# An Evolutionary Game Theory Approach to Rational Expectations

William R. Parke[1]
Department of Economics
University of North Carolina
Chapel Hill, NC 27599

George A. Waters[2]
Department of Economics
Illinois State University
Normal, IL 61790-4200

August 1, 2003

[1] parke@email.unc.edu
[2] gawater@ilstu.edu

**Abstract**

Multiple martingale, or bubble, solutions exist when expectations about future values of endogenous variables enter a model under rational expectations. Such solutions reproduce bubble behavior, as seen in asset markets for example, and have implications about the nature of rational expectations in many contexts. Economists often select a unique rational expectations solution from the class of martingale solutions using a principle such as minimum state variables. In contrast, we allow agents to choose from different forecasts based on bubble solutions. Evolutionary game theory selection dynamics describe agents' choices of forecasts using squared forecast errors as payoffs, and we study whether agents' expectations converge to a unique, fundamental solution. Under the commonly studied replicator dynamic, agents come to agree on a fundamental forecast. However, under convex monotonic selection dynamics, when agents switch more rapidly to forecasts with lower errors, bubble solutions can play an important role in the evolution of the model. These results show how bubbles can arise and raise doubts about focusing on a unique rational expectations solution.

*Metaphysics, or the attempt to conceive the world as a whole by means of thought, has been developed, from the first, by the union and conflict of two very different human impulses, the one urging men towards mysticism, the other urging them towards science.*
Bertrand Russell, **Mysticism and Logic**, 1917.

## 1. Introduction

The existence of the martingale solutions for models with forward-looking expectations reflects the fact that the rational expectations assumption alone does not identify a unique solution. Uniqueness in applications involving forward-looking expectations is generally achieved by assuming that agents use *deeper theorizing* to construct a principle that singles out a fundamental solution.[1] The candidates for this principle include expectational stability [Evans (1986), Evans and Honkapohja (2001)], stationarity [Gourieroux, Laffont, and Monfort (1982)], minimal use of state variables [McCallum (1983,1997)], and minimum variance [Taylor (1977)]. Although the multiplicity of these proposed principles raises the issue of how a single principle becomes common knowledge, the various principles in fact agree on the same fundamental solution in many cases.[2] The real challenge the martingale solutions pose is how agents acquire the sophisticated analytic ability it takes to understand the mathematical subtleties of the problem they face. If the agents are not endowed with the necessary sophisticated analytic ability, then the question becomes how they come to behave as if they are.[3]

Using an evolutionary game theory framework, we establish conditions under which practical

[1] Pesaran (1987, p. 96) uses the term *deeper theorizing* and surveys the various principles.

[2] We use the term common knowledge here in the sense that not only do all agents use the same forecasting strategy, but all agents know that they all use the same forecasting strategy, all agents know that all agents know they all use the same forecasting strategy, and so forth [Osborne and Rubinstein (1994, p. 73)]. If any link in this chain is broken by some free-spirited, inquisitive agent, then the original theory does not specify what any of the agents will do.

[3] The *as if* postulate has a long history in economics. Friedman (1953) supports the standard analytics of price theory with analogies to how leaves on a tree solve the problem of maximizing collective exposure to sunlight and how expert billiards players make shots as if they had spent years in the classroom learning to solve mathematical formulas. Lucas (1978, p. 1429) explicitly states that the sophisticated rational behavior he assumes is the outcome of an unspecified *as if* process. Here, we investigate how agents might learn to behave as if they had thought through the martingale solutions problem.

experience will lead agents holding heterogeneous beliefs to agree on the fundamental solution.[4] The agents choose among three forecasting strategies that all satisfy the principle of rational expectations.[5] A fraction $\gamma_t$ of the agents follow the *fundamentalist* forecast $e_{\gamma,t}$ that agrees with the deeper theorizing analysis. Another fraction $\lambda_t$ follow a *mystical* forecast $e_{\lambda,t}$ that adds a martingale to the fundamentalist forecast. The remaining fraction $\beta_t = 1 - \gamma_t - \lambda_t$ follow the *reflective* forecast $e_{\beta,t}$ that equals the mathematical expectation of $y_t$ taking into account $\gamma_t$, $e_{\gamma,t}$, $\lambda_t$, and $e_{\lambda,t}$.

The evolutionary dynamics of the fractions $\gamma_t$, $\lambda_t$, and $\beta_t$ provide some insights into the merits of assuming that a unique fundamental solution is common knowledge. If a small fraction of the agents experiment with the mystical forecast, there are two possible outcomes. The agents might return fairly quickly to unanimous belief in the fundamental solution, lending support to the notion of imposing the fundamental solution by assumption. Mysticism might, on the other hand, grow in popularity and persist for a long period of time, giving rise to an episode that could be characterized as a bubble. Occurrences of such episodes with a nontrivial probability would weaken the basis for simply imposing the fundamental solution by assumption.

Evolutionary game theory is well suited to studying the dynamics of beliefs and the robustness of an outcome.[6] It is based on the notion that agents might arrive at sophisticated behavior as the result of selection among heterogeneous, less sophisticated behavioral rules. Samuelson (1998, p. 15) cites the parallel to *the process by which competitive markets are typically described as reaching equilibrium, with high-profit behavior being rewarded at the expense of low-profit behavior.* Theoretical developments have been aimed at understanding how agents might narrow the set of interesting equilibria for models with large sets of Nash equilibria without necessarily understanding the subtle mathematical details of game theory. Binmore, Gale, and Samuelson (1995), for example, use drift, which is an injection of small fractions of agents using each strategy at each period,

[4] To focus attention on the issues raised by the martingale solutions, we assume that all agents know the model's parameters. The extensive literature on convergence to rational expectations detailing how agents might attempt to learn the parameters of a model with expectations includes Bray (1982) and Marcet and Sargent (1989a,1989b). The sunspot equilibria considered in that context by Woodford (1990) and Howitt and McAfee (1992) share the extraneous nature of the mysticism considered here, but our agents do know the parameters of the model. Evans and Honkapohja (2001) and Grandmont (1998) survey this literature.

[5] Others studies of heterogeneous expectations, including Brock and Hommes (1997,1998), Chiarella (2002), Degrauwe (1993), DeLong, Shleifer, Summers, and Waldmann (1990), and LeBaron, Arthur, and Palmer (1999), pit agents with various naive forecasting strategies against agents who have perfect foresight or rational expectations. All our agents understand rational expectations.

[6] Samuelson (1998) and Weibull (1997) survey evolutionary game theory.

to evaluate and select among Nash equilibria. This is similar to our efforts to study whether nearly universal belief in the fundamentalist forecast is robust to some agents experimenting with mysticism.

Implementing an evolutionary game theory approach to the contest among fundamentalism, mysticism, and reflectivism requires some specifics of how agents choose among the three strategies. In a general notation for $n$ strategies, let $x_t = (x_{1,t}, ..., x_{n,t})$ denote the population shares at time $t$ for strategies $(s_1..., s_n)$. Let $u(s_i, x_t)$ denote the payoff to following strategy $s_i$ given that the overall distribution across strategies at time $t$ is $x_t$. The payoffs are taken here to be the negatives of the squared forecast errors. We assume that the distribution of agents across strategies evolves according to

$$\frac{x_{i,t+1} - x_{i,t}}{x_{i,t}} = \delta(x_t)\ f(u(s_i, x_t)) + \mu(x_t)\ , \quad i = 1, ..., n, \tag{1.1}$$

where $f(\cdot)$ is a possibly nonlinear transformation of the payoff, $\delta(x_t)$ is a scale factor, and $\mu(x_t)$ ensures that the population shares sum to one. This is a discrete time version of Hofbauer and Weibull's (1996) equation (6).[7]

We consider two possibilities that we refer to as imitation of successful agents (ISA) and review of unsuccessful forecasts (RUF).[8] Imitation of successful agents emphasizes the choice of a new forecasting strategy, assuming that agents reconsider their strategies each period and that agents tend to adopt new strategies that have small realized forecasting errors and are popular in the overall population. Review of unsuccessful forecasts focuses on the decision to review a forecasting strategy, assuming that the decision to review depends on a forecast's current performance. Poorly performing forecasts are more likely to be reviewed and replaced. Both principles can be expressed as algebraic special cases of (1.1).

For both ISA and RUF, we are able to demonstrate monotone convergence to universal belief in the fundamental solution under certain conditions. Central among these is the *replicator dynamic*, where $f(\cdot)$ in (1.1) is linear in the payoff $u(s_i, x_t)$ and the growth rate for $x_{i,t}$ is proportional to the *fitness* of $x_{i,t}$, which is the difference between its payoff $u(s_i, x_t)$ and the population average payoff

[7] Bjornerstedt and Weibull (1993) and Weibull (1997) develop evolutionary game theory as a model of social interaction rather than biological evolution. Cabrales and Sobel (1992) provide some theoretical background for discrete time selection dynamics. Hofbauer and Weibull (1996) show that (1.1) ecompasses dynamics approaching both best-reply and worst-reply (when almost all agents abandon the strategy with the worst payoff) behavior, with the replicator dynamic being an intermediate case.

