

Endogenous Rational Bubbles

George A. Waters¹
Department of Economics
Illinois State University
Normal, IL 61790-4200

August 29, 2012

¹gawater@gmail.com

Abstract

This paper presents an asset pricing model that allows for heterogeneous forecasting strategies that satisfy rational expectations. An exponentially weighted replicator dynamic describes how agents switch between a forecast based on fundamentals, a rational bubble forecast that uses extraneous information and a reflective forecast, which is a weighted average of the former two. If the innovations to the extraneous martingale have a similar magnitude to those of the dividend process and agents are sufficiently aggressive in switching forecasting strategies, a significant portion of the population may adopt the rational bubble forecast. Tests on simulated data show excess variance in the price and persistence and volatility in the price-dividend ratio that is not explained by representative agent models satisfying the efficient markets hypothesis. The data also matches the stylized fact that returns are unpredictable in the short run. Conditions determining the frequency and duration of episodes where a significant fraction of agents adopt the rational bubble forecast leading to large deviations in the price-dividend ratio are discussed.

JEL Classification: C22, C73, G12, D84

Keywords: evolutionary game theory, rational bubble, heterogeneous forecasts, return predictability, excess variance

1 Introduction

Though the efficient markets hypothesis (EMH) has been a dominant paradigm in asset pricing for decades, the assumption that there is a representative expectation for an asset price contradicts the observed heterogeneity of forecasts, and is at odds with the popular perception that bubbles are a common phenomena in asset markets. Further, the strong version of the EMH, meaning asset prices are determined solely by expectations of fundamental information, cannot explain important features of the data, such as the volatility of the price-dividend ratio.

There are models that allow for bubbles. Models of rational bubbles (Blanchard (1979), Evans (1991)) are appealing, since expectations are unbiased, and the model matches the stylized fact that prices and returns are unpredictable in the short run. However, it is unclear how agents could coordinate on a single forecast based on extraneous information when an alternative forecast based on the strong EMH is available. Models with heterogenous behavioral forecasting strategies, such as Brock and Hommes (1998) and LeBaron (2010), can also produce large deviations in the asset price and price-dividend ratio from the predicted values of the strong EMH, though such approaches involve strategies do not satisfy rationality.

The present paper explains how agents with a choice of forecasting strategies could adopt a rational bubble forecast leading large deviations in the asset price from the predictions of the EMH. The outbreak of such bubbles can explain the persistence and volatility of the price-dividend ratio and the excess volatility in the price while also producing unforecastable returns. Further, the model provides a description of factors underlying the frequency and duration of bubbles.

An evolutionary game theory dynamic describes how agents switch between forecasting strategies based on their past performance, given by payoffs based on forecast errors, as in Parke and Waters (2007). Agents choose from a fundamental forecast, which corresponds to the strong EMH prediction, a mystic forecast, which uses an extraneous martingale as in the rational bubble model, and a reflective forecast, which is a weighted average of the former two forecasts. The reflective forecast embodies all the information available to the agents including the other forecasts and their relative popularity and is the unique unbiased forecast in an environment with heterogeneous forecasting strategies. The behavior of all agents in the present work satisfies the *cognitive consistency principle*, described in Evans and Honkapohja (2011), which specifies that agents in a model are as

smart as economists. More precisely, agents form expectations using reasonable models according to economic theory.

In the spirit of evolutionary stability, small fraction of agents experiments with mysticism, and this strategy can gain a significant following if agents are sufficiently aggressive in switching to better performing strategies and the shocks to the martingale are similar in magnitude to those of the dividend process¹. Bubbles in the asset price are associated with such outbreaks of mysticism, but endogenously collapse given the assumption that a small fraction of agents do not abandon the fundamental forecast.

Simulation results demonstrate that bubble episodes can explain a number of empirical features of the data. The underlying dividend process is calibrated to the annual data used in Shiller (2005). Simulations where heterogeneity in the forecasting strategies is common demonstrate persistence and volatility in the price-dividend ratio similar to that in the data, but not explained by the EMH. The simulated asset price also shows excess variance documented in Shiller (1981). In addition, the returns are unpredictable in the simulated data for any parameter choices. Parke and Waters (2007) demonstrate that heterogeneity in forecasting strategies in this model can explain ARCH and GARCH effects and excess kurtosis in returns, which models with a unique rational bubble forecast cannot. While these are important implications, this paper focuses on annual data, so these issues are not given detailed attention.

There are a number interesting alternative approaches to asset pricing that involve deviations from the strong EMH. Adam, Marcet and Niccolini (2008) and Lansing (2010) are able to match a number of the features of the U.S. stock market data. In the model in the former paper, a representative agent updates its estimate of the long run growth rate of the asset price, which is used for forecasting. In Lansing (2010), the forecasting model (perceived law of motion) includes a geometric random walk, making bubbles a possibility, and agents update a parameter in the forecasting model that determines the impact of the bubble. The agent in Branch and Evans' (2011) model of bubbles updates an estimate of the conditional variance of the return using a linear model. The time series implications of this approach have yet to be explored in detail.

The representative agent models referenced above all update parameters in a linear model used for forecasting, deviating from rational expectations but satisfying the cognitive consistency

¹These results are demonstrated analytically in Parke and Waters (2012).

principle. However, these models do not account for the observed heterogeneity of forecasts in asset markets, and how their forecasting strategies would perform in the presence of alternatives is an open question. Furthermore, to rule out explosive behavior, these models must restrict the parameters in the forecasting model used by the agent.

There are a number of models with heterogeneous forecasting strategies. As noted above, these approaches usually involved strategies that do not satisfy rationality. In LeBaron (2010), some agents use a "buy and hold" strategy, which has intuitive appeal but does not satisfy cognitive consistency. The cognitive consistency of agents in the asset pricing models of Brock and Hommes (1998) and Branch and Evans (2007) is open to interpretation². In the former paper, some agents have perfect foresight but must pay a cost. In contrast, the reflective forecast in the present work is constructed using all information available to the agents, and does not require a cost. In Branch and Evans (2007), agents use underparameterized models, which exclude available information that affects the asset price regardless of the choice of forecasting strategies, unlike the martingale in the present model.

The paper is organized as follows. Section 2 give details about the asset price model with heterogeneous expectation, while section 3 presents the dynamic describing the evolution of the forecasting strategies. Section 4 describes the simulations and the conditions for the formation of bubbles. Section 5 gives the results of formal tests on the simulated data, and Section 6 concludes.

2 Asset Pricing

This section specifies the three forecasts and the resulting realization of the asset price, which thereby determines the forecast errors for each strategy. The underlying motivation is the standard asset pricing equation

$$p_t = \alpha p_{t+1}^e + d_t, \tag{1}$$

where the asset price is p_t , the dividend is d_t and the parameter α is the discount factor. This model is not fully sufficient for our purpose, since there is a unique representative forecast of the price. Brock and Hommes (1998) develop a model with mean-variance optimization where investors

²These papers are part of a large literature using the multinomial logit dynamic to describe the evolution of heterogeneous forecasts. See Hommes (2006) for a survey.

choose between a riskless and risky asset in constant supply. With risk neutral agents who have a common belief about the variance of the returns, the model with heterogeneous forecasts can be written as

$$p_t = \alpha \sum_{h=1}^n x_{h,t} e_{h,t} + d_t + C \quad (2)$$

where the vectors $e_t = (e_{1,t}, \dots, e_{n,t})$ and $x_t = (x_{1,t}, \dots, x_{n,t})$ are the different forecasts of p_{t+1} and the fractions of agents using the forecasts, respectively. The constant C is a risk premium.

