

ON THE EVOLUTIONARY STABILITY OF RATIONAL EXPECTATIONS

WILLIAM R. PARKE

University of North Carolina

GEORGE A. WATERS

Illinois State University

Evolutionary game theory provides a fresh perspective on the prospect that agents with heterogeneous expectations might eventually come to agree on a single expectation corresponding to the efficient markets hypothesis. We establish conditions under which agreement on a unique forecast is stable, but also show that persistent heterogeneous expectations can arise if those conditions do not hold. The critical element is the degree of curvature in the payoff weighting functions agents use to value forecasting performance. We illustrate our results in the context of an asset pricing model where a martingale solution competes with the fundamental solution for agents' attention.

Keywords: Rational Expectations, Heterogeneous Expectations, Evolutionary Game Theory, Asset Pricing, Efficient Markets Hypothesis

1. INTRODUCTION

Models assuming a representative agent with rational expectations are ubiquitous in macroeconomics, though they implicitly make strong assumptions about the availability and use of information. Such approaches also stand in contrast to the variety of forecasting techniques (ARIMA, VAR) in the econometric literature.¹ In an asset-pricing environment, the assumption of a unique rational forecast is embodied in the strong version of the efficient markets hypothesis (EMH), where asset prices depend solely on expected future dividends.

The EMH is the basis for much asset pricing theory [Cochrane (2001)] and has empirical support based on the unpredictability of returns [Fama (1991)]. However, some recent work suggests that returns are predictable [Fama and French (1989)] and that there are a number of other features of the data, such as ARCH effects [Engle (2001)], excess kurtosis in returns, and excess variance in asset prices [Shiller (1981), LeRoy and Parke (1992)], that cannot be explained by models satisfying the EMH.

There are many alternative forecasting strategies that do not satisfy the EMH, including “technical” trading and behavioral forecasting strategies.² Although

Address correspondence to: George A. Waters, Department of Economics, Illinois State University, Normal, IL 61790-4200, USA; e-mail: gawater@ilstu.edu.

these approaches can be criticized for violating rational expectations, there does exist a continuum of martingale or rational bubble solutions that satisfy rationality.³ How agents are able to coordinate on the EMH equilibrium in the presence of such a multiplicity of forecasts is an open question.

The present work proposes to allow agents a choice. The model presented here is a variant of Parke and Waters (2007), where agents could switch between forecasting strategies based on the EMH and martingale solutions. Simulation results in that work show that the model explains empirical features of asset-pricing data, such as ARCH effects in returns, that models based on the EMH or representative agent rational bubbles cannot. The focus here is to examine the underlying conditions necessary for agents to adopt forecasts based on martingale solutions.

Evolutionary game theory provides a mechanism describing how agents might switch forecasting strategies based on performance, so their ability to coordinate becomes an outcome, not an assumption, of the model. The primary result is that coordination on a unique forecast based on the EMH is possible but requires very strong assumptions. The stochastic innovations in the models must be sufficiently bounded and agents must not be overly aggressive in switching strategies. Hence, the assumption of a representative forecast based on the EMH is potentially misleading.

All agents have the same information about the underlying model, and the available forecasts satisfy rational expectations.⁴ Differences in agents' choices of forecasting strategies arise because of differences of opinion about the information that should be used. We focus on three forecasting strategies: the *fundamental* forecast based solely on expected future dividends according to the EMH, the *mystic* forecast that uses an extraneous martingale, and the *reflective* forecast, which is a weighted average of the other two. The reflective forecast provides an unbiased forecast, satisfying rational expectations even in the presence of heterogeneity of forecasting strategies in the population. The reflective forecast follows the literature that sets out an econometric view of the merits of combining forecasts.⁵

The evolution of the fractions of agents using the three strategies is determined by a weighted replicator similar to that in Hofbauer and Weibull (1996), which is an example of an imitative dynamic from the evolutionary game theory literature; see Sandholm (2011). The payoffs are based on forecast errors, following common methods for evaluating forecasts in the time series literature; see Elliot and Timmerman (2008). Imitative dynamics allows analysis of a situation where one strategy is eliminated, in contrast to multinomial logit dynamics. A weighting function transforms the payoffs and allows the modeling of varying degrees of aggressiveness in switching between strategies over time. Under a linear weighting function, the dynamic is equivalent to the replicator, but, for a convex weighting function, agents switch to better-performing strategies faster.

The primary goal of the analysis is to determine conditions under which a small fraction of agents experimenting with mysticism can gain a significant following. When no agents use the mystic forecast, the fundamental and reflective forecasts

coincide and the model satisfies the EMH. However, if a significant fraction of the population adopts the mystic forecasting strategy, the extraneous martingale has an impact on the asset price. Now, we can state the results more precisely. If the shocks to the dividends and the martingale innovations are sufficiently large and the weighting function is sufficiently convex, implying aggressive switching, then persistent heterogeneity in the forecasting strategies is likely, which would violate the EMH. The results are determined through both formal stability analysis and simulations of the model.

The present approach has overlaps with the literatures on learning, finance, and the evolution of forecasting strategies. Within the learning literature, expectations are often homogeneous as well, as with the analysis of least squares learning in Marcet and Sargent (1989b) or studies on gradient learning [Evans et al. (2005)]. There are examples of heterogeneous expectations with fixed fractions of agents having idiosyncratic information both in the finance literature, Constanides and Duffie (1996), for example, and in the learning literature [Marcet and Sargent (1989a) and Evans and Guesnerie (2005)], including studies of agents learning sunspot solutions [Branch and McGough (2004)].

There is a substantial literature with dynamic switching of forecasting strategies using the multinomial logit model, particularly with the cobweb model [Brock and Hommes (1997), Hommes (2006)] and with asset pricing⁶ [Brock and Hommes (1998) and Föllmer et al. (2005)]. Anufriev and Hommes (2012a) estimate a multinomial logit model with multiple forecasting strategies using experimental asset market data, arguing that the data can only be explained through the consideration of heterogeneous forecasts.

There are applications of imitative dynamics. In a cobweb model, Sethi and Franke (1995) and Branch and McGough (2008) use alternative versions of the replicator to describe the evolution of agents' choices between adaptive and perfect foresight forecasts and show conditions where chaos arises. As with the multinomial logit approach in Brock and Hommes (1997), agents using the perfect foresight forecast must pay a fixed cost. In the present approach, all agents have the same choice of forecasting strategies available to them, and we do not impose a cost. Fundamentalists ignore the martingale because they believe it is extraneous, not because of the cost of information.⁷

Branch and Evans (2007) and Guse (2010) describe the evolution of models with two costless forecasts in a small macro model and an asset-pricing environment, respectively. The paper by Guse (2010) uses an imitative dynamic but varies from the present approach in a number of ways. Agents update parameter estimates and their choice of forecasting model each period; there are only two available forecasting strategies, based on fundamentals or a sunspot; and agents pay a cost for using a forecasting model that admits extraneous information. In this environment, results from simulations show that persistent heterogeneity cannot arise if the cost is sufficiently large.

Blume and Easley (1992) is a prominent example from a related literature that studies the long-run survival of investment strategies. These models do not specify

the fractions of followers of different strategies in each period but focus on the limiting ratios of payoffs between strategies. Hens and Schenk-Hoppé (2009) review recent developments.

The experimental literature confirms that asset pricing with heterogeneous agents can produce a wide range of behaviors. Anufriev and Hommes (2012b) provide results from recent studies.

The paper is organized as follows. Section 2 describes the asset pricing model, the three forecasts, and their payoffs, and Section 3 discusses the evolutionary dynamics. Section 4 contains the formal stability results, and Section 5 describes the simulation results. Section 6 concludes.

2. ASSET PRICING

A simple asset pricing model motivates a number of the important concepts. For an asset price y_t , the notional model is a basic recursion,

$$y_t = \alpha y_{t+1}^e + u_t, \tag{1}$$

where $\alpha < 1$ is a discount factor, y_{t+1}^e is the expected price next period, and the dividends term u_t is a stochastic income flow.⁸ This model has a solution under rational expectations based only on fundamentals (current and expected future dividends here), but (1) also admits rational bubble solutions that depend on extraneous variables. Convergence to agreement on a single expectation y_{t+1}^e is viewed here as a possible conclusion rather than as an assumption. Given heterogeneous expectations and mean–variance optimizing agents, Brock and Hommes (1998) discuss conditions⁹ under which the realized security price depends on the weighted average of agents’ expectations.¹⁰ They show that

$$y_t = \alpha x_t \cdot e_t + u_t, \tag{2}$$

where e_t is a vector of forecasts of y_{t+1} and x_t is a vector of fractions of the population using each forecast. Here, the vector $e_t = (e_{1,t}, e_{2,t}, e_{3,t})$ represents the three forecasts of reflectivism $e_{1,t}$, fundamentalism $e_{2,t}$, and mysticism $e_{3,t}$, and the fractions of followers are given by the vector $x_t = (x_{1,t}, x_{2,t}, x_{3,t})$.