[8] Both cases are discussed by Bjornerstedt and Weibull (1993) and Weibull (1997).

$\sum_{j=1}^{n} x_{j,t} u(s_j, x_t)$.[9] We also examine the dynamics for a convex $f(\cdot)$, where agents are switching to superior strategies more aggressively than under the replicator.[10] For convex functions $f(\cdot)$, we find that convergence occurs if the errors driving the process are sufficiently small relative to the degree of convexity of $f(\cdot)$. We conclude, therefore, that, if $f(\cdot)$ is linear, or minimally convex, mysticism will quickly disappear and agents will follow some combination of fundamentalism and reflectivism. The exact mix will not matter because, if the following for the mystical forecast is zero, the fundamentalist and reflective forecasts are identical.

We also demonstrate, however, that this convergence may not be robust with respect to the introduction of mysticism. Robustness tends to fail when the error variances are large, indicating substantial uncertainty about the fundamentals, and when $f(\cdot)$ is convex because agents are aggressively pursuing superior strategies. Under those conditions, there can be a nontrivial probability that mysticism has a significant following for an extended period. If so and if agents occasionally experiment with mysticism, then mysticism can be an important feature of the overall process.

Other evidence exists that the long run might not be dominated by the fundamental rational expectations solution. Brock and Hommes (1998) analyze heterogeneous expectations in an asset pricing model under the assumption that the level of $x_{i,t}$ (rather than the growth rate) is proportional to $\exp(u(s_i, x_t))$. Using trading profits as the payoffs, they come to a conclusion similar to ours in that they find *persistent deviations from the fundamental price, and highly irregular, chaotic asset price fluctuations, when the intensity of choice to switch prediction strategies becomes high* (p. 1265). This finding echoes their conclusions for a cobweb model [Brock and Hommes (1997)]. Branch (forthcoming) emphasizes that the variety of available forecasting strategies can be an important determinant of stability by showing that the introduction of agents following adaptive expectations can dampen the cobweb model price oscillations. Blume and Easley (1992) use an evolutionary approach to study asymptotic behavior in an asset market, showing that ad hoc behavior can survive asymptotically. LeBaron, Arthur, and Palmer (1999) conduct simulations using a genetic learning algorithm and are able to replicate many of the time series features of real markets. They find [p. 1513] that the agents *concentrate on rules using those pieces of forecasting*

[9] Branch (2001) studies a version of the Brock and Hommes (1997) cobweb model using the replicator dynamic.

[10] Such dynamics fall into the class of *convex monotonic* dynamics [Hofbauer and Weibull(1996)]. Extreme convexity implies that agents are approaching *best-reply* behavior where nearly all agents switch to the strategy with the highest current payoff.

*information which appear significant in the time series tests, but should not be important inside a homogeneous linear rational expectations equilibrium.*

Our results are unique in that, while earlier studies with heterogeneity explore the implications of misspecified beliefs, all our agents are attempting to implement behavior satisfying the concept of rational expectations. We find that disagreements about how to implement rational expectations can lead to an outcome that is not robust with respect to experimentation with mysticism. If some of the agents adopt the fundamentalist forecast and some adopt the reflective forecast, then the circumstances are fertile for the emergence of bubble-like episodes of mysticism. Such episodes are particularly likely if agents are aggressively pursuing superior forecasting strategies.

The outline of our discussion is as follows. In Section 2 we lay out the strategies and payoffs available to agents facing a model with potential martingale solutions. Section 3 puts forth a general structure of evolutionary dynamics. Sections 4 and 5 present convergence results for two specific forms of evolutionary dynamics, imitation of successful agents and review of unsuccessful forecasts. Section 6 explores the robustness of these results with respect to experimentation with mysticism. Section 7 summarizes the implications of our results for the existence of periodic episodes of bubbles in processes that admit martingale solutions.

## 2. The Model

Lucas's (1978) study of asset pricing shows that optimizing behavior can lead to an equation of the form

$$y_t = \alpha(y^e_{t+1} + E(u_{t+1}|\Omega_t)), \quad \alpha < 1. \tag{2.1}$$

In his setting, $y^e_{t+1}$ is a representative agent's expectation of next period's value for a stock price $y_t$, and $E(u_{t+1}|\Omega_t)$ is the mathematical expectation of next period's dividend for that stock given an information set $\Omega_t$. We assume that agents know the parameter $\alpha$ and the mathematical expectations $E(u_{t+k}|\Omega_t)$, $k = 1, 2, \ldots$ . Agents will determine the realization $y_t$ on the basis of $\Omega_t$ so we do not include $y_t$ in $\Omega_t$.

If we let

$$y^*_t = \sum_{s=1}^{\infty} \alpha^s u_{t+s} \tag{2.2}$$

denote the perfect foresight fundamental solution that would arise if agents knew the entire future path of a stationary $u_t$, then the solution set for (2.1) includes

$$y_t = E(y_t^*|\Omega_t) + \alpha^{-t} m_t, \tag{2.3}$$

for any martingale $m_t = m_{t-1} + \eta_t$, where $\eta_t$ is a serially independent martingale innovation. Such solutions arise because of the self-fulfilling nature of expectations in (2.1).

We apply evolutionary game theory to study the outcome when agents do not agree on a single expectation $y^e_{t+1}$, but do attempt to adhere to the rational expectations philosophy underlying Lucas's paper. From the equations in Lucas (1978), we construct three forecasts.

*The Fundamentalist Forecast.* A fraction $\gamma_t$ of our agents read Lucas (1978, p.1439) and adopt the fundamentalist one period ahead forecast

$$e^{+1}_{\gamma,t} = E(y^*_{t+1}|\Omega_t), \tag{2.4}$$

which is the equation immediately following Lucas' equation (14). We use the term fundamentalist because this forecast corresponds to the fundamental (bubble-free) solution in the deeper theorizing literature.

*The Mystical Forecast.* Another fraction $\lambda_t$ of the agents opt to follow the mystical forecast

$$e^{+1}_{\lambda,t} = E(y^*_{t+1}|\Omega_t) + \alpha^{-t-1} m_t. \tag{2.5}$$

Sargent (1987, equation (3.8)) confirms to them that (2.5) is a valid solution to Lucas' equation (14).

The results in this paper do not depend on the specific origin of the martingale $m_t$. While the mystical forecast could be constructed from an intentionally extraneous martingale, it could also arise from some disagreement about what is fundamental. The followers of the mystical forecast might envision an uncertain, but forecastable lump-sum return $\alpha^{-T} M_T$ at a date $T$ in the future. Although the fundamentalists would insist that the mystical forecast is based on a pot of gold at the end of a rainbow, the true nature of $\alpha^{-T} M_T$ might not be so obvious. For example, $\alpha^{-T} M_T$ might be the ultimate return envisioned for a hedge fund founded by Nobel Prize winning economists or $\alpha^{-T} M_T$ might be the ultimate return for a company with a near monopoly on hardware needed for the vastly expanding Internet. Given belief in a return $\alpha^{-T}M_T$ at a fixed date $T$ that remains in the future, calculating the discounted expected value $\alpha^{-t}m_t = \alpha^{T-t}E(\alpha^{-T}M_T|\Omega_t)$ will generate the martingale term in (2.5).

*The Reflective Forecast.* A third fraction $\beta_t$ of the agents consider (2.1) to be the essential conclusion of Lucas (1978) and adopt the reflective forecast $e^{+1}_{\beta,t}$, which is defined by the property

$$y_t = \alpha(e^{+1}_{\beta,t} + E(u_{t+1}|\Omega_t)). \tag{2.6}$$

Followers of the reflective forecast understand that their choice of $e^{+1}_{\beta,t}$ affects $y_t$ and they use the algebra given below and their knowledge of $\gamma_t$, $e^{+1}_{\gamma,t}$, $\lambda_t$, and $e^{+1}_{\lambda,t}$ to find the value for $e^{+1}_{\beta,t}$ that achieves (2.6).

Given these three heterogeneous forecasting strategies, determining $y_t$ requires a basis for aggregating expectations. For agents optimizing over the mean and variance of returns, Brock and Hommes (1998) extend Lucas's model of representative agent optimizing behavior to heterogeneous expectations. We adopt their conclusion (their equation (2.7)), which can be written in our notation as[11]

$$y_t = \alpha(\gamma_t e^{+1}_{\gamma,t} + \lambda_t e^{+1}_{\lambda,t} + \beta_t e^{+1}_{\beta,t}) + \alpha E(u_{t+1}|\Omega_t)). \tag{2.7}$$

The quantity $y^e_{t+1} = \gamma_t e^{+1}_{\gamma,t} + \lambda_t e^{+1}_{\lambda,t} + \beta_t e^{+1}_{\beta,t}$ functions as the aggregate expectation. Combining (2.6) and (2.7) yields

$$e^{+1}_{\beta,t} = (1 - n_t)e^{+1}_{\gamma,t} + n_t e^{+1}_{\lambda,t}, \tag{2.8}$$

where

$$n_t = \frac{\lambda_t}{\lambda_t + \gamma_t}.$$

The reflective forecast is thus aptly named because it achieves (2.6) by adopting a weighted average of the other two forecasts.[12]

The reflective forecast differs in one important regard from the fundamentalist and mystical forecasts. Followers of the reflective forecast take into account the beliefs of others, regardless of whether or not they agree with the basis for those beliefs. Followers of the fundamentalist and

[11] Evans and Honkapohja (2001, p. 46 and p. 223) introduce a similar equation by assumption.

