A Real Minsky Moment in an Artificial Stock Market

Blake LeBaron *
International Business School
Brandeis University

March 2012: Preliminary

Abstract

Many authors have contributed to the idea that financial markets are dynamically unstable. Much of this line of thinking suggests that bubbles and crashes will be a generic feature of most any speculative market, and that removing them would be difficult, if not impossible. Hyman Minsky is probably the one of the earlier contributors to this area, and recently his work has been looked at with new respect. This paper shows that some of his basic thinking about risk and extreme portfolio positions hold in agent-based financial markets in a way that appears close to his thinking. The advantage of this is that it presents a fully operational computer generated model with testable Minsky like effects which can be used to connect our understanding of Minsky to more modern, and rigorous approaches to macroeconomics.

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

The generic instability of financial markets has been discussed by many economists, but there is still no accepted theoretical framework for what makes a financial market special. One of the most important approaches to financial instability comes from the work of Hyman Minsky.[1] This paper presents an agent-based financial market that is shown to share many of the features of the classic Minsky style of instability. This adds to the debate on this subject by providing a runnable structural models which can be carefully explored to better understand the how the Minsky structure plays out in terms of dynamics.

Minksy conjectures that financial markets begin to build up bubbles as investors become increasingly overconfident about markets. They begin to take more aggressive positions, and can often start to increase their leverage as financial prices rise. Prices eventually reach levels which cannot be sustained either by correct, or any reasonable forecast of future income streams on assets. Markets reach a point of instability, and the over extended investors must now begin to sell, and are forced to quickly deleverage in a fire sale like situation. As prices fall market volatility increases, and investors further reduce risky positions. The story that Minsky tells seems compelling, but we have no agreed on approach for how to model this, or whether all the pieces of the story will actually fit together. The model presented in this paper tries to bridge this gap.

Heterogeneous agent-based models have been applied to financial markets for quite some time.[2] 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.[3] 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

[1] His classic book on the subject Minsky (1986) presents most of the basic ideas. Echoes of his framework can be felt in other work such as Kindleberger & Aliber (2005), or Reinhart & Rogoff (2011).

[2] 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.

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

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.

Several agent-based financial markets have highlighted the possibility for heterogeneity in the processing of past information by learning agents.[4] 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 paper begins by demonstrating that the model generates reasonable dynamics in terms of financial time series. This section is short since most of these results are presented in greater detail in LeBaron (2010). The critical dynamics around large price drops are then shown to exhibit very regular patterns that bear some resemblance to the Minksy agenda. The final sections look under the hood of the demand functions for the various strategies to show how they form the core for general financial behavior (with or without leverage), which is critical to generating instabilities.[5]

2 Model Structure

This section describes the 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 follow-

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

[5] Instability in all these cases drives from periods where, regardless of strategies, many agents are forced to liquidate positions of risky assets, in a kind of fire-sale situation. Examples of this can be seen in the agent-based world in Ussher (2008), and also Shin (2008) or Brunnermeir & Sannikov (2010).

ing 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.[6] Dividends are calibrated to the trend and volatility of real dividend movements from U.S. aggregate equity markets.[7] 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 nonlinear volatility generator for actual price series. The market structure also is important in that outside resources arrive only through the dividend flows entering the economy, and are used up only through consumption. The consumption levels are set to be proportional to wealth which, though unrealistic, captures the general notion that consumption and wealth must be cointegrated in the long run. Finally, prices are set to clear the market for the fixed supply of equity shares. The market clearing procedure allows for the price to be included in expectations of future returns, so an equilibrium price level is a form of temporary equilibrium for a given state of wealth spread across the current forecasting rules.

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

The empirical features of the market are emergent in that none of these are prewired into the individual trading algorithms. Some features from financial data that this market

[6] 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.

[7] Dividend calibration uses the annual Shiller dividend series available at Robert Shiller's 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.

replicates are very interesting. This would include the simple and basic feature of low return autocorrelations. In this market traders using short range autoregressive models play the role of short run arbitragers who successfully eliminate short run autocorrelations. They continually adapt to changing correlations in the data, and their adaptation and competition with others drives return correlations to near zero. This simple mechanism of competitive near term market efficiency seems consistent with most stories we think about occurring in real markets.

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

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

The constant d_g is the growth rate, or drift, for the log dividend process.[8] Time will be incremented in units of weeks. 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 $e^{d_g+(1/2)\sigma_d^2}$ which is approximately $D_g = d_g + (1/2)\sigma_d^2$.

[8] Lower case variables will represent logs of the corresponding variables.

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

$$R_t = \frac{P_t + D_t - P_{t-1}}{P_{t-1}}, \tag{2}$$

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,

$$\sum_{i=1}^{I} S_{t,i} = 1. \tag{3}$$

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 all 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,

$$W_{t,i} = P_t S_{t,i} + B_{t,i} + C_{t,i} = (P_t + D_t) S_{t-1,i} + (1 + r_f) B_{t-1,i}, \tag{4}$$

where $W_{t,i}$ represents the wealth at time t for agent i. Consumption at time t by agent i is given by $C_{t,i}$.

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,

$$\max_{\alpha_{t,i}} \frac{E_t^i W_{t+1,i}^{1-\gamma}}{1-\gamma}, \tag{5}$$

$$st. \quad W_{t+1,i} = (1 + R_{t+1,i}^p)(W_{t,i} - C_{t,i}), \tag{6}$$

$$R_{t+1,i}^p = \alpha_{t,i} R_{t+1} + (1 - \alpha_{t,i}) R_f. \tag{7}$$

$\alpha_{t,i}$ 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,

$$\alpha^*_{t,i} = \frac{E^i_t(r_{t+1}) - r_f + \frac{1}{2}\sigma^2_{t,i}}{\gamma\sigma^2_{t,i}}, \tag{8}$$

with $r_t = \log(1 + R_t)$, $r_f = \log(1 + R_f)$, and $\sigma^2_{t,i}$ is agent i's estimate of the conditional variance at time t. This is perturbed by a small amount of individual noise to give the actual portfolio choice for agent i,

$$\alpha_{t,i} = \alpha^*_{t,i} + \epsilon_{t,i}. \tag{9}$$

Where $\epsilon_{t,i}$ is an individual shock designed to make sure that there is some small amount of heterogeneity to keep trade operating.[9] It is normally distributed with variance, σ^2_ϵ.

In the current version of the model neither leverage nor short sales are allowed. The fractional demand is restricted to $\alpha_{t,i}$ with $\alpha_L \leq \alpha_{t,i} \leq \alpha_H$ with $\alpha_L = 0.05$ and $\alpha_H = 0.95$. 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

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

for intertemporal hedging.[10]

2.3 Expected Return Forecasts

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

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

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

$$\overline{(p-d)}_t = (1-g_L)\overline{(p-d)}_{t-1} + g_L(p-d)_{t-1} \tag{13}$$

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

$$\bar{\sigma}^2_{pd,t} = (1-g_L)\bar{\sigma}^2_{pd,t-1} + g_L((p-d)_t - \overline{(p-d)}_t)^2 \tag{15}$$

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

The forecasts used will combine four linear forecasts drawn from well known forecast families.[11] The first of these is an adaptive linear forecast which corresponds to,

$$f^j_t = f^j_{t-1} + g_j(r_t - f^j_{t-1}). \tag{16}$$

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

[10] 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.

