

An Evolutionary Route to Rational Expectations

William R. Parke¹

Department of Economics
University of North Carolina
Chapel Hill, NC 27599-3305

George A. Waters²

Department of Economics
Illinois State University
Normal, IL 61790-4200

July 24, 2008

¹parke@email.unc.edu

²gawater@ilstu.edu, phone: 312-307-0661, fax: 309-438-5228

Abstract

Evolutionary game theory provides a fresh perspective on the prospects that agents with heterogeneous expectations might eventually come to agree on a single expectation. We establish conditions under which convergence of beliefs could occur, but also show that persistent heterogeneous expectations can arise if those conditions do not hold. The critical element is the degree of curvature in 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

JEL codes: C73, D84, G12, C22

1 Introduction

One of the challenges to models with forward-looking expectations is explaining how all agents come to hold a single expectation. A lengthy debate has taken place, for example, over how agents seeking to adopt rational expectations might single out the fundamental solution from among the infinite number of martingale solutions.¹ More generally, econometrics texts describe a wide range of forecasting strategies (ARIMA models, for example) that may or may not exemplify rational expectations.² Agents in the financial markets pursue an even wider range of “technical” forecasting strategies.³ This diversity of empirical forecasting procedures is at odds with the rational expectations assumption, which compels theorists to embrace representative agent models admitting the existence of only one expectation.

Using evolutionary game theory, we attempt to reconcile the theoretical assumption of a single expectation with the empirical reality of multiple forecasting procedures. We establish conditions under which agents with heterogeneous expectations will tend to converge to agreement on a single fundamental expectation. This result is not entirely conclusive in justifying representative agent models, however, because we also show that alternative conditions lead to persistent heterogeneous expectations.

The game theoretic basis for our approach has several key features: (i) expectations can be heterogeneous, (ii) agents are allowed to switch strategies based on observed forecasting performance, (iii) convergence to a single expectation is possible, but not necessary, (iv) agents consider a strategy we refer to as reflective that is a weighted average of other forecasts, and (v) one forecast is regarded as fundamental on the basis of economic theory by at least some small fraction of the agents.

The majority of the macroeconomic literature, including the work on rational bubbles (Blanchard (1979), Evans (1991), Charemza and Deadman (1995)), assumes all agents share a unique rational expectation. Within the learning literature, expectations are often homogeneous as well, as with the analysis of least squares learning in Marcet and Sargent (1989a) or studies on gradient learning (Evans, Honkapohja and Williams 2005). There are examples of heterogeneous expecta-

¹See, for example, Cochrane (2001) and McCallum (1983, 1997).

²The econometrics literature in the tradition of Box and Jenkins (1970) specifically focuses on the advantages of empirical models over theory-based models.

³Covel (2004), for example, puts forward the merits of “trend following.”

tions with fixed fractions of agents having idiosyncratic information in both the finance literature, Constanides and Duffie (1996) for example, and in the learning literature (Marcet and Sargent (1989b) and Evans and Guesnerie (2005)) including studies of agents learning sunspot solutions (Branch and McGough (2004)).

Our approach includes heterogeneous forecasts and allows for dynamic switching between forecasting strategies based on past performance. As an example, we present an asset pricing model where payoffs to forecast strategies are based on negative squared forecast errors, following common methods for evaluation of forecasts in the time series literature, see Elliot and Timmerman (2008). Results in a more general model of heterogeneous expectations rely on payoffs that are concave in the forecast.

An evolutionary game theory mechanism describes switching between strategies while allowing for the possibility that agents will come to agree on a single forecast. 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 pricing⁴ (Brock and Hommes (1998) and Föllmer, Horst, and Kirman (2005)). Suppose fractions $x_{1,t}, \dots, x_{n,t}$ of the agents follow n strategies with forecasts $e_{1,t}, \dots, e_{n,t}$ that produce payoffs $\pi_{1,t}, \dots, \pi_{n,t}$. The multinomial logit mechanism describes the evolution of $x_{i,t}$ as follows,

$$x_{i,t+1} = \frac{\exp(\beta\pi_{i,t})}{\exp(\beta\pi_{1,t}) + \dots + \exp(\beta\pi_{n,t})}, \quad (1)$$

where β is called the *search intensity parameter*. Such an approach is not appropriate for our purpose since agreement on a single forecast is not possible with this model. In fact, unless there are infinite payoffs or search intensity, all strategies retain followers regardless of their payoff performance.

To allow for the possibility that agents agree on a single forecast, we study dynamics of the general form

$$x_{i,t+1} - x_{i,t} = x_{i,t-1} \frac{w(\pi_{i,t}) - \bar{w}_t}{\bar{w}_t}, \quad (2)$$

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

⁴Branch and Evans (2007) study dynamic switching with multinomial logit in a Lucas-style macro model.

$\bar{w}_t = x_{1,t-1}w(\pi_{1,t}) + x_{2,t-1}w(\pi_{2,t}) + \dots x_{n,t-1}w(\pi_{n,t})$. Strategies with above average payoffs gain adherents. Dynamics such as (2) with x_i on the right-hand side are also *imitative* in that popular strategies tend to remain popular. In contrast to the multinomial logit model (1), the outcome where $x_{j,t} = 0$ is a distinct possibility under the evolutionary game theory dynamic (2), so poorly performing strategies could be driven out. Sandholm (2007) provides a detailed discussion comparing dynamics.

If $w(\cdot)$ is linear, equation (2) corresponds to the replicator dynamic, which has been analyzed in detail in the game theory literature, see Weibull (1997). Sethi and Franke (1995) and Branch and McGough (2005) apply alternative versions of the replicator dynamic to the cobweb model and show conditions where chaos arises. All the referenced models studying the cobweb model have agents using a naive forecast based on a simple backward looking forecast and a rational forecast for which they must pay a cost. In contrast, in the present approach, all agents have the same information and there is no cost differential between forecasting strategies⁵.

The specification of a linear weighting in (2) is restrictive, so we consider convex $w(\cdot)$'s, following Hofbauer and Weibull (1996). Compared to the linear case, under convex weighting agents switch more aggressively to strategies that yield small squared errors. The convexity of the weighting function turns out to be a key factor determining whether persistent heterogeneity of forecasts is a possibility.

Though a number of results apply to a general model with heterogeneous forecasts, a standard asset pricing model motivates some key concepts. The analysis focuses on three forecasts that satisfy rationality. The fundamentalist forecast determined solely by expected future dividends has special status in that there is always a small fraction of agents using this forecasts. Agents may also choose a rational bubble or *mystic* forecast.

The *reflective* forecast is an essential element in both the asset pricing model and the general model. This forecast follows the literature⁶ that sets out an econometric view of the merits of combining forecasts. Furthermore, in an environment with heterogeneous expectations, the reflective forecast is the one that embodies the information about the alternative forecasts and the

⁵Branch and Evans (2006) study a cobweb model with two costless mis-specified forecasts.