The forecasts considered are motivated by the multiplicity of solutions to the model (1) in the homogeneous case. According to the strong efficient markets hypothesis (EMH), the price is given by the discounted expected future dividends as given by the following solution to the model (1).

$$p_t^* = d_t + \sum_{j=1}^{\infty} \alpha^j E_t(d_{t+j})$$

Agents referred to as *fundamentalists* adopt the forecast $e_{2,t}$ determined by the above solution.

$$e_{2,t} = E_t(p_{t+1}^*) = \sum_{j=1}^{\infty} \alpha^{j-1} E_t(d_{t+j-1}) \quad (3)$$

However, this solution is not unique. As discussed in the rational bubble literature, Evans (1991) for example, there is a continuum of solutions to (1) of the form

$$p_t^m = p_t^* + \alpha^{-t} m_t$$

where the stochastic variable m_t is a martingale such that $m_t = m_{t-1} + \eta_t$, for *iid*, mean zero shocks η_t . Though the information contained in the martingale m_t may be extraneous, if agents believe that information is important, it does affect the asset price. Agents that adopt the forecast $e_{3,t}$ based on the rational bubble solution above are called *mystics*, and their forecast is as follows.

$$e_{3,t} = E_t(p_{t+1}^m) = E_t(p_{t+1}^*) + \alpha^{-t-1} m_t \quad (4)$$

A primary objection to such a solution is that it violates a transversality condition, see Lundqvist and Sargent (2004, section 13.6). As pointed out by Lansing (2010), an agent could profitably

short the risky asset if the prices follows such a path. However, this hypothetical agent would need to be infinitely lived with unlimited resources or ability to borrow. Furthermore, agents in the present model can adopt or abandon the forecast at any time so this objection to the mystic forecast is not a major concern.

Both the mystic and fundamental forecasts satisfy rational expectations in that they are unbiased in the homogeneous case. However, our goal is to allow for possible heterogeneity in forecasting strategies, so we introduce the *reflective* forecast, which satisfies rational expectations even in the presence of heterogeneity. The reflective forecast $e_{1,t}$ is an average of the alternative forecasts used in the population weighted according to the relative popularity.

$$e_{1,t} = (1 - n_t) e_{2,t} + n_t e_{3,t} \quad (5)$$

where

$$n_t = \frac{x_{3,t}}{x_{2,t} + x_{3,t}}$$

The variable n_t is the relative popularity of mysticism among agents using mysticism or reflectivism.

Reflectivism depends on alternative strategies, so to ensure its existence, we make the following key assumption.

Assumption: The fraction of fundamentalists $x_{2,t}$ never falls below some minimum $\delta_2 > 0$.

This assumption is not particularly restrictive, considering that in most asset pricing models, all investors are fundamentalists. Given these three forecasting strategies (3), (4) and (5) and the asset pricing model allowing for heterogeneity (2), the realization of the asset price is

$$p_t = p_t^* + \alpha^{-t} n_t m_t. \quad (6)$$

One can verify that the reflective forecast has the same form as the realization of the price such that $e_{1,t} = E_t p_{t+1}$. The reflective forecast embodies the "beauty contest" characterization (Keynes 1935) of asset markets in that agents use the martingale in their forecast only to the extent that other agents use it, not because they regard it as inherently important.

Agents evaluate the performance of the forecasting strategies by comparing payoffs based on squared forecast errors. Hommes (2001) shows that the mean-variance optimization underpinning

the model (2) is equivalent to minimizing squared forecast errors. Payoffs are defined as follows.

$$\pi_{i,t} = -(p_t - e_{i,t-1})^2 \quad (7)$$

The reflective forecast error U_t plays an important role in the payoffs to all three forecasting strategies, and is comprised of two terms.

$$U_t = (p_t^* - E_{t-1}(p_t^*)) + \alpha^{-t}(n_t m_t - n_{t-1} m_{t-1}) \quad (8)$$

The first term is the current period dividend payment, which is the new fundamental information. The second term embodies the new information about the martingale's impact on the asset price. The part in parentheses can be written as $n_t \eta_t - \Delta n_t m_t$ so both the innovation in the martingale and its relative use for forecasting within the population affect the reflective forecast error. The representation of U_t shows that the reflective forecast is unbiased, under the assumption that agents are unable to forecast n_t . The innovations to the dividend (d_t) and the martingale (η_t) and the change in n_t are all independent, mean zero, so the reflective forecast is unbiased. The reflective forecast satisfies rational expectations in the presence of heterogeneity.

The fundamental and mystic forecasts satisfy the weak efficient markets hypothesis in that the forecasts are unbiased, but their forecast errors are affected by the level of the martingale in the presence of heterogeneity. A key term in the payoffs is the weighted martingale $A_{t-1} = \alpha^{-t} m_{t-1}$. The reflective forecast depends only on U_t and, using (7) and (8), has payoff

$$\pi_{1,t} = -U_t^2. \quad (9)$$

Fundamentalism has forecast error $U_t + n_{t-1} A_{t-1}$, so its payoff is

$$\pi_{2,t} = -U_t^2 - 2n_{t-1} U_t A_{t-1} - n_{t-1}^2 A_{t-1}^2. \quad (10)$$

Similarly, the payoff to mysticism is as follows.

$$\pi_{3,t} = -U_t^2 + 2(1 - n_{t-1}) U_t A_{t-1} - (1 - n_{t-1})^2 A_{t-1}^2 \quad (11)$$

Much of the intuition behind the potential for mysticism to gain a following can be observed in the above three payoffs. The third terms in the payoffs to mysticism (11) and fundamentalism (10) are unambiguously damaging to those payoffs in comparison with the payoff to reflectivism (9). If there is heterogeneity in the choice of forecasting strategies ($0 < n_{t-1} < 1$), then mysticism and fundamentalism over- and under-react to the martingale. In expectation, the covariance $U_t A_{t-1}$ is zero, so reflectivism outperforms the other two strategies.

However, mysticism can outperform the other strategies in a given period. If the realization of the covariance $U_t A_{t-1}$ is positive and sufficiently large, the second term in (11) may outweigh the third term so that $\pi_{3,t} > \pi_{1,t} > \pi_{2,t}$. Such a positive covariance corresponds to a fortunate (for the mystic) correlation between the martingale and the innovations in the model. In distribution, dividends are uncorrelated with the martingale, but over a number of periods, such correlations are likely to occur. For mysticism to have a chance of success, the level of A_t must be large enough to that the covariance is significant, but not so large that the martingale terms dominate. Intuitively, a forecast like "Dow 36K" might attract a significant following, but "Dow 36b" would be dismissed.