[11] 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).

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

compete against each other in the market. Finally, this forecast will be trimmed in that it is restricted to stay between the values of $[-h_j, h_j]$. These will be set to relatively large values, and are randomly distributed across the j rules.

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

$$f_t^j = \bar{r}_t + \beta_t^j((p-d)_t - \overline{(p-d)}_t). \tag{17}$$

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

The third forecast rule will be based on linear regressions, and is referred to as a "noise trader" strategy. It is a predictor of returns at time t given by

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

This strategy works to eliminate short range autocorrelations in returns series through its behavior, and $M_{AR} = 3$ for all runs in this paper. In this way it plays the role of a noise trader, albeit a purposeful one, from other models.

The previous two rules will be estimated each period using recursive least squares. There are many examples of this for financial market learning.[13] 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.[14]

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, $\bar{\sigma}^2_{r,t}$, 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.

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

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

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,

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

The fundamental regression is updated according to,

$$\begin{aligned} \beta^j_{t+1} &= \beta^j_t + \frac{g_j}{\bar{\sigma}^2_{pd,t}}((p-d)_{t-M_{PD}} u_{t,j}) \\ u_{t,j} &= (\tilde{r}_t - f_{j,t-M_{PD}}). \end{aligned} \tag{20}$$

For the lagged return regression this would be,

$$\begin{aligned} \beta^j_{t+1,i} &= \beta^j_{t,i} + \frac{g_j}{\bar{\sigma}^2_{r,t}}(r_{t-i} u_{t,j}) \qquad i = 1,2,3, \\ u_{t,j} &= (r_t - f^j_t) \end{aligned} \tag{21}$$

where g_j is again the critical gain parameter, and it varies across forecast rules.[15] In both forecast regressions the forecast error, $u_{t,j}$, is trimmed. If $u_{t,j} > h_j$ it is set to h_j, and if $u_{t,j} < -h_j$ it is set to $-h_j$. This dampens the impact of large price moves on the forecast estimation process.

2.5 Variance Forecasts

The optimal portfolio choice demands a forecast of the conditional variance as well as the conditional mean.[16] The variance forecasts will be generated from adaptive expectations

[15] 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 probably reasonable. This is done to avoid performing many costly matrix inversions.

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

as in,

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

$$e^2_{t,j} = (r_t - f^j_{t-1})^2, \quad (23)$$

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.[17] 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.[18] Finally, the gain level for the variance in a forecast rule, $g_{j,\sigma}$, is allowed to be different from that used in the mean expectations, g_j. This allows for rules to have a different time series perspective on returns and volatility.

There is one further aspect of heterogeneity that is important to the market dynamics. Agents do not update their variance estimates immediately. They do it with a lag using a stochastic updating processes. Agent i will update to the current variance estimate for rule j, $\hat{\sigma}^2_{t+1,j}$, with probability p_σ.[19] This allows for a greater amount of heterogeneity in variance forecasts, and mitigates some extreme moves in price which can be caused by a simultaneous readjustment in market risk forecasts. This is a form of simulating more heterogeneity in the variance forecasting process, but in a stochastic fashion.

2.6 Market Clearing

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

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

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

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

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

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

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

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

This market is cleared for the current price level P_t. This needs to be done numerically given the complexities of the various demand functions and forecasts, and also the boundary conditions on $\alpha_{t,i}$.[20] 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.[21]

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

$$f^j_{t+1} = (1-g_j)f^j_t + g_j e_{t+1}. \tag{27}$$

This easily maps to the simple exponential forecast rule,

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

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

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

or

$$g_j = 1 - 2^{-1/m_h}. \tag{30}$$

The distribution of m_h then is the key object of choice here. It is chosen so that $\log_2(m_h)$

[20] 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 the appendix.

[21] 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.

is distributed uniformly between a given minimum and maximum value. The gain levels are further simplified to use only 5 discrete values. These are given in table 1, and are $[1, 2.5, 7, 18, 50]$ years respectively. In the long memory (low gain) experiments these five values will be distributed between 45 and 50 years.

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

2.8 Adaptive rule selection

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

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

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

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

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

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

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

3 Results and Experiments

3.1 Calibration and parameter settings

Table 1 presents the key parameters used in the simulation. As mentioned the dividend series is set to a geometric random walk with drift. The drift level, and annual standard deviation are set to match those from the real dividend series in Shiller's annual data set. This gives a recognizable real growth rate for dividends of 2 percent per year. The level of risk aversion, γ will be fixed at 3.5 for all runs. This is a reasonable level for standard constant relative risk aversion.[22] The gain range for the learning models is set to 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 $I = 10000$ agents, and $J = 10000$ 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.

3.2 Time series features

3.2.1 Weekly Series

This section will give a brief picture of the time series features generated by the model. As with many agent-based models most of the common empirical features of financial time series are replicated.

Figure 1 displays a 100 year snapshot of the end of a simulation. The top panel shows the variability of the price/dividend ratio. It is clear that the price moves around the dividend value, and does not appear to converge or settle down. There is also some indication of asymmetry in how it moves over time, with steep drops, and long slow returns. The middle panel shows the corresponding return series which clearly shows several extreme returns, and pockets of high and low volatility as is common in most financial series.

Table 2 presents some simple summary statistics for the baseline simulation shown in figure 1. Table 2 presents weekly summary statistics which reflect most of these early

[22]Many of the results can be replicated for a range of γ from $3 - 4$. The value of 3.5 gives some of the most realistic looking series while still being a reasonable level.

graphical features. The two comparison series are weekly returns with dividends for the CRSP value weighted index and IBM, covering all years from 1926-2009. The returns show a reasonable level of volatility in the row labeled Std. They are more variable than the CRSP index returns, but slightly less than those generated by weekly IBM returns. Although they are calibrated to aggregate returns, the agents trade this as an individual stock, so it is not clear which is the best comparison.

The next two rows report skewness and kurtosis estimates. These are all comparable across the 3 different return series. Kurtosis levels for the simulation are a little higher than those in the data. The last two rows in the table present the tail exponent which is another measure of the shape of the tail in a distribution. This estimate uses a modified version of the Hill estimator as developed in Huisman, Koedijk, Kool & Palm (2001), and further explored in LeBaron (2008) who shows it gives a very reliable estimate of this tail shape parameter. Values in the neighborhood of $2 - 4$ are quite common for weekly asset return series, so the results here are all within reasonable ranges.[23]

Figure 2 displays the return autocorrelations for the two series. The top panel displays autocorrelations for returns on the baseline simulation and the weekly CRSP series. They reveal the common result of very low autocorrelations in returns. The correlation at 2 days for the simulation is near 0.12 which is a little high, but not completely out of line with actual correlations. It should be noted that the autocorrelation at lag 1 is near zero for the simulation. This is interesting since one strategy, the noise trader, is designed to try to take advantage of just this type of predictability.

