An Evolutionary Game Theory Explanation of ARCH Effects

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

June 14, 2006

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

Abstract

While ARCH/GARCH equations have been widely used to model financial market data, formal explanations for the sources of conditional volatility are scarce. This paper presents a model with the property that standard econometric tests detect ARCH/GARCH effects similar to those found in asset returns. We use evolutionary game theory to describe how agents endogenously switch among different forecasting strategies. The agents evaluate past forecast errors in the context of an optimizing model of asset pricing given heterogeneous agents. We show that the prospects for divergent expectations depend on the relative variances of fundamental and extraneous variables and on how aggressively agents are pursuing the optimal forecast. Divergent expectations are the driving force leading to the appearance of ARCH/GARCH in the data.

JEL Classification: C22, C73, G12, D84

Keywords:ARCH, autoregressive conditional heteroskedasticity, evolutionary game theory, rational expectations,

Corresponding Author: William R. Parke, Department of Economics, University of North Carolina, Chapel Hill, North Carolina, 27516. 919-966-5393 (voice) 919-966-2383 (messages), 919-966-4986 (fax), parke@email.unc.edu.

The goal of volatility analysis must ultimately be to explain the causes of volatility. While time series structure is valuable for forecasting, it does not satisfy our need to explain volatility. .... Thus far, attempts to find the ultimate cause of volatility are not very satisfactory.

- Robert Engle (2001)

Few models are capable of generating the type of ARCH one sees in the data. .... Most of these studies are best summarized with the adage that "to get GARCH you need to begin with GARCH."

- Adrian Pagan (1996)

1 Introduction

ARCH/GARCH models have been used to describe the behavior of inflation, interest rates and exchange rates[1], and they have become the standard tool for analyzing returns in financial markets[2]. Despite the widespread empirical successes of ARCH/GARCH models, discovering underlying mechanisms that lead to time-varying volatility has proved to be an elusive goal.

We propose that time-varying volatility is a natural feature of models with forward-looking agents. Our key condition is that agents are not constrained by assumption to agree on a single expectation. Instead, we apply recent developments in evolutionary game theory to explain how forward-looking agents might choose among differing forecasts. We assume that information arrives uniformly over time so that changes in volatility are due entirely to agents' behavior. Furthermore, our agents do not set out simply to invent ARCH. They use ideas drawn from the literature on rational expectations to choose among forecasts based on what they perceive to be fundamentals. We establish conditions under which ARCH effects will be a normal feature of the resulting data.

The mechanism generating time-varying volatility has a general formulation. Evolutionary game theory describes how fractions $x_t = (x_{1,t}, ..., x_{k,t})$ of the population using forecasting strategies $s_t = (s_{1,t}, ..., s_{k,t})$ evolve according to the performances of the strategies. An asset pricing model then shows how the price y_t depends on the fractions x_t and other information Θ_t so that

[1] Bollerslev (1986) examines inflation dynamics with a GARCH model. Engle, Lilien and Robins (1987) use the ARCH in mean model to study yield curve issues. Diebold and Nerlove (1989) use a multivariate ARCH model to analyze exchange rates.

[2] Engle (2001) provides a recent account of the methodology.

$y_t = y(x_t, \Theta_t)$. The fractions x_t are taken to be known and fixed before y_t is realized. We divide Θ_t into information Ω_t available to agents before y_t is realized and all other information ε_t. The model generating y_t can then be written in the form

$$y_t = y(x_t, \Omega_t, \varepsilon_t). \tag{1}$$

After agents observe y_t, they choose strategies for period $t+1$, updating x_t using a procedure of the form

$$x_{t+1} = g(x_t, y_t, \Omega_t). \tag{2}$$

Given this structure, the variance of the asset price can generally be written as

$$V(y_t | x_t, \Omega_t, \Sigma_t) = h(x_t, \Omega_t, \Sigma_t), \tag{3}$$

where Σ_t is a measure of the volatility of ε_t. To clarify the source of conditional heteroskedasticity, we assume that Σ_t is constant. The variance of y_t will not, however, be constant if it depends on the mix x_t of agents' strategies.

We demonstrate that ARCH effects can appear to be important in an empirical model for y_t if the econometrician does not take into account heterogeneity. A typical representative agent model for y_t, for example, would omit x_t from (1) leaving

$$y_t = y'(\Omega_t, \varepsilon_t).$$

The corresponding representative agent model of the conditional variance (3) would reduce to

$$V'(y_t | \Omega_t, \Sigma_t) = h'(\Omega_t, \Sigma_t). \tag{4}$$

Standard econometric tests may well diagnose that the representative agent model (4) leaving out the fractions x_t has ARCH, making it appear to be necessary to account for changes over time in Σ_t. We provide specific simulations that illustrate how apparent ARCH can be an artifact of ignoring heterogeneous expectations.