The payoffs also display a herding effect, since the variable n_{t-1} representing the state of the choices of forecasting strategies in the population enters the payoffs. For mysticism to outperform reflectivism, the following condition must hold, $U_t A_{t-1} > \frac{1}{2}(1 - n_{t-1})A_{t-1}^2$. As noted above, the covariance term $U_t A_{t-1}$ must be positive and sufficiently large for this condition to be met, but the closer n_{t-1} gets to one, the less strict the requirement on this term becomes. If there is any inertia in the the choice of forecasting strategies, more agents adopting mysticism leads to a greater impact of the martingale on the asset price and improves the performance of the mystic forecast. Hence, the adoption of mysticism is partially self-fulfilling.

3 Evolutionary Dynamics

A generalization of the replicator dynamic, a workhorse in the evolutionary game theory literature, describes the evolution of the vector x_t of the fractions of agents using the different forecasting strategies. This dynamic allows for the parameterization of agents' aggressiveness in switching to better performing strategies, which is a key determinant for the potential adoption of mysticism. This section discusses the necessary conditions for the resulting endogenous emergence of rational

bubbles and the reasons for their eventual collapse.

Let the weighting function $w(\pi)$ be a positive, increasing function of the payoffs. The general replicator dynamic³ is

$$x_{i,t+1} - x_{i,t} = x_{i,t} \frac{w(\pi_{i,t}) - \bar{w}_t}{\bar{w}_t}, \quad (12)$$

where the expression \bar{w}_t is the weighted population average $\bar{w}_t = x_{1,t}w(\pi_{1,t}) + \dots + x_{n,t}w(\pi_{n,t})$. A strategy gains followers if its weighted payoff above the weighted population average, i.e. has positive fitness, in the language of evolutionary game theory. Such a dynamic is said to be *imitative* since strategies that are popular today, larger $x_{i,t}$, tend to gain more adherents if they are successful.

A general form for the dynamic (12) allows for a range of behavior of the agents. Compared to a linear weighting function $w(\pi)$, under a convex $w(\pi)$, agents switch to better performing strategies more quickly, see Hofbauer and Weibull (1996). A linear weighting function in the dynamic (12) corresponds to the special case of the replicator dynamic studied in Weibull (1998) and Samuelson (1997). Sandholm (2011) gives a thorough comparison of the features of a number of evolutionary dynamics. Waters (2009) discusses discrete time dynamics used in macroeconomic applications.

Using a version of the dynamic (12) with an alternate timing, Parke and Waters (2012) demonstrate that, for bounded dividends, the payoff to reflectivism is always above the population average. Therefore, under the replicator (linear $w(\pi)$), mysticism cannot take followers away from reflectivism. Under linear weighting, the covariance (second) terms in the payoffs to mysticism and fundamentalism, (11) and (10), cancel in the population average payoff⁴, but the third terms with A_{t-1}^2 do not. Since the payoff to reflectivism is unaffected by the martingale, it is larger than the population average, so reflectivism gains followers.

Reflectivism's dominance is weaker in the case of a convex weighting function. Here, a positive covariance term $U_t A_{t-1} > 0$ has greater benefit to mysticism than harm to fundamentalism, so it enters the population average payoff and, if it is large enough, mysticism can gain a following. The model used for simulations focuses on the exponential weighting function

$$w(\pi) = e^{\theta^2 \pi}, \quad (13)$$

³Parke and Waters (2012) focus on a dynamic with the same form, but slightly altered timing and perform simulations with the present form as a robustness check.

⁴Given the timing of the present version of the model, the covariance terms may not cancel out to zero, but their impact is minimal.

so θ parameterizes the aggressiveness of the agents. An increase in θ means that agents are switching more quickly to the best strategy, but as θ decreases the dynamic approaches the linear weighting case.

One drawback to imitative dynamics such as the generalized replicator (12) is their lack of inventiveness, see Waters (2009) for a discussion. If a strategy has no followers ($x_i = 0$) then it cannot gain any. Hence, game theorists usually focus on equilibria that are evolutionarily stable, meaning they are robust to the introduction of a small fraction of deviating agents. Similarly, the focus of the present class of models is whether the fundamental forecast is robust to the introduction of a small fraction of mystics.

It is possible for mysticism to gain a following given the following conditions. *i)* Some agents believe that extraneous information may be important to the value of an asset. *ii)* In some periods, the extraneous martingale is correlated with fundamentals. *iii)* Agents must be sufficiently aggressive in switching to superior performing strategies.

Mysticism cannot maintain a following indefinitely given the existence of a minimum fraction of fundamentalists δ_2 . If fundamentalism could be eliminated from the population, then $n_t = 1$ and the payoff to mysticism (10) is identical to the payoff to reflectivism (9), and the model collapses to a representative agent rational bubble model. However, the presence of a minimum fraction of fundamentalists implies that $n_t < 1$ and that the reflective and mystic forecasts are not identical. Since the expected value of the covariance term $U_t A_{t-1}$ in (10) is zero, reflectivism outperforms mysticism in the long run. Further, the magnitude of A_t (a submartingale) grows over time, so the third term in the payoff to mysticism (10) dominates and the performance of mysticism deteriorates over time. While mysticism can gain a following temporarily whereby the martingale affects the asset price, eventually agents abandon mysticism, so bubbles endogenously form and collapse. The goal of the simulations is to determine the quantitative importance of such outbreaks of mysticism.

Since it limits the life of bubbles, the minimum fraction of fundamentalists plays a similar role as the *projection facility* used with least squares learning as in Adam, Marcet and Niccolini (2008). Similarly, Lansing (2010) limits parameters so that agents focus on the one bubble of a continuum of solutions, that leads to stationarity in the first difference of the endogenous variable being forecast. In these models, a representative agent updates the estimate of the parameters in a forecasting rule, but the projection facility limits the acceptable estimates. The projection facility

places stronger restrictions on agents beliefs than the minimum fraction of fundamentalists. In the present model, a small fraction of agents rejects extraneous information, but under the projection facility, all agents have a sophisticated understanding of the long run implications of the choice of forecasting rules.

The present model represents a minimal departure from rationality when mystics are introduced into the population. Mysticism appears due to a disagreement about what constitutes fundamental information, but all agents form expectations with a reasonable economic model, i.e. agents meet the *cognitive consistency* principle described in Evans and Honkapohja (2011). Further, both mysticism and fundamentalism satisfy rationality in the homogeneous case, and reflectivism satisfies rationality when there is heterogeneity in the forecasting strategies, and this forecasting strategy is available to agents at all times. When mystics are eliminated from the population, the reflective and fundamental forecasts coincide. Only when mystics are introduced do the mystic and fundamental forecasts deviate from rationality, but mystics believe that the extraneous information in the martingale is relevant to the forecast of the asset price, and that other agents will eventually realize it. All agents believe that they are making efficient use of the available information.

4 Simulations

Simulation results of the model with the three forecasting strategies described above verify that the potential for outbreaks of mysticism depends on the aggressiveness of the agents in switching to better performing strategies and the magnitude of the shocks to the dividends and the martingale. Furthermore, for reasonable parameterizations, when a significant portion of the population adopts the mystic forecasting strategy, there can be bubble-like deviations in the asset price and price-dividend ratio.

