persistence of volatility is one of the most important features in financial data. Figure 13 shows a good alignment between the data, and the multiple gain simulation. In both cases volatility is very persistent, but correlations drop off quickly at first, and then linger, still positive even out to lags of over two years. This is often attributed to long memory or potentially multiple factors in the volatility process. Here, it is endogenously generated by the simulated traders. Figure 14 repeats this estimation for the other two simulation runs. Both show some persistence, but none appear to look much like the actual data. The low gain case shows some persistence in volatility, but it is generally pretty small, starting at only slightly above 0.1. It eventually drops to be very close to zero. The single gain case is the most interesting. It shows very large autocorrelations, but they drop off fast. The rate appears almost exponential. This is interesting since it is the pattern one would see if the volatility process were a simple autoregressive process with one lag (an AR(1)). This makes sense since volatility forecasting is concentrated on one time horizon or frequency based on the single gain parameter. This forecasting behavior seems to be endogenously appearing in the time series in a manner that reflects the construction of the volatility forecaster. This difference between the multiple and single tgain forecast here is quite stark, and suggests a place where the multiple gain forecasting set will be necessary for volatility time series replication.

Table 4 replicates results from one of the most reported features of long range price/dividend ratios, simple linear return predictions. Future excess quarterly returns are regressed on the logged current price/dividend ratio. Predictability in these regressions is often significant, but weak in terms of R^2.[31] The regressions are estimated both for the long range stock return series, and the three different simulations. Surprisingly, the results to not vary that much across the different treatments. In all cases the estimated parameter, β, is significant, but R^2's show that the linear model is not capable of capturing much of the variability in quarterly returns.[32] Even the low gain case shows some evidence for predictability, although the R^2 of 0.01 suggests it is of minor importance. It is somewhat surprising given the previous results that this feature is so similar across the different runs. The linear forecasting model would appear to have little power to discern features which are even clear to an eyeball test. It is still important that this basic aspect of long range financial series is replicated well by this model.

[30] See LeBaron (2008) and Huisman, Koedijk, Kool & Palm (2001) for information on the modified Hill estimators that are used here.

[31] See Cochrane (2008) for a discussion and references to the vast literature. Also, see Goyal & Welch (2003) for results on how the out of sample predictability can be very weak from these regressions.

[32] Another well known feature of these regressions is that the R^2 increases as the forecasting horizon increases.

4 Conclusions

This paper has shown in a simple agent-based financial market that heterogeneous gains across agents are essential to generating realistic market dynamics. In terms of agent-based modeling this is an important message, and suggests that models built using single learning gains may never be completely adequate for financial modeling.

The assumption that market participants do not settle down to a common belief about the use of past data seems very reasonable. Market participants are often interested in performance measures over many different horizons. Also, their is no common agreement on the use of different predictive tools such as trend following, or momentum strategies. Finally, financial data does not help much, given that its irregularities and noise levels do not point to any single past horizon to use for future forecasting. This allows the rich heterogeneous population of forecast rules to thrive, contributing to the uncertain data environment in the future.

A world where agents use only a single gain level still generates some of the basic stylized facts of financial time series. This includes fat tailed distributions, pockets of large volatility, and persistent changes in the price/dividend ratio. However, the dynamics with a single gain are extreme, and could not be recognized as appearing like any real world financial time series. This reminds one of some of the simpler agent-based financial markets with a small set of trading strategies. They generate interesting dynamics which are qualitatively replicating various features, but are far off the mark in terms of quantitative features.

Differences between homogeneous and heterogeneous learning models show up in more ways than just basic time series features. Distribution of wealth across the different strategies changes dramatically. Relative wealth shifts from the passive buy and hold strategies to the adaptive trend following strategies. This is probably driven by the fact that in a single forecast gain world, the trends are easier to find and coordinate on by the different agents. A world with multiple gain levels yields much more confusion on the "right" length for trends, and makes coordination much messier. This messy coordination seems more realistic for observed market behavior.

Some aspects of financial time series clearly indicate hints of a multiple gain perspective. This is clearest in the world of volatility where time series models with either long memory or multiple horizons are standard. The results in this paper are supportive of the fact that multiple gain forecasters are critical to these long memory like features in terms of volatility persistence. Single gain learners generate persistent volatility, but it does not appear very close to the long memory, highly persistent features seen in actual returns data.

If heterogeneous gain learners are essential in terms of modeling financial markets, then it is important in terms of modeling strategies. A single representative gain agent offers the possibility of a much simpler market setting, and a tantalizing hope for some better analytic models.[33] The model used here would then be composed of four forecast types which is still a difficult task for dynamic analysis, but is possible. The presence of multiple gain levels in the learning process makes this analytic simplification much less likely. Technology for doing this still is not very well understood. These more complicated models may have to continue in the realm of simple computational experiments for some time to come.

[33] A good thought experiment on what might be possible are the heterogeneous agent DSGE models of which Krusell & Anthony A. Smith (1998) is an early example. The dynamics here is suggestive that describing wealth distributions with a small set of summary statistics may not be sufficient.

References

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, pp. 778–878.

Anufriev, M. & Dindo, P. (2010), 'Wealth-driven selection in a financial market with heterogeneous agents', *Journal of Economic Behavior and Organization* **73**, 327–358.

Bollerslev, T., Engle, R. F. & Nelson, D. B. (1995), ARCH models, *in* R. F. Engle & D. McFadden, eds, 'Handbook of Econometrics', Vol. 4, North-Holland, New York, NY, pp. 2959–3038.

Branch, W. A. & Evans, G. W. (2011), 'Learning about risk and return: A simple model of bubbles and crashes', *American Economic Journal: Macroeconomics* **3**, 159–191.

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.

Cagan, P. (1956), The monetary dynamics of hyperinflation, *in* M. Friedman, ed., 'Studies in the quantity theory of money', University of Chicago Press, Chicago, IL, pp. 25–120.

Campbell, J. Y., Lo, A. W. & MacKinlay, A. C. (1996), *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.

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.

Cochrane, J. (2008), 'The dog that did not bark: A defense of return predictability', *Review of Financial Studies* **21**(4), 1533–1575.

Corsi, F. (2009), 'A simple approximate long-memory model of realized volatility', *Journal of Financial Econometrics* **7**(2), 174–196.

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.