The discounted present value of the expected income stream,

$$y_t^* = u_t + \sum_{j=1}^{\infty} \alpha^j E_t(u_{t+j} | \Omega_t), \tag{3}$$

where Ω_t is an information set available to the agents, will serve as a point of reference as it satisfies the strong version of the EMH, though it is not the unique solution to (1). The natural candidate for the fundamental forecast of y_{t+1} is thus

$$e_{2,t} = E(y_{t+1}^* | \Omega_t) = \sum_{j=1}^{\infty} \alpha^{j-1} E(u_{t+j} | \Omega_t). \tag{4}$$

Although it is common to assume that all agents somehow recognize $e_{2,t}$ as the appropriate forecast, we make the following less restrictive assumption:

Assumption 1. The fraction $x_{2,t}$ of the agents using the fundamental forecast $e_{2,t}$ is bounded from below by $\delta_2 > 0$ for every t .

These unyielding fundamentalists might well be impressed by the fact that (4) is touted by a large fraction of the academic literature in economics and finance. They do not need to take a position on, for example, the merits of transversality conditions vs. minimum state variables as a basis for (4) to recognize that (4) is prominently featured in McCallum (1983, 1997) and Cochrane (2001).

We select a challenger to the fundamentalist forecast from among the martingale or rational bubble solutions to the model (1). This alternative forecast will be

$$e_{3,t} = e_{2,t} + \alpha^{-t-1}m_t = E(y_{t+1}^*|\Omega_t) + \alpha^{-t-1}m_t, \quad (5)$$

where $m_t = m_{t-1} + \eta_t$ is a martingale. We label this forecast *mysticism* because, although $\alpha^{-t}m_t$ is thought by economic theorists to be extraneous, agents believing in (1) cannot rule out a martingale solution on the basis of that mathematical model.¹¹ Followers of the mystical forecast might, for example, sincerely believe that $\alpha^{-t-1}m_t$ is a valid addition to the fundamentals for y_t . In fact, the martingale does influence y_t if the mystical forecast attracts followers.

Both the fundamental forecast (4) and mysticism (5) satisfy rational expectations in the homogeneous case, but this observation may not hold if there is heterogeneity in the choice of forecasting strategies in the population. Hence, the reflective forecast is postulated to be an average of the fundamental and mystic forecasts weighted according to their relative popularity,

$$e_{1,t} = (1 - n_t)e_{2,t} + n_t e_{3,t}, \quad (6)$$

where

$$n_t = \frac{x_{3,t}}{x_{2,t} + x_{3,t}} \quad (7)$$

is the proportion of followers of the mystical forecast $x_{3,t}$ among those not following the reflective forecast $1 - x_{1,t} = x_{2,t} + x_{3,t}$. Such a forecast follows the literature on the benefits of combining forecasts. See Elliot and Timmerman (2008) for multiple references. This ratio shows one reason for the imposition of a minimum fraction δ_2 using fundamentalism is to avoid dividing by zero. Because reflectivism is based on other forecasts, there must be at least one other forecast in the population for reflectivism to be well specified.

The reflective forecast can be written as

$$e_{1,t} = E(y_{t+1}^*|\Omega_t) + \alpha^{-t-1}n_t m_t, \quad (8)$$

which can be verified with the preceding four equations. The martingale affects the reflective forecast according to the relative popularity of mysticism and fundamentalism. The realization for y_t can be obtained by substituting the expectations

(4), (5), and (8) into (2), yielding

$$y_t = y_t^* + \alpha^{-t} n_t m_t. \tag{9}$$

The realization is thus the fundamentalist forecast plus, to the extent that some of the agents are following mysticism, a martingale term. Unlike the fundamental and mystic forecasts, both the realization (9) and the reflective forecast embody available information about the fractions of the population using the different strategies. By construction, the aggregate expectation $x_t \cdot e_t$ in (2) coincides with the reflective forecast. For the three given forecasting strategies, the aggregate expectation is $x_{1,t}e_{1,t} + x_{2,t}e_{2,t} + x_{3,t}e_{3,t}$. Substituting for $e_{1,t}$ and n_t using (6) and (7) demonstrates that the aggregate expectation and the reflective forecast are equivalent.

Payoffs are given by the negative of the squared forecast error,

$$\pi_{i,t} = -(y_t - e_{i,t-1})^2. \tag{10}$$

Evaluating forecasts using squared errors has a long tradition in econometrics.¹² In an asset pricing context, there are a number of choices for payoffs, such as realized profits, excess returns, or Sharpe ratios. For the present asset pricing model, Hommes (2001) shows that (10) is the natural objective function for mean–variance maximizing agents. If agents adjust for risk based on the variance of the profits, then forecast errors are a more appropriate payoff than realized profits.

The reflective forecast error,

$$U_t = y_t - e_{1,t-1},$$

includes two components:

$$U_t = [y_t^* - E(y_t^*|\Omega_{t-1})] + \alpha^{-t}(n_t m_t - n_{t-1} m_{t-1}). \tag{11}$$

The first term on the right is the innovation in fundamentals, and the second term on the right is the weighted martingale innovation.

The reflective forecast is unbiased, given the information assumed to be available to the agents. The innovation to the dividends has mathematical expectation zero. The expression in parentheses in the second term can be written $\Delta n_t m_{t-1} + n_t \Delta m_t$, where Δn_t is the change in n_t and $\Delta m_t = \eta_t$ is the martingale innovation. If the expectation $E(\Delta n_t|\Omega_{t-1})$ is zero, then the expected reflective forecast error $E(U_t|\Omega_{t-1})$ is zero, because the innovations to the dividend and martingale are also unforecastable. Agents are assumed to be unable to forecast changes in the choices of forecasting strategies across the population. Although they can observe fractions of followers of the forecasting strategies, they do not understand the dynamics describing the evolution of these fractions. (See note 13, cited in the next section.)

TABLE 1. Payoff ordering if $A_{t-1}, n_{t-1} > 0$

Realized U_t	Best payoff	Middle payoff	Worst payoff
$U_t < -n_{t-1}A_{t-1}/2$	Fund.	Refl.	Myst.
$-n_{t-1}A_{t-1}/2 < U_t < (1 - 2n_{t-1})A_{t-1}/2$	Refl.	Fund.	Myst.
$(1 - 2n_{t-1})A_{t-1}/2 < U_t < (1 - n_{t-1})A_{t-1}/2$	Refl.	Myst.	Fund.
$(1 - n_{t-1})A_{t-1}/2 < U_t$	Myst.	Refl.	Fund.

The payoff (10) to reflectivism is

$$\pi_{1,t} = -U_t^2. \tag{12}$$

The payoffs to mysticism and fundamentalism also depend on $A_{t-1} = \alpha^{-t}m_{t-1}$. Intuitively, U_t depends primarily on innovations and A_{t-1} depends on the level of the martingale and consequently is nonstationary. The fundamentalist forecast error from (9) and (4),

$$y_t - e_{2,t-1} = U_t + n_{t-1}A_{t-1},$$

includes a fraction of the martingale term because, to the extent that some of the agents are following the mystical forecast, the realization (9) is affected by the martingale term. Note that if mysticism is driven out of the population, so that $n_{t-1} = 0$, then the fundamental forecast coincides with the reflective forecast. The fundamentalist payoff is

$$\pi_{2,t} = -U_t^2 - 2n_{t-1}U_tA_{t-1} - n_{t-1}^2A_{t-1}^2. \tag{13}$$

The mystic forecast error is

$$y_t - e_{3,t-1} = U_t - (1 - n_{t-1})A_{t-1},$$

from (9) and (5), and the resulting payoff is

$$\pi_{3,t} = -U_t^2 + 2(1 - n_{t-1})U_tA_{t-1} - (1 - n_{t-1})^2A_{t-1}^2. \tag{14}$$

Any of the three forecasting strategies could have the best payoff depending on the realizations of U_t and A_{t-1} , as detailed in Table 1. If A_{t-1} is large relative to U_t , then the third terms, referred to as martingale terms, in the payoffs to fundamentalism (13) and mysticism (14) are the dominating feature, causing both payoffs to underperform reflectivism. The reflective forecast is constructed so that the martingale does not affect its payoff; thus, when the martingale term is large, reflectivism is best. However, if A_{t-1} is not large and the ‘‘covariance’’ U_tA_{t-1} is large and positive, then mysticism could have the best payoff. Such an outcome corresponds to a fortunate (for the mystic) correlation between the martingale and the innovations in the model. Similarly, a large and negative covariance favors fundamentalism.