These ARCH effects take place in a standard mean-variance optimization model of asset prices extended to an environment with heterogeneous agents. Brock and Hommes (1998) develop the theoretical basis for this model. They apply the model in an environment with multiple trader types who use an assortment of linear forecasting rules. Brock, Hommes and Wagener (2005) and Gaunersdorfer, Hommes, and Wagener (2003) extend these results.[3] Other studies with heterogeneous expectations include Chiarella and He (2002), DeGrauwe (1993) and DeLong, Shleifer, Summers and Waldmann (1990). Our study differs in that we focus on agents pursuing goals set forth in the literature on rational expectations.

Our specific example features agents choosing among three forecasting strategies. A *fundamentalist* uses only expected future dividends to form his forecast. A *mystic* uses fundamentals, but may also experiment with other extraneous information because, perhaps, there is some uncertainty about what belongs in the fundamentals. A *reflectivist* incorporates all available information about agents' expectations, calculating the average expectation using population share weights.

We implement the procedure for updating agents' choices of forecasting strategies (2) by having agents switch to strategies that have exhibited lower squared forecast errors. The switching probabilities are determined by the evolutionary dynamic of Hofbauer and Weibull (1996). That dynamic allows for a nonlinear weighting function in the forecast evaluation. We add an important dimension to our analysis by using the nonlinearity to parameterize how aggressively agents switch forecasts.

Standard econometric tests applied to simulated data confirm that the extent of ARCH effects depends on agent aggressiveness and on the variance of the potential extraneous element that might enter the mystical forecast. If the latter variance is small relative to the variance of the fundamentals or if agents are not very aggressive, then the asset price tends to follow fundamentals nearly all the time. If the variance of the extraneous element is larger and agents are more aggressive, then asset prices show occasional bubble behavior and both Engle's (1982) test for ARCH and estimates of a GARCH(1,1) model support the conclusion that the data can be described as ARCH/GARCH for many of the simulations.

The role of heterogeneity has been noted in other related contexts with boundedly rational

[3] Hommes (2006) surveys these and related models.

forecasting strategies. Lux and Marchesi (2000) construct an asset market with fundamentalists and two different types of chartists, who respond to trends in the data. The switching probabilities between the two strategies are determined by a modified discrete choice model that allows for sluggish adjustment. Their model shows that switching between strategies can produce ARCH effects for certain parameter values. Föllmer, Horst and Kirman (2005) show the existence of bubbles, but also show the existence of limiting distributions of asset prices in a discrete choice model with forecasting strategies put forward by 'gurus' that could include chartists and fundamentalists. In a similar framework, Gaunersdorfer and Hommes (2005) derive theoretical results indicating the presence of bubbles and volatility clustering.

LeBaron, Arthur and Palmer (1999) study the time series features of a simulated asset market and show the existence of ARCH effects and many other features of financial market data. They use a computational approach with many trader types introduced throughout the simulation according to a genetic algorithm. They find that the evidence of ARCH effects is much stronger in an environment they term fast learning than it is given slow learning.

The focus in this paper on the possibility of heterogeneous forecasts stands in contrast to those who argue that martingale solutions should be ruled out according to criteria such as transversality (see Cochrane (2001) p. 27), minimum state variables (McCallum (1983, 1997)), and expectational stability (Evans and Honkapohja (2001)).[4] In our context, the importance of the martingale solutions depends on the parameter values and, while convergence to a single expectation is possible under some conditions, we establish the range of parameter values for which the martingale solutions are an important feature. Using another approach that defines stability according to *rationalizability* of strategies, Evans and Guesnerie (2003) show convergence to the minimum state variables solution with homogeneous agents, but they also show instability in the case of heterogeneous agents.

The literature on convergence to rational expectations along the lines of least squares learning (Grandmont (1998) and Marcet and Sargent (1989), for example) is also concerned with multiple rational expectations equilibria. Woodford (1990) and Howitt and McAfee (1992) establish the

[4] Evans and Honkapohja (2001) also discuss the possibility that bubble solutions might be learnable if agents use a sufficiently complicated model to form expectations.

possibility of learning sunspot equilibria given accidental correlations between sunspots and fundamentals.[5] While these papers focus on agents learning model parameters over time, our agents know the parameters and are choosing among forecasts constructed from the multiple solutions to the model.

den Haan and Spear (1998) explain conditional volatility in real interest rate fluctuations. They construct an optimizing model where agents hold savings in the form of bonds. Agents are heterogeneous, as in the present paper, and receive idiosyncratic shocks. Volatility clustering arises due to borrowing constraints that vary across the business cycle.