The underlying dividend process is calibrated to the annual S&P 500 data used by Shiller (2005)⁵. Given the dividend d_t and the martingale m_t , the model is determined by the dynamic (12) along with the exponential weighting function (13), the payoffs (9), (10) and (11), and the realization of the asset price (6). The dividend process is specified as a stationary process with parameter choices below⁶.

⁵The data used is updated to included years up to 2011.

⁶LeBaron, Arthur and Palmer (1999) and Branch and Evans (2011) use stationary dividends. Adam, Marcet and

$$d_t = \bar{d} + \rho (d_{t-1} - \bar{d}) + v_t$$

\bar{d}	ρ	σ_v
0.1166	0.465	0.203

The constant \bar{d} is chosen so that for $\alpha = 0.95$, the steady state price-dividend ratio (log difference) is 2.66, which is close to the long run average for the S&P 500 from the Shiller data. The persistence parameter ρ and shocks $v_t \sim N(0, \sigma_v)$ are chosen to match values from the H-P detrended earnings series. Using a linear trend on post-war data gives a similar estimate. Since not all firms pay dividends, earnings are used as a proxy for d_t to allow for the inclusion of a larger number of firms.

Two other fixed parameters are the minimum fraction of fundamentalists $\delta_2 = 0.01$ and the fraction of mystics introduced into the population 0.001. The minimum fraction of mystics is set much smaller so that the introduction of mystics on its own does not have a quantitatively significant effect on the asset price (6) since the initial n_0 is small. If the dynamic used in the simulations is specified so that if the unconstrained dynamic (12) sets one of the fractions below its minimum, that fraction is set to its minimum, and the other two strategies split the remaining followers in the same proportion they would in the unconstrained case. If mysticism falls below its minimum, that level of followers is reintroduced and the martingale is restarted at $m_t = 0$.

4.1 Bubble frequency and duration

The free parameters θ , which measures agent aggressiveness, and σ_η , the standard deviation of the martingale innovations, play a large role in determining the potential for outbreaks of mysticism and bubbles. For such events to occur, agents must be sufficiently aggressive, meaning θ is sufficiently large, and the magnitude of the martingale innovations must be large enough to have a noticeable impact on the payoffs and the asset price, but not so large so that the third term in the payoff to mysticism (11) dominates.

Figures 1-5 demonstrate the role of the parameters θ and σ_η in determining the frequency and duration of bubbles. Figures 1 and 2 show the evolution of the price dividend ratio, the

Niccolini and Lansing (2010) both model dividends as a random walk with drift, which would complicate the present model and is left as a possibility for future work.

forecast errors of the reflective and fundamental forecasts and the fraction of followers of the three forecasting strategies for two different choices of θ . In figure 1, this parameter is set to $\theta = 5/8$, a relatively low value indicating sluggish adjustment to strategies with superior performance. The simulations are initiated at a point where the fraction of followers of reflectivism, the potentially dominant strategy, is at its maximum. For the low level of θ , the introduction of a small fraction of mystics does not induce others to adopt the strategy and has no appreciable impact on the evolution of the asset price. Again, for smaller θ 's, the dynamic (12) approaches the linear weighting case where reflectivism dominates.

Figure 2 shows the same variables as Figure 1, but for a higher level of θ at $\theta = 5.0$. Here, agents are sufficiently aggressive for mysticism to gain a following for significant stretches of time. There are a number of instances where well over half of the the population is using mysticism and some of these are associated with large and persistent deviations in both the fundamentalist forecast errors and the steady state value in the price-dividend ratio. Note that the martingale does not damage the reflectivist forecast error in such a persistent way, since reflectivists use information about the martingale and the relative popularity of mysticism in their forecast. Reflective forecast errors are large only when there is a large change in the popularity of mysticism coinciding with a large value for the martingale.

Figures 3 and 4 illustrate the role of the standard deviation of the martingale innovations σ_η in the formation and duration of bubbles. The agent aggression parameter is set to $\theta = 5.0$ as in Figure 2, but the parameter σ_η is lower at $\sigma_\eta = 0.25\sigma_v$. Hence, though mysticism often gains a following, it is more difficult for the martingale to attain a sufficient magnitude to noticeably affect the asset price. However, when they do occur, bubbles in the asset price tend to last longer, since the martingale grows relatively slowly and more time is required for the martingale (third) term in the mystic payoff (11) to overwhelm the covariance term. Conversely, a higher magnitude for martingale innovation, as in Figure 4 with $\sigma_\eta = 2.0\sigma_v$, shows that bubble outbreaks become rare and short-lived as the martingale quickly, if not immediately, grows too large for mysticism to dominate.

Finally, Figure 5 shows the case where the agent aggression parameter is large at $\theta = 10.0$, while the parameter $\sigma_\eta = \sigma_v$ as in Figures 1 and 2. Here, the martingale innovations are at a magnitude where mysticism can dominate and the agents are quickly switching to superior strategies means

that the dynamics are dominated by the covariance term in the payoffs to mysticism (11) and fundamentalism (10), and there is little inertia in the evolution of the $x_{i,t}$'s. Hence, mysticism quickly gains a following with a positive covariance, but quickly loses with the opposite. There are some occurrences of bubble-like behavior in the price-dividend ratio, but the primary impact of the martingale is an increase in the volatility of the asset price. We proceed by examining more formal econometric features of the data to support these qualitative observations.

5 Time Series Tests

The simulated data matches econometric features of asset market data in multiple respects. In the presence of bubbles, the price-dividend ratio has greater persistence and volatility, and there is excess variance in the price-dividend ratio. Returns are unpredictable in the short run for any parameter choices.

5.1 Mystic dominance and bubbles

Simple measures to detect mysticism and bubbles allow a demonstration of the correspondence between the impact of the martingale and formal econometric features of the data such as excess variance. We run 10,000 trials of 100 periods, roughly the size of the sample in the Shiller data, with 50 periods for initiation. Table 1 reports the fraction of periods (across all trials) where mysticism dominated, i.e. when the fraction of followers of mysticism is greater than 0.5. Table 2 reports the fraction of trials with an occurrence of a bubble in the asset price, defined as a price-dividend ratio (in levels⁷) greater than double its steady state value. This is a necessarily arbitrary but rather strict interpretation of a bubble. Observing the major U.S. stock market averages and using a steady state ratio of 20, the price-earnings ratio in the Shiller data only exceeded 40 in 2009. In the present model of a bubble, a negative bubble, when prices fall below their fundamental value, are just as likely as positive bubbles. If both classes of bubbles are included, the values in Table 2 should be doubled.

Tables 1 and 2 verify that outbreaks of mysticism and bubbles require sufficiently large choices for the parameter θ , the measure of agent aggressiveness, and the parameter σ_η the standard

⁷As a log difference, the bubble condition is $p_t - d_t > \bar{p} - \bar{d} + \ln 2$.

deviation of the shocks to the martingale. For low values of these parameters, there are no occurrences of bubbles or mystic dominance.