Ding, Z., Granger, C. W. & Engle, R. F. (1993), 'A long memory property of stock market returns and a new model', *Journal of Empirical Finance* **1**, 83–106.

Evans, G. W. & Honkapohja, S. (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.

Frankel, J. A. & Froot, K. A. (1987), 'Using survey data to test standard propositions regarding exchange rate expectations', *American Economic Review* **77**(1), 133–153.

Frankel, J. A. & Froot, K. A. (1988), Explaining the demand for dollars: International rates of return and the expectations of chartists and fundamentalists, *in* R. Chambers & P. Paarlberg, eds, 'Agriculture, Macroeconomics, and the Exchange Rate', Westview Press, Boulder, CO, pp. 73–126.

French, K. R., Schwert, G. W. & Stambaugh, R. F. (1987), 'Expected stock returns and volatility', *Journal of Financial Economics* **19**, 3–29.

Friedman, M. (1956), *A Theory of the Consumption Function*, Princeton University Press, Princeton, NJ.

Fuster, A., Laibson, D. & Mendel, B. (2010), 'Natural expectations and macroeconomic fluctuations', *Journal of Economic Perspecitives* **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.

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.

Goyal, A. & Welch, I. (2003), 'Predicting the equity premium with dividend ratios', *Management Science* **49**(5), 639–654.

Granger, C. W. (1980), 'Long memory relationships and the aggregation of dynamic models', *Journal of Econometrics* **14**, 227–238.

Hommes, C. H. (2011), 'The heterogeneous expectations hypotheis: Some evidence from the lab', *Journal of Economic Dynamcis and Control* **35**, 1–24.

Hommes, C. H. & Wagener, F. (2009), Complex evolutionary systems in behavioral finance, *in* T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', North-Holland, pp. 217–276.

Hong, H. & Stein, J. C. (1999), 'A unified theory of underreaction, momentum trading, and overreaction in asset markets', *Journal of Finance* **54**, 2143–2184.

Honkapohja, S. & Mitra, K. (2006), 'Learning stability in economies with heterogeneous agents', *Review of Economic Dynamics* **9**, 284–309.

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.

Kirman, A. P. (1991), Epidemics of opinion and speculative bubbles in financial markets, *in* M. Taylor, ed., 'Money and Financial Markets', Macmillan, London, pp. 354–368.

Kocherlakota, N. (2010), 'Modern macroeconomic models and tools for economic policy', *The Region: Federal Reserve Bank of Minneapolis* .

Krusell, P. & Anthony A. Smith, J. (1998), 'Income and wealth heterogeneity in the macroeconomy', *Journal of Political Economy* **106**(5), 867–896.

LeBaron, B. (2001*a*), 'Evolution and time horizons in an agent based stock market', *Macroeconomic Dynamics* **5**(2), 225–254.

LeBaron, B. (2001*b*), 'Stochastic volatility as a simple generator of apparent financial power laws and long memory', *Quantitative Finance* **1**, 621–631.

LeBaron, B. (2002), 'Short-memory traders and their impact on group learning in financial markets', *Proceedings of the National Academy of Science: Colloquium* **99**(Supplement 3), 7201–7206.

LeBaron, B. (2006), Agent-based financial markets: Matching stylized facts with style, *in* D. Colander, ed., 'Post Walrasian Macroeconomics', Cambridge University Press, pp. 221–235.

LeBaron, B. (2008), Robust properties of stock return tails, Technical report, International Business School, Brandeis University.

LeBaron, B. (2012), 'Heterogeneous gain learning and the dynamics of asset prices', *Journal of Economic Behavior and Organization* **83**(3), 424–445.

LeBaron, B. (2013), Heterogeneous gain learning and long swings in asset prices, *in* R. Frydman & E. S. Phelps, eds, 'Rethinking Expecations: The Way Forward For Macroeconomics', Princeton University Press, pp. 169–206.

LeBaron, B., Arthur, W. B. & Palmer, R. (1999), 'Time series properties of an artificial stock market', *Journal of Economic Dynamics and Control* **23**, 1487–1516.

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

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.

Lux, T. (2009), Stochastic behavioral asset pricing models and the stylized facts, *in* T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', North-Holland, pp. 161–215.

Mitra, K. (2005), 'Is more data better?', *Journal of Economic Behavior and Organization* **56**, 263–272.

Muth, J. (1960), 'Optimal properties of exponentially weighted forecasts', *Journal of the American Statistical Association* **55**, 299–306.

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.

Pastor, L. & Stambaugh, R. F. (2009), 'Predictive systems: Living with imperfect predictors', *Journal of Finance* **64**, 1583–1628.

Phelps, E. S. (1967), 'Phillips curves, expectations of inflation and optimal unemployment over time', *Economica* **34**(135), 254–281.

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

Schwert, G. W. (1990), 'Indexes of United States stock prices from 1802-1987', *Journal of Business* **63**, 399–426.

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

Taylor, M. & Allen, H. (1992), 'The use of technical analysis in the foreign exchange market', *Journal of International Money and Finance* **11**(3), 304–14.

Tesfatsion, L. & Judd, K. L., eds (2006), *Handbook of Computational Economics: Agent-based computational economics*, North-Holland, Amsterdam.

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.

Zschischang, E. & Lux, T. (2001), 'Some new results on the Levy, Levy, and Solomon microscopic stock market model', *Physica A* **291**, 563–573.

Table 1: *Parameter Definitions*

Parameter	Value
d_g	0.0126
σ_d	0.12
r_f	0
γ	3.5
λ	0.0005
I	16000
J	4000
g_j	[100,20,10,5,2] years
g_L	200 years
g_u	[100,20,10,5,2] years
L	5 percent/year
$[\alpha_L, \alpha_H]$	[0.01,0.99]
σ_ϵ	0.01
M_{AR}	3
h_j	[0.025,0.10]

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: **Annual Statistics**

Series	U.S. 1872-2012	Multiple gain 960 years	Single gain 960 years	Low gain 960 years
Mean excess return	6.5	8.4	16.4	4.3
Std.	16.3	19.2	31.8	5.0
Sharpe ratio	0.33	0.44	0.52	0.87
Mean(P/D)	22.21	20.87	21.11	35.93
Std(P/D)	5.77	5.16	11.70	1.25

Simulations use samples of the final 50,000 weeks or about 950 years. The U.S. series are annual from 1871-2012. P/D refers to the price dividend ratio. All returns are real including dividends. For P/D ratios only, the U.S. is stopped at 1992.