The organization of the paper is as follows. Section 2 develops the asset pricing model with heterogeneous agents. Section 3 and 4 describe the different forecasting strategies and their squared forecast errors. Section 5 shows how agents' choices of strategies evolve over time. Section 6 establishes how convergence of expectations depends on assumptions about agents' beliefs and willingness to consider new information. Section 7 describes the simulation methodology used in the remainder of the paper. Section 8 provides examples and compares the simulations to some stylized facts. Section 9 presents econometric analysis of the simulations using a GARCH(1,1) model. Section 10 concludes.

2 Asset Pricing with Heterogeneous Agents

This section develops an optimizing model of asset pricing with heterogeneous agents based on Brock and Hommes (1998). They extend a standard asset pricing model to the situation where agents can have heterogeneous beliefs about future prices. Agents are myopic, mean-variance optimizers who can choose between a risky asset and a riskless asset with gross rate of return R. An agent's wealth W_t evolves according to

$$W_{t+1} = RW_t + (y_{t+1} + d_{t+1} - Ry_t)\, z_t,$$

where y_t is the price of the risky asset, d_t is the dividend payment, and z_t is the number of shares of the risky asset purchased by the agent at time t. The asset price and dividend process are

[5] Timmermann (1994) argues that feedback from lagged variables into dividends rules out rational bubbles.

stochastic so agents do not have precise knowledge of y_{t+1} or d_{t+1} when making decisions about z_t.

Agents may have heterogeneous expectations. For an agent of type j, let $f_{j,t}(W_{t+1})$ be the expectation of wealth conditional on the information available at time t, and let $V_{j,t}(W_{t+1})$ be the perceived variance of W_{t+1}. Assume that the demand for shares $z_{j,t}$ by agents of type j maximizes

$$f_{j,t}(W_{t+1}) - \frac{a}{2} V_{j,t}(W_{t+1}), \tag{5}$$

where the parameter a denotes the level of risk aversion. If the perceived conditional variance $V_{j,t}(y_{t+1} + d_{t+1} - Ry_t)$ for per share excess returns is the same constant σ_r^2 for all agent types, then the demand for shares by agents of type j is given by

$$z_{j,t} = \left(\frac{1}{a\sigma_r^2}\right) f_{j,t}(y_{t+1} + d_{t+1} - Ry_t).$$

Summing the demand over the n types of agents, where the fraction of the population of type j is $x_{j,t}$, and equating total demand to a constant supply z^s per investor yields the following condition for the price of the risky asset

$$Ry_t = \sum_{j=1}^{n} x_{j,t} f_{j,t}(y_{t+1} + d_{t+1}) - C, \tag{6}$$

where $C = a\sigma_r^2 z^s$ functions as a risk premium. This equation is analogous to Brock and Hommes' (1998) equation (2.7).

We retain their assumption that the conditional variance σ_r^2 of the per share excess returns is constant. If agents were to expect bubbles, σ_r^2 could be higher, but the results remain the same for any constant σ_r^2. More sophisticated agents could try to estimate an ARCH model every period to anticipate changes in volatility, causing σ_r^2 to vary over time and across agents, but in order to focus exclusively on ARCH effects that arise from heterogeneity in the forecasting strategies, we rule out such behavior. It seems likely that variation in σ_r^2 would add to rather than reduce conditional volatility.

The pricing condition for agents with homogeneous rational expectations provides a useful benchmark for asset pricing given heterogeneous expectations (6). If all agents were of a single type, then y_t would follow

$$Ry_t = E_t\left(y_{t+1} + d_{t+1}\right) - C \tag{7}$$

The bubble-free solution to this equation is $y_t = y_t^*$, where

$$y_t^* = \sum_{s=1}^{\infty} \alpha^s E_t d_{t+s} - \left(\frac{\alpha}{1-\alpha}\right) C \tag{8}$$

and $\alpha = R^{-1}$. The self-fulfilling nature of expectations, however, admits the class of solutions

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

where m_t is a martingale that we will express as $m_t = m_{t-1} + \eta_t$ for an *i.i.d.*, mean zero stochastic process η_t with variance σ_η^2. Requiring the variance of η_t to be constant is not necessary to satisfy (7), but we want to rule out conditional volatility in the martingale innovation or in the dividend innovation as sources of conditional volatility in the asset price.

To focus attention on the effects of choosing among forecasting strategies, we retain the key underlying assumption of Brock and Hommes (1998) that agents have common beliefs about the dividend process so that y_t^* is common knowledge. We also take the realized martingale m_{t-1} to be common knowledge although agents will disagree about whether it should be used to predict y_t.