As the choices of the parameters θ and σ_η become very large, the occurrences of bubbles and mystic dominance fall from their maximum values. For example in Table 1 given $\theta = 3/4$, the fraction of mystic dominance initially rises with σ_η to a maximum of 0.188 at $\sigma_\eta = 1.0\sigma_v$, corresponding to Figure 2, but falls for larger magnitudes of the shock to the martingale for two reasons. For large σ_η , the martingale (third) term in the payoff to mysticism (11) dominates, diminishing the payoff and making the emergence of mysticism more difficult, as shown in Figure 4. Second, for large θ and σ_η , bubbles rise and collapse faster, lowering the number of periods satisfying the criteria for mystic dominance and bubbles, as in Figure 5.

If agents are sufficiently aggressive about switching to superior strategies, $\theta \geq 3/4$, the role of the martingale becomes significant. In these cases, the fraction of periods showing mystic dominance is always greater than the fraction with bubbles. Even if mysticism has a large following, the magnitude of the martingale may not be large enough to have a dramatic effect on the asset price, pointing up the difficulty identifying bubbles. It is possible that agents are always using extraneous information to value assets, but that such information only drives asset prices away from their fundamental values on rare occasions.

5.2 Persistence and volatility

For parameter settings that produce outbreaks of mysticism and bubbles, the simulated price-dividend series displays greater persistence than the dividend series and matches the volatility observed in the Shiller data. Table 3 reports the average autocorrelation coefficient across the 10,000 trials and demonstrates higher persistence for values of θ and σ_η where bubbles can arise. As with the succeeding tables, the values in the upper left, where both θ and σ_η are small, replicate the predictions of the EMH. The standard deviation and autocorrelation coefficient of $p_t - d_t$ of 0.18 and 0.45 under the EMH are low compared to the Shiller data values of 0.38 and 0.8. The standard deviation is matched closely by many of the simulations (Table 3) where bubbles are prevalent. While the highest value in the table of 0.64 does not show the persistence in the annual data of 0.8, Table 4 reports the standard deviation of the autocorrelation coefficients over the trials and shows that such levels of persistence do occur in a number of trials. Interestingly, for high

values of θ and σ_η , the persistence falls to low levels as mysticism is adopted and abandoned very quickly, as in Figure 5.

5.3 Return Predictability

Returns are not predictable in the short run, an implication of the strong and weak versions of the EMH that is found in stock market data. We examine whether the price-dividend ratio is informative about per share excess returns

$$Z_t = d_t + p_t - \alpha^{-1}p_{t-1}, \quad (14)$$

which is the part of the optimization problem underlying the asset pricing model (2), see Brock and Hommes (1998). Similar to the reflective forecast error (8), the excess return depends on the innovation to the dividend and the term $n_t\eta_t - \Delta n_tm_t$ representing new information about the martingale. As long as agents are unable to forecast n_t , the reflective forecast is unbiased and excess returns are unpredictable.

To test predictability, the following equation to test whether lagged price-dividend ratios contain information about current returns, similar to those used in Fama and French (1988), is estimated on simulated data with 100 observations.

$$Z_t = \beta_0 + \beta_1(p_{t-k} - d_{t-k}) + \varepsilon_t,$$

where k is the lead time for the prediction. If the R^2 from the estimation is over 0.1, returns are defined to be predictable. For the annual Shiller data, the value of the R^2 for any specification of the data⁸ never exceeds 0.12 so 0.1 is a conservative threshold. For any lead time, none of the simulated series series had predictable returns. Returns at longer horizons are also unpredictable, which is unsurprising given the stationarity of the dividend process. Tests on the gross returns give identical results. For $P_t = \exp p_t$ and $D_t = \exp d_t$, the gross return is $(P_t + D_t - P_t)/P_t$, which is regressed on the price dividend ratio P_{t-k}/D_{t-k} .

In the related literature on learning in asset markets, the results on return predictability vary.

⁸Some research, see Cochrane (2001), reports predictable returns over longer horizons for different samples or frequencies of the data.

Adam, Marcet and Niccolini (2008) try to match the features of quarterly data for the U.S. stocks for 1927-2005 and succeed quite well with long run return predictability, but for horizons shorter than 6 quarters the R^2 from the simulated data is too high, roughly the opposite of the results of the present work. Lansing (2010) reports autocorrelation coefficients on the returns⁹ from simulated data that are slightly higher than the annual Shiller data 1871-2005.

In the model with heterogeneous forecasts of Blake (2010), return predictability depends on the available choice of forecasting strategies. For some simulations, returns show serial correlation, but the inclusion of a small fraction of agents using a simple AR1 forecasting model makes returns unpredictable. Though such a strategy was not particularly successful, its presence seems to arbitrage away the predictability of returns in a way reminiscent of arguments in support of efficient markets. Whether forecasts of returns using an AR1 model satisfy cognitive consistency is open to interpretation, but the result is a possible explanation for weakly efficient markets and points up the need to consider heterogeneous forecasts.

5.4 Excess Variance

Studies such as Shiller (1981) demonstrate that asset prices fluctuate more than predicted by the EMH, and endogenous rational bubbles can explain such excess variance. Simulations determine a ratio of the realized variance and the predicted variance based on the variance of the dividends and the EMH, though the statistical significance is difficult to assess. A statistical test of the variance of the price-dividend ratio provides more definitive evidence.

In the absence of mysticism ($n_t = 0$), the asset price behaves according to the strong version of the EMH and depends only on the dividend process.

$$y_t^* = \bar{d} \left(\frac{\alpha}{1 - \alpha} - \frac{\alpha\rho}{1 - \alpha\rho} \right) + d_t \left(1 + \frac{\alpha\rho}{1 - \alpha\rho} \right)$$

Hence, the variance of the asset price should be $\sigma_{y^*}^2 = (1 - \alpha\rho)^{-2} \sigma_d^2$. Table 6 reports the ratio $\sigma_y^2/\sigma_{y^*}^2$ of the variance of the simulated asset prices and the predicted variance using the variance of the simulated dividends. Under the strong EMH, the ratio is unity, which occurs for very low levels of θ and σ_η . For higher levels, the ratio rises above one, and, for one pair of parameter

⁹ The associated values of the R^2 are not reported.

values well over three. This level is smaller than Shiller’s initial estimate, but other research¹⁰ has found smaller estimated values.

To examine the statistical significance of the observed excess volatility in the asset prices, we conduct a test on a transformation of the price-dividend ratio. Let the notation \hat{x} denote the deviation of x from its steady state value. The series $\widehat{pd}_t = \frac{1 - \alpha\rho}{\alpha\rho\sigma_d} (\hat{y}_t - \hat{d}_t)$ has the standard normal distribution so the variance of \widehat{pd}_t is distributed $\chi^2(n)/n$ where n is the number of periods. Table 7 reports the fraction of runs such that the variance of the realized \widehat{pd}_t is excessive at a significance level of 0.05. The results demonstrate that the excess variance shown in Table 6 is statistically significant and corresponds to outbreaks of bubbles. For small values of θ and σ_η (not reported), the simulations correspond to the EMH and the probability of excess variance corresponds to the significance level 0.05. The pattern of the excess return probabilities in Table 7 follows that of the probability an occurrence of a bubble in Table 2. For example, for a sufficiently large choice of θ such that $\theta \geq 3/2$, both probabilities rise with the magnitude of the martingale innovation σ_η for all values reported, but for smaller choice of θ , the probabilities both peak at a choice of σ_η less than the maximum value $\sigma_v \times 8$ in the tables.