Despite these observations, reflectivism does have an inherent advantage over the other strategies, as one might expect given that the reflective forecast embodies extra information about the other forecasts and their fractions of supporters. Many evolutionary game theory dynamics depend on the *fitness* of the strategies, the difference between payoffs and the population average payoff $\bar{\pi}_t$, given by $\bar{\pi}_t = x_{1,t-1}\pi_{1,t} + x_{2,t-1}\pi_{2,t} + x_{3,t-1}\pi_{3,t}$. In particular, the replicator dynamic with a linear weighting function $w(\cdot)$ in (16) has the property that the fraction of agents using a strategy adjusts proportionally to the fitness of that strategy. Here, the population average is

$$\bar{\pi}_t = -U_t^2 - \frac{x_{2,t-1}x_{3,t-1}}{x_{2,t-1} + x_{3,t-1}}A_{t-1}^2. \tag{15}$$

Although A_{t-1} does not enter the payoff $\pi_{1,t} = -U_t^2$ to reflectivism, if there is any heterogeneity in the population, it does affect the population average. Under the linear weighting according to population shares for $\bar{\pi}_t$, the covariance terms in the mystic and fundamental payoffs cancel, so the population average payoff cannot be superior to the reflective payoff. Hence, the following holds:

Remark. The fitness of reflectivism $\pi_{1,t} - \bar{\pi}_t$ is always nonnegative.

Under the standard replicator dynamic, the evolution of the vector x_t describing the strategy choices in a population is driven by the fitness of the strategies. The Remark suggests that reflectivism has an ever-increasing following under the replicator, and the point where the maximum fraction of the population is using reflectivism is stable, though much work remains to be done to formalize this idea. First of all, the standard replicator may be problematic in an environment with payoffs that are negative and stochastic. Furthermore, the analysis should allow for a wider range of behavior than specified by the replicator.

3. CURVATURE AND SELECTION DYNAMICS

The evolutionary dynamic studied here, called the weighted replicator, allows the study of stability of points where one strategy is eliminated. Further, the speed at which agents switch to better performing strategies is naturally parameterized within the weighted replicator. The goal is to examine the relationship between agent aggressiveness and the stability of a point where mysticism is not adopted by any agents. Given the previously defined notation for payoffs and strategy choices, the general form of the dynamic is

$$x_{i,t+1} - x_{i,t} = x_{i,t-1} \frac{w(\pi_{i,t}) - \bar{w}_t}{\bar{w}_t}, \tag{16}$$

where the *weighting function* $w(\cdot)$ is increasing in the payoffs, and the expression \bar{w}_t is the population average weighted according to the popularity of the strategies, so that $\bar{w}_t = x_{1,t-1}w(\pi_{1,t}) + x_{2,t-1}w(\pi_{2,t}) + x_{3,t-1}w(\pi_{3,t})$. Strategies with above average payoffs gain adherents. Dynamics such as (16) with $x_{i,t-1}$ on the

right-hand side are also *imitative* in that the popularity of a strategy affects its popularity in the following periods.¹³ The timing in (16) differs from the customary replicator dynamic. We make this choice because the time- t payoffs depend on a forecast made at time $t - 1$, which in turn depends on the time- $(t - 1)$ fractions of followers of the strategies $x_{i,t-1}$. Simulations with an alternative timing are conducted as a robustness check.¹⁴

For a linear weighting function, the dynamic (16) corresponds to the replicator—see Weibull (1998)—but this is not the only possibility. Hofbauer and Weibull (1996, p. 563) note that curvature in $w(\cdot)$ can change agent behavior in an evolutionary game, somewhat analogously to how curvature in utility functions can affect behavior toward risk in other settings. The degree of nonlinearity parameterizes a property that we interpret as agent aggressiveness. We interpret a weighting function that places a relatively high weight only on small squared errors as symptomatic of agent aggressiveness in pursuing accurate forecasting strategies. A weighting function that makes a more moderate distinction between large and small squared errors characterizes less aggressive agents.

A second consideration, enforcing the nonnegativity condition $w(\pi_{i,t}) \geq 0$, also motivates using nonlinear weighting functions. Nonnegativity, which is common in static games, is desirable because \bar{w}_t appears in the denominator of (16) and $\bar{w}_t \leq 0$ would be a problem. Weighting functions that achieve nonnegativity are inherently nonlinear because no linear transformation of the squared error payoffs will be unambiguously nonnegative.

We consider two particular nonlinear payoff weighting functions that enforce $w(\pi_{i,t}) \geq 0$ and parameterize agent aggressiveness for the general dynamic (16).

Truncation weighting. A simple way to achieve nonnegativity is to work with the weighting $w(\pi_{i,t}) = C + \pi_{i,t}$, where C is a constant chosen so that $C + \pi_{i,t} > 0$ for all strategies and all periods. The revised replicator dynamic then becomes

$$x_{i,t+1} - x_{i,t} = x_{i,t-1} \frac{\pi_{i,t} - \bar{\pi}_t}{C + \bar{\pi}_t}. \quad (17)$$

In static game theory, the form of the replicator (17) is quite sufficient, as it is easy to choose C to be larger than the biggest payoff.¹⁵ Here, there is no lower bound for $-U_i^2$ in the payoffs, but the truncation function,

$$w(\pi) = \begin{cases} C + \pi & \text{if } C + \pi \geq 0 \\ 0 & \text{if } C + \pi < 0 \end{cases}, \quad (18)$$

guarantees nonnegativity without requiring $C + \pi_{i,t} \geq 0$ for all strategies and all periods.

The parameter C can be viewed as parameterizing agent aggressiveness. If C is small, then $w(\pi) = 0$ for all but the smallest forecast errors because agents regard strategies with larger forecast errors as worthless. Smaller values for C increase the ratio $(\pi_{i,t} - \bar{\pi}_t)/(C + \bar{\pi}_t)$, causing a bigger change $x_{i,t+1} - x_{i,t}$ for a given $\pi_{i,t} - \bar{\pi}_t$ and $\bar{\pi}_t$.

Exponential weighting. The exponential transformation

$$w(\pi) = e^{\pi/(2\sigma^2)} \tag{19}$$

achieves nonnegativity without sacrificing smoothness. The resulting dynamic is an example of a convex monotonic dynamic as discussed in Hofbauer and Weibull (1996).

Applying the exponential transformation to the squared forecast error $\pi_{i,t} = -(y_t - e_{i,t-1})^2$ produces a familiar functional form,

$$w(\pi_{i,t}) = e^{-(y_t - e_{i,t-1})^2/(2\sigma^2)}. \tag{20}$$

This is a normal probability density. The “mean” of the forecast error $y_t - e_{i,t-1}$ is zero and the “variance” is σ^2 . The parameter σ^2 has no necessary relation to the statistical properties of the forecast errors. It instead determines how agents react to large and small squared forecast errors.

In the present context, σ^2 parameterizes agent aggressiveness. If σ^2 is large, $w(\pi_{i,t})$ is not very sensitive to the magnitudes of the forecast errors. Equivalently, agents are not very aggressive about pursuing the best forecasting strategy. If σ^2 is small, agents assign appreciable value only to the smallest forecast errors. Hofbauer and Weibull (1996, p. 563) describe the effect of convexity as “individuals react over-proportionally to higher payoffs,” as opposed to the replicator, where the population shares adjust proportionally to fitness.

A more common choice of dynamic for the modeling of heterogeneous forecasts is the multinomial logit (MNL),

$$x_{i,t+1} = \frac{e^{\pi_{i,t}/\gamma}}{\sum_{j=1}^n e^{\pi_{j,t}/\gamma}}. \tag{21}$$

The *search intensity* parameter γ is assumed to be positive. The primary reason we opt for an imitative dynamic is that under the MNL dynamic, all strategies maintain a fraction of followers even if their performance is consistently poor. See Waters (2009) for a detailed discussion.

The parameterization of agent aggressiveness in the present context with imitative dynamics is related to the search intensity parameter. Compare the weighted replicator (16) with exponential weighting (19) and the multinomial logit dynamic (21), where there are two strategies to choose from. A lower σ and γ indicate higher aggressiveness and search intensity in the respective dynamic, leading to more followers for the better performing strategy under both. However, because the multinomial logit dynamic determines the level, as opposed to the change, in $x_{i,t}$, the search intensity parameter should not be interpreted as a speed-of-adjustment parameter as with the aggressiveness parameter in the present approach. To parameterize the speed of adjustment with a multinomial logit approach, one can

use the weighting on an average of past payoffs, as in equation 4.5 of Brock and Hommes (1998).

4. STABILITY ANALYSIS

We are interested in the stability properties of a system that we can summarize as follows. In period t , the population fractions $x_{i,t}$, $i = 1, \dots, k$, determine the choices among forecasts $e_{i,t}$, $i = 1, \dots, k$, and the realization of the asset price

$$y_t = \alpha x_t \cdot e_t + u_t.$$

This is (9). Agents choose among the reflective forecast (8), the fundamentalist forecast (4), and the mystical forecast (5). The payoffs are determined by squared errors of the forecasts:

$$\pi_{i,t} = -(y_t - e_{i,t-1})^2.$$

This is (10). It takes the specific forms (12), (13), and (14) for reflectivism, mysticism, and fundamentalism, respectively. The selection dynamic,

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

produces the population fractions for period $t + 1$. This is (16). We consider two choices for $w(\pi_{i,t})$, truncation weighting (18) and exponential weighting (19).