It is well known that the absolute magnitude of returns is 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 2 shows that this positive correlation holds for the simulation. Positive correlations in returns continue out well beyond 1 year (52 weeks) which is also common for many financial series.[24] 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 3 presents both the S&P price/earnings (P/E) ratio, and the price/dividend (P/D) ratio from the simulation. The simulation contains only one fundamental for the stock, and it can be viewed as

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

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

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 4 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 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 0.67 and 0.72 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 slightly high, real return of 11.1 percent as compared to 7.95 percent for the S&P. One conjecture is that much of this is driven by the increase in volatility and its contribution to the adjustment from geometric to arithmetic returns. This is examined in the next row, which reports the annual log real return for the two series. This shows a much closer relationship. The higher volatility reported earlier is repeated in the annual volatility estimates. The simulation gives an annual standard deviation of 0.27 which compares to a value of only 0.17 for the S&P. The last row reports the annual Sharpe ratio for the two series.[25]

4 Agent composition

This section turns to analyzing the properties of the agents after the market has stabilized into a steady state. It is important to know how wealth is distributed both across strategy types, and across gain levels within each strategy. This distribution of wealth is the key state of the market that controls dynamics as it moves through time.

Figure 5 shows the time series behavior of the fractions of wealth in the different strate-

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

gies. Wealth has some periods in which it fluctuates erratically, but over most of the simulations the proportions remain relatively stable. They reach a steady state after about 20,000 weeks have gone by. There still is a some extensive variability of wealth across rules even in this steady state. The ordering of the wealth series is important. The adaptive rules often control the largest share with about 40-45 percent. The next largest level goes to the buy and hold strategy with close to 40 percent. The fundamental strategy controls only about 10 percent of the wealth on average. The noise traders are the smallest with only about 5 percent of wealth.

The large amount of wealth in the adaptive strategy relative to the fundamental is important. The fundamental traders will be a stabilizing force in a falling market. If there is not enough wealth in that strategy, then it will be unable to hold back sharp market declines. This is similar to a limits to arbitrage argument. In this market without borrowing the fundamental strategy will not have sufficient wealth to hold back a wave of self-reinforcing selling coming from the adaptive strategies.

Agents are distributed across gain levels as well as over different strategy types. This form of heterogeneity is also important to the market dynamics. Figure figure 6 shows a snapshot of the distribution of gains for the different strategies. As noted before, the 5 discrete gain levels range from 2 years through 50 years in terms of half-lives. The fundamental strategy displays a relatively uniform pattern across gains. There is no tendency to select any particular range of adaptive learning which is interesting and important. The adaptive rules are different with a some selection toward higher gain (short memory) trading strategies. This does suggests that some fraction of wealth is using relatively short range momentum behavior in their trading decision making. Finally, the noise traders, who are running short horizon return auto regressions, concentrate on relatively low gain strategies. This suggests that they are locking down on relatively effective, and stable, use of their adaptive learning recursive least squares forecasts.

Figure figure 7 repeats the previous wealth distributions for the variance forecasting gain levels. Remember that agents are allowed to forecast volatility using a different gain level from what they use to forecast returns.[26] The distributions in figure 7 show a uniform pattern across all three forecasting rules. There is no tendency to converge to any specific horizon on these. This will be important to the dynamics of the market as will be shown in the next section.

[26]This may seem strange at first, but it is important to remember that the gain level may be driven by perceptions of signal/noise ratios as well as expectations about the usefulness of past data.

5 Crash dynamics

5.1 Dynamics around crashes

This dynamics of the market as shown in figure 1 appears to be defined by long run increases in the price of the asset or P/D ratio, followed by brief dramatic price declines, or market crashes. Figure 8 presents an event study summarizing these comovements over many periods. It uses the last 50,000 weeks of data. A crash is defined as a return which is below the 0.005 quantile in the return distribution. Once a day is determined as a crash day, then the next year (52 weeks) of returns are blocked. This is an attempt to make the classification detect the beginning of a period of distress in the market.

The upper left panel displays the results for trading volume. The pattern shows low volume before the crash hits, then a spike up, and persistent high volume after the crash. There is some indication that volume may start to increase before the actual date of the crash, but this may also be due to misdating of the event. A reverse pattern is given in the upper right panel for the equity fraction which falls dramatically after a crash. Again, this process begins slightly before the event date which could again be a dating issue.

The lower left panel displays the cross sectional mean for the coefficient used in the P/D regression equation. In this case the value tends to be closer to zero before the crash, and then it swings dramatically negative after. There is a 52 week lag which corresponds to the length of the regression. Since this forecast looks 52 weeks into the future, the crash, and the corresponding P/D ratio, will remain in the regression until 52 weeks after the crash. After this point the value begins to increase. The relatively large values prior to the crash are indicative of learning algorithms which have not observed a crash in the recent data, and are beginning to forget what a large P/D ratio implies for conditional return forecasts.

The lower right panel is the most important of the figures in terms of price dynamics. This figure estimates the actual demand elasticity estimated from the aggregate demand curve using all the individual agent forecasts. The sign of this indicates whether the demand is well behaved around a period of financial instability. A negative sign corresponds to a well behaved, downward sloping demand curve. The sign, and demand slope, is negative on average, but the values are quite different around a crisis. It appears to turn positive about 5 weeks before a crises starts as shown by the dark vertical line in the figure. It also continues positive for several weeks after the crises. If one were looking for a reliable "early warning" indicator, then this would work for this model. Unfortunately, in actual financial data we don't get to observe the slope of the demand curve, so it would not be feasible.

5.2 Demands and mechanisms

In order to get a better sense for the mechanism behind market instability figure 9 plots the excess demand curves in both a normal and a crash period. The prices in both plots are normalized so that the current price is 1. The latter is defined as a period where the price fall is larger than 10 percent. The first demand curve is taken as the last period in the benchmark simulation, and the crash is the last price fall in the simulation run. This plot is typical of what the two different demand situations look like in many simulations runs. In the normal case the demand curve is well behaved, and it is clear that small wiggles in this curve will generate only small changes in the equilibrium price. The crash market demand shows how it contributes to instability. The demand curve is near vertical around the current price leading to a dramatic drop in price after a small shift in the demand curve.

Figure 10 splits the demand curves apart by strategy family, and also decomposes the demands as follows. Total demand for shares is given by,

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

Split this into,

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

While both of these aggregate demands are complicated expressions, the first is more likely to be destabilizing, since it only involves the function $\alpha_i(P_t)$. The fraction of wealth for the risky asset can be increasing in P_t since this would be a characteristic of the adaptive strategy. The first part is a kind of active strategy response to a price change. The second component is a portfolio composition component, since it will change even for a fixed equity fraction. It corresponds to the need to adjust the portfolio in response to a price change in order to maintain appropriate strategy fractions. For example, if the equity price were to fall, investors holding a constant wealth fraction in equity, would need to respond by purchasing more equity to maintain current wealth fractions. This buying activity in response to a falling price is a major stabilizing force.[27]

Figure 10's four panels correspond to the four different strategies. Demands are now

[27] One could carefully work out the components of price impact on demand both through the portfolio composition and α, but this will be complicated by the fact that the final analysis depends critically on the overall wealth weights on different agents in the market.

in units of actual shares on the x-axis. Comparisons across strategies give an idea of each strategy's size in the market. The solid blue line corresponds to total demand, the dashed red line is the active strategy component of demand, and the dot-solid green line corresponds to the portfolio composition part.