⁶Bates and Granger (1969) and Granger and Raathan (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.

fractions of the agents using them.

The central question of the paper is whether the fundamental forecast is robust to the introduction of agents experimenting with alternative forecasts such as mysticism when the reflective forecast is an alternative strategy. In both the asset pricing and general models, if all strategies besides fundamentalism and reflectivism are driven from the model, these two forecasts will coincide, and all agents use the fundamental forecast.

To analyze conditions yielding such an outcome, we define *reflective dominance* where the fraction following reflectivism is monotone increasing, which implies that the point with the maximal following for reflectivism is stable. In the asset pricing model, for appropriately bounded payoffs and linear weighting, which is associated with the replicator dynamic, reflective dominance obtains.

However, even in this framework favorable to reflectivism, the elimination of alternative strategies such as mysticism is not a foregone conclusion. Given sufficient convexity of the weighting function $w(\cdot)$ relative to the magnitudes of the payoffs in (2), when agents are more aggressively pursuing better performing strategies, reflectivism is not robust to the introduction of followers of non-fundamental forecasts. The potential for persistent heterogeneity is established by both analytic and simulation results for the asset pricing model.

The organization of the paper is as follows. The next section introduces the asset pricing model to motivate the different forecasts and provide an intuitive interpretation of reflective dominance. Section 3 describes the general model of heterogeneous forecasts, and Section 4 shows results for the replicator dynamic. Section 5 introduces and gives results for convex weighting. Sections 6 and 7 respectively demonstrate theoretical and simulation results for the asset pricing model. Section 8 concludes.

2 A Specific Game: Asset Pricing

Though our theoretical results are relevant for a large class of models with heterogeneous forecasts, 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

$$y_t = \alpha y_{t+1}^e + u_t, \tag{3}$$

where $\alpha < 1$ is a discount factor and u_t is an income flow.⁷ Convergence to agreement on a single expectation y_{t+1}^e is viewed here, however, as a possible conclusion rather than as an assumption. Given heterogeneous expectations and mean-variance optimizing agents, Brock and Hommes⁸ (1998) show that the realized security price depends on the weighted average of agents' expectations

$$y_t = \alpha x_t \cdot e_t + u_t, \quad (4)$$

where e_t is a vector of the forecasts of y_{t+1} and x_t is a vector of the fractions of the population using each forecast. The discounted present value of the expected income stream

$$y_t^* = u_t + \sum_{j=1}^{\infty} \alpha^j E_t(u_{t+j} | \Omega_t), \quad (5)$$

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 efficient markets hypothesis, though it is not the unique solution to (3). The natural candidate for the fundamental forecast of y_{t+1} is thereby

$$e_{2,t} = E(y_{t+1}^* | \Omega_t) = \sum_{j=1}^{\infty} \alpha^{j-1} E(u_{t+j} | \Omega_t). \quad (6)$$

While it is common to assume that all agents somehow recognize $e_{2,t}$ as the appropriate forecast, we assume only that the fraction δ_2 of the agents always do so. These unyielding fundamentalists might well be impressed by the fact that (6) 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 (6) to recognize that (6) is prominently featured in Cochrane (2001) and McCallum (1983, 1997).

Our theoretical results allow for an arbitrary number of alternative forecasting strategies, but for the model of asset pricing we select a challenger to the fundamentalist forecast from among the

⁷This equation could equally apply to aggregate prices, exchange rates, etc..

⁸Brock and Hommes (1998) study a similar model, but there are several important differences between their analysis and ours. 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 introduction, does not allow for convergence to a single expectation.

martingale or rational bubble solutions to the model (3). This alternative forecast will be

$$e_{3,t} = e_{2,t} + \alpha^{-t-1}m_t = E(y_{t+1}^*|\Omega_t) + \alpha^{-t-1}m_t. \quad (7)$$

where $m_t = m_{t-1} + \eta_t$ is a martingale. We label this forecast *mysticism* because, while $\alpha^{-t}m_t$ is thought by economic theorists to be extraneous, agents believing in (3) cannot rule out a martingale solution on the basis of that mathematical model.⁹ Followers of the mystical forecast might, for example, sincerely believe that $\alpha^{-t}m_t$ is a valid addition to the fundamentals for y_t . In fact, the martingale does influence y_t if the mystical forecast attracts followers.

The reflective forecast is postulated to be an average of the fundamental and mystic forecasts weighted according to their relative popularity.

$$e_{1,t} = (1 - n_t)e_{2,t} + n_te_{3,t}$$

where

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

is the proportion of followers of the mystical forecast among those not following the reflective forecast. Such a forecast follows the literature on the benefits of combining forecasts. See Elliot and Timmerman (2008) for multiple references. The reflective forecast can be written as

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

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 (6), (7), and (8) into (4), yielding

$$y_t = y_t^* + \alpha^{-t}n_tm_t. \quad (9)$$

The realization is thus the fundamentalist forecast plus, to the extent that some of the agents are

⁹Even 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.

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.

Payoffs are given by the negative of the squared forecast error

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

Evaluating forecasts using squared errors has a long tradition in econometrics¹⁰. For the present asset pricing model, Hommes (2001) shows that (10) is the natural objective function for mean-variance maximizing agents.

The reflective forecast error $U_t = y_t - e_{1,t-1}$ includes two components:

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

The first term on the right is the innovation in fundamentals. The second term on the right is the weighted martingale innovation. Using this approach, the payoff to reflectivism is

$$\pi_{1,t} = -U_t^2. \quad (12)$$

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 depends on the martingale and consequently is non-stationary. The fundamentalist forecast error $y_t - e_{2,t-1} = U_t + n_{t-1} A_{t-1}$, from (9) and (6), 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. The fundamentalist payoff is

$$\pi_{2,t} = -U_t^2 - 2n_{t-1}U_t A_{t-1} - n_{t-1}^2 A_{t-1}^2. \quad (13)$$

The mystical forecast error is $y_t - e_{3,t-1} = U_t - (1 - n_{t-1})A_{t-1}$, from (9) and (7), and the resulting payoff is

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

¹⁰Elliott and Timmerman (2008) discuss the role of mean squared prediction error and forecast combination in the literature on forecasting.

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 mysticism (13) and fundamentalism (14) are the dominating feature causing both payoffs to under-perform reflectivism. The reflective forecast is constructed such that the martingale does not affect its payoff, so, when the martingale term is large, reflectivism is best. However, if A_{t-1} is not large and the "covariance" $U_t A_{t-1}$ 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*¹¹ 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 (2) has the property that the fraction of agents using a strategy adjusts proportionally with the fitness of that strategy. Here, the population average is

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

While A_{t-1} does not enter the payoff to reflectivism $\pi_{1,t} = -U_t^2$, if there is any heterogeneity in the population, it does affect the population average.