Table 3: **Weekly Statistics**

Series	U.S.	Multiple gain	Single gain	Low gain
Skewness	-0.5	-0.7	-7.9	-0.0
Kurtosis	8.4	10.5	164.1	3.1
Left tail index	3.9	4.7	1.5	9.7
Right tail index	4.8	8.5	2.5	10.1

$$r_{t+1} - r_{f,t+1} = \alpha + \beta \log(P_t / D_t)$$

Table 4: **Price/dividend Return Forecasts**

Series	β	R-squared
U.S. Merged (1872-2012)	-0.047	0.02
	(0.010)	
All gain	-0.078	0.04
	(0.006)	
Single gain	-0.077	0.06
	(0.01)	
Low gain	-0.054	0.01
	(0.01)	

All regressions are run at the quarterly frequency with OLS. Numbers in parentheses are standard errors.

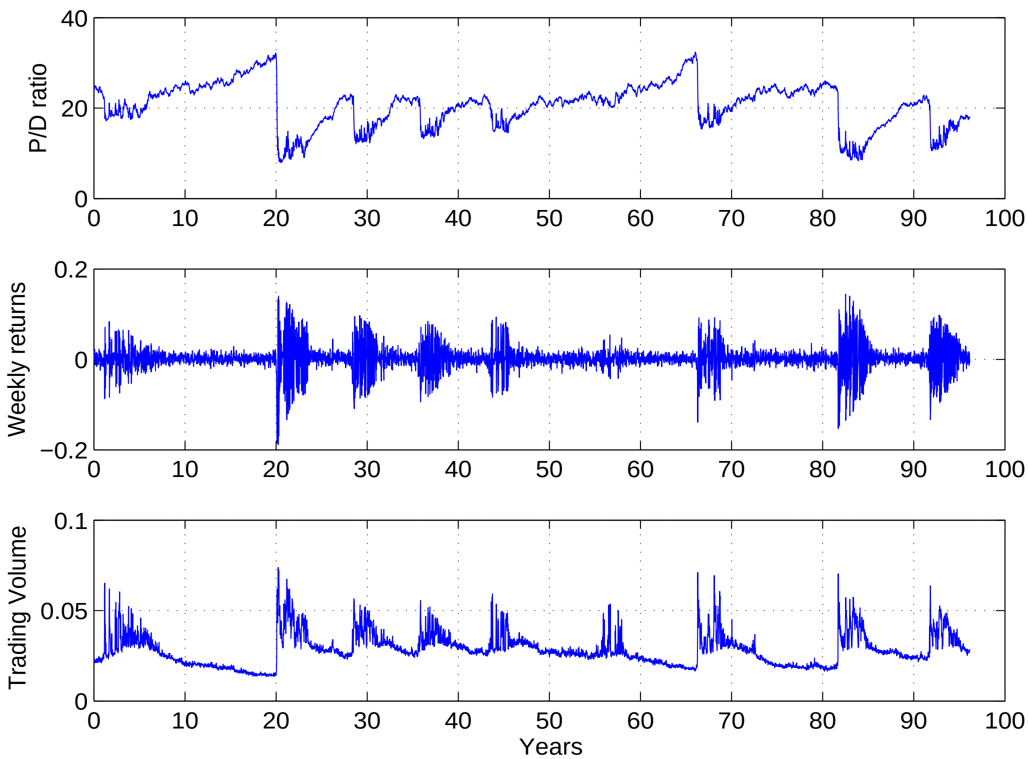