4.1. Truncation Weighting

The stability of the system depends on the parameter C that characterizes agent aggressiveness compared to the magnitude of the shocks. In Section 5 we explore some unstable regions using simulations. In this section, we establish a rigorous extension of the Remark following equation (15) for truncation weighting (18). The proof of stability shows that, given a bound on U_t^2 and a minimum fraction of agents following fundamentalism, for every t either the fraction following reflectivism increases or mysticism is eliminated from the population. In either case, all agents end up following either the fundamental forecast or the reflective forecast, which are then identical. Intuitively, under such a bound, the covariance terms in the payoffs to mysticism and fundamentalism cannot have sufficient influence on the dynamics for mysticism to attract a following.

The proof can be described informally. If the payoffs to all three strategies are greater than $-C$, then there is no truncation and the logic of the Remark is straightforward. The covariance terms do not appear in the population average payoff, so reflectivism's fitness is always nonnegative, and its share cannot fall. For mysticism to attract followers, the payoff to fundamentalism must be below $-C$ and the payoff to mysticism must be best. However, for fundamentalism to perform so badly, A_{t-1}^2 must be large; but if it is too large, the reflective payoff is greater than the mystic payoff. Therefore, for U_t sufficiently bounded, the effect

of the covariance term on the mystic payoff is limited, and there is no A_{t-1}^2 that is simultaneously large enough to force the fundamentalist payoff below $-C$ and small enough so that mysticism outperforms reflectivism.

DEFINITION 1. Let $\underline{\Delta} = \{x | x_1 + x_2 + x_3 = 1, x_2 \geq \delta_2\}$ be the set of admissible points in the simplex, and let $B(y, \varepsilon) = \{z \in \underline{\Delta} | \|y - z\| < \varepsilon\}$ be the open ball around a point y within $\underline{\Delta}$.

Note that if x_t is in the set of admissible points $\underline{\Delta}$, then $n_t < 1 - \delta_2$, which is necessary for the proof of the following.

PROPOSITION 2. There exists a constant $\varphi \in (0, 1)$ such that for the dynamics given by (16) and (18) and for the payoffs (12), (13), and (14), if U_t satisfies the condition, for any $x_t \in \underline{\Delta}$,

$$U_t^2 < \varphi C,$$

then $x_{1,t+1} \geq x_{1,t}$ or $x_{3,t+1} = 0$.

Proof. See the Appendix. ■

Given this bound, either the fraction of followers of reflectivism increases or mysticism is eliminated from the population. This proposition implies stability of the point where the maximum number of agents are using reflectivism. To be more precise, consider the following definition, adapted from Lakshmikantham and Trigiante (2002, Section 4.1).

DEFINITION 3. A point \tilde{x} is uniformly stable if for any $\varepsilon > 0$ there exists a $\delta > 0$ such that $x_t \in B(\tilde{x}, \delta)$ implies that $x_{t+k} \in B(\tilde{x}, \varepsilon)$ for any positive integer k .

Although this version of stability does not apply directly to the present model, a slightly modified version captures the implications of Proposition 2 in the following corollary:

COROLLARY 4. For the model in Proposition 2, given that the condition on U_t holds for all t , the point where the maximum fraction of agents are using reflectivism, $\tilde{x} = (1 - \delta_2, \delta_2, 0,)$ has the following property. For any $\varepsilon > 0$, there exists a $\delta > 0$ such that if $x_t \in B(\tilde{x}, \delta)$, then for every $k > 0$, $x_{t+k} \in B(\tilde{x}, \varepsilon)$ or $x_{3,t+k} = 0$, meaning that mysticism is eliminated.

Proof. See the Appendix. ■

Proposition 2 implies that the point with the maximum fraction of agents using reflectivism satisfies a weak version of stability, in the sense that if x_t is in a neighborhood of \tilde{x} , then, for an appropriate bound on U_t from (11), it will remain there for future periods or mysticism will be eliminated. The logic of the proof of Proposition 2 is relevant for any point in the interior of the simplex for an appropriately chosen bound. The bound on U_t implies bounds on both the

innovations in the dividends and the martingale, which are necessary conditions¹⁶ for Proposition 2 to hold.

Proposition 2 shows that mysticism, which does not conform to the EMH, cannot achieve a significant following if a sufficiently tight bound on the reflective forecast error can be established, relative to the aggressiveness of the agents represented by C . However, persistent heterogeneity of forecasts is also a viable possibility. It is notable that a result such as Proposition 2 can be established in an environment allowing heterogeneous expectations, but it is questionable that bounding stochastic elements such as dividend innovations is realistic. “Black swans” in asset markets would violate such bounds, for example, as would some realizations under normally distributed dividends, though such events could be rare. Simulation results in the next section help to clarify the quantitative importance of the formal analysis.

The stability concept arising from Proposition 2 is adapted from the theory of deterministic difference equations to an environment with stochastic payoffs. Stability theory for deterministic difference equations is well developed; see Lakshmikantham and Trigiante (2002), for example. Even so, they identify 11 different versions of stability (Definition 4.1.2), depending on whether the variable in question remains within a neighborhood of a point or converges to a point in the limit and depending on the speed of convergence. The stability theory for stochastic games has few general results because there are a number of ways to introduce randomness. For example, Hofbauer and Sandholm (2007) study the evolutionary dynamics of simple games with random payoffs and examine the limiting behavior as the randomness is eliminated. In Kandori et al. (1993), agents make occasional mistakes in their strategy choices, and the authors show that there are absorbing sets for the population’s strategy choices.

The stability concept in Proposition 2 is “weak” in the sense that it does not guarantee convergence to \tilde{x} either locally or globally. However, although the present model is stochastic, the stability concept is deterministic. With more restrictive model assumptions, it is possible to demonstrate a stronger version of stability, but the primary goal of the paper is to demonstrate that both adherence to the EMH and outbreaks of heterogeneity of the forecasts are possible outcomes. The model presented here provides a formal representation of the view that the EMH is an important benchmark for understanding asset prices, but it does not fully describe the underlying formation of expectations or the resulting dynamics.

4.2. Exponential Weighting

For agents using the exponential weighting function (20), we present an informal argument that a restriction on U_i in relation to agent aggressiveness again guarantees that the fraction of agents following reflectivism is increasing over time. For reflectivism to increase under the dynamic (16), the weighted payoff to reflectivism $w(\pi_{1,t})$ must be larger than the weighted average payoff \bar{w}_t , so

$\bar{w}_t/w(\pi_{1,t}) < 1$. Using exponential weighting (19), this fraction may be written

$$\frac{\bar{w}_t}{w(\pi_{1,t})} = x_{1,t-1} + x_{2,t-1} \exp(\pi_{2,t} - \pi_{1,t}) + x_{3,t} \exp(\pi_{3,t} - \pi_{1,t}).$$

To examine the dynamics around the introduction of the mystic, assume the martingale is small and the payoff differences in the equation are close to zero. The following equation uses Taylor approximations of the exponential functions around $A_{t-1} = 0$ in the preceding equation with payoffs (12), (13), and (14), eliminating powers of A_{t-1} higher than two because they are small:

$$\frac{\bar{w}_t}{w(\pi_{1,t})} \cong 1 - (2\sigma^4)^{-1} \left(\frac{x_{2,t-1}x_{3,t-1}}{1 - x_{1,t-1}} \right) A_{t-1}^2 (\sigma^2 - U_t^2). \tag{22}$$

If we restrict the Taylor expansion to linear terms, then U_t^2 does not appear, the ratio $\bar{w}_t/w(\pi_{1,t}) < 1$, and reflectivism’s share increases over time as with linear weighting. When the second-order term of the Taylor approximation is included, the sign of $\bar{w}_t/w(\pi_{1,t})$ depends on the sign of $\sigma^2 - U_t^2$. Hence, if U_t is appropriately bounded, then $\sigma^2 - U_t^2$ is positive, but for $U_t^2 > \sigma^2$, $x_{1,t}$ decreases, opening the door to persistent heterogeneity in forecasting strategies. So, if the stochastic innovations in the model are sufficiently large relative to agents’ aggressiveness in switching to better performing strategies, reflective monotonicity does not hold and mysticism has an opportunity to gain adherents. Again, stability and persistent heterogeneity are both possible, depending on whether there is a sufficient bound on the dividends relative to the aggressiveness of the agents. To determine the quantitative effects of changes in model parameters on the likelihood for such instability, we examine simulations of the asset pricing model.

5. ROBUSTNESS

The results in the preceding sections show the potential for either agreement on the fundamental forecast within the population or persistent heterogeneity in forecasting strategies. This section describes simulation results that quantitatively characterize the conditions that might lead to heterogeneous expectations by considering how the following for mysticism might increase from the initial minimal following¹⁷ $n_{t-1} \approx 0$. We then discuss the interpretation of the model, using it to describe bubbles and other features of financial market data.

5.1. Simulations

The simulation results confirm the intuition from the formal discussion. For sufficiently bounded dividends and shocks to the martingale and for sluggish switching between forecasting strategies, mysticism cannot attract a following. With mysticism playing little or no role in the dynamics, the reflective and fundamental

forecasts coincide, and the model reduces to the homogeneous environment corresponding to the EMF. However, for sufficiently large shocks, mysticism can attract a significant following, and the extraneous martingale can have an impact on the asset price, as seen in equation (9) for $n_t > 0$. In such a situation, the asset price still satisfies a weak version of the efficient markets hypothesis in that prices and returns are not forecastable, but the strong version is not satisfied because information besides expected future dividends impacts the asset price.