3 Strategies

The organization of this section and the next reflects the sequential process agents follow to determine y_t. They begin period t with the vector of fractions x_t already determined. Information Ω_t then arrives, and they compute forecasts for period $t+1$. These forecasts, given the asset pricing model in the previous section, determine their share positions and, hence, the price y_t. Once y_t is realized, they can compute forecast errors for the forecasts made in period $t-1$. They then use an evolutionary game theory mechanism to decide on fractions x_{t+1} for period $t+1$.

This section postulates three possible strategies for forecasting the asset price y_t. The vector of fractions following the three strategies will be denoted $x_t = (\beta_t, \gamma_t, \lambda_t)$. The forecasting strategies differ in two ways. Agents are not immediately certain whether newly proposed types of information are fundamental or extraneous, and they differ in the extent to which they attempt to exploit knowledge of other agents' expectations.

A fraction γ_t of the population uses (8) to form the *fundamentalist forecast*

$$f_{\gamma,t}(y_{t+1}) = E_t y^*_{t+1} \tag{10}$$

Agents following fundamentals base their forecast solely on expected future dividends, ruling out forecasts that involve the martingale m_t. A fraction λ_t of the population uses a *mystical forecast*

$$f_{\lambda,t}(y_{t+1}) = E_t y^*_{t+1} + \alpha^{-t-1} m_t \tag{11}$$

based on a martingale solution of the form (9). The basis for allowing agents to follow a martingale solution is, of course, a critical issue.

Some researchers exclude such solutions on the basis of a transversality condition or similar criterion (see Cochrane (2001) or other references in the introduction). These arguments require strong assumptions about the information available to the agents. Recognizing that m_t is, in fact, a martingale and not stationary data is a classic econometric problem with the characteristic that agents cannot know with certainty, on the basis of a finite data sample, that m_t is nonstationary and that a transversality condition is violated. Furthermore, the predictive properties of m_t may well be uncertain in finite samples. Later results in this paper show that, given the randomness of sample correlations, forecasting strategies based on extraneous martingales can outperform other strategies over the short term. Therefore, while agents may indeed rule out martingales in the long run, ruling out martingales in finite data samples is another matter.

A martingale might well appeal to agents trying to follow fundamentals. Suppose m_t is a new idea (the impact of the Internet on commerce or the influence of sunspots[6], for example) that might or might not be a revision to the set of fundamentals. If m_t is in fact a spurious martingale,

[6] Typically, sunspot equilibria can be constructed in a model with martingale solutions. See Farmer (1999, Chapter 10) for a full discussion.

agents attempting to choose between (10) and (11) will eventually reject (11), but the process of evaluating and rejecting (11) on the basis of empirical evidence may take some time. We diverge here from the sunspot literature in that we assume that only a portion of the agents adopt the sunspot forecast while others stick with the fundamentalist forecast.

For our third forecasting strategy, agents give primary attention to anticipating the choices of others in a manner reminiscent of Keynes' (1935) "beauty contest" interpretation of predicting stock prices. We refer to this third forecast as the *reflective forecast* because it simply reflects the views of others without attempting to consider the intrinsic value of the asset. Formally, we assume that the reflective forecast is an average of the other two forecasts, weighted according to their relative popularity. That is,

$$f_{\beta,t}(y_{t+1}) = n_t f_{\lambda,t}(y_{t+1}) + (1 - n_t)\, f_{\gamma,t}(y_{t+1}), \tag{12}$$

where

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

The reflective forecast can be expressed in terms of fundamentals and the martingale by substituting (10) and (11) into (12) to obtain

$$f_{\beta,t}(y_{t+1}) = E_t y^*_{t+1} + \alpha^{-t-1} n_t m_t. \tag{13}$$

The martingale thus influences the reflective forecast to the extent that the fraction n_t of the other agents are following the mystical forecast. If $n_t = 0$ or $n_t = 1$, then the reflective forecast will numerically match the dominating forecast.[7] The reflective forecast (13) differs from perfect foresight because $y^*_{t+1} - E_t y^*_{t+1}$ and $\alpha^{-t-1} n_{t+1} m_{t+1} - \alpha^{-t-1} n_t m_t$ are not known in period t. The reflective forecast brings to the system the notion of taking into account the expectations of other agents, which is an important element of rational expectations.