[12] If $n_t$ equals zero or one, then there is effectively no disagreement about forecasts because $e^{+1}_{\beta,t}$ equals either $e^{+1}_{\gamma,t}$ or $e^{+1}_{\lambda,t}$.

mystical forecasts, on the other hand, do not react to the other forecasts. Considerably before the advent of the rational expectations principle, Keynes' *beauty contest* view of investment recognized the growing importance of the reflective forecast: *We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. ... Investment based on genuine long-term expectation is so difficult today as to be scarcely practicable. He who attempts it must surely lead much more laborious days and run greater risks than he who tries to guess better than the crowd how the crowd will behave; and, given equal intelligence, he may make more disastrous mistakes.*[13] We will quantify this natural advantage accruing to the reflective forecast.

The first step is computing the forecast errors that agents consider in choosing strategies. At time $t-1$, the forecasts of $y_t$ are

$$e^{+1}_{\gamma,t-1} = E(y^*_t|\Omega_{t-1}), \tag{2.9a}$$

$$e^{+1}_{\lambda,t-1} = E(y^*_t|\Omega_{t-1}) + \alpha^{-t} m_{t-1}, \tag{2.9b}$$

$$e^{+1}_{\beta,t-1} = E(y^*_t|\Omega_{t-1}) + \alpha^{-t} n_t m_{t-1}. \tag{2.9c}$$

The reflective forecast includes a martingale term $\alpha^{-t} n_t m_{t-1}$, where the weight $n_t$ depends on the fraction of other agents following mysticism. The fractions $\gamma_t$, $\lambda_t$, $\beta_t$, and $n_t$ are set after $y_{t-1}$ is determined and held constant until after $y_t$ is determined. At time $t$, $E(y^*_t|\Omega_t)$ and $m_t$ become available, resulting in the three forecasts of $y_{t+1}$[14]

$$e^{+1}_{\gamma,t} = E(y^*_{t+1}|\Omega_t), \tag{2.10a}$$

$$e^{+1}_{\lambda,t} = E(y^*_{t+1}|\Omega_t) + \alpha^{-t-1} m_t, \tag{2.10b}$$

$$e^{+1}_{\beta,t} = E(y^*_{t+1}|\Omega_t) + \alpha^{-t-1} n_t m_t. \tag{2.10c}$$

These forecasts enter (2.7) to yield the realization

$$y_t = E(y^*_t|\Omega_t) + \alpha^{-t} n_t m_t.$$

The negatives of the three period $t$ squared forecast errors can be written as

$$\pi_{\beta,t} = -U^2_t, \tag{2.11a}$$

[13] Keynes (1935, p. 156-157)

[14] We assume that agents commit to using a specific forecasting strategy at time $t$ before $m_t$ and, therefore, $y_t$ is known. This ordering does not allow agents to revise their choices of strategies given the realizations of the forecasts. Brock and Hommes (1998) and LeBaron, Arthur and Palmer (1999) use the same assumption about $y_t$.

$$\pi_{\gamma,t} = -U_t^2 - 2n_t U_t A_t - n_t^2 A_t^2, \tag{2.11b}$$

$$\pi_{\lambda,t} = -U_t^2 + 2(1-n_t) U_t A_t - (1-n_t)^2 A_t^2, \tag{2.11c}$$

where $A_t = \alpha^{-t} m_{t-1}$ is the level of the martingale term and the reflective forecast error $U_t = F_t + G_t$ includes the innovation in the fundamentals for $y_t$

$$F_t = E(y_t^* | \Omega_t) - E(y_t^* | \Omega_{t-1})$$

and a term involving the innovation in the martingale multiplied by the weight $n_t$ measuring the importance of mysticism

$$G_t = \alpha^{-t} n_t (m_t - m_{t-1}).$$

Given these squared forecast errors at time $t$, agents choose $(\gamma_{t+1}, \lambda_{t+1}, \beta_{t+1})$, which will remain fixed until $y_{t+1}$ is determined. Note that the reflective forecast (2.10c) becomes

$$e_{\beta,t}^{+1} = E(y_{t+1}^* | \Omega_t) + \alpha^{-t-1} n_{t+1} m_t \tag{2.10c*}$$

because $n_{t+1}$ is now known just as $n_t$ was known for (2.9c).

Three overall properties of the model are important to the evolution of beliefs. First, the results are invariant to the actual process determining the fundamentals in the following sense. The payoffs (2.11) involve the fundamentals only through the innovation $F_t$, which is serially independent regardless of the time series properties of $E(y_t^* | \Omega_t)$. A given realization for the series $F_t$, $t = 1, 2, ...$ is consistent with a wide range of dynamic properties for $y_t^*$.

Second, we can now quantify the natural advantage that Keynes recognized for the reflective forecast. The payoff for the reflective forecast is given by

$$\pi_{\beta,t} = \bar{\pi}_t + \frac{\gamma_t \lambda_t}{\gamma_t + \lambda_t} A_t^2, \tag{2.12}$$

where $\bar{\pi}_t = \gamma_t \pi_{\gamma,t} + \lambda_t \pi_{\lambda,t} + \beta_t \pi_{\beta,t}$ is the population average payoff and $\pi_{\beta,t} \geq \bar{\pi}_t$. This relation does not involve $U_t A_t$ because the terms involving $U_t A_t$ in (2.11b) and (2.11c) cancel in the weighted average. The covariance terms $-2U_t A_t n_t$ and $2U_t A_t (1-n_t)$ do, on the other hand, determine the ranking of $\pi_{\beta,t}$ relative to $\pi_{\gamma,t}$ and the ranking of $\pi_{\beta,t}$ relative to $\pi_{\lambda,t}$. In particular, the mystical forecast can have the highest payoff if $2(1-n_t) U_t A_t > A_t^2$. This can happen if $U_t$ differs significantly from zero and $A_t$ happens to be of the same sign.

Third, the contest between fundamentalism and mysticism is symmetric, but biased in favor of the currently favored alternative. This follows from

$$\pi_{\lambda,t} - \pi_{\gamma,t} = 2U_t A_t + 2(n_t - \tfrac{1}{2})A_t^2. \tag{2.13}$$

If $n_t = \frac{1}{2}$, then (2.13) does not favor either fundamentalism or mysticism because $E(U_t A_t) = 0$. If $n_t > \frac{1}{2}$, then the term $2(n_t - \frac{1}{2})A_t^2$ favors a further increase in $\lambda_t$ and, hence, $n_t$. A further decrease in $n_t$ is, on the other hand, made more likely if $n_t < \frac{1}{2}$.

These general observations provide some intuition about the properties of the complete model. In Section 3 we present a general model of evolutionary dynamics for $(\gamma_t, \lambda_t, \beta_t)$, and in Sections 4 and 5 we give specific behavioral equations that complete the specification of agents' behavior.

## 3. Evolutionary Dynamics

Studying the evolutionary dynamics of the model requires a specification of how agents choose forecasts given the payoffs. We consider two alternative models of agents' behavior that we refer to as *imitation of successful agents* and *review of unsuccessful forecasts*. These principles feature agents choosing among possible forecasts, all of which attempt to implement rational expectations, using plausible behavioral rules without making the final step of adopting one of the deeper theorizing principles as common knowledge.

We use a discrete time version of evolutionary dynamics. Let $x_{i,t}$ be the fraction of agents using forecast $i$ at time $t$. Let $r_{i,t}$ be the fraction of agents using forecast $i$ who review their choice of strategy at time $t$, and let $p^i_{j,t}$ be the probability that an agent using forecast $j$ in period $t$ who reviews switches to forecast $i$ in the next period. If there are $n$ available forecasts, then the change in $x_{i,t}$ is given by

$$x_{i,t+1} - x_{i,t} = \sum_{j=1}^{n} r_{j,t} x_{j,t} p^i_{j,t} - r_{i,t} x_{i,t}. \tag{3.1}$$

This is the discrete time version of equation (4.24) in Weibull (1997). Hofbauer and Weibull (1996, pp. 564-6) consider two specific behavioral models contained within this general framework.