The magnitude of the stochastic elements of the model is crucial. For mysticism to succeed, the term based on the innovations U_t —see equation (11)—must be large relative to the curvature of $w(\pi)$, which we interpret as aggressiveness of the agents. Furthermore, the term A_t , which is based on the martingale, must be in a range where the covariance term $U_t A_{t-1}$ in the mystic payoff (14) outweighs the impact of the third term with A_{t-1}^2 , as indicated in the proof of Proposition 2.

Some normalization is necessary, and we set the standard deviation of the fluctuation in fundamentals, $y_t^* - E(y_t^* | \Omega_{t-1})$ in (11), to $\sigma^* = 1$. We set the discount factor to $\alpha = 0.99$. The other two parameters are either C or σ^2 and the standard deviation σ_η of the martingale innovation $\eta_t = m_t - m_{t-1}$. The innovation in fundamentals and the innovation in the martingale are both taken to be normally distributed, which means that the conditions for stability in Proposition 2 and equation (22) may be violated for some t . Whether this is quantitatively important is a key question to be examined through simulations.

The initial population share for fundamentalism is set to the minimum $x_{2,0} = 0.05$, but mysticism starts at $x_{3,0} = 0.0001$, which is 500 times smaller. At the start, then, $n_0 = 0.002$ and agents are very nearly following the fundamentalist forecast. Therefore, if mysticism cannot attract a much greater following, then the martingale term has little impact and the asset price is governed by the EMH.

The unconstrained dynamic given by the weighted replicator (16) must be augmented to account for behavior at or near the boundaries of the simplex. Assumption 1 ensures that the fraction of followers of fundamentalism cannot fall below the minimum δ_2 . Similarly, if $x_{3,t}$ falls below 0.0001 in a given period, we reset the fraction of agents following mysticism to that starting value. When the fractions given by the unconstrained dynamics break these bounds, we set those fractions to their minima and allocate the other fractions so that $x_{1,t} + x_{2,t} + x_{3,t} = 1$ and the unconstrained fractions are in proportion to their weighted payoffs. For example, to make Assumption 1 operational, if equation (16) sets $x_{2,t+1} < \delta_2$, then we let $x_{2,t+1} = \delta_2$ and use equation (16) with only the other two forecasting strategies to determine the division of the remaining fraction¹⁸ of agents $1 - \delta_2$. Last, when the fraction of mystics is reset to its minimum $x_{3,t+1} = 0.0001$, the martingale is restarted at zero $m_{t+1} = 0$.

We define robustness in terms of the probability that mysticism attains a specific percentage following. The simulations start at $x_{3,0} = 0.0001$. We calculate the probability that $x_{3,t} \geq 0.20$ at any time within the first 100 periods. If we frequently observe $x_{3,t} \geq 0.20$ within the first 100 periods, we conclude that, for the given




FIGURE 1. $\sigma^2 = 1$.

parameter values, the tendency to converge to a single forecast is not robust to one agent in 10,000 experimenting with mysticism.

Figures 1 and 2 show sample simulations with $\sigma_\eta = 1.0$ for exponential weighting with a low level of aggressiveness, $\sigma^2 = 1$ in Figure 1, and a high level, $\sigma^2 = 1/4$ in Figure 2. Along with the fractions of followers for the three strategies, the figures show the resulting asset price deviations from its steady state. Following Branch and Evans (2011), we assume that dividends are iid, $u_t \sim N(1.0, 1.0)$, which implies that the asset price under the EMH is a constant, so any variation




FIGURE 2. $\sigma^2 = 1/4$.

TABLE 2. Truncated squared errors (18): Probability $x_{3,t} \geq 0.20$ for some $t \leq 100$, 10,000 trials

	$C = 16$	$C = 8$	$C = 4$	$C = 2$	$C = 1$
$\sigma_\eta = 1/16$	0.001	0.065	0.357	0.642	0.776
$\sigma_\eta = 1/8$	0.002	0.088	0.507	0.829	0.939
$\sigma_\eta = 1/4$	0.002	0.122	0.649	0.951	0.994
$\sigma_\eta = 1/2$	0.002	0.151	0.752	0.988	1.000
$\sigma_\eta = 1$	0.003	0.173	0.834	0.997	1.000
$\sigma_\eta = 2$	0.003	0.183	0.846	0.998	1.000
$\sigma_\eta = 4$	0.003	0.172	0.774	0.979	0.997
$\sigma_\eta = 8$	0.002	0.110	0.570	0.882	0.953
$\sigma_\eta = 16$	0.001	0.064	0.359	0.662	0.791

in the asset price must be the result of heterogeneity in the forecasting strategies. Parke and Waters (2007) show simulations of a related model with persistence in the dividend process. Persistence of the dividends affects the asset price, but the robustness analysis presented here depends only on the magnitude of the innovations.

The simulation for a low level of aggressiveness in switching between strategies in Figure 1 shows minimal variation in the asset price, and mysticism never exceeds the threshold 0.2. In contrast, the simulation in Figure 2 has multiple significant outbreaks where mysticism is the dominant strategy and the asset price shows deviations from its steady state value.

Table 2 (the truncation weighting function) and Table 3 (the exponential weighting function) report results for these probabilities and show rather forcefully that mysticism will be an important factor if agents are sufficiently aggressive and the martingale innovations are sufficiently large. Although convergence in Table 2 to a single expectation is common for $C = 16$ and likely for $C = 8$, the probability of episodes of mysticism approaches one for smaller values of C . In Table 3,

TABLE 3. Exponential weighting (19): Probability $x_{3,t} \geq 0.20$ for some $t \leq 100$, 10,000 trials

	$\sigma^2 = 1$	$\sigma^2 = 1/2$	$\sigma^2 = 1/4$	$\sigma^2 = 1/8$	$\sigma^2 = 1/16$
$\sigma_\eta = 1/16$	0.001	0.015	0.157	0.444	0.615
$\sigma_\eta = 1/8$	0.001	0.051	0.382	0.802	0.948
$\sigma_\eta = 1/4$	0.003	0.128	0.671	0.970	0.998
$\sigma_\eta = 1/2$	0.006	0.235	0.848	0.996	1.000
$\sigma_\eta = 1$	0.011	0.327	0.910	0.999	1.000
$\sigma_\eta = 2$	0.018	0.341	0.890	0.998	1.000
$\sigma_\eta = 4$	0.015	0.258	0.762	0.973	0.998
$\sigma_\eta = 8$	0.010	0.153	0.530	0.861	0.953
$\sigma_\eta = 16$	0.006	0.070	0.305	0.613	0.780

TABLE 4. Truncated squared errors (18): mean of $x_{3,100}$, 10,000 trials

	$C = 16$	$C = 8$	$C = 4$	$C = 2$	$C = 1$
$\sigma_\eta = 1/16$	0.0011	0.0428	0.0863	0.0655	0.0504
$\sigma_\eta = 1/8$	0.0013	0.0599	0.1118	0.1023	0.0880
$\sigma_\eta = 1/4$	0.0017	0.0580	0.1361	0.1406	0.1425
$\sigma_\eta = 1/2$	0.0012	0.0485	0.1259	0.1650	0.1765
$\sigma_\eta = 1$	0.0013	0.0257	0.0871	0.1591	0.1936
$\sigma_\eta = 2$	0.0005	0.0115	0.0528	0.1145	0.1540
$\sigma_\eta = 4$	0.0002	0.0051	0.0303	0.0661	0.0913
$\sigma_\eta = 8$	0.0001	0.0024	0.0154	0.0344	0.0465
$\sigma_\eta = 16$	0.0001	0.0012	0.0066	0.0179	0.0253

if the convexity parameter σ^2 is $1/4$ or smaller and the martingale innovation standard deviation σ_η is at least $1/8$, then the probability of significant episodes of mysticism ranges from over one-half to 1.000. Together, Tables 2 and 3 tell a consistent story. If agents are aggressive, which we can identify as $\sigma^2 \leq 1/4$ or $C \leq 4$, then the outcome is not well characterized as convergence to homogeneous expectations. The probability of repeated episodes of mysticism approaches 1.000 for some parameter combinations.

For martingale innovations with standard deviation above $\sigma_\eta = 2$, however, the frequency of mystic success declines. For high levels of A_{t-1} , the reflective payoff (12) is superior because the large martingale terms involving A_{t-1}^2 in the payoffs to fundamentalism (13) and mysticism (14) overwhelm the covariance terms involving $U_t A_{t-1}$. This intuition is also the reason that only a restriction on U_t , and not A_{t-1} or the payoffs, is necessary for stability in Proposition 2.