Remark 1 *The fitness of reflectivism $\pi_{1,t} - \bar{\pi}_t$ is positive so, under the replicator, the point where the maximum fraction of agents is using reflectivism is stable.*

Our primary goal is to turn this remark into a rigorous proposition and to determine the conditions under which agents agree on a unique forecast given an arbitrary number of available strategies and a wider class of dynamics, so we now examine a more general framework.

¹¹The word fitness reflects the biological origins of evolutionary game theory, but there are social-economic foundations to these dynamics as well, see Weibull (1997).

3 An Evolutionary Framework for Expectations

We begin with a schematic description of an evolutionary model of expectation formation. More specific assumptions in later sections determine whether the mechanism we describe here supports agreement on a unique forecast.

Assume that agents choose among n forecasting strategies. In period t , the fraction $x_{i,t}$ of the agents follows strategy i , which provides them with the forecast $e_{i,t}$ for y_{t+1} . Let $x_t = (x_{1,t}, \dots, x_{n,t})$ and let $e_t = (e_{1,t}, \dots, e_{n,t})$ so that the average forecast is given by the inner product $\bar{e}_t = x_t \cdot e_t$. Similarly, let $\pi_t = (\pi_{1,t}, \dots, \pi_{n,t})$ denote the vector of payoffs.

The model for y_t is given by

$$y_t = g(x_t, e_t, \Omega_t, \varepsilon_t), \quad (15)$$

where the information set Ω_t is observable in period t , but the information set ε_t is not. We assume y_t takes on values in a set Y . The forecast for strategy i is given by

$$e_{i,t} = f_i(x_t, \Omega_t, \Gamma_{i,t}), \quad i = 1, \dots, n, \quad (16)$$

where the functions f_i , $i = 1, \dots, n$, differ across strategies. The information in Ω_t is common to all forecasts and corresponds to the expected future dividends in the asset pricing model in section 2. $\Gamma_{i,t}$ is a generic notation for the ingredients unique to a given forecasting strategy. These could include data, parameters, estimates of g , or even other forecasts. $\Gamma_{i,t}$ might well be extraneous in the sense that the realization of y_t depends on $\Gamma_{i,t}$ only because agents adopting $e_{i,t}$ think it does, as with the martingale in the asset pricing model.

Nothing about this structure rules out forecasting strategies that attempt to improve over time in the manner of, perhaps, least squares learning (Marcet and Sargent (1989a, 1989b)). Our basic notion of convergence, however, is not that all the forecasting strategies converge to a single strategy, but rather that the agents abandon unsuccessful forecasting strategies in favor of strategies with better forecasting records. If the forecast strategies are based on different values of some key parameter, then agents could come to agree on the value for that parameter by discarding unsuccessful strategies. For example, Guse (2005) applies the replicator dynamic to a model where

forecaster use different versions of a linear forecast.

We describe agents' efforts to find the best forecasting strategy in terms of an evolutionary game theory mechanism. In period t , agents calculate the payoffs based on a comparison of the forecasts $e_{i,t-1}$ of y_t with its realization,

$$\pi_{i,t} = h(y_t, e_{i,t-1}), \quad i = 1, \dots, n, \quad (17)$$

the squared error payoffs (10) being a particular example. Given these payoffs, they use the updating mechanism

$$x_{i,t+1} = k_i(\pi_t, x_t, \Psi_t), \quad i = 1, \dots, n, \quad (18)$$

to choose the forecasting strategies they will use in period $t+1$. The argument Ψ_t denotes whatever information agents take into account beyond the current payoffs π_t and the current fractions x_t .

While it is not necessary for our analysis, it is convenient to make a distinction between agents and forecasters.¹² The information and calculations (16) needed to produce a forecast could be quite sophisticated, but these calculations need be performed by only one forecaster for each strategy. The agents have only to choose a among the forecasters.

This distinction helps to justify our construction of a forecast such as the reflective that assumes knowledge of the populations shares $x_{i,t}$. While it might be difficult to argue that each agent knows all other agents' strategies, only one forecaster needs that information. That forecaster could bear the costs of surveying the agents and calculating the sample average forecast, but that is not the only possibility.¹³ The current x_t could be calculated using the recursion (18) and the histories of π_t and Ψ_t , which must be observable if agents are to implement (18). While either option has its complications, the agents do not conduct a survey or do the recursive calculations. They simply pick a forecaster.

Now we define the reflective and fundamental forecasts for the general model (15). The reflective

¹²In there discussion of bubbles and heterogeneous expectations, Föllmer, Horst, and Kirman (2005) make a similar distinction between agents and "gurus."

¹³In this paper, we do not rely on the costs of calculating forecasts to obtain results. If only the forecasters bear the costs of constructing forecasts and then spread that cost across their clients, the practical importance of cost might not be great. An issue of the Wall Street Journal, for example, costs \$1.00 at the newsstand. For an investor with a \$1 million portfolio, the problem is not paying for the newspaper, it is choosing among the differing outlooks expressed therein.

forecast highlights the role of agents who take into account the expectations of other agents. To discuss this strategy without cluttering later results with summations and subscripts, we will use the notation $\underline{e}_t = (e_{2,t}, \dots, e_{n,t})$ and $\underline{x}_t = (x_{2,t}, \dots, x_{n,t})/(1 - x_{1,t})$. For the general model, the reflective forecast

$$e_{1,t} = \underline{x}_t \cdot \underline{e}_t \tag{19}$$

is the population weighted average of the other forecasts. We refer to $e_{1,t}$ as reflective because, rather than putting forward an independent forecast of its own, it simply reflects the views of others.

As noted above, such forecast combination frequently outperform more sophisticated techniques for forming expectations. Elliot and Timmerman (2008) discuss possible reasons for this including model mis-specification and non-stationarity of explanatory variables, both of which could be significant factors affecting the dynamics of the models in the present work. Furthermore, rational agents should have the option of using a strategy that embodies available information about the alternative forecasts and the prevalence of their use in the population.

As with the appeal to efficient markets following (6), economic theory will often supply a *fundamentalist forecast*. For example, for many models that might admit the possibility of heterogeneous expectations, economic theorists derive a unique expectation by imposing a representative agent assumption.

An important characteristic of the fundamentalist forecasting strategy is that it has an unyielding core of followers. Some small percentage of the agents stick with the fundamentalist forecast regardless of the current payoffs because they have faith in economic theorizing. We express this as:

Assumption 1: Minimum Percentage Fundamentalism.

$$x_{2,t} \geq \delta_2, \quad \forall t \geq 0, \text{ for some constant } \delta_2 > 0.$$

We use a simple notation for this restriction. The population fractions have to be contained in the simplex

$$\Delta = \{x | x_1 + \dots + x_n = 1, x_1 \geq 0, \dots, x_n \geq 0\}.$$

The assumptions about the fundamentalist and reflective forecasting strategies further restrict the fractions to the simplex

$$\underline{\Delta} = \{x \in \Delta | x_i \geq \delta_i, i = 1, \dots, n\}, \quad \text{where } \delta_i = 0 \text{ for } i \neq 2.$$

4 Reflective Dominance

We can now set down conditions for reflectivism to outperform all non-fundamental strategies. Such conditions are analogous to the inherent advantage of reflectivism in the asset pricing model expressed in Remark 1. Note that in both the asset pricing and general frameworks, if reflectivism and fundamentalism are the only two strategies in the population then both forecasts are identical and expectations converge to the fundamentalist forecast.