The motive for agents to include the reflective forecast among the strategies they consider might be some instinctive belief in the possible merit of averaging across available forecasts. An average

[7] The reflective forecast is not well-defined for $\lambda_t = \gamma_t = 0$, but, as we explain in Section 6, that case is not the focus of this paper.

is an obvious statistic to compute so it would not be be surprising if agents did the calculation. In the next section we give this intuition a mathematical foundation by showing that, in every period, the payoff to the reflective forecast is greater than or equal to the average payoff to the other forecasts. Unless the fundamentalist and mystical forecasts happen to be numerically equal (because $m_t = 0$), the reflective forecast is guaranteed to have an above average payoff.[8]

The information requirements for these strategies differ somewhat, but to make switching strategies possible we assume that all agents have access to the same information, including expected dividends and the martingale. While the knowledge of the fractions β_t, γ_t, and λ_t needed to implement the reflective forecast might come from direct observation of other agents, that is not necessary. The fractions change over time according to updating equations, given below as (24), that are functions of the payoffs. Given that the agents know the payoffs and a starting point for the fractions, they can recursively calculate the fractions over time just as we do in our simulations later in this paper.[9]

Knowledge of the fractions β_t, γ_t, and λ_t is easier to justify under the alternative interpretation that there are many agents, but only three forecasters. Only the forecasters need to know how to calculate a forecast. The agents simply choose among the available forecasts. Föllmer, Horst, and Kirman (2005) make a similar distinction between agents and "gurus."

The realized asset price (6) can be written as

$$y_t = \alpha(\gamma_t f_{\gamma,t}(y_{t+1}) + \lambda_t f_{\lambda,t}(y_{t+1}) + \beta_t f_{\beta,t}(y_{t+1}) + E_t d_{t+1} - C).$$

Using (12) to express $\gamma_t f_{\gamma,t}(y_{t+1}) + \lambda_t f_{\lambda,t}(y_{t+1})$ in terms of $f_{\beta,t}(y_{t+1})$ yields

$$y_t = \alpha(f_{\beta,t}(y_{t+1}) + E_t d_{t+1} - C),$$

which emphasizes that the reflective forecast does summarize the available information in deter-

[8] If $m_t = 0$, then $A_t = 0$ in equation (20) below.

[9] It is beyond the scope of this paper, but one might consider the situation where the reflective forecast is based on agent fractions estimated with some imprecision. For example, a loose interpretation of the updating equations (24) is that the fraction following a strategy can be estimated from its recent payoffs.

mining y_t. Substituting (13) and noting that y_t^* satisfies (7) yields

$$y_t = y_t^* + \alpha^{-t} n_t m_t. \tag{14}$$

The market price y_t deviates from the price y_t^* implied by the true fundamentals to the extent that a nonzero fraction n_t of the agents not following the reflective forecast are in fact following the mystical forecast.

4 Evaluating Forecasts

Once y_t is realized, agents are in a position to evaluate their current strategies, and there are several possible criteria.[10] We adopt the squared error criterion in light of its long tradition in the econometrics literature for evaluating forecasts. LeBaron et. al. (1999) make the same choice. Brock and Hommes (1998), on the other hand, use current realized trading profits. They also suggest a more general weighted average of trading profits that includes accumulated wealth as another case. Hommes (2001, p. 156) shows that the trading profits criterion does not take into account risk, but that the squared error criterion can be derived from the utility function (5) for a mean-variance optimizing risk averse agent. Gaunersdorfer, Hommes, and Wagener (2003) replace the trading profits payoff in Brock and Hommes (1998) with a squared error payoff function and derive qualitatively similar results.

We assume that agents attempt to forecast $d_t + y_t$ rather than just y_t because $d_t + y_t$ is the total return to holding an asset with price y_{t-1} in period $t-1$. The agents share the common expectation $E_{t-1}d_t$ for d_t so their forecasts for d_t are identical. Using (10), (11), and (13), the forecasts of y_t from period $t-1$ are given by:

$$\begin{aligned} f_{\gamma,t-1}(y_t) &= E_{t-1}y_t^*, \\ f_{\lambda,t-1}(y_t) &= E_{t-1}y_t^* + \alpha^{-t} m_{t-1}, \\ f_{\beta,t-1}(y_t) &= E_{t-1}y_t^* + \alpha^{-t} n_{t-1} m_{t-1}. \end{aligned}$$

[10] Blume and Easley (1992) discuss some issues involved in choosing payoff functions in the context of an evolutionary study of asset pricing. They are concerned with long run survival of strategies.

The reflective forecast error can be decomposed using (14) as $U_t = F_t + G_t$, where

$$F_t = d_t + y_t^* - E_{t-1}(d_t + y_t^*) \tag{15}$$

is the innovation in the fundamentals and

$$G_t = \alpha^{-t}(n_t m_t - n_{t-1} m_{t-1}) \tag{16}$$

is the innovation in the weighted martingale. Using the notation $A_{t-1} = \alpha^{-t} m_{t-1}$, the fundamentalist and mystical forecast errors can be expressed as $U_t + n_{t-1}A_{t-1}$ and $U_t - (1 - n_{t-1})A_{t-1}$. We let σ_F^2 denote the variance of F_t, and, to focus unambiguously on other sources of heteroskedasticity, we assume this variance is constant over time. The process F_t will be serially independent regardless of the structure of the dividend process.