The behavior of the simulated time series of the endogenous rational bubbles model matches multiple features found with stock price and dividend data. Simulations with moderate values of the agents aggressiveness parameter and the standard deviation of the martingale innovation such as $\theta = 5$ and $\sigma_\eta = 1/2$ produce data in line with the Shiller data, and these are values where more than half the runs have bubbles associated with mysticism.

6 Conclusion

Models of asset pricing where a representative agent forms expectations according to the strong efficient markets hypothesis are at odds with the observed heterogeneity of forecasts and cannot explain bubbles or related formal econometric features of asset market data. The model with mysticism introduces heterogeneous forecasting strategies that satisfy rationality but puts few other restrictions on agents beliefs. Agent aggressiveness in switching strategies is parameterized by the convexity of the weighting function in the generalized replicator dynamic. While the model is

¹⁰Some examples are LeRoy and Porter (1981), Campbell and Shiller (1989) and LeRoy and Parke (1992). The issue is complicated since some of these models account for a time varying interest rate or discount factor.

capable of mimicking the behavior under the strong efficient markets hypothesis, for sufficiently aggressive agents and martingale innovations of a magnitude similar to the dividend innovations, bubbles can arise. Though returns remain unpredictable in the short run for any parameter settings, outbreaks of mysticism explain the observed volatility and persistence of the price-dividend ratio and excess volatility in asset prices that the efficient markets hypothesis cannot.

There are implications for policies that could minimize bubbles in asset markets. Bubbles arise due to herding behavior and require agents who switch aggressively between strategies. Hence, policies that slow trading may limit how quickly agents change strategies and short-circuit coordination on forecasts based on extraneous information. Some such policies such as limits on single day movements of an index are already in place, and others options such as a Tobin tax should be considered.

That rational bubbles can arise endogenously and extraneous information can have an impact on asset prices and returns is a positive result in that it explains a number of important aspects of the data, but it also points up the limitations on the information that can be obtained through observations of asset prices and returns. If extraneous information has an impact on the asset price, then that price does not reflect the fundamental value of the asset. While the strong efficient markets hypothesis remains an important benchmark for defining bubbles, it is not a full description of the behavior of asset markets. The present model provides a tool for assessing whether prices and returns are truly informative about the value of the underlying asset.

References

- Adam, Klaus, Albert Marcet and Juan Pablo Nicolini (2008) Stock market volatility and learning. CEPR Discussion paper 6518.
- Branch, William. and George Evans (2007) Model uncertainty and endogenous volatility. *Review of Economic Dynamics* 10, 207-237.
- Branch, William. and George Evans (2011) Learning about risk and return: A Simple model of bubbles and crashes. *American Economic Review: Macroeconomics* 3(1), 159-191.
- Blanchard, Olivier (1979) Backward and forward solutions for economies with rational expectations, *American Economic Review* 69(2), 114-118.

Brock, William A. and Cars H. Hommes (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22, 1235-1274.

Campbell, Jeffrey Y. and Robert J. Shiller (1989) The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1(3), p.195-228.

Cochrane, John (2001) *Asset Pricing*. Princeton University Press, Princeton, NJ.

Evans, George (1991) Pitfalls in testing for explosive bubbles in asset prices, *American Economic Review* 81(4), 922-930.

Evans, George and Seppo Honkapohja (2011) Learning as a rational foundation for macroeconomics and finance, manuscript.

Fama, Eugene F. and French, Kenneth R. (1989) Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics*, Elsevier 25(1), p. 23-49.

Hofbauer, Josef and Jorgen Weibull (1996) Evolutionary selection against dominated strategies. *Journal of Economic Theory* 71, 558-573.

Hommes, Cars H. (2001) Financial markets as nonlinear adaptive evolutionary systems. *Quantitative Finance* 1, 149-167.

Hommes, C. 2006, Heterogeneous agent models in economics and finance. in Tesfatsion, L., Judd, K.I. (Eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. Elsevier Science BV, Amsterdam

Keynes, J. M., 1935. *The General Theory of Employment, Interest and Money*. Harcourt Brace, New York, NY.

Lansing, K. (2010) Rational and near-rational bubbles without drift. *Economic Journal* 120, 1149-1174.

LeBaron, Blake (2010) Heterogeneous gain learning and long swings in asset prices. *Journal of Economic Behavior and Organization*, forthcoming.

- LeBaron, B., Arthur, B., Palmer, R. (1999) Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23, 1487-1516.
- LeRoy, Stephen F. and William R. Parke (1992) Stock price volatility: Tests based on the geometric random walk. *American Economic Review* 82(4), 981-992.
- LeRoy, Stephen F. and Richard D. Porter (1981) The present value relation: Tests based on implied variance bounds. *Econometrica* 49(3), p. 555-574.
- Lundqvist, Lars and Thomas J. Sargent (2004) *Recursive Macroeconomics Theory*, The MIT Press, Cambridge, Massachusetts.
- Parke, William R. and George A. Waters (2007) An evolutionary game theory explanation of arch effects. *Journal of Economic Dynamics and Control* 31(7), 2234-2262.
- Parke, William R. and George A. Waters (2012) On the evolutionary stability of rational expectations, manuscript.
- Samuelson, Lawrence (1997) *Evolutionary games and equilibrium selection*, The MIT Press, Cambridge, Massachusetts.
- Sandholm, William (2011) *Population Games and Evolutionary Dynamics*, The MIT Press, Cambridge, Massachusetts.
- Shiller, Robert J. (1981) Do stock prices move too much to be justified to expected changes in dividends? *American Economics Review* 71(3), 421-426.
- Shiller, Robert J. (2005) *Irrational Exuberance*. Princeton University Press, Princeton, N.J.
- Waters, George A. (2009) Chaos in the cobweb model with a new learning dynamic. *Journal of Economic Dynamics and Control* 33(6), 1201-1216.
- Weibull, Jorgen (1998) *Evolutionary Game Theory*. MIT Press, Cambridge, MA.