The process that we refer to as *imitation of successful agents* assumes that agents review at a rate invariant to the payoffs (we use $r_{i,t} \equiv 1$, assuming all agents review every period), but that the transition probabilities $p^i_{j,t}$ are functions of forecast performance in that agents tend to switch to strategies with lower forecast errors. Agents use payoff weighting functions $w\left(\pi_{i,t}\right)$ to arrive at the transition probabilities

$$p^i_{j,t} = \frac{w\left(\pi_{i,t}\right) x_{i,t}}{\overline{w}_t}, \tag{3.2}$$

where $\overline{w}_t = \sum\limits_{h=1}^{n} w\left(\pi_{h,t}\right) x_{h,t}$. The transition probability $p^i_{j,t}$ into strategy $i$ depends on its current weighted payoff $w\left(\pi_{i,t}\right)$ relative to the population average $\overline{w}_t$ and its current popularity $x_{i,t}$. The latter factor can be thought of as a measure of past forecast performance. Substituting (3.2) into (3.1) yields

$$x_{i,t+1} - x_{i,t} = x_{i,t}\ \frac{w\left(\pi_{i,t}\right) - \overline{w}_t}{\overline{w}_t}, \tag{3.3}$$

which is of the form (1.1) with $f(u(s_i, x_t)) = w\left(\pi_{i,t}\right)$, $\delta(x_t) = 1/\overline{w}_t$, and $\mu(x_t) = -1$.

The process we refer to as *review of unsuccessful forecasts* assumes that the review rates $r_{i,t}$ depend on forecast performance, with agents more likely to review forecasts with large squared errors. The review rates can then be given by a function

$$r_{i,t} = r(\pi_{i,t}), \tag{3.4}$$

where $r(\pi_{i,t})$ is decreasing in its argument. To focus on the effects of conditional review rates, we assume reviewing agents choose a new strategy using $p^i_{j,t} = x_{i,t}$. This is equivalent to adopting the strategy of a randomly selected member of the population, which could be motivated by a belief that popular strategies are the more successful ones. Substituting (3.4) into (3.1) yields

$$x_{i,t+1} - x_{i,t} = x_{i,t}\ (\overline{r}_t - r(\pi_{i,t})), \tag{3.5}$$

where $\overline{r}_t = \sum_{j=1}^{n} r(\pi_{j,t}) x_{j,t}$. This is of the form (1.1) with $f(u(s_i, x_t)) = -r\left(\pi_{i,t}\right)$, $\delta(x_t) = 1$, and $\mu(x_t) = \overline{r}_t$. Convexity of $f(\cdot)$ in (1.1) in this case takes the form of concavity in $r(\cdot)$ because $r(\pi_{i,t})$ enters (3.5) with a negative sign.

The evolutionary dynamics given by (3.3) or (3.5) along with the payoffs (2.11) determine the movement of $(\gamma_t, \lambda_t, \beta_t)$ within a simplex

$$\Delta = \{(\gamma_t, \lambda_t, \beta_t) \mid \gamma_t, \lambda_t, \beta_t \geq 0,\ \lambda_t + \gamma_t + \beta_t = 1\}.$$

The two forces described by (2.12) and (2.13) will generally push $(\gamma_t, \lambda_t, \beta_t)$ toward one of two edges, $\{(\gamma_t, \lambda_t, \beta_t) \mid \lambda_t = 0,\ \gamma_t + \beta_t = 1\}$ and $\{(\gamma_t, \lambda_t, \beta_t) \mid \gamma_t = 0,\ \lambda_t + \beta_t = 1\}$. Along either edge there is no tendency for further change because agents agree on a single forecast. If $\lambda_t = 0$, the reflective forecast equals the fundamentalist forecast so that fundamentalists and reflectivists follow different procedures, but in fact agree on the same forecast. A similar agreement on the mystical forecast occurs if $\gamma_t = 0$.

The fraction of the agents following the fundamentalist forecast is not, however, likely to ever reach zero. While fundamentalism is, in a sense, just one of the martingale solutions, it is the only specific solution that is written up in numerous scholarly publications. We assume, therefore, that a percentage $\gamma_{min} > 0$ of the agents have read and firmly agree with the deeper theorizing literature.[15] This restricts $(\gamma_t, \lambda_t, \beta_t)$ to the simplex

$$\underline{\Delta} = \{(\gamma_t, \lambda_t, \beta_t) \mid \gamma_t \geq \gamma_{min},\ \lambda_t, \beta_t \geq 0,\ \lambda_t + \gamma_t + \beta_t = 1\}. \tag{3.6}$$

Restricting $(\lambda_t, \gamma_t, \beta_t)$ to the simplex $\underline{\Delta}$ eliminates the paradox that, if $\beta_t$ were to equal 1 because all agents adopt reflectivism, then the reflective forecast, which is a weighted average of other forecasts, would not exist.

The actual mechanism restricting $(\gamma_{t+1}, \lambda_{t+1}, \beta_{t+1})$ to the simplex $\underline{\Delta}$ takes the following form. Let $D[(\gamma_t, \lambda_t, \beta_t)] = (\gamma'_{t+1}, \lambda'_{t+1}, \beta'_{t+1})$ denote the dynamics described by (3.3) or (3.5), and let $H[(\gamma'_{t+1}, \lambda'_{t+1}, \beta'_{t+1})] = (\gamma_{t+1}, \lambda_{t+1}, \beta_{t+1})$ denote a projection onto $\underline{\Delta}$. That is, $H[(\gamma, \lambda, \beta)] = (\gamma, \lambda, \beta)$ for $(\gamma, \lambda, \beta) \in \underline{\Delta}$, and $H[(\gamma, \lambda, \beta)] \in \underline{\Delta}$ for any $(\gamma, \lambda, \beta)$ with $\gamma + \lambda + \beta = 1$. To avoid $\gamma_{t+1} < \gamma_{min}$ or, more generally, to avoid $(\gamma_{t+1}, \lambda_{t+1}, \beta_{t+1}) \notin \underline{\Delta}$ we assume that the evolutionary dynamics are given by the composition $H[D[(\gamma_t, \lambda_t, \beta_t)]] = (\gamma_{t+1}, \lambda_{t+1}, \beta_{t+1})$. In Section 6, we consider a specific $H$, but our convergence results in Sections 4 and 5 require only a general

[15] It is also possible to implement this constraint by creating a fourth forecast strategy (unshakeable fundamentalism) that uses the fundamental solution with a review rate of zero. This would complicate the notation, but lead to very similar conclusions.

assumption about $H$. We assume that $H$ does not change the signs of the changes given by the mapping $D$. That is, $H$ must satisfy

$$
\begin{aligned}
sgn(\gamma_{t+1} - \gamma_t) &= sgn(\gamma'_{t+1} - \gamma_t), \\
sgn(\lambda_{t+1} - \lambda_t) &= sgn(\lambda'_{t+1} - \lambda_t), \\
sgn(\beta_{t+1} - \beta_t) &= sgn(\beta'_{t+1} - \beta_t),
\end{aligned}
$$

where $sgn(x) = 1$ for $x > 0$, $sgn(0) = 0$, and $sgn(x) = -1$ for $x < 0$. This property is easily attained for any projection $H$ by adding the condition that $(\gamma_{t+1}, \lambda_{t+1}, \beta_{t+1}) = (\gamma_t, \lambda_t, \beta_t)$ if the projection would have violated the assumption.

Our convergence results in Sections 4 and 5 will focus on conditions under which $\beta_t$ is monotone nondecreasing, forcing $\lambda_t$ to zero. Once $\lambda_t$ reaches zero, the reflective and fundamentalist forecasts are identical. In Section 6, we consider cases where this convergence is not monotone, leaving open the possibility that mysticism might survive long enough and gain enough popularity to make the effects of agents experimenting with that strategy an important consideration.

## 4. Imitation of Successful Agents

While some intuitive observations on the imitation of successful agents (ISA) can be gained using general equations and Taylor series approximations, concrete results on convergence require specific functional forms. We present two, a linear function and an exponential function.

### 4.1 The General Model

We will describe the evolution of the populations shares $(\gamma_t, \lambda_t, \beta_t)$ in terms of $\beta_t$ and $n_t = \lambda_t/(\gamma_t + \lambda_t)$. It will prove convenient to write the equation of motion for $\beta_t$ (3.3) with $\beta_t/\beta_{t+1}$ on the left, keeping in mind that a ratio less than one implies that $\beta_t$ is increasing. This yields

$$\frac{\beta_t}{\beta_{t+1}} = 1 + (1-\beta_t)\left(\frac{n_t w(\pi_{\lambda,t}) + (1-n_t)w(\pi_{\gamma,t})}{w(\pi_{\beta,t})} - 1\right). \tag{4.1}$$

The following for the reflective forecast $\beta_t$ increases if $w(\pi_{\beta,t})$ is greater than the weighted average $n_t w(\pi_{\lambda,t}) + (1-n_t)w(\pi_{\gamma,t})$. For a differentiable payoff weighting function, a second-order Taylor series approximation to (4.1) is

$$\frac{\beta_t}{\beta_{t+1}} \cong 1 - \frac{\gamma_t \lambda_t}{\gamma_t + \lambda_t}\left[\frac{w'(\pi_{\beta,t})}{w(\pi_{\beta,t})}A_t^2 + 2\frac{w''(\pi_{\beta,t})}{w(\pi_{\beta,t})}(U_t A_t)^2\right],$$

where we omit powers of $A_t$ higher than two on the grounds that $A_t$ is initially small. The term involving $A_t^2$ unambiguously works in favor of an increase in $\beta_t$, but the term involving $(U_t A_t)^2$ can cause a decrease if $U_t A_t$ is sufficiently large in absolute value. Large absolute values for $U_t A_t$ are, from (2.11b) and (2.11c), associated with success or failure for the fundamentalist and mystical forecasts.