The payoffs are the negatives of the squared forecast errors:

$$\pi_{\beta,t} = -U_t^2, \tag{17}$$

$$\pi_{\gamma,t} = -U_t^2 - 2n_{t-1}A_{t-1}U_t - n_{t-1}^2 A_{t-1}^2, \tag{18}$$

$$\pi_{\lambda,t} = -U_t^2 + 2\left(1 - n_{t-1}\right) A_{t-1} U_t - \left(1 - n_{t-1}\right)^2 A_{t-1}^2. \tag{19}$$

The terms involving A_{t-1} appear because the realized price depends on A_{t-1} to the extent that some agents follow the mystical forecast. The differences in the squared forecast errors depend on the product of A_{t-1} and U_t and the square of A_{t-1}.

Any of the three forecasts could have the best payoff in a particular period, and the ordering depends on the realizations of A_{t-1} and U_t. The ordering $\pi_{\lambda,t} > \pi_{\beta,t} > \pi_{\gamma,t}$ can occur if the term $2\left(1 - n_{t-1}\right) A_{t-1} U_t$ in the payoff to mysticism is positive (the mystic has conjured a fortuitously accurate forecast) and sufficiently large. Similarly, $\pi_{\gamma,t} > \pi_{\beta,t} > \pi_{\lambda,t}$ can occur if the term $2n_{t-1}A_{t-1}U_t$ is sufficiently negative. If $A_{t-1}U_t$ is sufficiently small in absolute value, then $\pi_{\beta,t}$ is greater than both $\pi_{\gamma,t}$ and $\pi_{\lambda,t}$.

The reflective forecast has a natural advantage that we can express in terms of the *fitness* of a given strategy, which is the difference between a given strategy's payoff and the population average payoff. Using the period $t-1$ population shares of agents following the forecasts that produce the realized payoffs in period t, the average payoff is $\bar{\pi}_t = \beta_{t-1}\pi_{\beta,t} + \gamma_{t-1}\pi_{\gamma,t} + \lambda_{t-1}\pi_{\lambda,t}$. The fitness of the reflectivist forecast is given by

$$\pi_{\beta,t} - \bar{\pi}_t = (1-\beta_{t-1})n_{t-1}\left(1-n_{t-1}\right)A_{t-1}^2 \geq 0. \tag{20}$$

The reflectivist payoff is thus unambiguously better than the average payoff.

The nature of the contest between fundamentalism and mysticism is captured by

$$\pi_{\lambda,t} - \pi_{\gamma,t} = 2A_{t-1}U_t - (1-2n_{t-1})A_{t-1}^2. \tag{21}$$

If fundamentalism has a large following (n_{t-1} is near zero), then the term involving A_{t-1}^2 favors continued domination by fundamentalism. A large positive $A_{t-1}U_t$ could, however, reverse this tendency. Mysticism has a symmetric advantage if n_{t-1} is near one, and that advantage could be reversed by a large negative $A_{t-1}U_t$.

5 Evolution

To translate the relative payoffs into changes in agents' beliefs we adopt a behavioral process from evolutionary game theory known as "imitation of successful agents."[11] The name emphasizes that the payoffs of other agents affect an agent's probability of switching strategies. This contrasts with strategies where agents focus only on their own payoffs.

Imitation of successful agents can be developed within a more general model of how agents review and change forecasting strategies. Let $r_{j,t}$ be the fraction of agents using forecast j who review their choice of strategy at time t, and let $p_{j,t}^i$ be the probability that a reviewing agent using forecast j in period t switches to forecast i in the next period. We let $x_t = (\beta_t, \gamma_t, \lambda_t)$ denote the vector of population shares, and we will use $x_{i,t}$ to reference the elements of this vector. If there

[11] Björnerstedt and Weibull (1996), Hofbauer and Weibull (1996), and Weibull (1997) develop imitation of successful agents. DeLong, Schleifer, Summer, and Waldman (1990) use another form of imitation in their noise trader model.

are k available forecasts, then the change in $x_{i,t}$ is given by

$$x_{i,t+1} - x_{i,t} = \sum_{j=1}^{k} r_{j,t} x_{j,t} p_{j,t}^{i} - r_{i,t} x_{i,t}. \tag{22}$$