Figure 1
 $\theta = 5/8, \sigma_\eta = \sigma_v \times 1.0$

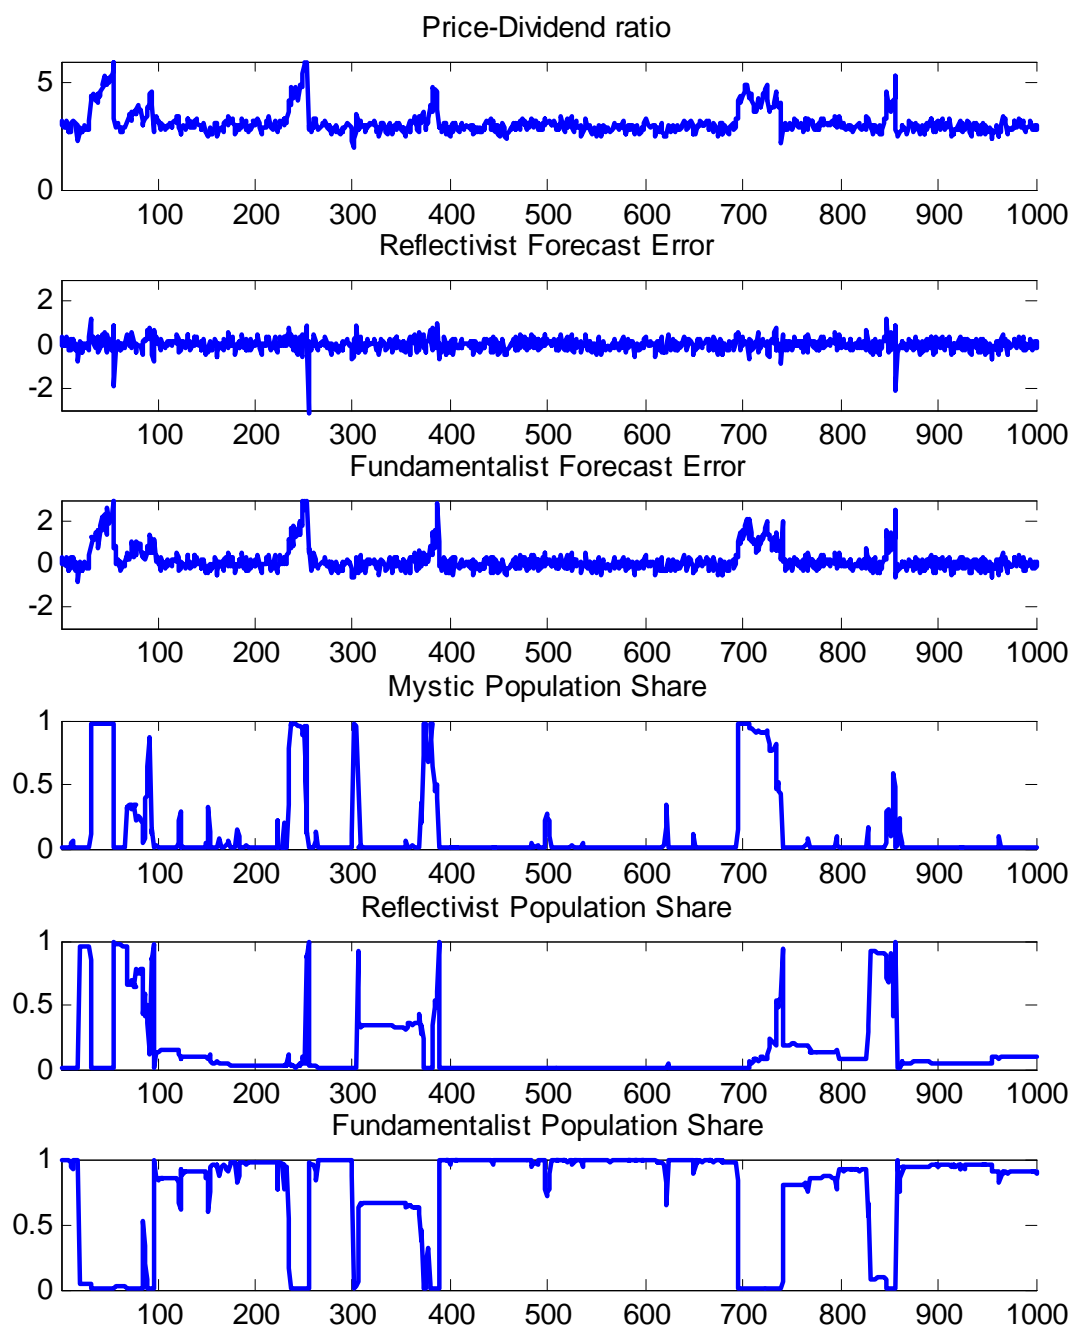


Figure 2
 $\theta = 5.0, \sigma_\eta = \sigma_v \times 1.0$

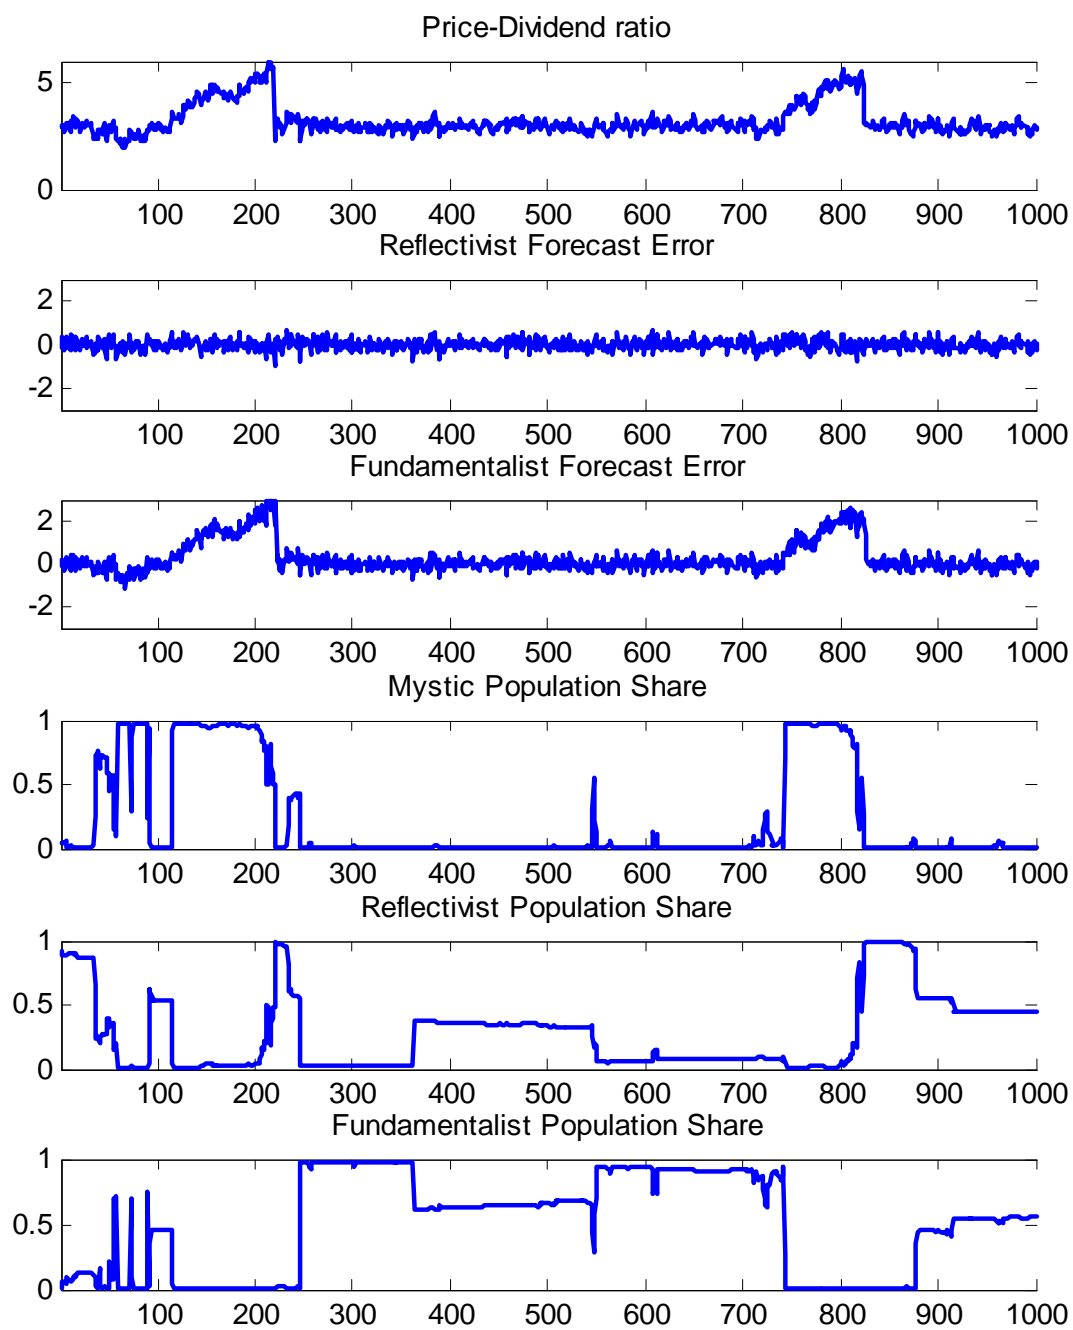




Figure 3
 $\theta = 5.0, \sigma_{\eta} = \sigma_v \times 0.25$



9

Figure 4
 $\theta = 5.0, \sigma_{\eta} = \sigma_v \times 2.0$



9

Figure 5
 $\theta = 10.0, \sigma_{\eta} = \sigma_v \times 1.0$

				$\sigma_\eta = \sigma_v \times$				
		1/8	1/4	1/2	1	2	4	8
	5/8	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	5/4	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	5/2	0.000	0.000	0.002	0.005	0.006	0.002	0.001
θ	5	0.031	0.053	0.064	0.043	0.017	0.007	0.002
	10	0.069	0.134	0.174	0.149	0.087	0.042	0.022
	20	0.037	0.095	0.152	0.146	0.096	0.054	0.030
	40	0.035	0.082	0.116	0.105	0.069	0.039	0.021

Table 1

The fraction of periods over all runs where the mysticism exceeds 50% ($x_2 > 0.5$)

				$\sigma_\eta = \sigma_v \times$				
		1/8	1/4	1/2	1	2	4	8
	5/8	0.008	0.006	0.008	0.014	0.023	0.059	0.323
	5/4	0.007	0.009	0.012	0.016	0.024	0.053	0.251
	5/2	0.014	0.019	0.035	0.070	0.110	0.109	0.131
θ	5	0.124	0.203	0.265	0.276	0.278	0.235	0.142
	10	0.093	0.265	0.539	0.739	0.839	0.837	0.750
	20	0.007	0.017	0.140	0.508	0.815	0.883	0.847
	40	0.007	0.007	0.027	0.245	0.625	0.750	0.741

Table 2

The fraction of runs with one period where $p_t - d_t > \ln 2 + (\bar{p} - \bar{d})$

				$\sigma_{\eta} = \sigma_v \times$				
		1/8	1/4	1/2	1	2	4	8
	5/8	0.445	0.446	0.442	0.431	0.408	0.348	0.222
	5/4	0.445	0.447	0.445	0.433	0.410	0.358	0.243
	5/2	0.454	0.459	0.465	0.465	0.448	0.402	0.360
θ	5	0.536	0.575	0.589	0.552	0.464	0.399	0.388
	10	0.529	0.599	0.635	0.553	0.369	0.218	0.148
	20	0.454	0.470	0.477	0.398	0.265	0.160	0.096
	40	0.442	0.440	0.424	0.359	0.267	0.185	0.117

Table 3

The average across all trials of the one lag autocorrelation coefficient for the price dividend ratio $p_t - d_t$.