Tables 4–7 give a snapshot of the behavior of the model after 100 periods. Tables 4 and 6 show the mean fraction of mystic followers $x_{3,100}$ for the linear and exponential weighting cases, whereas Tables 5 and 7 show the standard deviations of those values over 10,000 trials for the two cases. In cases where mysticism

TABLE 5. Truncated squared errors (18): standard deviation of $x_{3,100}$, 10,000 trials

	$C = 16$	$C = 8$	$C = 4$	$C = 2$	$C = 1$
$\sigma_\eta = 1/16$	0.0296	0.1880	0.2593	0.2295	0.2032
$\sigma_\eta = 1/8$	0.0326	0.2186	0.2888	0.2792	0.2614
$\sigma_\eta = 1/4$	0.0371	0.2128	0.3128	0.3179	0.3195
$\sigma_\eta = 1/2$	0.0295	0.1963	0.3030	0.3377	0.3464
$\sigma_\eta = 1$	0.0310	0.1452	0.2590	0.3337	0.3574
$\sigma_\eta = 2$	0.0194	0.0993	0.2063	0.2917	0.3275
$\sigma_\eta = 4$	0.0087	0.0658	0.1581	0.2283	0.2639
$\sigma_\eta = 8$	0.0000	0.0450	0.1142	0.1678	0.1948
$\sigma_\eta = 16$	0.0000	0.0304	0.0746	0.1230	0.1457

TABLE 6. Exponential weighting (19): mean of $x_{3,100}$, 10,000 trials

	$\sigma^2 = 1$	$\sigma^2 = 1/2$	$\sigma^2 = 1/4$	$\sigma^2 = 1/8$	$\sigma^2 = 1/16$
$\sigma_\eta = 1/16$	0.0002	0.0032	0.0218	0.0459	0.0405
$\sigma_\eta = 1/8$	0.0004	0.0067	0.0464	0.0955	0.0957
$\sigma_\eta = 1/4$	0.0005	0.0148	0.0819	0.1459	0.1621
$\sigma_\eta = 1/2$	0.0010	0.0262	0.1124	0.1952	0.2124
$\sigma_\eta = 1$	0.0011	0.0285	0.1022	0.1783	0.2058
$\sigma_\eta = 2$	0.0010	0.0173	0.0613	0.1150	0.1529
$\sigma_\eta = 4$	0.0005	0.0062	0.0249	0.0631	0.0883
$\sigma_\eta = 8$	0.0003	0.0029	0.0116	0.0300	0.0459
$\sigma_\eta = 16$	0.0003	0.0012	0.0053	0.0146	0.0246

could not gain followers, the mean values are at or close to the starting value of 0.0001 with minimal variation. However, when mysticism has a chance, the means are well above the minimum and show great variation in the possible outcomes, implying that it is not hard to find instances where mysticism is the dominant strategy.

As noted in the Introduction, the timing of the general dynamic (16) is different from that in some standard developments of the replicator. For example, Samuelson (1997, p. 64) discusses a dynamic of the form

$$x_{i,t+1} - x_{i,t} = x_{i,t} \frac{w(\pi_{i,t}) - \bar{w}_t}{\bar{w}_t}, \tag{23}$$

where $\bar{w}_t = x_{1,t}w(\pi_{1,t}) + x_{2,t}w(\pi_{2,t}) + \dots + x_{n,t}w(\pi_{n,t})$, whereas the dynamic (16) lags the right-hand-side population shares $x_{i,t-1}$, a natural step given that the time t payoffs $\pi_{i,t}$ depend on time- $(t - 1)$ strategy choices represented by n_{t-1} in the payoffs to fundamentalism (13) and mysticism (14). Simulations of the model using the more standard dynamic¹⁹ (23) give outcomes almost identical to those presented here, so the results are not dependent on the choice of timing.

TABLE 7. Exponential weighting (19): standard deviation of $x_{3,100}$, 10,000 trials

	$\sigma^2 = 1$	$\sigma^2 = 1/2$	$\sigma^2 = 1/4$	$\sigma^2 = 1/8$	$\sigma^2 = 1/16$
$\sigma_\eta = 1/16$	0.0012	0.0307	0.1025	0.1506	0.1368
$\sigma_\eta = 1/8$	0.0056	0.0524	0.1583	0.2263	0.2208
$\sigma_\eta = 1/4$	0.0088	0.0864	0.2202	0.2807	0.2893
$\sigma_\eta = 1/2$	0.0169	0.1249	0.2595	0.3221	0.3305
$\sigma_\eta = 1$	0.0217	0.1341	0.2512	0.3160	0.3316
$\sigma_\eta = 2$	0.0215	0.1065	0.2000	0.2663	0.3029
$\sigma_\eta = 4$	0.0143	0.0620	0.1282	0.2064	0.2433
$\sigma_\eta = 8$	0.0091	0.0442	0.0861	0.1427	0.1816
$\sigma_\eta = 16$	0.0127	0.0273	0.0587	0.1001	0.1357

5.2. Interpretation

The model with outbreaks of mysticism represents an appealing model of a bubble. Although mysticism can gain a following for certain parameter values, it cannot last indefinitely. If mysticism is at its maximum, the presence of the minimum fraction following fundamentalism ensures that the mystic and reflective forecasts are not identical. Although the mystic forecast can outperform the reflective forecast for given periods, the expected payoff of the reflective forecast is superior to the expected payoff to mysticism so agents eventually abandon mysticism. Hence, bubbles arise and collapse endogenously, as shown in Figure 2, in contrast to the models on rational bubbles.²⁰ Furthermore, the reflective forecast is an apt description of bubble psychology. Agents adopt a forecast based on information they believe to be extraneous because other agents are using that information.

In practice there could be many mystic forecasts, making persistent heterogeneity more likely. There is an abundance of variables—exchange rates and commodity prices being two of many examples—that might be considered for use in forecasting asset prices. Hence, for a given asset, it is quite possible that an extraneous variable exists that has the characteristics necessary to serve the role of the martingale in the present model. We focus on a single martingale alternative to clarify the conditions necessary for mysticism to arise. Furthermore, when one mystic forecast does gain a following, agents will tend to coordinate on that strategy. Note that when mysticism gains a large following, i.e., n_t is close to one, the mystic payoff (14) is close to the reflective payoff because the impact of the martingale (third) term in (14) is minimized. Although an extension of the model to include more forecasts might be desirable to calibrate to financial market data, the intuition of the present model should hold because agents will tend to coordinate on a successful mystic.

It might be argued that the existence of a minimum fraction of fundamentalists is unlikely, given that mysticism can have significant periods of success. However, if all agents adopt the mystic forecast, the model collapses to a never-ending rational bubble, which is economically implausible. As we note following Assumption 1, this minimum fraction of fundamentalists exists as long as some agents believe standard asset pricing textbooks. The small fraction of unyielding fundamentalists is all that is necessary to ensure that the asset price remains connected to the dividends in the long run.

The conclusion that extraneous information embodied in the mystic forecast can affect asset prices depending on the magnitude of the shocks and the aggressiveness of the agents has both theoretical and quantitative support. Hence, asset price bubbles arising from the presence of mysticism are a definite possibility, but so are stretches of time in which the EMH is satisfied. The shocks to the fundamentals are closely related to the uncertainty about future dividends, suggesting that bubbles are more likely to arise for the stocks of firms in new industries or for recently developed asset classes, where judgements about future profits and performance are difficult.

Parke and Waters (2007) show that, when conditions are favorable for repeated episodes of mysticism, the observed excess returns will exhibit volatility clustering. Standard tests for autoregressive conditional heteroskedasticity (ARCH) reject the null of homoskedasticity. The widespread finding of ARCH effects in empirical work could thus be taken as supporting the notion that aggressive agents are precluding convergence to a single expectation.

6. SUMMARY AND CONCLUSIONS

Our evolutionary game theory approach demonstrates that agreement on a unique rational expectation and the persistent presence of heterogeneous forecasts based on extraneous information are both distinct possibilities. The primary determining factor is the aggressiveness of the agents in switching forecasts relative to the magnitude of the stochastic elements. Convexity in the weighting function of the payoffs in the evolutionary game theory dynamic is directly related to the switching speed. Sluggish adjustment and small shocks correspond to stability of the point where mysticism is eliminated and agents coordinate on a forecast corresponding to the efficient markets hypothesis. Higher aggressiveness gives rise to the possibility that mysticism could play a significant role in asset price dynamics, potentially leading to persistent deviations from the fundamental forecast.

These results have a direct bearing on the merits of assuming that all agents agree on a single forecast. The theoretical results imply that such an assumption may be a reasonable abstraction for agents not overly aggressive in pursuing the best forecast, but also imply that an environment populated by aggressive agents may be fertile ground for emergence of heterogeneous expectations. The simulation results confirm that, for our asset pricing example, sufficient agent aggressiveness can lead to persistent heterogeneous expectations.

NOTES

1. The econometrics literature in the tradition of Box and Jenkins (1970) specifically focuses on the advantages of empirical models over theory-based models.

2. Covel (2004) describes the benefits of trend following. Hommes (2006) surveys models with heterogeneous, boundedly rational forecasts.

3. Blanchard (1979), Evans (1991), and Charemza and Deadman (1995) are examples of rational bubble models. Pesaran (1987) discusses the implications of the multiplicity of solutions.