The only assumptions we impose on the model (15) are that the realization of y_t is linear in the vector of forecasts e_t and the influence of forecast $e_{i,t}$ in determining y_t depends on the percentage $x_{i,t}$ of the population following that forecast. That is,

Assumption 2: Linear Expectations.

$$y_t = \alpha x_t \cdot e_t + z_t, \tag{20}$$

where α is a parameter and z_t is a convenient shorthand notation for some function of Ω_t and ε_t that constitutes the remainder of the equation. Again, note that a model of this form corresponds to the asset pricing equation (3) that has foundations based on the choice between risky versus riskless assets, see Brock and Hommes (1998).

We introduce a technical assumption designed to rule out a convergence failure that could result if the competing strategies actually had numerically identical forecasts.

Assumption 3: Differing Forecasts. If $x_{i,t} > 0$ and $x_{j,t} > 0$ for $i > 1$ and $j > 1$ and $i \neq j$, then with probability one $e_{i,t} \neq e_{j,t}$.

We exclude strategy 1 from this assumption because the whole point of the reflective forecast is to mimic other forecasts. The shape of the payoff function (17) is an important element of our proof that expectations converge.

Assumption 4: Concave Payoff Function. The payoff function (17) is globally concave in its second argument and has a finite maximum.

The payoffs based on squared forecast errors in Section 2 are concave in the forecast error with the maximum realized when the forecast coincides with the realized value. Elliot and Timmerman (2002) give other criteria for evaluating forecasts satisfying concavity such as those responding asymmetrically to square errors, though they discuss other methods that do not.

We next impose a minimal structure on the updating functions (18). Let $\bar{\pi}_t = x_{t-1} \cdot \pi_t$ denote the average payoff across agents. The time subscripts differ because $\pi_{i,t}$ is the payoff received in period t by the $x_{i,t-1}$ agents following strategy i in period $t-1$. We revisit this point in Section 7.

*Assumption 5: Payoff Positivity.*¹⁴ The updating function (18) satisfies

$$\begin{aligned} x_{i,t+1} &> x_{i,t} \text{ if } \pi_{i,t} > \bar{\pi}_t \\ x_{i,t+1} &= x_{i,t} \text{ if } \pi_{i,t} = \bar{\pi}_t \\ x_{i,t+1} &< x_{i,t} \text{ if } \pi_{i,t} < \bar{\pi}_t \end{aligned}$$

The replicator is the primary example of a dynamic satisfying payoff positivity.

These four assumptions set out a general, fairly unconstrained forecasting environment. Linear expectations are analytically convenient. A concave payoff function with a finite maximum rules out payoffs running off to infinity for certain forecasting results. Payoff positivity puts a weak restriction on the form of the updating functions without specifying those functions explicitly.

Definition 2 *We will say that the system (15), (16), (17), and (18) has **reflective dominance** if, for every x_t interior to $\underline{\Delta}$, the only fraction increasing with probability one is $x_{1,t}$.*

The practical implication of reflective dominance is that $x_t \cdot e_t$ in (20) converges to the fundamentalist forecast $e_{2,t}$. Any individual $x_{i,t}$ for $i = 2, \dots, n$ could increase in a given period t , but if $x_{1,t}$ is increasing in the interior of $\underline{\Delta}$, then $x_{2,t} + \dots + x_{n,t}$ must be decreasing. This process will stop only when x_t is on the boundary of $\underline{\Delta}$ containing points x_t of the form $(1-a, a, 0, \dots, 0)$ where $\delta_2 \leq a \leq 1 - \delta_1$, and the reflective and fundamentalist forecasts are numerically identical. Agents will have effectively arrived at a single forecast.

¹⁴Weibull (1997, p. 197) discusses payoff positivity in a different context.

Establishing conditions such that elements $x_{2,t}, \dots, x_{n,t}$ decrease with some nonzero probability should be fairly straightforward in the context of a particular model. We establish this point for the asset pricing model in Section 6. We focus here on proving that $x_{1,t}$ is increasing with probability one under general assumptions.

Proposition 3 *Assumptions 2 through 5 yield reflective dominance.*

Proof. Let $\underline{h}(y_t, \underline{e}_{t-1}) = (h(y_t, e_{2,t-1}), \dots, h(y_t, e_{n,t-1}))$ denote the vector of payoffs $(\pi_{2,t}, \dots, \pi_{n,t})$. Assumption 1 guarantees that $x_{2,t} > 0$, and, for x_t in the interior of $\underline{\Delta}$, it must be the case that $x_{i,t} > 0$ for at least one $i > 2$. Given that h is concave in its second argument and the forecasts satisfy Assumption 4, and using form of the reflective forecast (19)

$$h(y_t, \underline{x}_{t-1} \cdot \underline{e}_{t-1}) > \underline{x}_{t-1} \cdot \underline{h}(y_t, \underline{e}_{t-1})$$

with probability one. We can restate this as

$$\pi_{1,t} > \underline{x}_{t-1} \cdot \underline{\pi}_t.$$

It follows immediately that $\pi_{1,t} > x_{1,t-1}\pi_{1,t} + (1 - x_{1,t-1})\underline{x}_{t-1} \cdot \underline{\pi}_t = \bar{\pi}_t$ so we have

$$\pi_{1,t} > \bar{\pi}_t.$$

Payoff positivity then implies $x_{1,t+1} > x_{1,t}$ with probability one. ■

Proposition 3 describes convergence of expectations under general conditions. It does not depend on the number of forecasts $e_{3,t}, \dots, e_{n,t}$ competing with reflectivism and fundamentalism and it does not depend on the details of those forecasts. Given that some nonzero fraction of the agents initially follows the reflective forecast, the pivotal assumption is that the payoff function is globally concave. This proposition applies to the asset pricing model in section 2 and formalizes Remark 1 with strict inequality since $x_{i,t} > 0$ on the interior of $\underline{\Delta}$ and the requirement for differing forecasts in Assumption 3 implies that $A_t \neq 0$.

To further understand the formal implications of reflective dominance, let us define stability as follows, where $\|\cdot\|$ is the Euclidean metric.

Definition 4 Define a neighborhood of $x \in \underline{\Delta}$ given $\varepsilon > 0$ as follows, $N(x, \varepsilon) = \{y \in \underline{\Delta} \mid \|y - x\| < \varepsilon\}$. A point $x_0 \in \underline{\Delta}$ is **stable** if for any neighborhood $N(x_0, \varepsilon)$, there exists a $\xi > 0$ such that if $x_t \in N(x_0, \xi)$ then $x_{t+s} \in N(x_0, \varepsilon)$ for any positive integer s .