This is a discrete time version of equation (4.25) in Weibull (1997). We assume that all agents review every period regardless of the payoff so $r_{j,t} \equiv 1$, but that the transition probabilities $p_{j,t}^{i}$ depend on the performances of the strategies. Agents will tend to switch to strategies with better payoffs, meaning lower squared forecast errors. We assume that agents arrive at the transition probabilities using payoff weighting functions $w(\pi_{i,t})$ to calculate

$$p_{j,t}^{i} = \frac{w(\pi_{i,t})\, x_{i,t}}{\overline{w}_t}, \tag{23}$$

where $\overline{w}_t = \sum_{h=1}^{n} w(\pi_{h,t})\, x_{h,t}$. The transition probability $p_{j,t}^{i}$ into strategy i depends on its current popularity $x_{i,t}$ and on its current payoff $w(\pi_{i,t})$ relative to the population weighted average $\overline{w}_t$. Substituting (23) into (22) with $r_{j,t} \equiv 1$ yields the equation of motion for the population shares

$$x_{i,t+1} = x_{i,t} \frac{w(\pi_{i,t})}{\overline{w}_t}. \tag{24}$$

The specific dynamics of the system will depend on the functional form of $w(\cdot)$.

In particular, the dynamics of the system depend on the convexity of the weighting function $w(\pi)$. If $w(\pi)$ is linear, the evolution of the fractions $x_{i,t}$ follows the *replicator dynamic* as the fractions change proportionally with the fitness (payoff relative to the population average) of a given strategy. For example, if $w(\pi) = \tau + \pi$ for a constant τ, then (24) can be written as

$$x_{i,t+1} = x_{i,t} \frac{\tau + \pi_{i,t}}{\tau + \overline{\pi}_t}. \tag{25}$$

Given (20), β_t is monotone increasing, leaving no opportunity for the emergence of mysticism because γ_t and λ_t will be forced to their minimum possible values. (For Case 2 and Case 3 discussed in the next section, $\gamma_t + \lambda_t$ is small, but positive.) No finite value for τ can guarantee that $\tau + \pi_{i,t}$ and $\tau + \overline{\pi}_t$ in (25) are positive, however, and setting $w(\pi)$ to some positive constant

κ if $\tau + \pi_{i,t} \leq \kappa$ does introduce a convexity into $w(\pi_{i,t})$. A convex weighting function provides conditions that can lead agents to adopt the mystical forecast. Simulations (not reported in this paper) show that the prospects for heterogeneous expectations depend on the value of τ, which determines how frequently the constraint $\tau + \pi \geq \kappa$ is binding.

To explore the relation between convexity in the payoff weighting function and convergence of expectations, we base our analysis in this paper on the exponential weighting function

$$w(\pi) = e^{\theta^2 \pi}, \tag{26}$$

where θ parameterizes the convexity of the function. Compared to the linear weighting function underlying the replicator dynamic, convexity of the weighting function means the population shares change *overproportionally* with the fitness of the strategies[12]. In economic terms, greater convexity of $w(\pi)$ implies that agents are seeking out the best performing strategy more aggressively.

We derive in the appendix approximations for the equations of motion for β_t and $n_t = \lambda_t/(\lambda_t + \gamma_t)$ that qualitatively characterize how the properties of the updating functions depend on the convexity θ. For β_t, when the accumulated martingale innovations have not yet made A_{t-1}^2 large, we have

$$\frac{\beta_t}{\beta_{t+1}} \cong 1 - \theta^2(1-\beta_t)n_{t-1}(1-n_{t-1})A_{t-1}^2(1-2\theta^2 U_t^2). \tag{27}$$

Note that $\dfrac{\beta_t}{\beta_{t+1}} < 1$ implies that reflectivism's share is increasing. The fraction β_t following reflectivism will fall, making it possible for fundamentalism or mysticism to gain followers if agent aggressiveness θ and the squared forecast error U_t^2 for the reflective forecast are sufficiently large.

The nature of the contest between the fraction λ_t following mysticism and the fraction γ_t following fundamentalism can be seen in an approximation to the equation of motion for n_t:

$$\frac{n_t}{n_{t+1}} \cong 1 - 2\theta^2(1-n_t)(A_{t-1}U_t + (n_{t-1} - \tfrac{1}{2})A_{t-1}^2). \tag{28}$$

This equation of motion inherits properties noted for the simple difference in the mystical and

[12] Hofbauer and Weibull (1996) examine the specification of the weighting function in detail.