We can reorganize this as

$$\frac{\beta_t}{\beta_{t+1}} \cong 1 - \frac{\gamma_t \lambda_t}{\gamma_t + \lambda_t}\frac{w'(\pi_{\beta,t})}{w(\pi_{\beta,t})}A_t^2\left[1 - 2\frac{w''(\pi_{\beta,t})}{w'(\pi_{\beta,t})}U_t^2\right].$$

The factor $w'(\pi_{\beta,t})/w(\pi_{\beta,t})$ measures the extent to which agents differentiate between forecasts. Convexity, measured by the factor $w''(\pi_{\beta,t})/w'(\pi_{\beta,t})$, determines whether large values of $U_t^2$ can have a negative impact on $\beta_t$. If $U_t^2$ is small or the convexity is small, there is a natural tendency for $\beta_t$ to increase toward one.[16]

The equation of motion for $n_t$ is

$$\frac{n_t}{n_{t+1}} = 1 + (1 - n_t)\left(\frac{w(\pi_{\gamma,t})}{w(\pi_{\lambda,t})} - 1\right) \tag{4.2}$$

if $n_{t+1} > 0$ and $w(\pi_{\lambda,t}) > 0$. (If $w(\pi_{\lambda,t}) = 0$, then $n_{t+1} = 0$.) If the payoff weight $w(\pi_{\gamma,t})$ for fundamentalism is less than the payoff weight $w(\pi_{\lambda,t})$ for mysticism, then $n_t = \lambda_t/(\gamma_t + \lambda_t)$ increases. The Taylor series approximation to (4.2), again using powers of $A_t$ of order 2 and lower, is

$$\frac{n_t}{n_{t+1}} \cong 1 - 2(1 - n_t)\frac{w'(\pi_{\beta,t})}{w(\pi_{\beta,t})}\left[U_t A_t + (n_t - \tfrac{1}{2})A_t^2\right].$$

If $n_t$ is near zero or one, then the factor $n_t - \frac{1}{2}$ acts to put the weight of the squared martingale $A_t^2$ toward reinforcing that value of $n_t$. The factor $w'/w$ again measures the extent to which agents differentiate between forecasts. This result is a more specific instance of the general property (2.13).

### 4.2 Linear Payoff Weighting Functions

The linear payoff weighting function $w(\pi) = c + \pi$ leads to the *replicator dynamic*. The growth rate for strategy $i$ is proportional to that strategy's fitness, which is the difference between its payoff and the population average payoff. The updating equation (3.3) takes the form[17]

$$x_{i,t+1} - x_{i,t} = x_{i,t}\,\frac{\pi_{i,t} - \bar{\pi}_t}{c + \bar{\pi}_t}.$$

The parameter $c$ is given various interpretations in the evolutionary game theory literature.[18] Larger values for $c$ decrease the change in $x_{i,t}$ caused by a given difference $\pi_{i,t} - \bar{\pi}_t$. In the present

[16] Given that $\pi_{\beta,t} = -U_t^2$, this statement formally requires that $x \cdot w''(-x)/w'(-x) \to \infty$ as $x \to \infty$.

[17] The general linear payoff weighting function would be $w(\pi) = a + b\pi$, but the results are invariant to a scalar multiple. We will thus work with a linear weighting function of the form $w(\pi) = c + \pi$.

[18] Samuelson (1998, pp. 63-67).

case, the payoffs are negative and we impose the constraint

$$\omega(\pi) = \left\{ \begin{array}{cc} c+\pi & \text{for } \pi > -c \\ 0 & \text{for } \pi \leq -c \end{array} \right\} \tag{4.3}$$

to guarantee nonnegative payoff weights.

**Proposition 1**. The dynamics given by (3.3) and (4.3) will lead to $\beta_{t+1} \geq \beta_t$ for $(\gamma_t, \lambda_t, \beta_t) \in \underline{\Delta}$ if $w(\pi_{\gamma,t}) > 0$ and $w(\pi_{\lambda,t}) > 0$. The latter bounds are satisfied if $(U_t + A_t)^2 < c$ and $(U_t - A_t)^2 < c$.

Proof: The first statement follows directly from (2.12) and (3.3) because all the payoff weights are positive. The second statement can be seen from (2.11). □

The conditions on the payoffs in the above proposition suggests that a large martingale term $A_t^2$ could cause nonrobustness. Our simulations in Section 6 show that the variance of the martingale is not a major determinant of the properties of the outcome, and we are able to establish the following result that does not involve the size of the martingale.

**Proposition 2**. If there exists a constant $K$ such that $U_t^2 < K$ for all $t \geq 0$ and if $(\gamma_0, \lambda_0, \beta_0) \in \underline{\Delta}$ with $n_0 \leq 1 - \dfrac{K}{c}$, then the dynamics given by (3.3) and (4.3) starting at $(\gamma_0, \lambda_0, \beta_0)$ will converge to a point where $\lambda = 0$ with $\beta_{t+1} \geq \beta_t$ until $\lambda_{t+1} = 0$.

Proof: Appendix A. □

Proposition 2 shows that bounding the uncertainty $U_t^2$ driving the reflective forecast error guarantees a region of monotone convergence given by $n_0 \leq 1 - K/c$.[19] While the value $K = \gamma_{min} \cdot c$ necessary to guarantee monotone convergence over the entire simplex $\underline{\Delta}$ is small, the conditions in the proposition are sufficient, but not necessary for convergence, which simulations show to occur from anywhere in $\underline{\Delta}$. In Section 6, we show that convergence can be nonmonotone if $U_t^2$ is large, making it possible for $\lambda_t$ to be significantly nonzero for a period of time.

[19] These bounds are reminiscent of the projection facility in least squares learning of Marcet and Sargent (1989b), which involves restricitions on agents' beliefs concerning model parameters. Grandmont (1998) discusses the interpretation of the projection facility in detail.

### 4.3 Exponential Payoff Weighting Functions

A convex payoff weighting function incorporates the notion that agents emphasize the importance of the best payoffs. The degree of convexity then determines how aggressively agents pursue this goal. The exponential function

$$w(\pi) = e^{\sigma\pi}$$

with $\sigma > 0$ provides a straightforward example.[20] The relative slope $w'(\pi)/w(\pi) = \sigma$ and the convexity $w''(\pi)/w'(\pi) = \sigma$ are constant over $\pi$. Discarding powers of $A_t$ greater than two (on the grounds that the martingale is small in its early stages) yields:

$$\frac{\beta_t}{\beta_{t+1}} \cong 1 - n_t(1-n_t)A_t^2\sigma\left[1 - 2\sigma U_t^2\right]. \tag{4.4}$$

The magnitude of $U_t^2$ relative to the convexity $\sigma$ determines the direction of change in $\beta_t$. We can extend this intuition with the following result that does not involve any Taylor series approximation.

**Proposition 3**: The dynamics given by (3.3) with $w(\pi) = e^{\sigma\pi}$ will lead to $\beta_{t+1} > \beta_t$ for $(\lambda_t, \gamma_t, \beta_t) \in \underline{\Delta}$ if

$$max\{\sigma(U_t + \tfrac{1}{2}A_t)^2, \sigma(U_t - \tfrac{1}{2}A_t)^2\} < a^*, \tag{4.5}$$

where $a^* = 0.351733711...$ satisfies $a^* = \frac{1}{2}exp(-a^*)$.

Proof: Appendix B. □

This condition is sufficient, but not necessary for $\beta_{t+1} > \beta_t$. In Section 6, we use simulations to explore the properties of the process when (4.5) is not satisfied.

### 5. Review of Unsuccessful Forecasts

The review of unsuccessful forecasts (RUF) provides an alternative to imitation of successful agents (ISA). The broad similarities in the results for the two principles provide evidence that our overall results are not unique to specific behavioral assumptions. We consider linear and exponential review rate functions.

[20] Bjornerstedt and Weibull (1993) and Weibull (1997) use this functional form.