The upper left panel corresponds to the adaptive forecast strategy. In this case there is a backward bending component to the active strategy component of demand (dashed line). A price increase for this strategy will lead to an increase in its forecast of expected returns, and therefore the destabilizing nature of the strategy is made clear in the demand for shares. Adding the composition component diminishes the destabilizing parts of the demand as was previously discussed.

The upper right panel displays the fundamental strategy. Expected returns in this case fall as the price rises, and this is realized through a downward sloping active component in the demand curve. The curve eventually goes vertical as the strategy reaches its minimum share demand. Again, the addition of the composition further smoothes the demand curve.

The lower left panel corresponds to the noise trader component. This can vary depending on the distribution of the regression parameter across noise traders which can be positive or negative. In this case it displays a backward bending part in both components of its demand, and therefore in the total demand. It should be noted that this type of trader corresponds to a small part of the market as shown by the low level of share demands. The final panel corresponds to the buy and hold strategy. These traders put little weight on current price changes in terms of their portfolio fractions. This gives an active component which is vertical. After the addition of the composition component these agents display a well behaved, downward sloping, demand curve.

Figure 11 gives another view of how the market dynamics is tied to agents' perceptions of volatility or risk as the market moves through its cycles. The upper panel displays the price/dividend ratio again. The middle panel shows the equity fraction for the fundamental strategy only. However, it is divided into two categories, one corresponding to low variance gain risk forecasts (blue solid line), and one corresponding to high variance gain forecasts (dashed red line). The low variance gain agents are less sensitive to short run swings in risk and are generally immune to Minsky like effects of misperceiving market risk. They follow a well defined strategy of backing off their positions as the P/D ratio increases, and they come back strong after a crash. The high gain strategies do not follow the "fundamental program" as well because of their added sensitivity to short run risk. When the market rises, they begin to increase their equity positions. They are respond-ing to their perceptions of risk which are displayed in the dashed red line in the lower panel. As the market reaches a top, they perceive risk as being low, even though a crash is eminent. The solid blue line in the bottom panel corresponds to the risk perceptions of the low gain forecasters who use a near constant long range risk measure. The high gain fundamental agents also back off their positions as when the P/D ratios are low. This is their response to their heightened risk after a crash. Crudely speaking, they don't have the guts to maintain their aggressive positions in the face of market turmoil.

Since figure 7 has shown that fundamental forecasts are evenly spread across all the variance gain levels, there is a good fraction of wealth that will not be implementing the fundamental strategy as well as the other agents, and contributes to both market instability, and probably their overall relative performance in the population. In terms of the entire market, this generates instability from a larger fraction fully invested in equity at the market top. It also contributes to the limits to arbitrage feature in that the fundamental strategy cannot gain enough of the overall wealth to be able to stabilize the market.

6 Conclusions

This paper has presented results of a stylized agent-based financial market. The model passes the usual hurdles of replicating most of the key features of financial time series including, long range fundamental deviations, volatility persistence, and fat tailed return distributions. The fact that this is possible in an agent-based market is not particularly new. However, for this market, the underlying structure reveals aspects of agent behavior which point to generic stabilizing and destabilizing properties that may cary over into real markets.

The dynamics are dominated by somewhat irregular swings around fundamentals, that show up as long persistent changes in the price/dividend ratio. Prices tend to rise slowly, and then crash fast and dramatically with high volatility and high trading volume. During the slow steady price rise, agents using similar volatility forecast models begin to lower their assessment of market risk. This drives them to be more aggressive in the market, and sets up a crash. All of this is reminiscent of the Minksy market instability dynamic, and other more modern approaches to financial instability.

Instability in this market is driven by agents steadily moving to more extreme portfolio positions. Much, but not all, of this movement is driven by risk assessments made by the traders. Many of them continue to use models with relatively short horizons for judging market volatility. These beliefs appear to be evolutionarily stable in the market. When short term volatility falls they extend their positions into the risky asset, and this eventually destabilizes the market. Portfolio composition varying from all cash to all equity yields very different dynamics in terms of forced sales in a falling market. As one moves more into cash, a market fall generates natural rebalancing and stabilizing purchases of the risky asset in a falling market. This disappears as agents move more of their wealth into the risky asset. It would reverse if they began to leverage this position with borrowed money. Here, a market fall will generate the typical destabilizing fire sale behavior shown in many models, and part of the classic Minsky story. Leverage can be added to this market in the future, but for now it is important that leverage per se is not necessary for market instability, and it is part of a continuum of destabilizing dynamics.

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.

Bollerslev, T., Engle, R. F. & Nelson, D. B. (1995), ARCH models, *in* 'Handbook of Econometrics', Vol. 4, North-Holland, New York, NY, pp. 2961–3038.

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

Brock, W. A. & Hommes, C. H. (1998), 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model', *Journal of Economic Dynamics and Control* **22**(8-9), 1235–1274.

Brunnermeir, M. K. & Sannikov, Y. (2010), A macroeconomic model with a financial sector, Technical report, Princeton University, Princeon, NJ.

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.

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

Farmer, J. D. & Geanakoplos, J. (2008), 'The virtues and vices of equilibrium and the future of financial economics', *Complexity* **14**, 11–38.

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* L. Tesfatsion & K. L. Judd, eds, 'Handbook of Computational Economics', Elsevier, pp. 1109–1186.

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.

Kindleberger, C. P. & Aliber, R. Z. (2005), *Manias, panics and crashes: A history of financial crises*, John Wiley.

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* L. Tesfatsion & K. L. Judd, eds, 'Handbook of Computational Economics', Elsevier, pp. 1187–1233.

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

LeBaron, B. (2010), Heterogeneous gain learning and the dynamics of asset prices, 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.

Minsky, H. P. (1986), *Stabilizing an Unstable Economy*, Yale University Press.

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.

Reinhart, C. M. & Rogoff, K. (2011), *This Time is Different: Eight Centuries of Financial Folly*, Princeton University Press.

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

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

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

Shin, H. S. (2008), *Risk and Liquidity*, Oxford University Press.

Ussher, L. (2008), 'A speculative futures market with zero-intelligence', *Eastern Economics Journal* **34**(4), 518–549.

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	0.06
r_f	0
γ	3.5
λ	0.0006
I	10000
J	10000
g_j	$[2, 50]$ years
g_u	17 years
L	10 percent/year
$[\alpha_L, \alpha_H]$	$[0.05, 0.95]$
σ_ϵ	0.03
M_{PD}	52 weeks
h_j	$[0.005, 0.050]$

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 Return Statistics*

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

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

Table 3: *Annual Statistics*

	Baseline	Shiller Earnings	Shiller Dividends
Mean(P/D)	20.42	15.32	26.62
Std (P/D)	7.27	5.97	13.81
Autocorrelation(1)	0.68	0.72	0.93
Mean(Return)	11.10		7.95
Mean(Log(Return))	6.82		6.22
Std(Return)	0.27		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 1871-2008. 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.