The following proposition relates stability and reflective dominance.

Proposition 5 Given Assumptions 2 through 5, the point $\tilde{x} = (1 - \delta_2, \delta_2, 0, \dots, 0)$ where reflectivism attains its maximum following is stable.

Proof. Given $\varepsilon > 0$, let $\xi = \min \{1 - \delta_2 - y_{1,t} \mid y_t \in \bar{N}(\tilde{x}, \varepsilon)\}$. If $x_t \in N(\tilde{x}, \xi)$, then reflective dominance ensures that $x_{t+s} > x_t$ and $x_{t+s} \in N(x_0, \varepsilon)$, so \tilde{x} is stable. ■

If x_t is in a neighborhood of \tilde{x} , it will stay there under the conditions of Propositions 3 and 5. At the point \tilde{x} where reflectivism achieves its maximum, the whole population agrees on the fundamental forecasts, which is also true at any point where reflectivism and fundamentalism are the only two forecasting strategies present in the population.

5 Payoff Weighting Functions

We can now discuss the above results on reflective dominance in light of the various specifications of the weighting function $w(\cdot)$ for the general dynamic (2). Curvature in $w(\cdot)$ can change agent behavior here somewhat analogously to how curvature in utility functions can affect behavior toward risk in other settings (Hofbauer and Weibull 1996, p. 563)

The replicator dynamic is a common benchmark in evolutionary game theory. A common form setting the weighting function in (2) to $w(\pi_{i,t}) = \pi_{i,t}$ has payoff positivity by construction, but makes sense only for positive payoffs. A simple alternative is to work with the revised weighting such that $w(\pi_{i,t}) = C + \pi_{i,t}$, 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

$$x_{i,t+1} - x_{i,t} = x_{i,t-1} \frac{\pi_{i,t} - \bar{\pi}_t}{C + \bar{\pi}_t}. \quad (21)$$

The factor on the right is $x_{i,t-1}$ rather than $x_{i,t}$ because the payoff $\pi_{i,t}$ accrues in period t to the fraction $x_{i,t-1}$ of agents choosing strategy i in period $t - 1$. Unlike simpler evolutionary games, a

model with forward expectations inherently requires payoffs to be realized at least one period after strategies are chosen. The timing in the general dynamic (2) could be interpreted in terms of the flow of information through word of mouth, published information or advertising.

Hofbauer and Sigmund (1988, p. 133) and Samuelson (1997, p. 66) discuss this version of the replicator dynamic with a constant (21). Since this dynamic also satisfies payoff positivity, Propositions 3 and 5 guarantee the stability of \tilde{x} where all expectations agree with the fundamental forecast. This result does not require us to specify the details of the model (20) or to even specify the number or nature of the alternative forecasts $e_{3,t}, \dots, e_{n,t}$. The next issue is whether reflective dominance and the stability of x_0 extend to the more general dynamic (2).

In static game theory, the form of the replicator (21) is quite sufficient as it is easy to choose C to be larger than the biggest payoff. For economic applications, such a choice may not be so obvious so we describe here two specific functional forms that achieve $w(\pi_{i,t}) \geq 0$ by construction. For a constant $C > 0$, the truncation function

$$w(\pi) = \left\{ \begin{array}{ll} C + \pi & \text{if } C + \pi \geq 0 \\ 0 & \text{if } C + \pi < 0 \end{array} \right\} \quad (22)$$

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

From a game theoretic point of view, 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. For the replicator dynamic, agent aggressiveness as measured by C determines the speed of adjustment of the fractions $x_{i,t}$. Smaller values for C increase the ratio $(\pi_{i,t} - \bar{\pi}_t)/(C + \bar{\pi}_t)$ driving $x_{i,t+1} - x_{i,t}$.

The exponential transformation

$$w(\pi) = e^{\pi/(2\sigma^2)} \quad (23)$$

achieves nonnegativity without sacrificing smoothness.¹⁵ Applying the exponential transformation to the squared forecast error $\pi_{i,t} = -(y_t - e_{i,t-1})^2$ produces a familiar functional form

$$w(h_t(y_t, e_{i,t-1})) = e^{-(y_t - e_{i,t-1})^2/(2\sigma^2)}. \quad (24)$$

¹⁵It has the properties that the relative slope $w'(\pi)/w(\pi) = (2\sigma^2)^{-1}$ and the convexity $w''(\pi)/w'(\pi) = (2\sigma^2)^{-1}$ are constant over π .

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 is chosen by the agents to reflect how they feel about large and small squared forecast errors.

For the exponential weighting function (23), agent aggressiveness takes the form of the parameter σ^2 . In terms of the normal density (24), if the notional variance σ^2 is large, $w(h_t(y_t - e_{i,t-1}))$ is not very sensitive to the magnitudes of the forecast errors. Equivalently, agents are not very aggressive about pursuing the best strategy. If σ^2 is small, agents assign appreciable value to only 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 share adjust proportionally to fitness.

Some element of convexity is unavoidable in the present framework. The constraint that $w(h(y_t, \cdot)) \geq 0$ with a finite maximum guarantees that $w(h(y_t, \cdot))$ will not be globally concave in its second argument. It will, however, have a concave region about its maximum. We formalize this point as:

Assumption 4': Concave Weighted Payoff Function. For every $y_t \in Y$, the function $w(h(y_t, \cdot))$ is concave in its second argument over an interval $\Theta(y_t)$ containing y_t .

In general, the size of the concave region $\Theta(y_t)$ of $w(h(y_t, \cdot))$ depends on the degree of convexity of $w(\cdot)$. For the normal density (24), the concave region $\Theta(y_t)$ is the interval $(y_t - \sigma, y_t + \sigma)$.¹⁶ For the truncation function (22), $\Theta(y_t) = (y_t - C^{\frac{1}{2}}, y_t + C^{\frac{1}{2}})$. In both cases, the size of the interval $\Theta(y_t)$ depends on the degree of agent aggressiveness. The more aggressive agents are, the shorter is $\Theta(y_t)$.

Even without specifying a particular $w(\cdot)$ we can give conditions implying $x_{1,t+1} > x_{1,t}$ in terms of a modification of payoff positivity.

Assumption 5': Weighted Payoff Positivity. The updating function (18) satisfies

$$\begin{aligned} x_{i,t+1} &> x_{i,t} \text{ if } w(\pi_{i,t}) > \bar{w}_t \\ x_{i,t+1} &= x_{i,t} \text{ if } w(\pi_{i,t}) = \bar{w}_t \\ x_{i,t+1} &< x_{i,t} \text{ if } w(\pi_{i,t}) < \bar{w}_t \end{aligned}$$

¹⁶The standard normal density has inflection points at -1 and +1.

where $\bar{w}_t = x_{t-1} \cdot w(\pi_t) = \sum_{i=1}^n x_{i,t-1} w(\pi_{i,t})$.