Figure 1: **Multiple Gain Summary:** P/D ratio, returns, volume

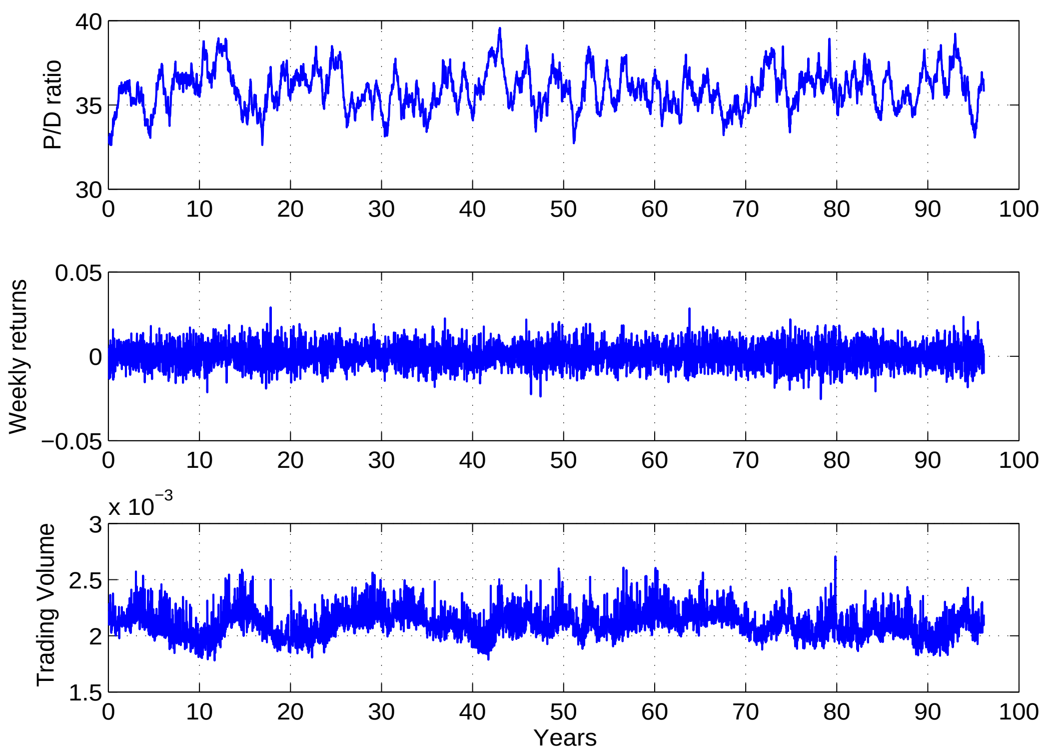

Figure 2: **Low Gain Summary:** P/D ratio, returns, volume

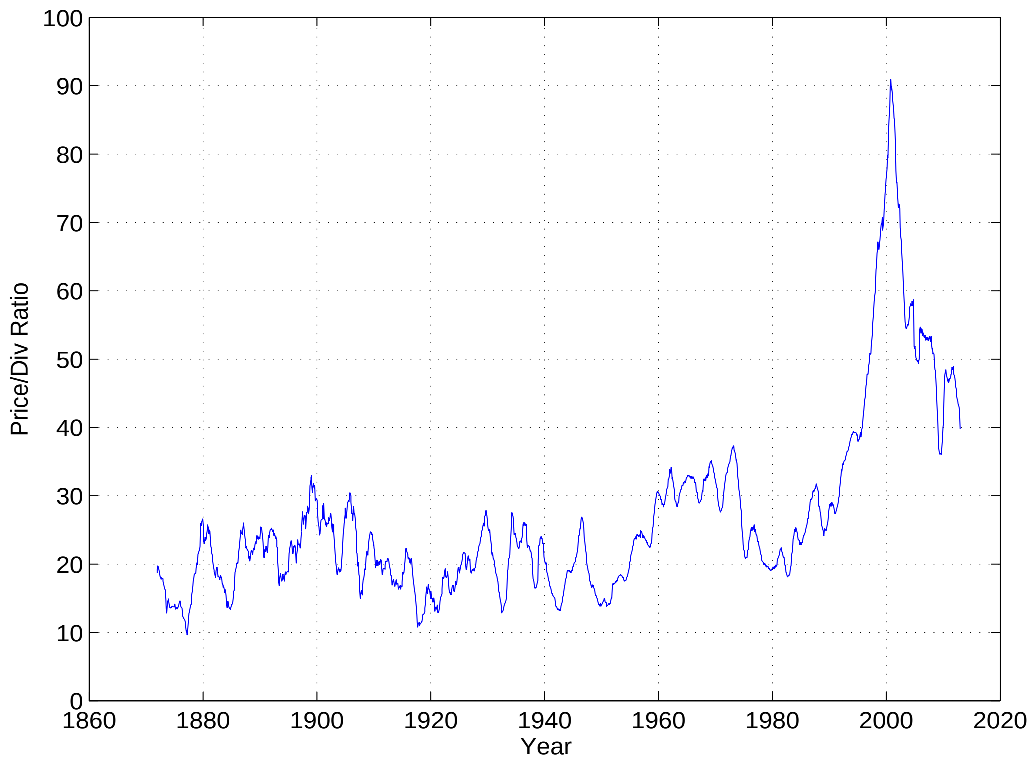

Figure 3: **U.S. Aggregate Price/Dividend Ratio**

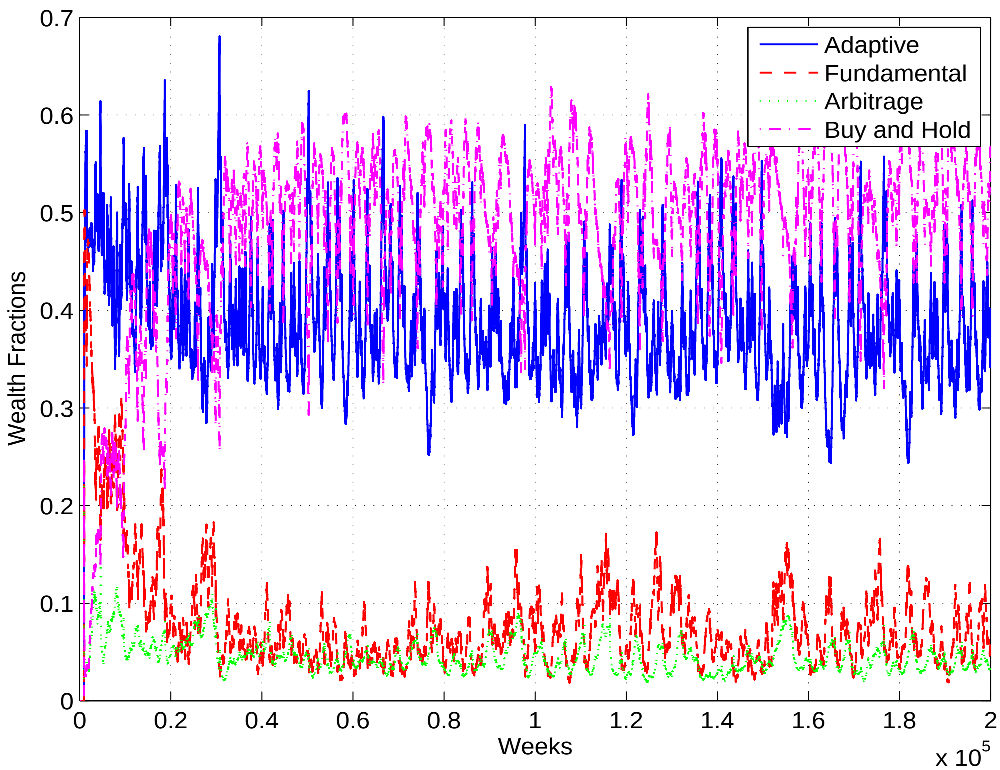