4. There are a number of boundedly rational or behavioral alternative forecasting strategies. Although some of these are appealing, we focus on rational bubble forecasts to introduce heterogeneity in a disciplined way.

5. Bates and Granger (1969) and Granger and Ramanathan (1984) discuss the potential benefits. Elliot and Timmerman (2008) has references such as Stock and Watson (1999) that empirically verify that combining forecasts can improve performance.

6. Branch and Evans (2007) study dynamic switching with multinomial logit in a Lucas-style macro model. Horst and Wenzelburger (2008) examine the long-run behavior of a related asset-pricing model.

7. One could impose a cost on reflectivists because they use more information, which would make the adoption of mysticism more likely.

8. This equation could apply equally to aggregate prices, exchange rates, etc..
9. In particular, there is a constant supply of a risky asset, and agents have a common belief about the variance of the returns.
10. There are several important differences between our analysis and that of Brock and Hommes (1998). Our choice of strategies differs from theirs. Brock and Hommes take payoffs to equal trading profits, which are a linear rather than concave function of forecast errors. Their discrete choice updating mechanism, as we note in the next section, does not allow convergence to a single expectation.
11. Even an auxiliary belief in a stationary solution does not rule out a martingale in a finite number of periods. It is not possible to know with certainty that m_t is nonstationary from a finite data sample. In fact, when mysticism grows in popularity, that growth will often occur in the first few periods, well before tests for nonstationarity have useful power.
12. Elliott and Timmerman (2008) discuss the role of mean squared prediction error and forecast combination in the literature on forecasting.
13. Such dynamics emerges from learning models where myopic agents repeatedly play a game and update their strategy choices over time. Hence they are not aware of the equation describing the evolution of x_t such as the one preceding, as noted in the discussion of the reflective forecast error following equation (11).
14. Parke and Waters (2007) use an alternative timing in a related model. As occurs for some discrete-time dynamics—see Weibull (1998, Sect. 4.1)—the present approach can produce overshooting, where $x_{i,t}$ could fall below zero or the fraction of fundamentalists fall below δ_2 . For the simulations, we specify that if $x_{i,t}$ falls below its minimum, it is reset to that value for the next period and the other fractions are set proportional to the fractions given (16). More details are given in Section 5.
15. Hofbauer and Sigmund (1988, p. 133), Samuelson (1997, p. 66), Hofbauer and Sigmund (1998, pp. 76–77), and Weibull (1998, pp. 122–123) discuss this version of the replicator dynamic.
16. It is not possible to show that bounds on the innovations imply a bound on the reflective forecast, but, for practical purposes, this is the case. Investigation of alternative versions of the model where this point may be made explicitly are left for future work.
17. Binmore et al. (1995) introduce drift, which is a similar approach that examines the effects of introducing a small fraction using a strategy.
18. A similar rule is used if $x_{1,t}$ becomes negative for the unconstrained dynamic, though the details do not affect our analysis of the robustness to the introduction of mysticism.
19. Note that this dynamic could be simplified so that only $x_{i,t+1}$ appeared on the left-hand side.
20. See, for example Evans (1991). See Parke and Waters (2007) for further discussion and simulation examples.

REFERENCES

- Anufriev, Mikhail and Cars H. Hommes (2012a) Evolution of market heuristics. *Knowledge of Engineering Review* 27, 255–271.
- Anufriev, Mikhail and Cars H. Hommes (2012b) Evolutionary selection of individual expectations and aggregate outcomes. *American Economic Review—Micro* 4, 35–64.
- Bates, J.M. and Clive W.J. Granger (1969) The combination of forecasts. *Operations Research Quarterly* 20, 451–468.
- Binmore, Kenneth, John Gale, and Lawrence Samuelson (1995) Learning to be imperfect: The ultimatum game. *Games and Economic Behavior* 8, 56–90.
- Blanchard, Olivier (1979) Backward and forward solutions for economies with rational expectations. *American Economic Review* 69(2), 114–118.
- Blume, Lawrence E. and David Easley (1992) Evolution and market behavior. *Journal of Economic Theory* 58, 9–40.
- Box, George E.P. and Gwilym M. Jenkins (1970) *Time Series Analysis: Forecasting and Control*. San Francisco: Holden-Day.

- 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.
- Branch, W. and B. McGough (2004) Multiple equilibria in heterogeneous expectations models. *Contributions to Macroeconomics* 4(1), Article 12.
- Branch, William and Bruce McGough (2008) Replicator dynamics in a cobweb model with rationally heterogeneous agents. *Journal of Economic Behavior and Organization* 65(2), 224–244.
- Brock, William A. and Cars H. Hommes (1997) A rational route to randomness. *Econometrica* 65, 1059–1095.
- 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.
- Charemza, Wojciech and Derek Deadman (1995) Speculative bubbles with stochastic explosive roots: The failure of unit root testing. *Journal of Empirical Finance* 2, 153–163.
- Cochrane, John (2001) *Asset Pricing*. Princeton, NJ: Princeton University Press.
- Constanides, George M. and Darrel Duffie (1996) Asset pricing with heterogeneous consumers. *Journal of Political Economy* 104, 219–240.
- Covel, Michael W. (2004) *Trend Following: How Great Traders Make Millions in Up or Down Markets*. Upper Saddle River, NJ: Prentice Hall.
- Elliot, Graham and Allan Timmerman (2008) Economic forecasting. *Journal of Economic Literature*, 46(1), 3–56.
- Engle, Charles (2001) GARCH 101: The use of ARCH/GARCH models in applied econometrics. *Journal of Economic Perspectives* 15(4), 157–168.
- Evans, George (1991) Pitfalls in testing for explosive bubbles in asset prices. *American Economic Review* 81(4), 922–930.
- Evans, George and Roger Guesnerie (2005) Coordination on saddle path solutions: The educative viewpoint—linear multivariate models. *Journal of Economic Theory* 124, 202–229.
- Evans, George, Seppo Honkapohja, and Noah Williams (2005) Generalized Stochastic Gradient Learning. NBER Working Paper 0317.
- Fama, Eugene F. (1991) Efficient markets II. *Journal of Finance* 46(5), 1575–1617.
- Fama, Eugene F. and Kenneth R. French (1989) Business condition and the expected returns on stocks and bonds. *Journal of Financial Economics* 25(1), 23–49.
- Föllmer, Hans, Ulrich Horst, and Alan Kirman (2005) Equilibria in financial markets with heterogeneous agents: A probabilistic perspective. *Journal of Mathematical Economics* 41, 123–155.
- Granger, Clive W.J. and Ramu Ramanathan (1984) Improved methods of combining forecasts. *Journal of Forecasting* 3, 197–204.
- Guse, Eran (2010) Heterogeneous expectations, adaptive learning and evolutionary dynamics. *Journal of Economic Behavior and Organization* 74, 42–57.
- Hens, Thorsten and Klaus R. Schenk-Hoppé (2009) *Handbook of Financial Markets: Dynamics and Evolution*. Amsterdam: North Holland.
- Hofbauer, Josef and William Sandholm (2007) Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory* 132, 47–69.
- Hofbauer, Josef and Karl Sigmund (1988) *The Theory of Evolution and Dynamical Systems*. Cambridge, UK: Cambridge University Press.
- 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, Cars H. (2006) Heterogeneous agent models in economics and finance. In L. Tesfatsion and K.I. Judd (eds.), *Agent-Based Computational Economics*, Handbook of Computational Economics, vol. 2, pp. 1109–1186. Amsterdam: Elsevier Science BV.

- Horst, Ulrich and Jan Wenzelburger (2008) On nonergodic asset prices. *Economic Theory* 34, 207–234.
- Kandori, Michihiro, George J. Mailath, and Rafael Rob (1993) Learning, mutation and long run equilibria in games. *Econometrica* 61(1), 29–56.
- Lakshmikantham, V. and D. Trigiante (2002) *Theory of Difference Equations: Numerical Methods and Applications*. New York: Marcel Dekker.
- Marcet, Albert and Thomas Sargent (1989a) Convergence of least squares learning in an environment with hidden state variables and private information. *Journal of Political Economy* 97(6), 1306–1322.
- Marcet, Albert and Thomas Sargent (1989b) Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, 337–368.
- McCallum, Bennett T. (1983) On nonuniqueness in rational expectations models: An attempt at perspective. *Journal of Monetary Economics* 11, 139–168.
- McCallum, Bennett T. (1997) The role of the minimum state variables criterion in rational expectations models. *International Journal of Tax and Finance* 6(4), 621–639.
- 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.
- Pesaran, Hashem (1987) *The Limits of Rational Expectations*. Oxford, UK: Blackwell Publishing.
- Samuelson, Lawrence (1997) *Evolutionary games and equilibrium selection*. Cambridge, MA: MIT Press.
- Sandholm, William (2011) *Population Games and Evolutionary Dynamics*. Cambridge, MA: MIT Press.
- Sethi, Rajiv and Reiner Franke (1995) Behavioural heterogeneity under evolutionary pressure: Macroeconomics implications of costly optimization. *Economics Journal* 105, 583–600.
- Shiller, Robert J. (1981) Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71(3), 421–436.
- Stock, James H. and Mark W. Watson (1999) A comparison of linear and nonlinear models for forecasting macroeconomic time series. In R.F. Engle and H. White (eds.), *Cointegration, Causality and Forecasting: A Festschrift in Honor of Clive W. Granger*, pp. 1–44. Oxford, UK: Oxford University Press.
- 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 (1997) *Evolutionary Game Theory*. Cambridge, MA: MIT Press.