Forecast strategy i gains in popularity if its weighted payoff $w(\pi_{i,t})$ is greater than the average weighted payoff \bar{w}_t . In the next section, we discuss a specific updating mechanism with weighted payoff positivity.

We can extend Propositions 3 and 5 to the case of weighted payoffs. Again, we do not have to specify the model, the strategies competing with fundamentalism and reflectivism, or even the number of strategies.

Proposition 6 *Given Assumptions 2, 3, 4', and 5', if $e_{i,t} \in \Theta(y_t)$ for all i , reflective dominance obtains, and \tilde{x} is stable.*

Proof. The proof of Proposition 3 applies with minor changes in notation. ■

The assumption that the forecasts are in a bounded region about y_t may or may not be plausible for a given model. In the asset pricing mode, if the dividends and innovations to the martingale are normally distributed, then the payoffs are unbounded, and we find that forecasts outside $\Theta(y_t)$ can lead to persistent heterogeneous expectations.

6 Asset Pricing Dynamics

Now, applying the results for the general model, we can extend Remark 1 and give formal results for the asset pricing model in section 2 by considering the class of evolutionary dynamics (2) that includes the replicator and convex monotonic dynamics as special cases. If the stochastic elements of the model, the dividends and martingale innovations, are appropriately bounded, reflective dominance (Definition 1) obtains as does the stability (Definition 4) of \tilde{x} , implying that mysticism is driven from the population and all agents are using the fundamental forecast. However, if agents are sufficiently aggressive, such bounds cannot be established and persistent heterogeneity is possible.

The asset pricing model discussed in this section is defined by the fundamental solution (5), the realization of the asset price (9) and the payoffs to reflectivism (12), mysticism (13) and fundamentalism (14). The payoffs are determined by squared errors of the forecasts as in (10).

If agents update their strategies using the replicator dynamic with the truncation function (22),

then Assumption 4' is satisfied. Also, all the transformed payoffs are positive if A_{t-1} and U_t satisfy

$$\max\{U_t^2, (U_t + n_{t-1}A_{t-1})^2, (U_t - (1 - n_{t-1})A_{t-1})^2\} < C. \quad (25)$$

This condition puts all the forecasts in the concave region $\Theta(y_t)$ of $w(h(y_t, \cdot))$ thereby bounding the payoffs to the three forecasting strategies, so the conditions of Proposition 6 are satisfied, which yields reflective dominance.

While (25) involves n_{t-1} , the following result emphasizes that bounds on combinations of A_{t-1} and U_t alone are sufficient to establish $x_{1,t+1} > x_{1,t}$.

Proposition 7 *The dynamics given by (22) and (2) yields reflective dominance if A_{t-1} and U_t satisfy*

$$\max\{(U_t + A_{t-1})^2, (U_t - A_{t-1})^2\} < C. \quad (26)$$

Proof. With probability one, $A_t \neq 0$, and Assumption 3 is satisfied. The quadratic function $(U_t + \theta A_{t-1})^2$ is convex. It attains its maximum over $-1 \leq \theta \leq 1$ at one of the two endpoints. ■

Condition (26) might suggest that a large martingale term A_{t-1} could cause nonrobustness. However, for large A_{t-1} , the payoffs to fundamentalism (13) and mysticism (14) are unambiguously worse than the payoff to reflectivism, which does not depend on A_{t-1} . The following result shows that, given a bound on n_t , a bound on U_t^2 guarantees $x_{1,t+1} > x_{1,t}$ for all possible values of A_{t-1} . The bound on n_t is not particularly restrictive because we are primarily interested in whether n_t is likely to grow from a small initial value.

Proposition 8 *Suppose that*

$$n_{t-1} \leq 1 - \varphi$$

for a constant $0 < \varphi < 1$. The dynamics given by (22) and (2) if U_t satisfies

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

then either $x_{1,t+1} > x_{1,t}$ or $x_{3,t+1} = 0$.

Proof. The bound on U_t^2 guarantees that $w(\pi_{1,t}) > 0$. If $w(\pi_{2,t}) > 0$ and $w(\pi_{3,t}) > 0$, Proposition 3 implies that $x_{1,t+1} > x_{1,t}$. The same inequality holds if $w(\pi_{2,t}) = 0$ and $w(\pi_{3,t}) = 0$.

If $w(\pi_{2,t}) > 0$ and $w(\pi_{3,t}) = 0$, then $x_{3,t+1} = 0$. Suppose then that $w(\pi_{2,t}) = 0$ and $w(\pi_{3,t}) > 0$.

We have $x_{1,t+1} > x_{1,t}$ if

$$w(\pi_{1,t}) > n_{t-1}w(\pi_{3,t}) + (1 - n_{t-1})w(\pi_{2,t}), \quad (27)$$

which can be rewritten as

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

The maximum value of the right-hand side of this inequality is achieved at $A_{t-1} = U_t/(1 - n_{t-1})$, which maximizes $\pi_{3,t}$. Making this substitution shows that (27) must hold if $U_t^2/(1 - n_{t-1}) < C$. Equivalently, (27) must hold if $n_{t-1} \leq 1 - \frac{U_t^2}{C}$. The latter inequality follows from

$$n_{t-1} \leq 1 - \varphi \leq 1 - \frac{U_t^2}{C}.$$

■

The logic of the above proof can be extended to any point in the interior of the simplex for an appropriate bound to U_t .

If agents use the exponential weighting function (24), Proposition 6 still applies and the conditions under which the forecast errors are within the concave region of $w(h(y_t, \cdot))$ are similar to those for Proposition 7 with σ^2 representing the aggressiveness of the agents.

Proposition 9 *The dynamics given by (23) and (2) will lead to $x_{1,t+1} > x_{1,t}$ with probability one for x_t in the interior of $\underline{\Delta}$ if A_{t-1} and U_t satisfy*

$$\max\{(U_t + A_{t-1})^2, (U_t - A_{t-1})^2\} < \sigma^2. \quad (28)$$

Proof. The standard normal density is concave between -1 and +1. All three forecast errors will be on the concave segment if

$$\max\{U_t^2, (U_t + n_{t-1}A_{t-1})^2, (U_t - (1 - n_{t-1})A_{t-1})^2\} < \sigma^2.$$

The quadratic function $(U_t + \theta A_{t-1})^2$ is convex. It attains its maximum over $-1 \leq \theta \leq 1$ at one of

the two endpoints. ■

Furthermore, for the exponential weighting case, we present an informal argument for a restriction solely on U_t leading to reflective dominance as in Proposition 8. The dynamic (2) shows that for reflectivism to increase the weighted payoff to reflectivism $w(\pi_{1,t})$ must be larger than the weighted average \bar{w}_t payoff so $\frac{\bar{w}_t}{w(\pi_{1,t})} < 1$. Using exponential weighting (23), 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 above equation are close to zero. The following approximation uses Taylor approximations of the exponential functions around zero, eliminating powers of A_{t-1} higher than two since they are small.