Figure 4: **Wealth Fractions Across Forecast Strategies**

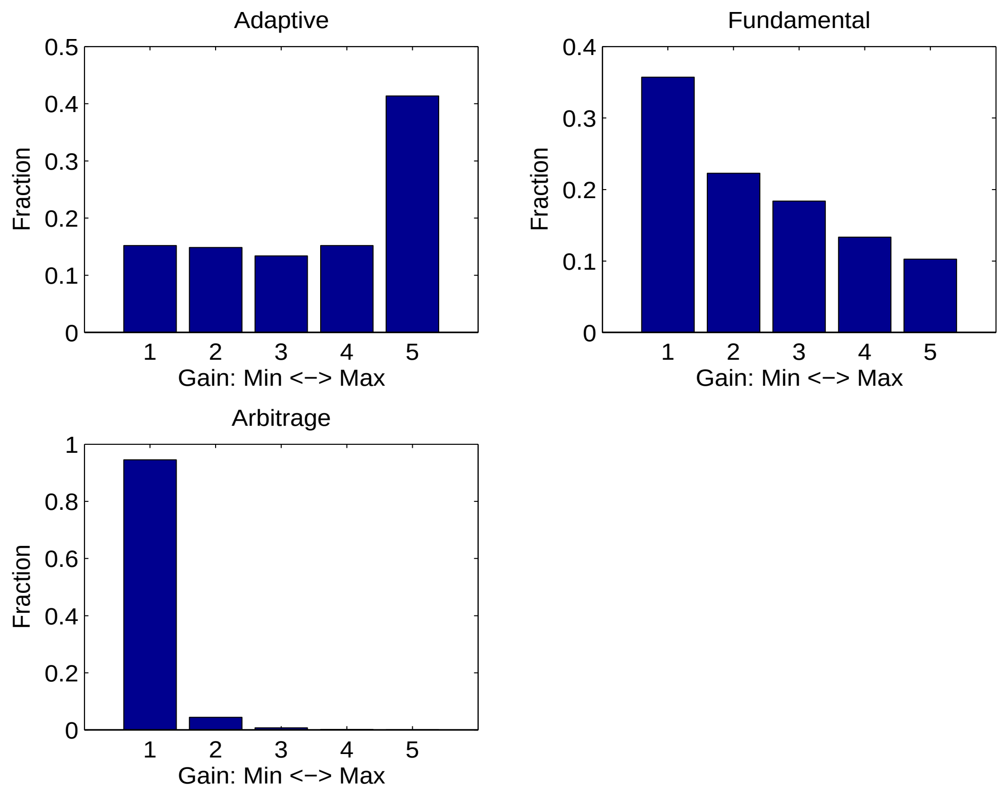

Figure 5: **Forecast Gain Distributions**

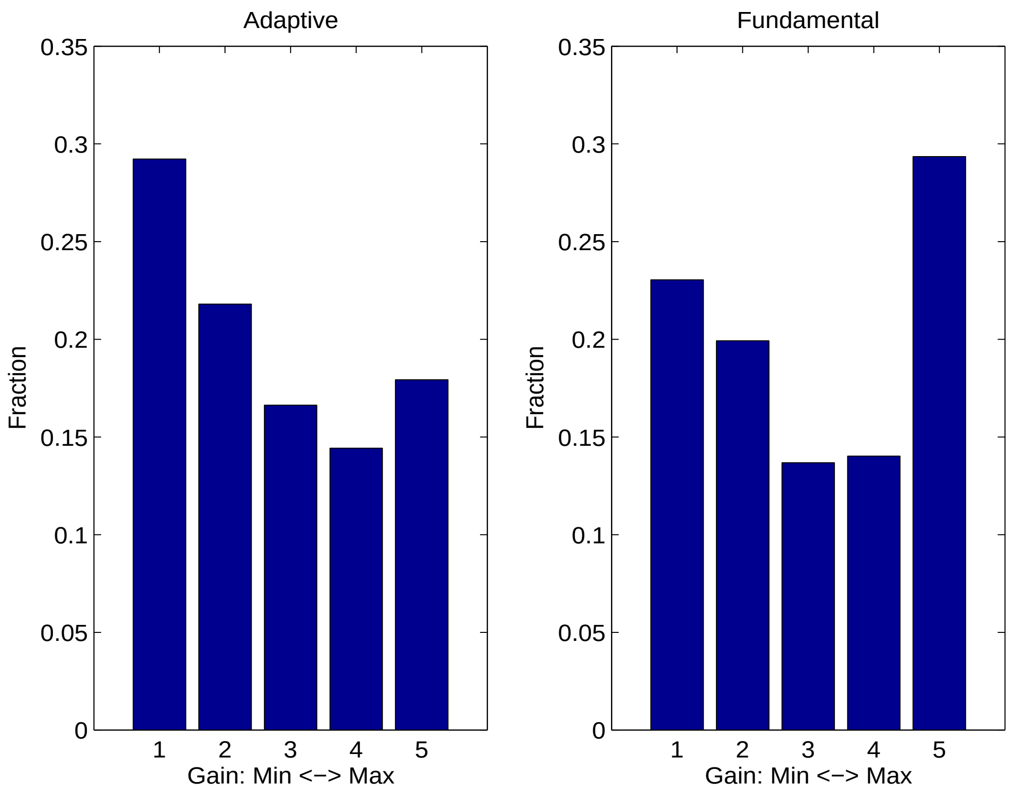

Figure 6: **Volatility Forecast Gain Distributions**

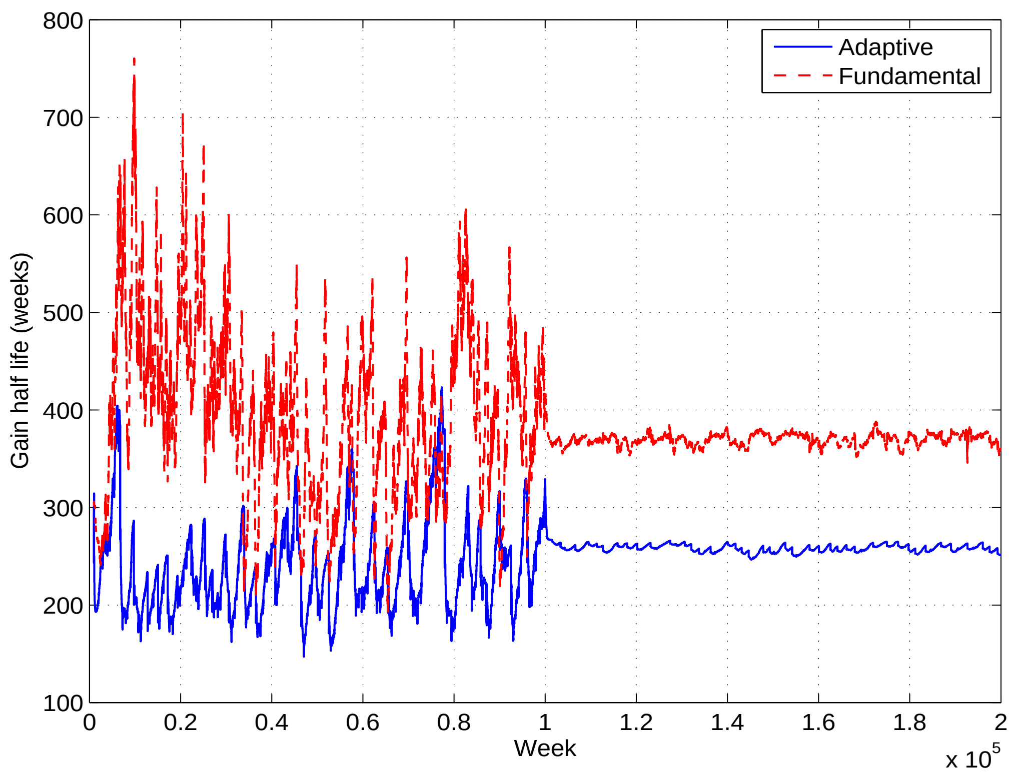

Figure 7: **Forecast Gain Convergence**