### 5.1 The General Model

The equation of motion (3.5) for $\beta_t$ is

$$\frac{\beta_{t+1}}{\beta_t} = 1 - (r(\pi_{\beta,t}) - \overline{r}_t), \tag{5.1}$$

where $\overline{r}_t = \gamma_t r(\pi_{\gamma,t}) + \lambda_t r(\pi_{\lambda,t}) + \beta_t r(\pi_{\beta,t})$. The reflective forecast gains weight if its review rate $r(\pi_{\beta,t})$ is less than the weighted average of the review rates for the other two forecasts. The basic properties of $\beta_t$ can again be seen from a Taylor series approximation

$$\frac{\beta_{t+1}}{\beta_t} \cong 1 - \frac{\gamma_t \lambda_t}{\gamma_t + \lambda_t} \left[ r'\left(\pi_{\beta,t}\right) A_t^2 + 2r''\left(\pi_{\beta,t}\right)\left(U_t A_t\right)^2 \right].$$

The term involving $A_t^2$ favors an increase in $\beta_t$ because $r'\left(\pi_\beta\right) < 0$. The term involving $U_t A_t$ only appears for a nonlinear review rate function (i.e. $r''\left(\pi_{\beta,t}\right) \neq 0$). We can reorganize this as

$$\frac{\beta_{t+1}}{\beta_t} \cong 1 - \frac{\gamma_t \lambda_t}{\gamma_t + \lambda_t} r'\left(\pi_{\beta,t}\right) A_t^2 \left[ 1 - 2\frac{r''\left(\pi_{\beta,t}\right)}{r'\left(\pi_{\beta,t}\right)} U_t^2 \right]. \tag{5.2}$$

The factor in square brackets can be negative if the product of $U_t^2$ and the curvature $r''\left(\pi_{\beta,t}\right)/r'\left(\pi_{\beta,t}\right)$ is large.[21] In that case, $\beta_t$ falls, leaving open the possibility that mysticism could gain popularity.

### 5.2 Linear Review Rate Functions

The linear case is again the replicator dynamic. The piecewise linear, but globally concave function

$$r\left(\pi\right) = \left\{ \begin{array}{cc} -b\pi & \text{for } -b\pi < 1 \\ 1 & \text{for } -b\pi \geq 1 \end{array} \right\} \tag{5.3}$$

where $b > 0$, imposes the constraint that the review rates must be no greater than unity.[22] An analysis very similar to Propositions 1 yields:

[21] We assume, using (2.11a), that $x \cdot r''(-x)/r'(-x) \to \infty$ as $x \to \infty$.

[22] The versions of the replicator dynamic following from (3.3) and (3.5) differ because (3.3) includes the denominator $c + \overline{\pi}_t$. Both are valid discrete time approximations to the continuous time replicator dynamic, and adding a denominator to (3.5) would unnecessarily complicate the notation for this section.

**Proposition 4**. The dynamics given by (3.5) and (5.3) lead to $\beta_{t+1} \geq \beta_t$ for $(\gamma_t, \lambda_t, \beta_t) \in \underline{\Delta}$ if $r(\pi_{\gamma,t}) < 1$ and $r(\pi_{\lambda,t}) < 1$. The latter bounds are satisfied if $b(U_t + A_t)^2 < 1$ and $b(U_t - A_t)^2 < 1$.

Proof: All the review rates are less than unity, and (2.12) applies directly. □

Again, although the term $A_t$ appears in these bounds, the magnitude of the martingale is not a primary determinant of whether $\beta_t$ converges monotonically. Using methods similar to Proposition 2, it is possible to show that, if there exists a constant $K$ such that the errors are bounded by $U_t^2 < K$ and if $n_t$ satisfies $bK < n_t < 1 - bK$, then $\beta_t$ is monotone nondecreasing for all values of $A_t$.

### 5.3 Exponential Review Rate Functions

A concave review rate function places extra emphasis on reviewing less than successful forecasting strategies. The degree of the concavity determines how aggressively agents pursue this goal. The concave review rate function

$$r(\pi) = 1 - e^{\sigma\pi}$$

naturally incorporates the nonnegativity constraint and has the property that $r''(\pi)/r'(\pi) = \sigma$ is constant. The approximation (5.2) becomes

$$\frac{\beta_{t+1}}{\beta_t} \cong 1 + \frac{\gamma\lambda}{\gamma + \lambda} A_t^2 \sigma exp\left(-\sigma U_t^2\right)\left[1 - 2\sigma U_t^2\right].$$

This shows that $\beta_t$ will tend to increase unless the product $\sigma U_t^2$ is large. In particular, we have

**Proposition 5**: The dynamics given by (3.5) and $r(\pi) = 1 - e^{\sigma\pi}$ will lead to $\beta_{t+1} > \beta_t$ for $(\lambda_t, \gamma_t, \beta_t) \in \underline{\Delta}$ if

$$max\{\sigma(U_t + \tfrac{1}{2}A_t)^2, \sigma(U_t - \tfrac{1}{2}A_t)^2\} < a^*, \tag{5.4}$$

where $a^* = 0.351733711...$ satisfies $a^* = \frac{1}{2}exp(-a^*)$.

Proof: Appendix B. □

## 6. Convergence and Robustness

Propositions 1-4 establish conditions under which $\beta_t$ is unambiguously monotone nondecreasing, leading the process toward the edge of $\underline{\Delta}$ given by $\gamma_t + \beta_t = 1$. Along that edge mysticism is eliminated and all agents agree numerically with the fundamentalist forecast. The conditions in the propositions are, of course, sufficient, but not necessary to have convergence to that edge of the simplex, and simulations from various starting points show that convergence to $\gamma_t + \beta_t = 1$ occurs from everywhere in $\underline{\Delta}$. The question of interest is thus not convergence, but the speed of convergence. If $\lambda_t$ remains significantly positive long enough to generate an episode where the martingale is important in determining $y_t$, then the behavior of $y_t$ might well be characterized as a bubble. Occasional experimentation with mysticism could then lead to periodic episodes of bubbles.

To quantify this possibility, we consider the path of the population shares when a small fraction $\lambda^*$ of the agents experiment with mysticism. We use $\lambda^* = \gamma_{min} = 0.05$ so that the simulations start from the point $(\gamma_0, \lambda_0, \beta_0) = (0.05, 0.05, 0.90)$. We start with $\beta_0$ large and $\gamma_0$ small on the grounds that previous experience would have tended to have driven out earlier episodes of mysticism as $\beta_t$ increased. For low curvature and small errors, the system will very seldom leave a small neighborhood of this starting point.

Tables 1 and 2 show simulation estimates of the survival probabilities for mysticism under imitation of successful agents (3.3) and review of unsuccessful forecasts (3.5) using the payoffs (2.11). Table 1 shows a small martingale innovation variance ($var(\eta_t) = 0.01$), and Table 2 shows a larger martingale innovation variance ($var(\eta_t) = 0.25$). In the linear cases (4.3) and (5.3), $c = 1$ and $b = 1$. In the exponential cases in Sections 4.3 and 5.3, $\sigma = 1$. In all cases, the discount factor in (2.1) is $\alpha = 0.95$. We ensure that the population fractions remain within the simplex $\underline{\Delta}$ by projecting points outside that simplex back to its boundary using the projection described in Appendix C. Each cell is based on 10,000 trials.

To quantify how long mysticism persists at a significant level, we define robustness in terms of the probability $Pr(\lambda_N > \lambda^{**})$ that mysticism has a following of at least $\lambda^{**}$ in period $N$. We use three values, 0.05, 0.10, and 0.20, for $\lambda^{**}$. We use six values, 0.01, 0.05, 0.1, 0.25, 1, and 5, for $var(F_t)$, where $F_t = E(y_t^*|\Omega_t) - E(y_t^*|\Omega_{t-1})$ is the innovation in fundamentals as the date of the information set changes, and we calculate the survival probabilities for $N = 10$ and $N = 20$.

As expected, the fate of mysticism depends on $var(F_t)$. For small values of $var(F_t)$, mysticism does not survive for 10 or 20 periods with any appreciable probability. For larger values of $var(F_t)$, however, the probability that mysticism survives and even has a following greater than 0.10 or 0.20 is large enough to be of practical importance. In some cases, the probability that $\lambda_t$ is greater than 0.20 is nearly as large as the probability it is greater than 0.05. To put these figures into perspective, if a process that survives 10 periods has a 0.03 probability of starting in a given period, then we would expect that process to be in existence about 1/3 of the time. Survival probabilities of even a few percent for 10 or 20 periods are thus potentially of practical importance.

While nonrobustness occurs for both ISA and RUF given a sufficiently large $var(F_t)$, RUF is noticeably more robust with respect to experimentation with mysticism than is ISA for a given $var(F_t)$. This difference is due to the RUF assumption about how reviewing agents choose a new strategy. Even when the outcome $U_tA_t$ in the payoffs (2.11) favors an increase in $\lambda_t$ for RUF, reviewing agents randomly select among strategies according to their current popularities. This does not tend to generate a rapid increase in a small $\lambda_t$. For ISA, on the other hand, agents are directly attracted to currently successful forecasts and a rapid increase in $\lambda_t$ can more easily occur.