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

If we restrict the Taylor approximation to linear terms, then the U_t^2 term would not appear, the ratio $\frac{\bar{w}_t}{w(\pi_{1,t})} < 1$, and reflective dominance holds as with linear weighting. When the second order term of the Taylor approximation is included, reflective dominance 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 and reflective dominance holds, 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 dominance does not hold and mysticism has an opportunity to gain adherents. To determine the quantitative effects of changes in model parameters on the likelihood for such instability, we examine simulations of the asset pricing model.

7 Robustness

The results in the previous sections show the potential for either agreement on the fundamental forecast within the population or persistent heterogeneity in forecasting strategies. If forecast errors are sometimes outside the concave region $\Theta(y_t)$ of $w(h(y_t, \cdot))$, alternative strategies such

as mysticism may not be driven from the population, leading to the instability of \tilde{x} . We can characterize the conditions that might lead to heterogeneous expectations by considering how the following for mysticism might increase from the initial minimal following¹⁷ $n_{t-1} \approx 0$.

When n_{t-1} is near zero, growth of mysticism requires two conditions. The payoff $\pi_{3,t} = -(U_t - (1 - n_{t-1})A_{t-1})^2$ is maximized over A_{t-1} by $A_{t-1} = U_t/(1 - n_{t-1})$, which in this case means $A_{t-1} \approx U_t$, and the term U_t is largely the innovation in fundamentals, see (11). The first condition is thus that the martingale term A_{t-1} predicts the innovation in fundamentals. The payoff $\pi_{3,t} \approx 0$ so the mystical forecast is not a candidate to violate the bounds (25) or (28). The fundamentalist and/or reflective forecasts must be outside $\Theta(y_t)$. In fact, $U_t + n_{t-1}A_{t-1} \approx U_t$ so the fundamentalist and reflective forecasts are nearly equal. The second condition is thus that the squared innovation in fundamentals is large enough ($U_t^2 > C$ or $U_t^2 > \sigma^2$) to push the fundamentalist and reflective forecasts out of the concave region $\Theta(y_t)$ of $w(h(y_t, \cdot))$.

We conduct simulations to explore the properties of the model when the conditions proving reflective dominance for all t do not hold. 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.

The initial population share for fundamentalism is set to the minimum $x_{2,0} = 0.05$, but mysticism starts at $x_{3,t} = 0.0001$, which is 500 times smaller. At the start then, $n_t = 0.002$ and agents are very nearly following the fundamentalist forecast. If $x_{3,t}$ falls below 0.0001 in a given period, we restart the martingale at $m_t = 0$ and reset the fraction of the agents following mysticism to $x_{3,t} = 0.0001$.¹⁸ The discrete time replicator dynamic does not guarantee $x_{i,t} \leq 1$, so we impose the minima $x_{1,t} \geq 0.05$, $x_{2,t} \geq 0.05$, and $x_{3,t} \geq 0.0001$. If 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. While this process is a bit tedious to describe, it is not of practical importance because

¹⁷Binmore, Gale, and Samuelson (1995) introduce drift, which is a similar approach examining the effects of introducing a small fraction using a strategy.

¹⁸It would be possible to have martingales starting more often, perhaps every period. That would complicate the simulations and probably work to favor more frequent episodes of heterogeneous expectations.

the outcome is determined by the behavior of x_t in the interior of the simplex $\underline{\Delta}$.

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

Table 2 (the truncation weighting function) and Table 3 (the exponential weighting function) show rather forcefully that mysticism will be an important factor if agents are sufficiently aggressive and the martingale innovation is large enough. While 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, 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.

There appear to be some differences in the two cases. For the truncation function (22), we find that convergence to a single expectation fails if the forecasts have an appreciable probability of being at all outside the concave region. The sharp change in curvature at the truncation point $C + \pi_{i,t} = 0$ apparently forces a sharp change in convergence behavior when $C + \pi_{i,t}$ is negative. The normal density (24), on the other hand, has a smoother transition from concavity to convexity, and the forecasts appear to have to be substantially outside $\Theta(y_t)$ for convergence to a single expectation to fail. For both cases, the size of the forecast errors relative to agent aggressiveness is the primary determinant of whether we observe convergence to a single expectation.

Though 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. Though the mystic can outperform the reflective forecast for given periods, over time the reflective forecast is superior so agents eventually abandon mysticism. Hence, bubbles arise and collapse endogenously¹⁹.

¹⁹This feature contrasts with the exogenous probability of a bursting bubble in the rational bubble literature, for

Parke and Waters (2007) show that, when conditions are favorable for repeated episodes of mysticism, the observed y_t 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.

8 Summary and Conclusions

Our evolutionary game theory approach provides two main insights into the issue of heterogeneity of expectations. The first is that agreement on a unique forecast follows from a symbiotic relationship between two forecasting strategies. A minimum percentage of belief in the fundamentalist strategy establishes the role of fundamentals in the eventual outcome, but the robustness of that outcome relies on the existence of agents following the reflective forecasting strategy, which takes into account the expectations of the other agents. The reflective forecasting strategy drives out all other strategies with the exception of the percentage δ_2 of the agents who simply will not abandon their belief in fundamentals. At that point, the fundamental and reflective forecasts agree numerically so that, in effect, all agents agree on fundamentals.

Our second insight is that the reflective dominance driving this stability depends on the curvature of agents' payoff weighting functions (relative to the variances of the stochastic elements of the model and the concavity of the payoff function). If the payoff weighting functions are linear over the relevant range or not too convex, reflective dominance drives the outcome to universal belief in the fundamentalist forecast. For larger degrees of convexity, however, the outcome is better described as persistent heterogeneous expectations.

We interpret convexity in the payoff weighting functions as agent aggressiveness. For a highly convex payoff weighting function, agents put little value on all but the best forecasts, which approaches best reply behavior. In a stochastic environment, aggressive pursuit of the best strategy can cause agents to switch to alternative forecasts $e_{3,t}, \dots, e_{n,t}$ based on a small number of lucky forecasts, creating the potential for persistent heterogeneity.

We demonstrate these results in the context of a basic asset pricing model where the standard example Evans (1991). See Parke and Waters (2007) for further discussion and simulation examples.

rational expectations solution is regarded as the fundamentalist forecast and a martingale solution is an alternative. If the martingale innovation variance is small and the payoff weighting function is not too convex, then agents agree on the standard rational expectations solution. If, on the other hand, the martingale innovation variance is larger and the payoff weighting function is more convex, then mysticism can gain a following. In the latter case, we find that if only one agent in 10,000 experiments with the martingale solution, the probability that at least 20% of the agents switch to that martingale forecast at some time within the first 100 periods approaches 100%. There is thus a decidedly nonzero probability that a martingale forecast with a very small following will become, for a time at least, a major factor if agents are aggressively pursuing the best forecast.

These results have a direct bearing on the merits of simply 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.