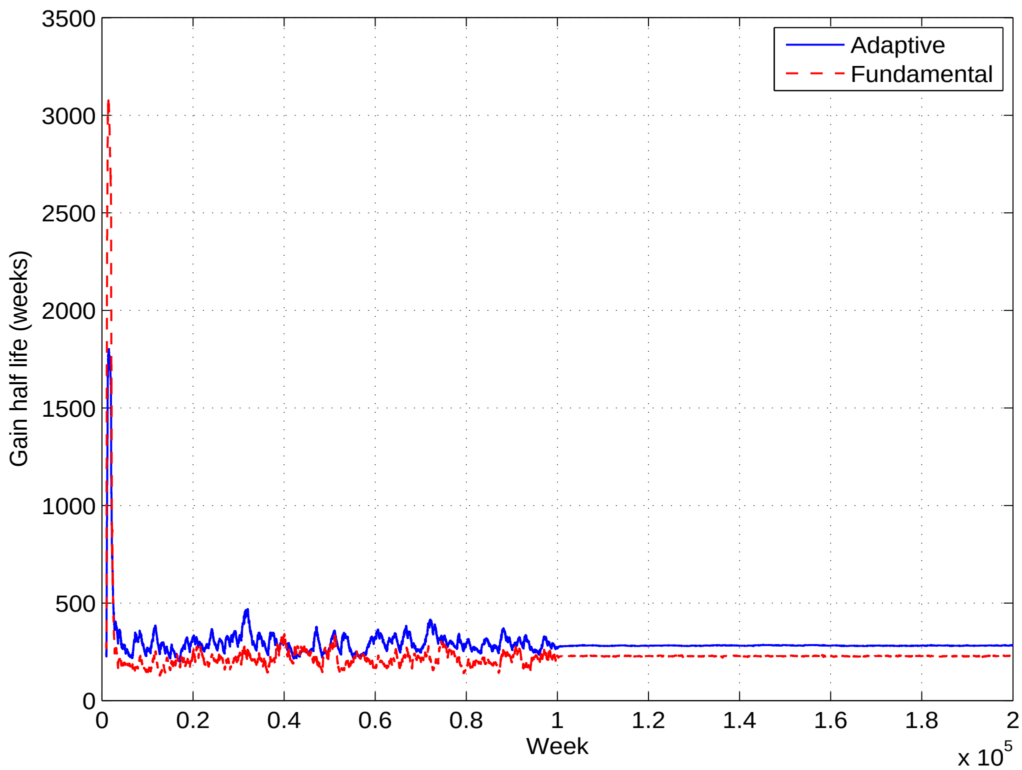

Figure 8: **Volatility Forecast Gain Convergence**

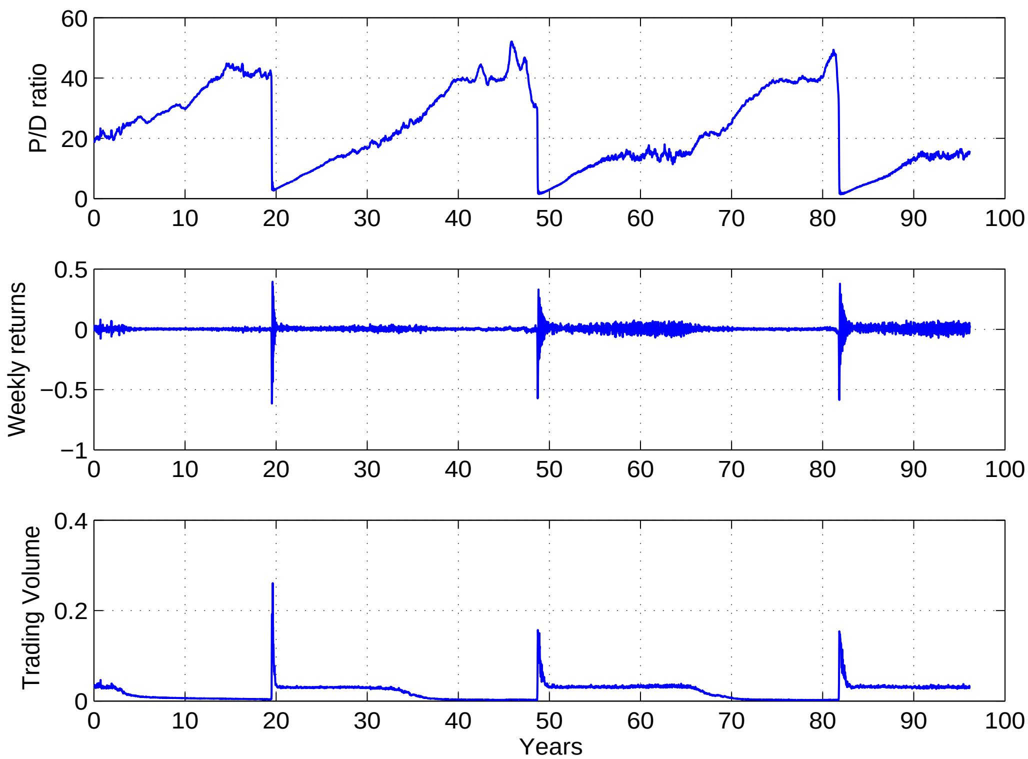

Figure 9: **Single Gain: p/d ratio, returns, volume**

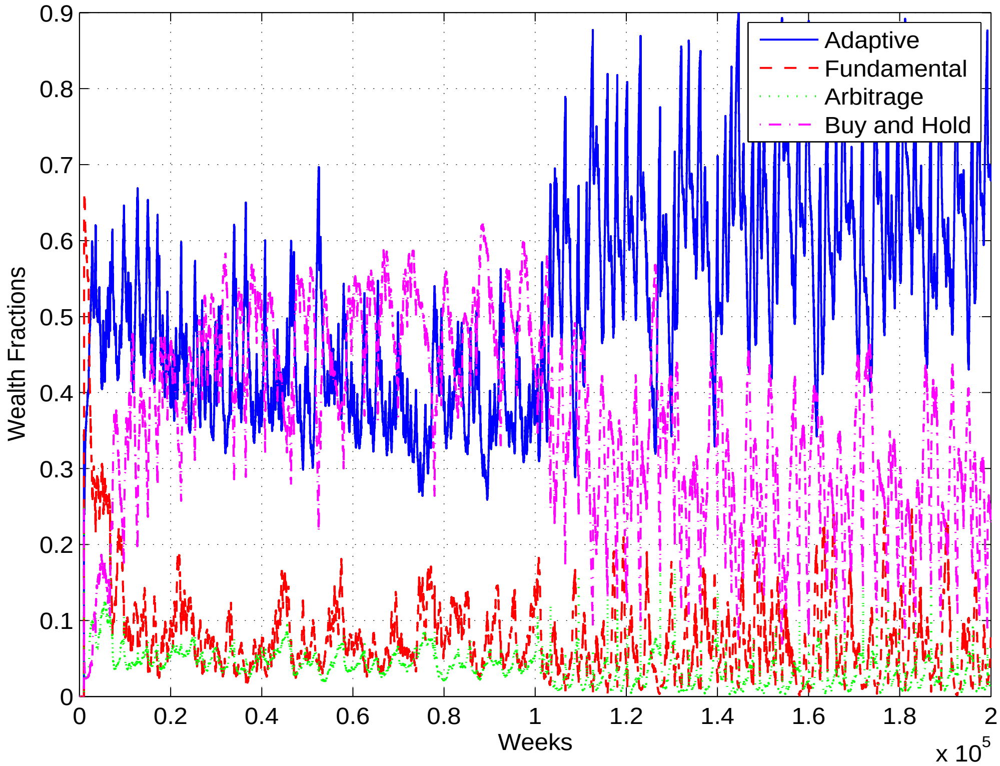

Figure 10: Wealth fractions across forecast strategies

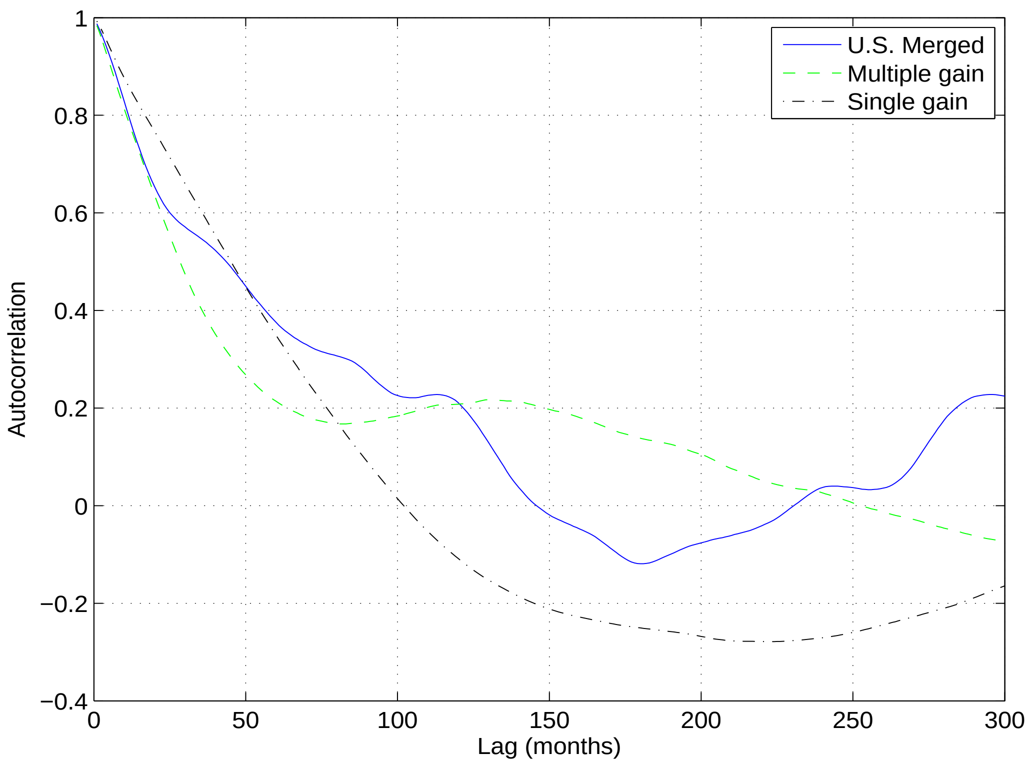