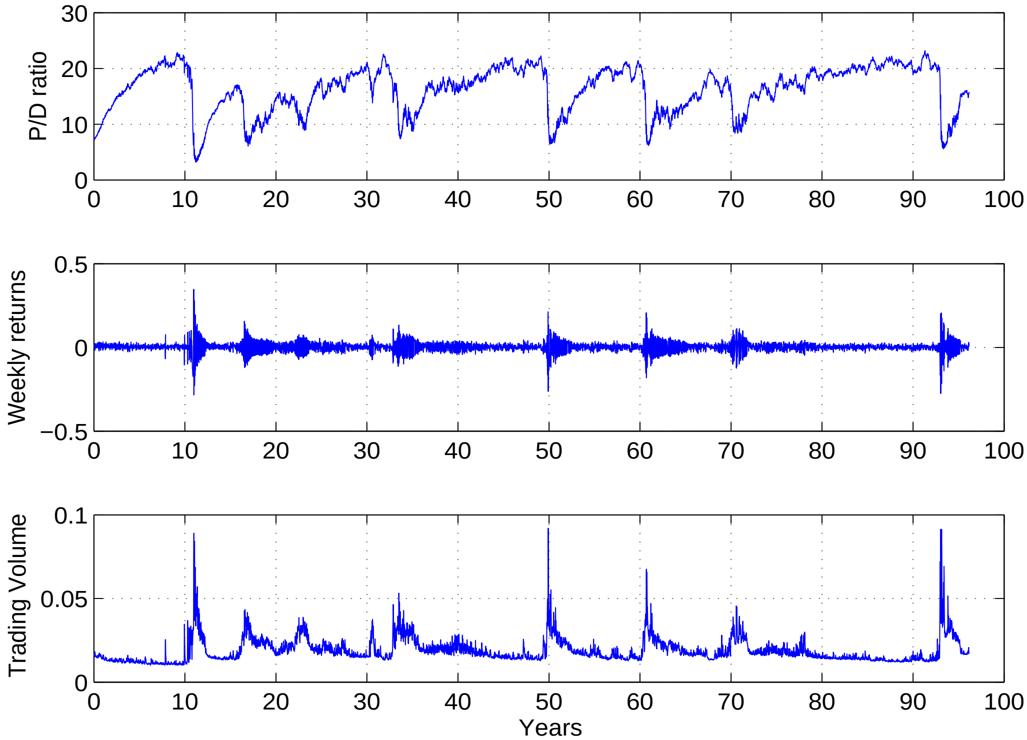

Figure 1: **Basic simulation dynamics**

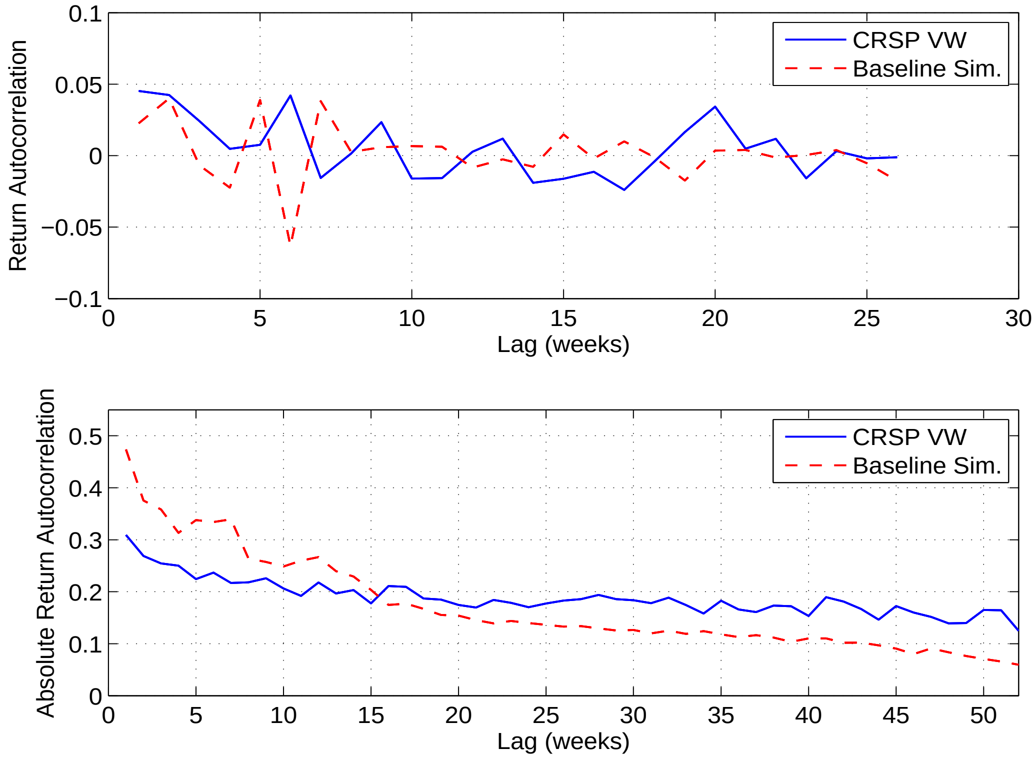

Figure 2: **Return Autocorrelations: Returns and absolute value of returns**

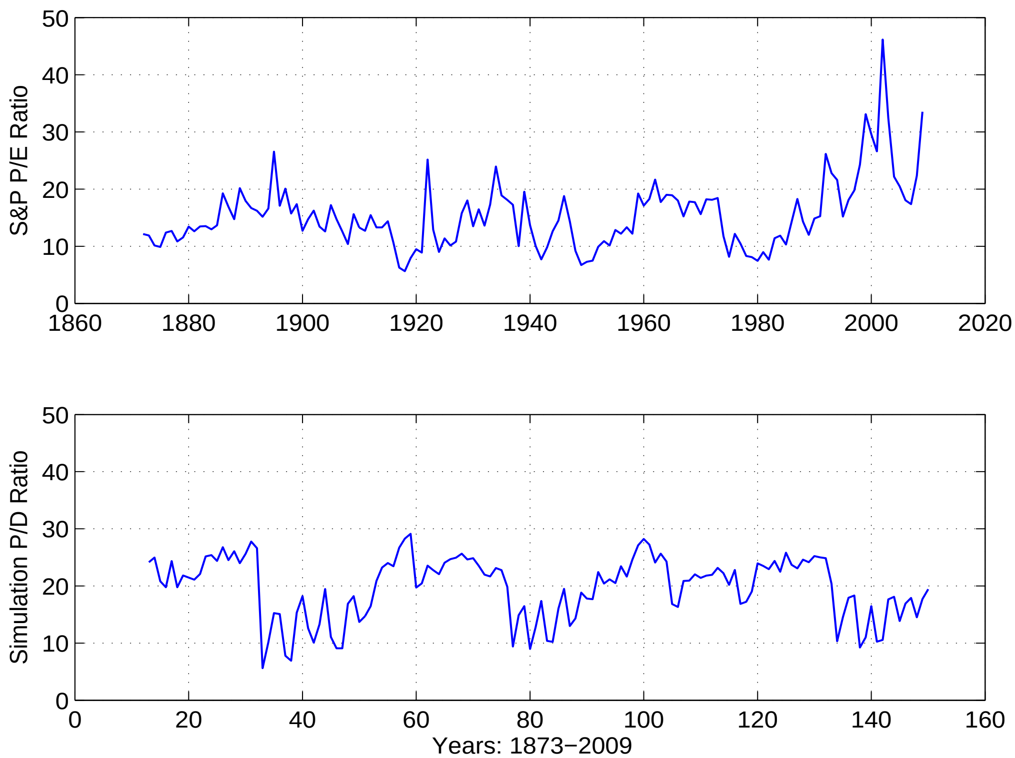