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.	Ref.	Myst.
$-n_{t-1}A_{t-1}/2 < U_t < (1 - 2n_{t-1})A_{t-1}/2$	Ref.	Fund.	Myst.
$(1 - 2n_{t-1})A_{t-1}/2 < U_t < (1 - n_{t-1})A_{t-1}/2$	Ref.	Myst.	Fund.
$(1 - n_{t-1})A_{t-1}/2 < U_t$	Myst.	Ref.	Fund.

Table 2					
Truncated Squared Errors (22)					
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	0.114	0.641	0.943	0.991
$\sigma_\eta = 1/8$	0.002	0.132	0.700	0.971	0.998
$\sigma_\eta = 1/4$	0.002	0.145	0.750	0.986	1.000
$\sigma_\eta = 1/2$	0.002	0.165	0.788	0.992	1.000
$\sigma_\eta = 1$	0.003	0.172	0.822	0.996	1.000
$\sigma_\eta = 2$	0.002	0.178	0.839	0.997	1.000

Table 3					
Exponential Weighting (23)					
Probability $x_{3,t} \geq 0.20$ for some $t \leq 100$					
10,000 trials					
	$\sigma^2 = 1$	$\sigma^2 = 1/2$	$\sigma^2 = 1/4$	$\sigma^2 = 1/8$	$\sigma^2 = 1/16$
$\sigma_\eta = 1/16$	0.000	0.026	0.407	0.884	0.985
$\sigma_\eta = 1/8$	0.000	0.044	0.538	0.954	0.998
$\sigma_\eta = 1/4$	0.000	0.070	0.662	0.984	1.000
$\sigma_\eta = 1/2$	0.000	0.112	0.766	0.995	1.000
$\sigma_\eta = 1$	0.001	0.148	0.825	0.997	1.000
$\sigma_\eta = 2$	0.003	0.198	0.850	0.999	1.000

References

- Bates, J.M. and C.W.J. Granger, 1969, The Combination of Forecasts, *Operations Research Quarterly* 20, 451-468.
- Binmore, K., Gale, J., Samuelson, L., 1995. Learning to be imperfect: the ultimatum game. *Games and Economic Behavior* 8, 56-90.
- Björnerstedt, J., Weibull, J.W., 1996. Nash equilibrium and evolution by imitation, in *The Rational Foundations of Economic Behavior*, K. Arrow, et. al., eds., Macmillan, London, 1996.
- Blanchard, O. 1979. Backward and forward solutions for economies with rational expectations, *American Economic Review* 69(2), 114-118.
- Box, G.E.P., Jenkins, G.M. , 1970, *Time Series Analysis: Forecasting and Control*, Holden-Day, San Francisco.
- Branch, W., Evans G., 2007, Model uncertainty and endogenous volatility. *Review of Economic Dynamics* 10, 207-237.
- Branch, W., McGough, B., 2004, Multiple equilibria in heterogeneous expectations models, *Contributions to Macroeconomics* 4(1), Article 12.
- Branch, W. and B. McGough, 2005, Replicator dynamics in a cobweb model with rationally heterogeneous expectations, Working Paper.
- Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. *Econometrica* 65, 1059-1095.
- Brock, W., Hommes, C., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22, 1235-1274.
- Brock, W.A., Hommes, C.H., Wagener, F.O.O., 2005. Evolutionary dynamics in financial markets with many trader types. *Journal of Mathematical Economics* 41, 7-42.
- Charemza, W., Deadman, D. 1995. Speculative bubbles with stochastic explosive roots: the failure of unit root testing, *Journal of Empirical Finance* 2, 153-163.
- Cochrane, J., 2001. *Asset Pricing*. Princeton University Press, Princeton, NJ.
- Constanides, G.M., Duffie, D. 1996. Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219-240.
- Covel, M , 2004, *Trend Following: How Great Traders Make Millions in Up or Down Markets*, (Upper Saddle River, New Jersey, Prentice Hall).
- Elliot, G. and Timmerman, A. 2004. Optimal Forecast Combinations under General Loss Functions and Forecast Error Distributions, *Journal of Econometrics* 122(1), 47-79.
- Elliot, G. and Timmerman, A. 2008. Economic Forecasting, *Journal of Economic Literature*, 46(1), 3-56.

- Evans, G.S. 1991. Pitfalls in testing for explosive bubbles in asset prices, *American Economic Review* 81(4), 922-930.
- Evans, G., Guesnerie, R., 2005. Coordination on saddle path solutions: the educative viewpoint – linear multivariate models. *Journal of Economic Theory* 124, 202-229.
- Evans, G., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.
- Evans, G., Honkapohja, S., Williams, N. 2005. Generalized stochastic gradient learning, NBER Working Paper 0317.
- Föllmer, H, Horst, U., Kirman, A., 2005. Equilibria in financial markets with heterogeneous agents: a probabilistic perspective. *Journal of Mathematical Economics* 41, 123-155.
- Granger, C.W.J. and R. Ramanathan, 1984, Improved methods of combining forecasts, *Journal of Forecasting* 3, 197-204.
- Guse, E., 2005, Learning with heterogeneous expectations in an evolutionary world, working paper.
- Hofbauer, J., Sigmund, K., 1988. *The theory of evolution and dynamical systems*. Cambridge University Press, Cambridge.
- Hofbauer, J., Weibull, J., 1996. Evolutionary selection against dominated strategies. *Journal of Economic Theory* 71, 558-573.
- Hommes, C. 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
- Marcet, A., Sargent, T., 1989a. Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, 337-368.
- Marcet, A. Sargent, T., 1989b. Convergence of least squares learning in an environment with hidden state variables and private information. *Journal of Political Economy* 97(6), 1306-1322.
- McCallum, B., 1983. On non-uniqueness in rational expectations models: an attempt at perspective. *Journal of Monetary Economics* 11, 139-168.
- McCallum, B., 1997. The role of the minimum state variables criterion in rational expectations models. *International Journal of Tax and Finance* 6(4), 621-639.
- Parke, W., Waters, G. 2007. An evolutionary game theory explanation of arch effects. *Journal of Economic Dynamics and Control* 31(7), 2234-2262.
- Samuelson, L. 1997. *Evolutionary games and equilibrium selection*, The MIT Press, Cambridge, Massachusetts.

Sethi, R. and Franke, R., 1995. Behavioural heterogeneity under evolutionary pressure: macroeconomics implications of costly optimization. *The Economics Journal* 105, 583-600.

Timmerman, A., 1994. Present value models with feedback. *Journal of Economic Dynamics and Control* 18, 1093-1119.

Stock, J., Watson M., 1999. A Comparison of linear and non-linear models for forecasting macroeconomic time series. in *Cointegration, Causality and Forecasting: A Festschrift in Honour of Clive W. Granger* ed. R.F. Engle and H. White, Oxford and New York, Oxford University Press, 1-44.

Weibull, J., 1997. *Evolutionary Game Theory*. MIT Press, Cambridge, MA.