				$\sigma_{\eta} = \sigma_v \times$				
		1/8	1/4	1/2	1	2	4	8
	5/8	0.090	0.090	0.090	0.090	0.092	0.095	0.100
	5/4	0.090	0.090	0.090	0.092	0.091	0.095	0.100
	5/2	0.097	0.102	0.118	0.136	0.146	0.139	0.146
θ	5	0.177	0.198	0.201	0.189	0.165	0.160	0.159
	10	0.130	0.145	0.142	0.147	0.170	0.178	0.181
	20	0.089	0.094	0.104	0.118	0.132	0.139	0.132
	40	0.089	0.090	0.092	0.102	0.118	0.126	0.122

Table 4

The standard deviation across all trials of the one lag autocorrelation coefficient.

				$\sigma_{\eta} = \sigma_v x$				
		1/8	1/4	1/2	1	2	4	8
	5/8	0.182	0.184	0.186	0.188	0.192	0.205	0.254
	5/4	0.182	0.184	0.186	0.188	0.191	0.203	0.253
	5/2	0.204	0.211	0.282	0.390	0.396	0.315	0.281
θ	5	0.390	0.493	0.499	0.403	0.295	0.249	0.233
	10	0.223	0.277	0.336	0.347	0.335	0.350	0.425
	20	0.185	0.196	0.225	0.259	0.292	0.352	0.464
	40	0.181	0.185	0.197	0.217	0.248	0.302	0.403

Table 5
The standard deviation of the price-dividend ratio.

				$\sigma_{\eta} = \sigma_v x$				
		1/8	1/4	1/2	1	2	4	8
	5/8	1.010	1.017	1.026	1.036	1.044	1.073	1.209
	5/4	1.011	1.018	1.028	1.036	1.042	1.066	1.203
	5/2	1.059	1.078	1.280	1.740	1.743	1.409	1.285
θ	5	1.732	2.273	2.304	1.792	1.334	1.180	1.133
	10	1.132	1.314	1.543	1.585	1.522	1.566	1.907
	20	1.017	1.052	1.135	1.236	1.345	1.575	2.136
	40	1.004	1.015	1.046	1.103	1.187	1.368	1.812

Table 6
The ratio $Var(y_t^*)/Var(y_t)$ for each run.s

				$\sigma_{\eta} = \sigma_v x$				
		1/8	1/4	1/2	1	2	4	8
	5/8	0.091	0.111	0.144	0.181	0.251	0.539	0.997
	5/4	0.098	0.116	0.158	0.185	0.226	0.481	0.966
	5/2	0.119	0.156	0.209	0.260	0.287	0.315	0.396
θ	5	0.330	0.445	0.503	0.490	0.431	0.367	0.263
	10	0.419	0.723	0.923	0.968	0.971	0.952	0.932
	20	0.118	0.315	0.782	0.961	0.989	0.990	0.997
	40	0.083	0.120	0.337	0.724	0.915	0.960	0.997

Table 7
The fraction of runs with excess variance