Table 2 is included to show that the effects of experimentation with mysticism do not greatly depend on the variance of the martingale.[23] In fact, many of the survival probabilities in Table 2 are smaller than the corresponding figures in Table 1. The explanation for this can be seen in the payoffs (2.11). Conditions favor an increase in $\lambda_t$ when $U_tA_t$ is large and positive, but a large $A_t^2$ works against increases in either $\lambda_t$ or $\gamma_t$.

Figures 1 and 2 illustrate the potential implications of mysticism for the ISA-Exponential model in Section 4.3.[24] Each figure shows a typical realization for a simulation over 1000 periods for $var(\eta_t) = 0.01$. We assume that a fraction $\lambda_t = 0.05$ of the agents experiment with mysticism in any period where $\lambda_t$ would otherwise be below 0.05. If $\lambda_t > 0.05$, we simply let the process run. The top graph in each figure shows the fundamentalist forecast error $U_t + n_tA_t$.[25] The vertical scale on the bottom graph shows $\gamma_t$ from the bottom and $\lambda_t$ from the top so that $\beta_t$ is the distance between the two lines.

[23] An extreme exception should be noted. It is possible for $var(\eta_t)$ to be large enough to cause instability with $var(F_t) = 0$.

[24] The figures for the other cases would be similar.

[25] As noted in Section 2, the results are invariant to the dynamic properties of the realized $y_t$. The fundamentalist forecast error is serially correlated only to the extent that the martingale component is sometimes significant. The realized $y_t$ is the sum of the fundamentalist forecast error and whatever process the fundamentalist forecast (2.4) follows.

Figure 1 with $var(F_t) = 0.25$ shows an outcome fairly robust with respect to the introduction of mysticism. Similar figures for smaller values of $Var(F_t)$ would show smaller and less frequent deviations from minimal values for $\gamma_t$ and $\lambda_t$. Figure 2 with $var(F_t) = 0.50$ shows the character of the nonrobustness that can result from a somewhat larger $var(F_t)$. The upper graph shows the occurrence of bubbles in the forecast error for the fundamentalist forecast. These bubbles occur during episodes when mysticism gains a significant following as shown in the lower graph. The magnitude, but not the frequency of the bubbles depends on the size of the discount factor $\alpha$, and values of $\alpha$ closer to unity would produce less dramatic bubbles.

Figures 1 and 2 illustrate another feature of the contest among mysticism, fundamentalism, and reflectivism. The shifting weight on the martingale induces ARCH effects in the outcome.[26] During periods when mysticism has an appreciable following, the extra variance in the fundamentalist forecast error $U_t + n_t A_t$ due to the martingale causes a variance increase that tends to persist as long as mysticism survives.

## 7. Summary and Conclusions

In models with future expectations, economists have been inclined to rule out all martingale solutions save one fundamental solution on the grounds that agents resort to deeper theorizing. We use evolutionary game theory to study how agents with heterogeneous expectations might attempt to choose among the martingale solutions. We assume that agents respond to squared forecast errors through either a payoff weighting function (imitation of successful agents) or a review rate function (review of unsuccessful forecasts).

Our main result is that the prospects for convergence to a single fundamental solution depend primarily on the curvatures of the payoff weighting or review rate functions relative to the variances of the underlying error processes. We show that small error variances lead to stability with little likelihood that incipient mysticism will survive for more than a few periods. This establishes that nonrobustness is not an inevitable outcome of our model's construction. In particular, the mere existence of agents flirting with mysticism is not sufficient to cause nonrobustness.

The curvature that can lead to nonrobustness is not an arbitrary construct. The payoff weighting functions must be convex because the payoff weights cannot be negative, and the review rate

[26] LeBaron, Arthur, and Palmer (1999) find a similar effect in simulations in a much more complex environment.

function must be concave because review rates cannot exceed unity. If the errors are large enough to make this curvature an important factor, then we observe nonrobustness characterized by mysticism that survives for prolonged periods of time.

Our results thus provide support for assuming convergence to the fundamental solution if the error variances are small relative to the curvature in agents' behavioral functions. If that is not the case, then our results provide a plausible mechanism for speculative bubbles. They require not that all agents believe in mysticism, but only that some fraction of agents follow mysticism when it has the smallest squared forecast errors. The latter event can happen even if only a small fraction of the agents experiment with mysticism. It is also not necessary that any individual bubble last forever for bubbles to be a permanent feature of the model. Most importantly, it is not necessary for the agents in this model to ever abandon their belief in rational expectations. The agents simply carry out a systematic empirical search among martingale solutions using squared forecast errors to compare alternative forecasting strategies.

One might argue that the two evolutionary game theory principles we examine, imitation of successful agents and review of unsuccessful forecasts, are too simple to guarantee the convergence to the fundamental solution that economists envision. Another view, however, is that the agents are trying to learn the *right* martingale solution when there simply is no underlying right answer. The more aggressively they pursue the optimal forecasting strategy, as measured by the curvature of their payoff weighting or review rate functions, the less robust the outcome is to the introduction of mysticism as a possible strategy. Indeed, these results suggest that sophisticated, aggressive searching for the optimal forecasting strategy will provide fertile ground for the growth of mysticism and bubbles. Recent events in the financial markets do not contradict these conclusions.

## Appendix A

## Proof of Propositions 2

**Proof of Proposition 2.** For $t \geq s$, the bound on $U_t^2$ guarantees that $w(\pi_{\beta,t}) > 0$. If $w(\pi_{\gamma,t}) > 0$ and $w(\pi_{\lambda,t}) > 0$, Proposition 1 implies that $\beta_{t+1} > \beta_t$. The same inequality holds if $w(\pi_{\gamma,t}) = 0$ and $w(\pi_{\lambda,t}) = 0$. If $w(\pi_{\gamma,t}) > 0$ and $w(\pi_{\lambda,t}) = 0$, then $\lambda_{t+1} = 0$ and we are done. Suppose then that $w(\pi_{\gamma,t}) = 0$ and $w(\pi_{\lambda,t}) > 0$. We have $\beta_{t+1} > \beta_t$ if

$$w(\pi_{\beta,t}) > n_t w(\pi_{\lambda,t}) + (1 - n_t) w(\pi_{\gamma,t}), \tag{A.1}$$

which can be rewritten as

$$c - U_t^2 > n_t(c - U_t^2 + 2\,(1 - n_t)\,U_t A_t - (1 - n_t)^2\,A_t^2).$$

Factoring out $1 - n_t$ yields

$$U_t^2 + 2 n_t U_t A_t - n_t\,(1 - n_t)\,A_t^2 < c.$$

The maximum value of the left hand side of this inequality is achieved at $A_t = U_t/(1 - n_t)$, and making this substitution shows that (A.1) must hold if $U_t^2/(1 - n_t) < c$ or $n_t \leq 1 - \dfrac{U_t^2}{c}$. The latter inequality follows from $n_t \leq 1 - \dfrac{K}{c} \leq 1 - \dfrac{U_t^2}{c}$. We thus have $\beta_t$ monotone nondecreasing, and the process stops when $\lambda_t = 0$. □

## Appendix B

## Proof of Propositions 3 and 5

**Lemma 1**: For $y < x \leq 0$ and $0 \leq \theta \leq 1$,

$$e^x + \theta(e^y - e^x) - e^{x+\theta(y-x)} \leq \tfrac{1}{2} e^x \theta(1 - \theta)(y - x)^2.$$

Proof: The function

$$f(z) = e^z - 1 - \tfrac{1}{2} z^2$$

has the properties that that $f(0) = 0$ and $f''(z) = e^z - 1 \leq 0$ for $z \leq 0$. The function $f$ is thus concave, and

$$\theta f(z) - f(\theta z) \leq 0$$

for $z \le 0$. This can be rearranged as

$$1 + \theta(e^z - 1) - e^{\theta z} \le \tfrac{1}{2}\theta(1-\theta)z^2.$$

Letting $z = y - x$ and multiplying by $e^x$ establishes the result. □

**Proof of Propositions 3 and 5:** In Propositions 3 and 5 $\beta_{t+1} \ge \beta_t$ if $n_t w(\pi_{\lambda,t}) + (1 - n_t)w(\pi_{\gamma,t}) \le w(\pi_{\beta,t})$ or $n_t r(\pi_{\lambda,t}) + (1-n_t)r(\pi_{\gamma,t}) \ge r(\pi_{\beta,t})$, respectively. Both conditions are equivalent to

$$n_t exp(\sigma\pi_{\lambda,t}) + (1-n_t)exp(\sigma\pi_{\gamma,t}) \le exp(\sigma\pi_{\beta,t}). \tag{B.1}$$