Figure 3: **P/E ratios: Annual 1871-2009**

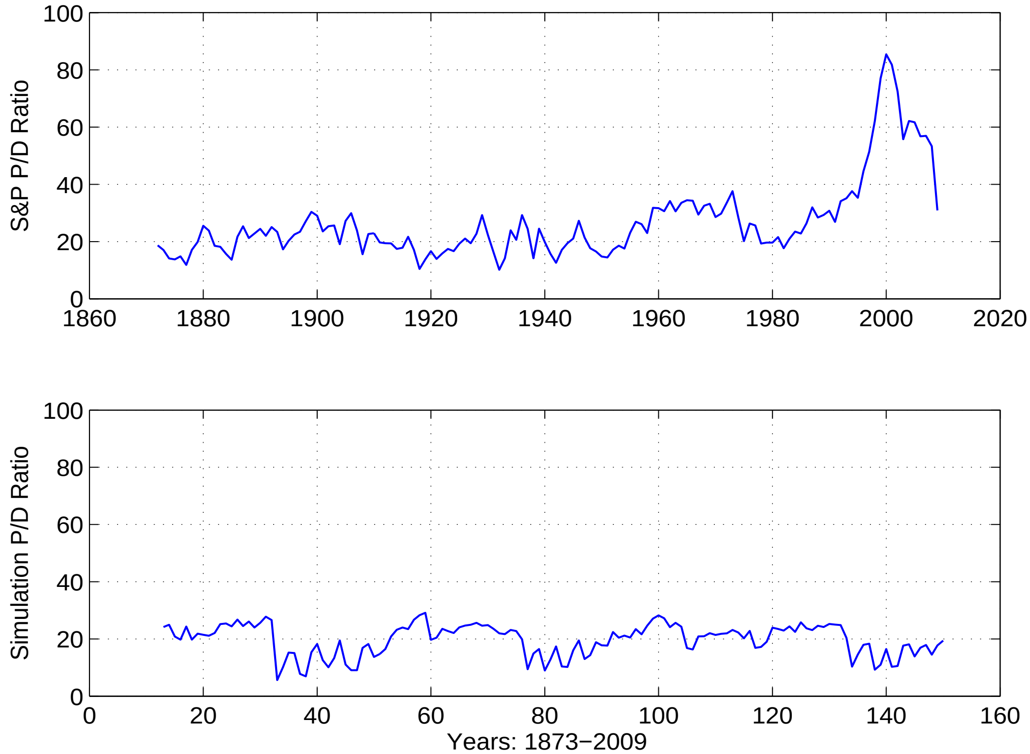

Figure 4: **P/D ratios: Annual 1871-2009**

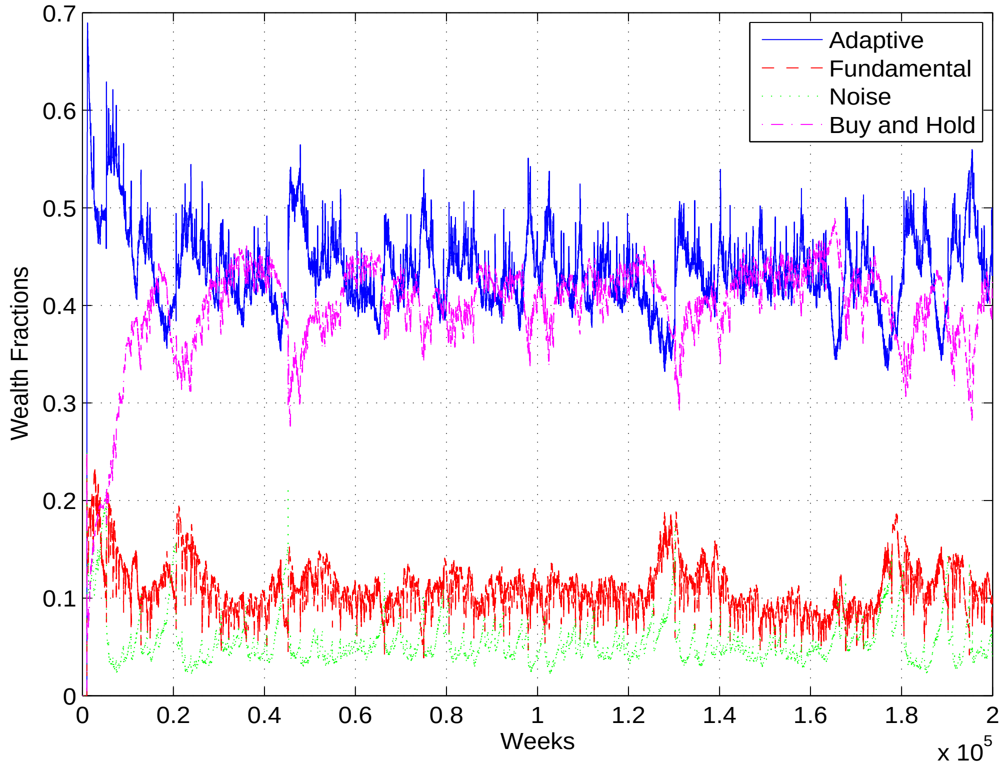

Figure 5: **Wealth Fraction Time Series**

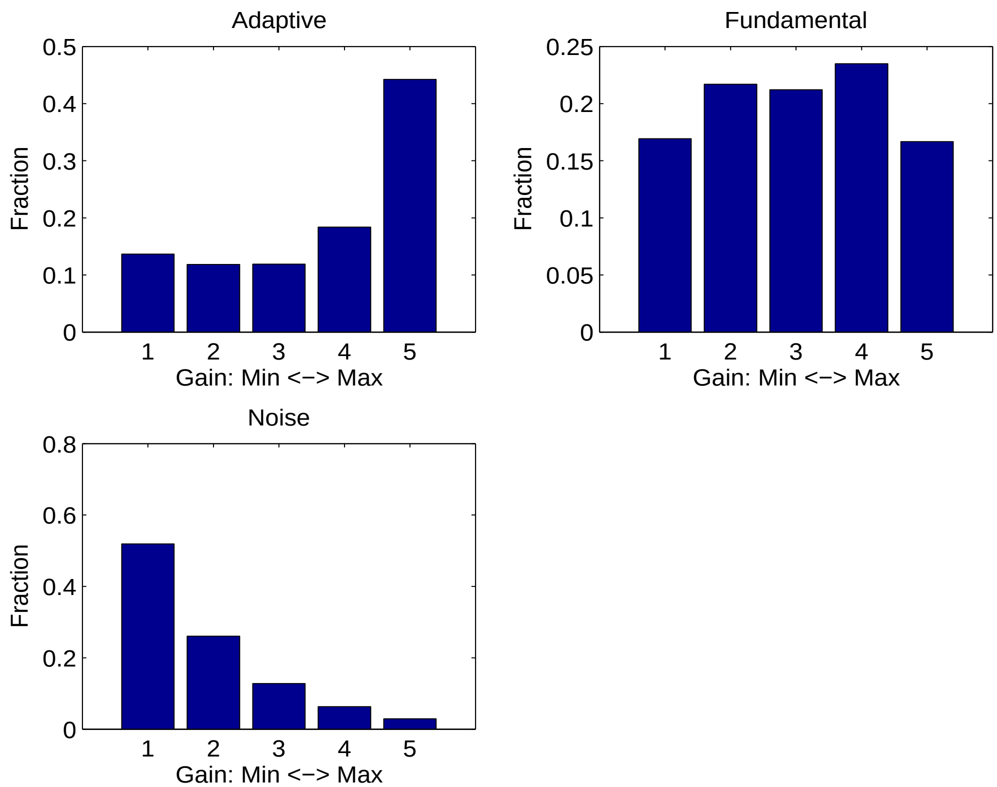