Figure 11: **Price/Dividend autocorrelations**

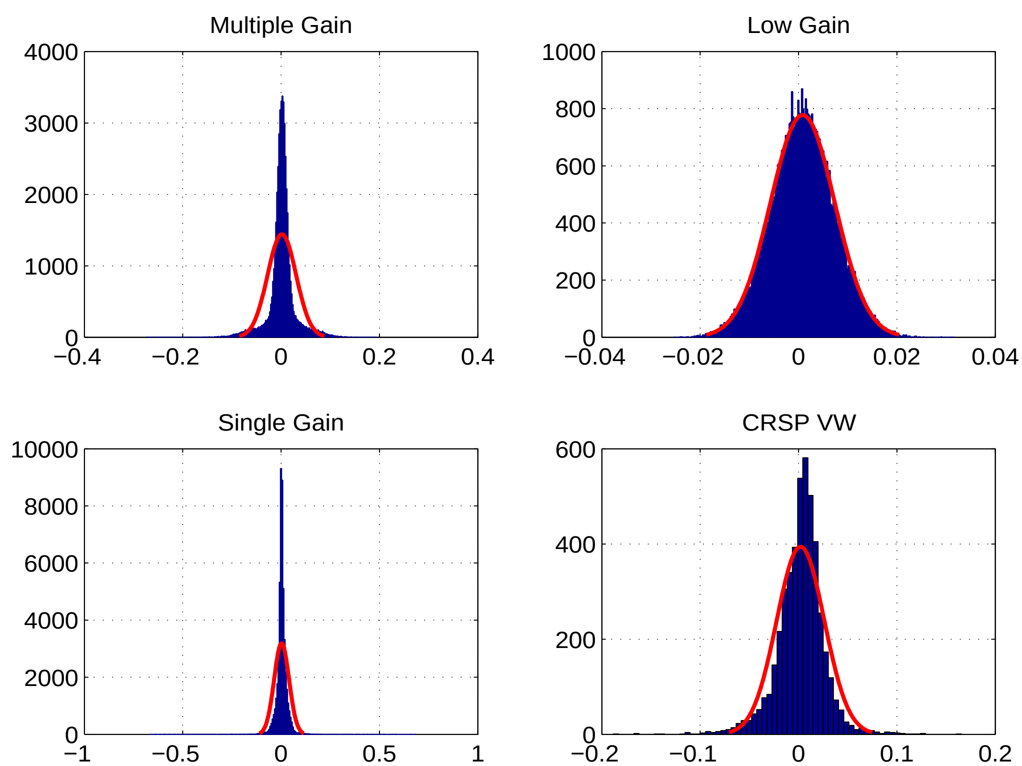

Figure 12: **Weekly return densities**

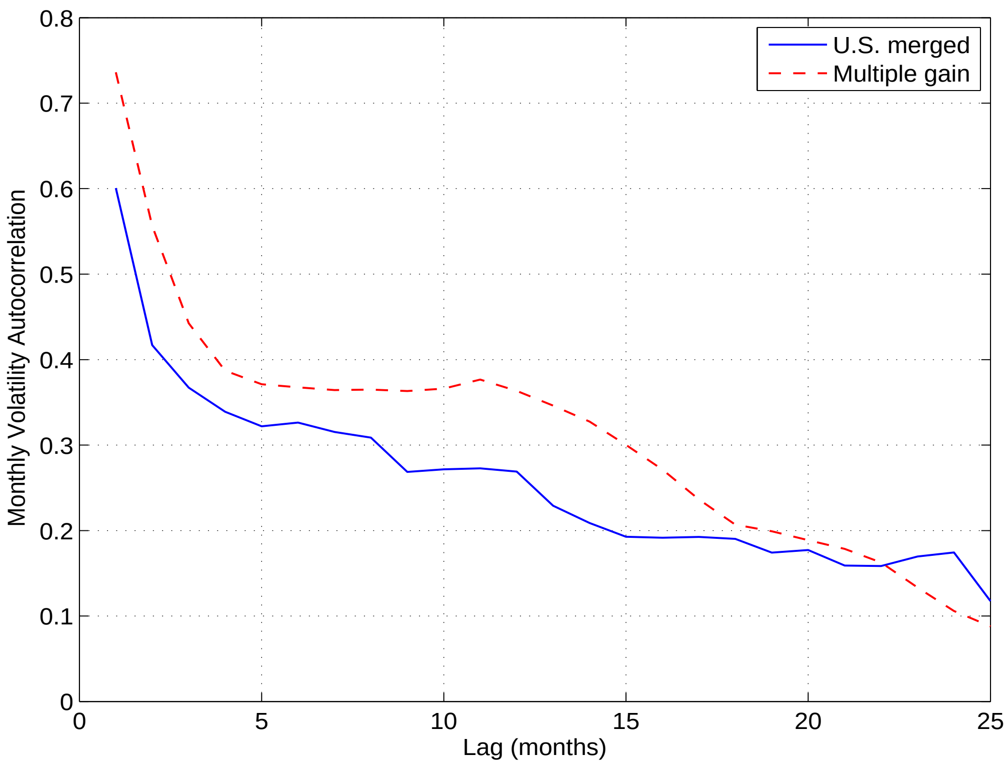

Figure 13: **Volatility Autocorrelations**

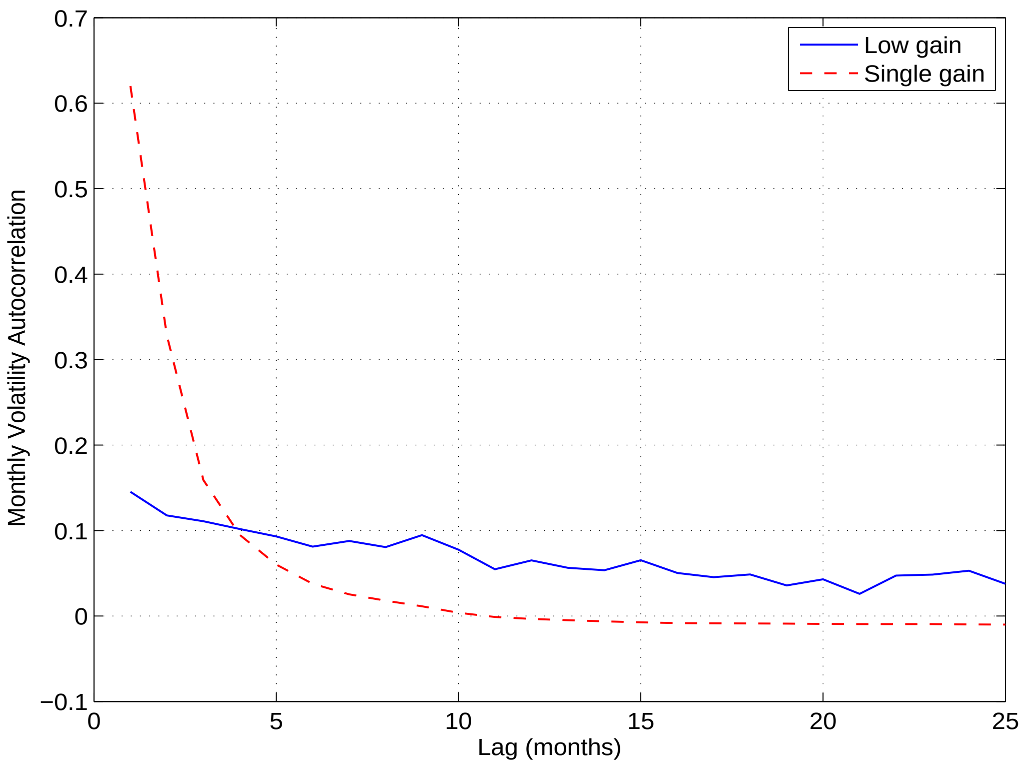

Figure 14: **Volatility Autocorrelations**