Let $b_1$ and $b_4$ denote the left-hand and right-hand sides of this inequality, which can be written as $b_1 \le b_4$. We will show that $b_1 \le b_2 \le b_3 \le b_4$, where

$$b_2 = exp(\sigma\pi_t^*) + \tfrac{1}{2}n_t(1-n_t)\sigma^2(\pi_{\lambda,t} - \pi_{\gamma,t})^2,$$

$$b_3 = exp(\sigma\pi_t^*) + \sigma\ exp(\sigma\pi_t^*)(\pi_{\beta,t} - \pi_t^*), \text{ and}$$

$$\pi_t^* = n_t\pi_{\lambda,t} + (1-n_t)\pi_{\gamma,t}.$$

Lemma 1 implies, using $\theta = n_t$, $x = \sigma\pi_{\gamma,t}$, and $y = \sigma\pi_{\lambda,t}$, that

$$n_t exp(\sigma\pi_{\lambda,t}) + (1-n_t)exp(\sigma\pi_{\gamma,t}) \le exp(\sigma\pi_t^*) + \tfrac{1}{2}exp(\sigma\pi_t^+)n_t(1-n_t)\sigma^2(\pi_{\lambda,t} - \pi_{\gamma,t})^2,$$

where $\pi_t^+$ is the greater of $\pi_{\gamma,t}$ and $\pi_{\lambda,t}$. Because $exp(\sigma\pi_t^+) \le 1$, this implies $b_1 \le b_2$. The inequality $b_3 \le b_4$ holds because $\pi_{\beta,t} - \pi_t^* = n_t(1-n_t)A_t^2$, which implies that $\sigma exp(\sigma\pi_t^*)$ is a lower bound on $\partial\exp(\sigma\pi)/\partial\pi$ for $\pi$ between $\pi_t^*$ and $\pi_{\beta,t}$. We confirm $b_2 \le b_3$ by substituting $\pi_t^* = -U_t^2 - n_t(1-n_t)A_t^2$ and $\pi_{\lambda,t} - \pi_{\gamma,t} = 2U_tA_t + 2(n_t - \frac{1}{2})A_t^2$ to produce the equivalent inequality

$$\sigma(U_t + (n_t - \tfrac{1}{2})A_t)^2 \le \tfrac{1}{2}exp(-\sigma(U_t^2 + n_t(1-n_t)A_t^2)). \tag{B.2}$$

Calculations show that $a^* = \frac{1}{2}exp(-a^*)$ for $a^* = 0.351733711\ldots$ and that $a_1 < \frac{1}{2}exp(-a_2)$ if $a_1 < a^*$ and $a_2 < a^*$. The bounds (4.5) and (5.4) imply that $(U_t + (n_t - \frac{1}{2})A_t)^2 < a^*$ because $(n_t - \frac{1}{2})^2 = \frac{1}{4} - n_t(1-n_t) \le \frac{1}{4}$ and $\left|2(n_t - \frac{1}{2})\right| \le 1$. The bounds (4.5) and (5.4) also imply that

$U_t^2 + n_t(1-n_t)A_t^2 < a^*$ because $n_t(1-n_t) \leq \frac{1}{4}$. The last two inequalities together establish (B.2), which is the last link in showing that $b_1 \leq b_2 \leq b_3 \leq b_4$. □

## Appendix C

## Constraining to the Simplex

The following algorithm achieves $x_{i,t+1} \geq x_{i,min}$, $i = 1, ..., 3$, given the point $(x'_{1,t+1}, x'_{2,t+1}, x'_{3,t+1})$ reached by the difference equations (3.3) or (3.5). If all three fractions satisfy $x'_{i,t+1} \geq x_{i,min}$, then no adjustment is necessary. If one fraction, say $x'_{1,t+1}$ is less than $x_{1,min}$, then set $x_{1,t+1} = x_{1,min}$ and set $x_{2,t+1}$ and $x_{3,t+1}$ using

$$x_{2,t+1} = (1 - x_{1,min})\frac{x_{2,t+1}}{x_{2,t+1} + x_{3,t+1}}$$

and

$$x_{3,t+1} = (1 - x_{1,min})\frac{x_{3,t+1}}{x_{2,t+1} + x_{3,t+1}}.$$

If two fractions, say $x'_{1,t+1}$ and $x'_{2,t+1}$ are less than $x_{1,min}$ and $x_{2,min}$, then set $x_{1,t+1} = x_{1,min}$, $x_{2,t+1} = x_{2,min}$, and $x_{3,t+1} = 1 - x_{1,min} - x_{2,min}$. All three fractions cannot violate $x'_{i,t+1} \geq x_{i,min}$ because $x'_{1,t+1} + x'_{2,t+1} + x'_{3,t+1} = 1$. If the outcome for these steps would cause $sgn(x_{i,t+1} - x_{i,t}) \neq sgn(x'_{i,t+1} - x_{i,t})$ for any $i$, set $x_{i,t+1} = x_{i,t}$ for $i = 1, ..., 3$.

## References


Binmore, K., J. Gale, and L. Samuelson (1995): "Learning to be Imperfect: The Ultimatum Game," *Games and Economic Behavior*, 8, 56–90.

Bjornerstedt, J., and J. W. Weibull (1993): "Nash Equilibrium and Evolution by Imitation," in *Rationality in Economics*, ed. by K. Arrow, and E. Colombatto. Macmillan, New York.

Blume, L., and D. Easley (1992): "Evolution and Market Behavior," *Journal of Economic Theory*, 58, 9–40.

Branch, W. A. (2001): "A Noisy Cobweb Model with Dynamic Predictor Selection," *Working paper, Department of Economics, College of William and Mary.*

——— (forthcoming): "Local Convergence Properties of a Cobweb Model with Rationally Heterogeneous Expectations," *Journal of Economic Dynamics and Control.*

Bray, M. (1982): "Learning, Estimation, and the Stability of Rational Expectations," *Journal of Economic Theory*, 26, 318–339.

Brock, W. A., and C. H. Hommes (1997): "A Rational Route to Randomness," *Econometrica*, 65, 1059–1095.

——— (1998): "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model," *Journal of Economic Dynamics and Control*, 22, 1235–1274.

Cabrales, A., and J. Sobel (1992): "On the Limit Points of Discrete Selection Dynamics," *Journal of Economic Theory*, 57, 407–419.

Chiarella, C., and X.-Z. He (2002): "Heterogeneous Beliefs, Risk and Learning in a Simple Asset Pricing Model," *Computational Economics*, 19, 95–132.

De-Grauwe, P. (1993): *Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets.* Blackwell, Cambridge, MA.

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990): "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98, 703–738.

Evans, G. W. (1986): "Selection Criteria for Models with Non-Uniqueness," *Journal of Monetary Economics*, 18, 147–157.

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

Friedman, M. (1953): "The Methodology of Positive Economics," in *Essays in Positive Economics.* University of Chicago Press, Chicago.

Gourieroux, C., J. Laffont, and A. Monfort (1982): "Rational Expectations in Dynamic Linear Models: Analysis of the Solutions," *Econometrica*, 50, 409–425.

Grandmont, J.-M. (1998): "Expectations Formation and Stability of Large Socioeconomic Systems," *Econometrica*, 66, 741–781.

Hofbauer, J., and J. W. Weibull (1996): "Evolutionary Selection against Dominated Strategies," *Journal of Economic Theory*, 71, 558–573.

Howitt, P., and R. P. McAfee (1992): "Animal Spirits," *American Economic Review*, 82, 493–507.

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

LeBaron, B., W. B. Arthur, and R. Palmer (1999): "Time Series Properties of an Artificial Stock Market," *Journal of Economic Dynamics and Control*, 23, 1487–1516.

Lucas, R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1426–1445.

Marcet, A., and T. J. Sargent (1989a): "Convergence of Least-Squares Learning in Environmen with Hidden State Variables and Private Information," *Journal of Political Economy*, 97, 1306–1322.

——— (1989b): "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.

McCallum, B. T. (1983): "On Non-Uniqueness in Rational Expectations Models: An Attempt at Perspective," *Journal of Monetary Economics*, 11, 139–168.

——— (1997): "The Role of the Minimum State Variables Criterion in Rational Expectations Models," *International Journal of Tax and Finance*, 6(4), 621–639.

Osbourne, M., and A. Rubinstein (1994): *A Course in Game Theory.* MIT Press, Cambridge, MA.

Pesaran, M. H. (1987): *The Limits to Rational Expectations.* Basil Blackwell, Oxford.

Russell, B. (1917): *Mysticism and Logic and Other Essays.* George Allen and Unwin, London.

Samuelson, L. (1998): *Evolutionary Games and Equilibrium Selection.* MIT Press, Cambridge, MA.

Sargent, T. J. (1987): *Dynamic Macroeconomic Theory.* Harvard Universtiy Press, Cambridge, MA.

Taylor, J. (1977): "Conditions for Unique solutions in Stochastic Macroeconomic Models with Rational Expectations," *Econometrica*, 45, 1377–1385.

Weibull, J. W. (1997): *Evolutionary Game Theory.* MIT Press, Cambridge, MA.

Woodford, M. (1990): "Learning to Believe in Sunspots," *Econometrica*, 58, 277–307.