APPENDIX

To prove Proposition 2, we first examine the one case where, given the restriction in the proposition, reflectivism could lose followers, but mysticism is eliminated. This situation arises for a small n when the covariance term UA is negative and fundamentalism gains followers. The details of the proof are rather involved, but the intuition is the following. If n is small, for fundamentalism to be the best, the term A^2 must be so large that the weighted payoff to mysticism must be zero in two successive periods, ensuring that mystic followers are eliminated from the population.

LEMMA 5. *If $w(\pi_{2,t}) > 0$ and $w(\pi_{3,t}) = 0$, then either $x_{1,t+1} > x_{1,t}$ or $x_{3,t+1} = 0$.*

Proof. Examining the dynamic (16) for mysticism, if $w(\pi_{3,t}) = 0$ and $x_{3,t-1} \geq x_{3,t}$, then $x_{3,t+1} = 0$ follows directly. We must show that if $x_{1,t+1} > x_{1,t}$ is not true, then it must be the case that $x_{3,t-1} \geq x_{3,t}$, which implies that mysticism is eliminated in period $t + 1$. This case is somewhat different from the case where the weighted payoff on fundamentalism is zero because we do not assume a minimum bound on mysticism.

If $x_{1,t+1} \leq x_{1,t}$ and $w(\pi_{3,t}) = 0$, then fundamentalism must gain adherents, so $w(\pi_{2,t}) > 0$. The condition $w(\pi_{1,t}) \leq \bar{w}_t$ becomes $(1 - x_{1,t-1})w(\pi_{1,t}) \leq x_{2,t-1}w(\pi_{2,t})$, which is equivalent to

$$C - U_t^2 < (1 - n_{t-1}) (-2U_t A_{t-1} - n_{t-1} A_{t-1}^2). \tag{A.1}$$

Because $w(\pi_{2,t}) > 0$, it must be the case that $U_t A_{t-1} < 0$. From the condition in the proposition, $C - U_t^2 > (1 - \varphi)C$ and $|U_t| < \sqrt{\varphi C}$. Therefore, because $n_{t-1} > 0$, for $A' = \frac{(1-\varphi)C}{\sqrt{\varphi C}}$, inequality (A.1) implies that $|A_{t-1}| \geq A'$. Note that as $\varphi \rightarrow 0$, $A' \rightarrow \infty$.

To show that $x_{3,t-1} \geq x_{3,t}$, we must demonstrate that $w(\pi_{3,t-1}) = 0$, which is equivalent to

$$C - U_{t-1}^2 < (1 - n_{t-2}) (-2U_{t-1} A_{t-2} + (1 - n_{t-2}) A_{t-2}^2). \tag{A.2}$$

If this inequality is true for $U_{t-1} A_{t-2} > 0$, it is true for $U_{t-1} A_{t-2} < 0$ as well. Because $C - U_{t-1}^2 > C$, $|U_{t-1}| < \sqrt{\varphi C}$, and $1 - n_{t-2} > \delta_2$, there exists an A'' such that $|A_{t-2}| > A''$ guarantees the inequality, and A'' is decreasing in φ . Here, for the inequality to hold, the A^2 term must be sufficiently large, in contrast to the inequality (A.1), where $|UA|$ must be sufficiently large.

However, a lower bound on A_{t-1} implies a lower bound on A_{t-2} . By definition (see the discussion after equation (12)), the term $A_{t-1} = \alpha^{-t-1} m_{t-1} = \alpha^{-1} A_{t-2} + \alpha^{-t} \eta_{t-1}$, recalling that η_t is the innovation to the martingale, or $A_{t-2} = \alpha A_{t-1} - \alpha^{-t} \eta_{t-1}$. The definition of U (11) in light of its bound $|U| < \sqrt{\varphi C}$ for any $n \in [0, 1]$ implies that $|\alpha^{-t} \eta_{t-1}| < \sqrt{\varphi C}$. Hence, if $x_{1,t+1} \leq x_{1,t}$ and $w(\pi_{3,t}) = 0$, implying that $A_{t-1} > A'$ from above, then $A_{t-2} > \alpha A' - \sqrt{\varphi C}$. Because $\alpha A' - \sqrt{\varphi C} \rightarrow \infty$ as $\varphi \rightarrow 0$ and A'' is decreasing in φ , there exists a φ' sufficiently small so that A' is large enough and A'' is small enough to ensure that $A_{t-2} > A''$, so the inequality (A.2) is satisfied, $w(\pi_{3,t-1}) = 0$, and $x_{3,t-1} \geq x_{3,t}$, as required. ■

The following proof of Proposition 2 covers the other cases.

Proof. The bound on U_t^2 guarantees that $w(\pi_{1,t}) > 0$. If $w(\pi_{2,t}) > 0$ and $w(\pi_{3,t}) > 0$, then the inequality in the Remark following equation (15) implies that $w(\pi_{1,t}) > \bar{w}_t$ and so $x_{1,t+1} > x_{1,t}$. The same inequality holds if both $w(\pi_{2,t}) = 0$ and $w(\pi_{3,t}) = 0$.

Suppose then that $w(\pi_{2,t}) = 0$ and $w(\pi_{3,t}) > 0$. The case where $x_{2,t+1}$ hits its minimum δ_2 is treated first as a separate case. If $x_{2,t+1}$ is fixed, then $x_{1,t+1} > x_{1,t}$ if and only if $w(\pi_{1,t}) > w(\pi_{3,t})$ for any reasonable dynamic. This condition is equivalent to

$$2(1 - n_{t-1}) U_t A_{t-1} - (1 - n_t)^2 A_{t-1}^2 < 0. \tag{A.3}$$

Because $w(\pi_{2,t}) = 0$ in this case, it must be true that

$$C - U_t^2 - 2n_{t-1} U_t A_{t-1} - n_{t-1}^2 A_{t-1}^2 < 0,$$

but the condition from the proposition $U_t^2 < \varphi C$ implies that $C - U_t^2 > (1 - \varphi)C$. Combining inequalities yields

$$2n_{t-1} U_t A_{t-1} + n_{t-1}^2 A_{t-1}^2 > (1 - \varphi) C.$$

For this inequality to be true, there must be a minimum $A''' > 0$ such that $|A_{t-1}| > A'''$.

Now, we can show the inequality (A.3), guaranteeing that $w(\pi_{1,t}) > w(\pi_{3,t})$ is satisfied in this case for some φ' sufficiently small. If U_t and A_{t-1} have different signs, (A.3) is

satisfied automatically. If they have the same signs, (A.3) is satisfied as long as

$$|U_t| < \frac{1}{2} (1 - n_t) |A_{t-1}|. \tag{A.4}$$

In this case, we have $|A_{t-1}| > A''' > 0$. Furthermore, because $x_{2,t+1} \geq \delta_2$, it is also the case that $(1 - n_t) > \delta_2$. Therefore, there exists a $\varphi'' \in (0, 1)$ such that for $U_t^2 < \varphi C$, the inequality (A.4) is satisfied and therefore so is $x_{1,t+1} > x_{1,t}$.

For the case where $w(\pi_{2,t}) = 0$ and $w(\pi_{3,t}) > 0$ when $x_{2,t+1}$ is greater than its minimum, the condition for $x_{1,t+1} > x_{1,t}$ is $w(\pi_{1,t}) > n_{t-1}w(\pi_{3,t})$, which is weaker than the condition $w(\pi_{1,t}) > w(\pi_{3,t})$ for the case when $x_{2,t+1} = \delta_2$, so the preceding argument applies.

So $x_{1,t+1} > x_{1,t}$, except in the case covered by the preceding lemma where $x_{3,t+1} = 0$. So for $\varphi = \min\{\varphi', \varphi''\}$, Proposition 2 holds. ■

The proof of Corollary 4 follows. The one minor complication is that $x_1 - \tilde{x}_1$ is not the same as the distance between x and \tilde{x} .

Proof. Given $\varepsilon > 0$, there exists a $\delta' > 0$ such that $x_{1,t} - \tilde{x}_1 < \delta'$ implies that $x_{1,t} \in B(\tilde{x}, \varepsilon)$. Similarly, given $\delta' > 0$, there exists a $\delta > 0$ such that $x_{1,t} \in B(\tilde{x}, \delta)$ implies that $x_{1,t} - \tilde{x} < \delta'$. Therefore, $x_{1,t} \in B(\tilde{x}, \delta)$ implies that $x_{1,t} - \tilde{x} < \delta'$ and Proposition 2 implies that $x_{1,t+k} - \tilde{x} < \delta'$ or $x_{3,t+k} = 0$. Hence, $x_{t+k} \in B(\tilde{x}, \varepsilon)$ or $x_{3,t+k} = 0$. ■