Figure 6: **Gain Wealth Distributions**

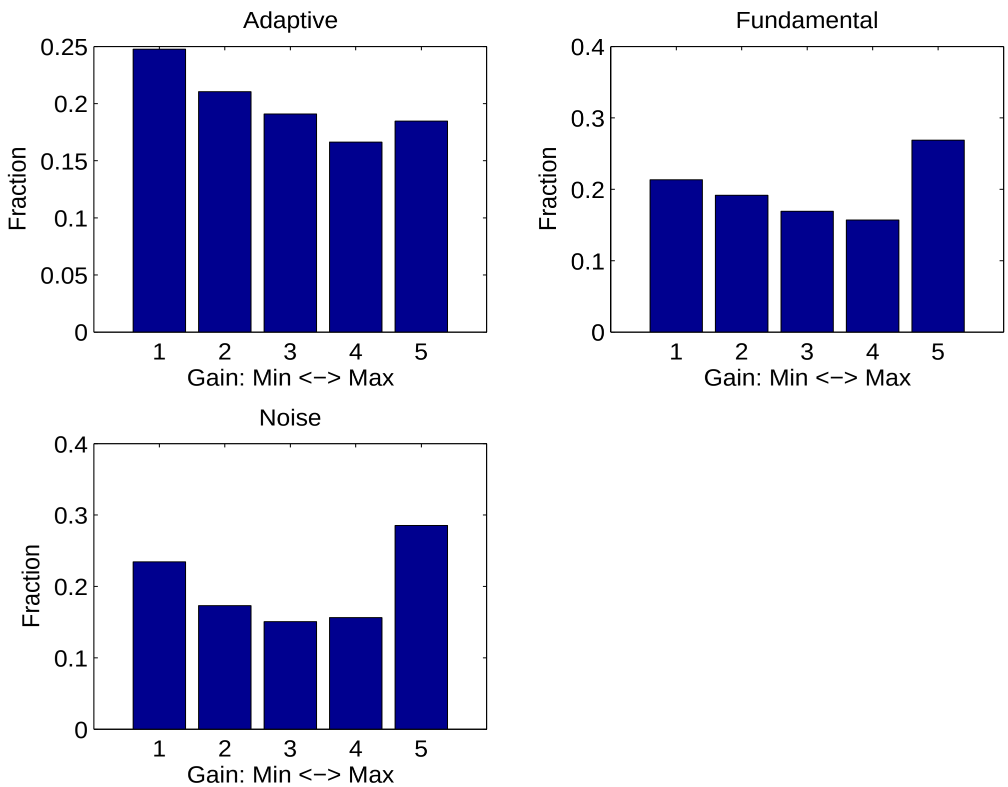

Figure 7: **Volatility Gain Wealth Distributions**

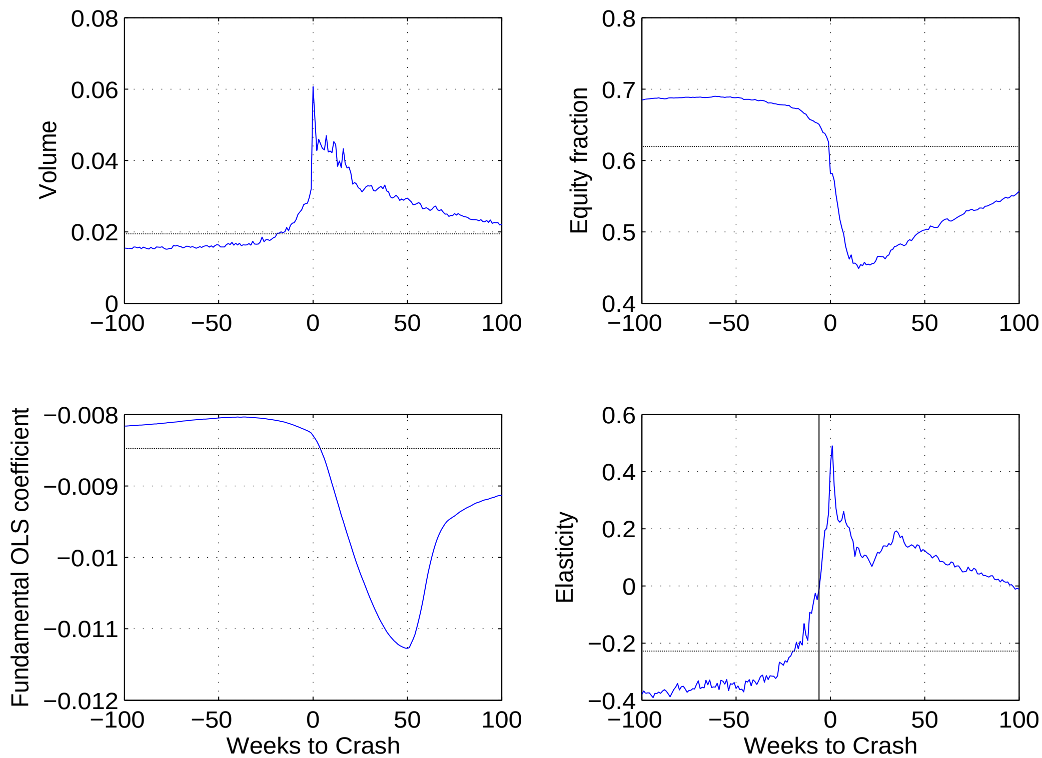

Figure 8: **Crash event dynamics**

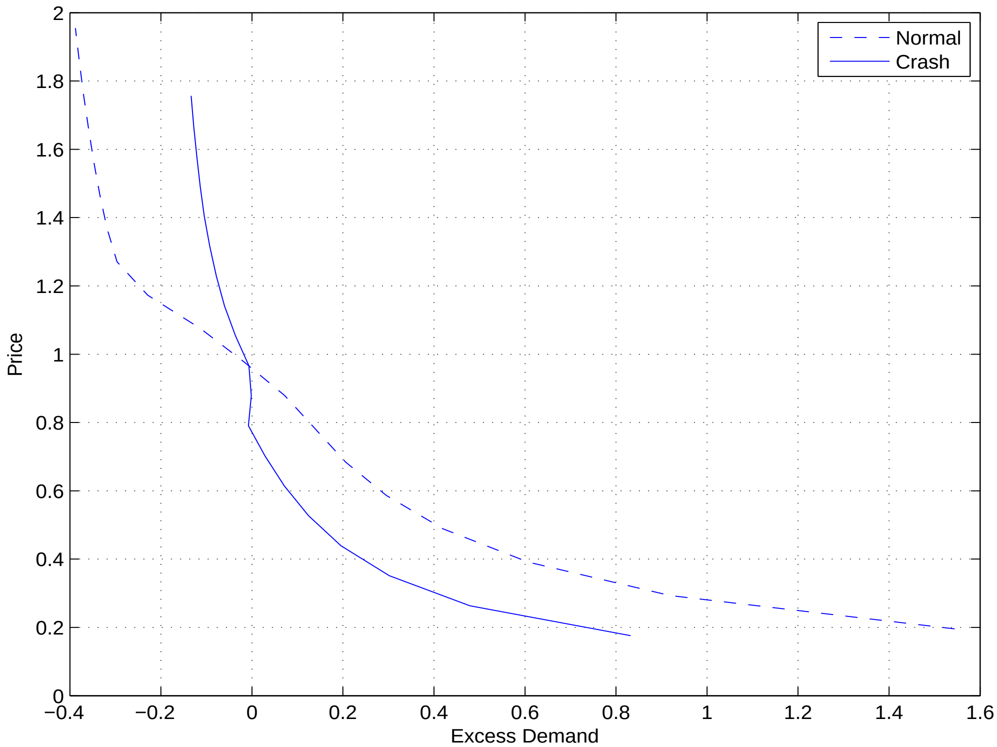

Figure 9: **Excess Demands**

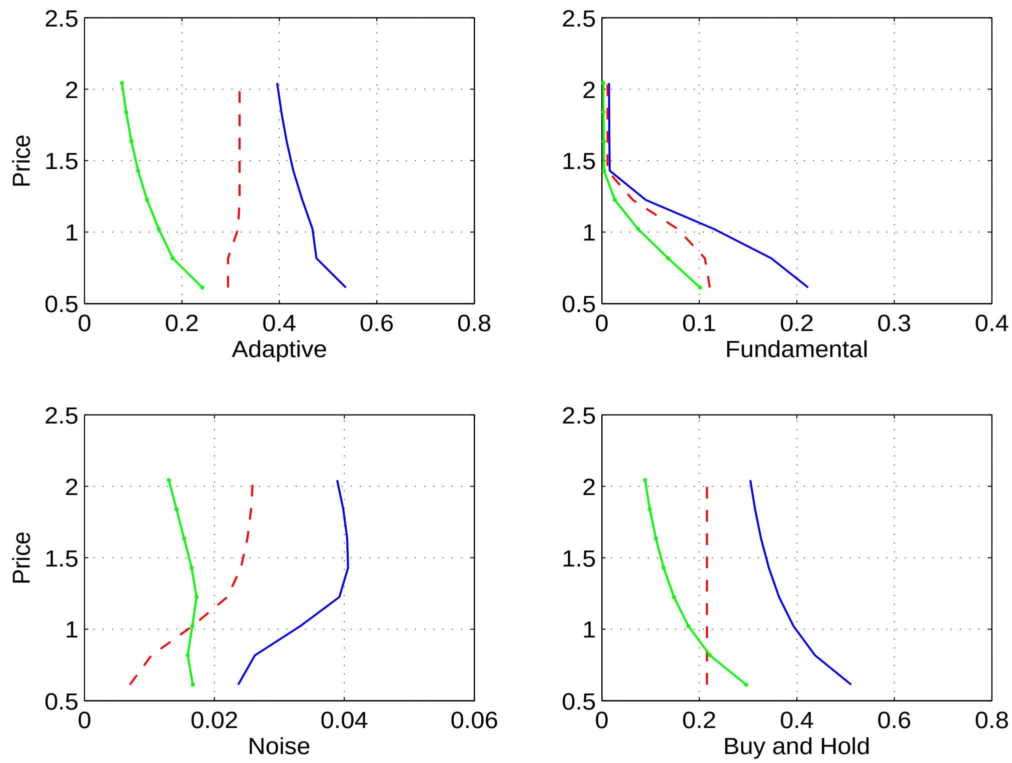

Figure 10: **Demand by Strategy**

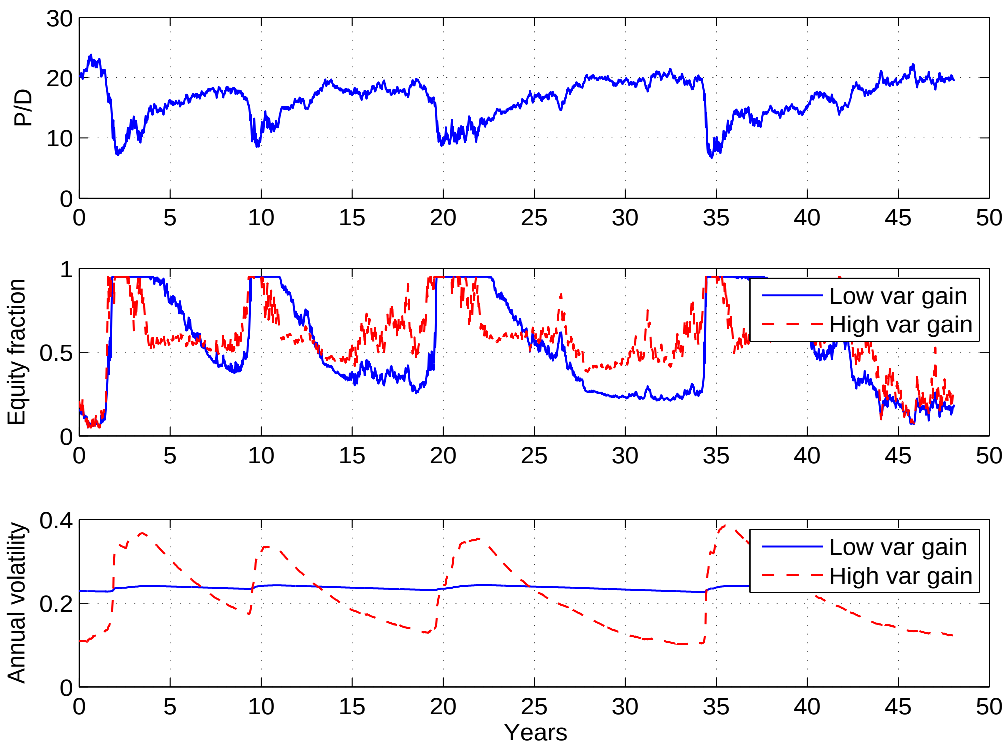

Figure 11: **Fundamental portfolio strategies and volatility**