fundamental forecast payoffs (21). If n_{t-1} is near zero or one, then the factor $n_{t-1} - \frac{1}{2}$ acts to put the weight of the squared martingale A_{t-1}^2 toward reinforcing that value of n_{t-1}. Reversing that trend requires a large value for $A_{t-1}U_t$ of the appropriate sign. For example, if n_{t-1} is less than $\frac{1}{2}$, then mysticism gains relative to reflectivism if $A_{t-1}U_t$ is a large positive number with $A_{t-1}U_t > -(n_{t-1} - \frac{1}{2})A_{t-1}^2$. (A symmetric result for gains in fundamentalism applies if n_{t-1} is greater than $\frac{1}{2}$.) Mysticism thus gains overall if a large squared forecast error U_t^2 in (27) causes a decrease in reflectivism and a large product $A_{t-1}U_t$ in (28) appears to show that the martingale predicts the forecast error. The factor θ^2 in both (27) and (28) causes the magnitude of the changes to be in proportion to the square of agent aggressiveness θ. A greater value for θ^2 reduces the number of consecutive fortuitous martingale realizations it would take to propel mysticism to a given popularity.

Imitation of successful agents differs from the evolutionary mechanism Brock and Hommes (1997) refer to as the discrete choice model.[13] That model, which is a close relative of the multinomial logit model, can be written as

$$x_{i,t+1} = \frac{w(\pi_{i,t})}{\sum_{j=1}^{k} w(\pi_{j,t})},$$

where $w(\pi_{i,t}) = e^{\beta\pi_{i,t}}$ and β is the "intensity of choice," which is similar to our measure θ of agent aggressiveness. A fraction $x_{i,t}$ cannot reach 0 under the discrete choice model although for high levels of β the fractions can become very small. The steady states for this model will, therefore, be interior to the simplex containing the vector x_t. Convergence to homogeneous expectations is not possible.

Imitation of successful agents (24), on the other hand, has the property that $x_{1,t}/x_{2,t}$ decreases if $w(\pi_{1,t}) < w(\pi_{2,t})$. Inferior strategies can be driven to zero popularity because $x_{i,t+1}$ depends on $x_{i,t}$ in (24). The factor $x_{i,t}$ appears because the transition probability in (23) depends on popularity, as measured by $x_{i,t}$. In fact, if $p_{j,t}^i = x_{i,t}$ in (23), then we would have a model of pure imitation, where agents choose a new strategy by randoming picking another agent and adopting

[13] Brock and Hommes (1997) introduced the discrete choice model as a method for studying the evolution of heterogeneous expectations in a cobweb model. They extend this approach to the asset pricing framework discussed here in Brock and Hommes (1998). Chiarella and He (2002) is one of many extensions.

that person's strategy. Imitation of successful agents, as generated by (23), assumes that agents switching strategies consider both current payoffs and popularities. The latter can be thought of as measuring the quality of a strategy's previous payoffs because a strategy gains in popularity to the extent it secures a series of favorable payoffs.

Considering the possibility of convergence to homogeneous expectations is thus reasonable under imitation of successful agents. The central question will be whether that outcome is robust to the introduction of small fractions of agents using alternative strategies.

6 Convergence of Expectations

We have constructed an evolutionary model of expectation formation in order to consider whether expectations converge and, if they do not, to characterize the nature of the resulting heterogeneous expectations. We will show that persistent heterogeneous expectations are more likely if agents are more aggressive. Two additional features of agents' behavior are important factors.
(i) Fundamentalism might have a special appeal because it is so widely cited by learned economists.
(ii) Agents might be willing to consider new information thought by some to help predict y_t. We analyze three cases.

Case 1: No Underlying Beliefs

If agents simply play the game as it is described to this point (without (i) and (ii) above), then $x_t = (\beta_t, \gamma_t, \lambda_t)$ will eventually cease changing when it reaches one of two edges of the simplex $\Delta = \{(\beta_t, \gamma_t, \lambda_t) | \beta_t \geq 0, \gamma_t \geq 0, \lambda_t \geq 0, \text{ and } \beta_t + \gamma_t + \lambda_t = 1\}$. The evolution equation (24) shows that $x_{i,t} = 0$ implies $x_{i,s} = 0$ for $s \geq t$. That is, if a strategy has no followers, it cannot acquire any. On the edge where $\lambda_t = 0$ and on the edge where $\gamma_t = 0$, nonzero weights apply to the reflective forecast and one other forecast. The two forecasts are numerically identical so there will be no further change in $x_t = (\beta_t, \gamma_t, \lambda_t)$. The edge where $\beta_t = 0$ is not an absorbing state because the fundamental and mystical forecasts will not be numerically equal and the evolution will continue until $\lambda_t = 0$ or $\gamma_t = 0$. The outcome will thus be some combination of reflectivism and one other strategy where the third strategy is extinct.[14]

[14] The reflective strategy is not well-defined at the point where $(\beta, \gamma, \lambda) = (1, 0, 0)$, but the fractions cannot reach

This first case does not provide a very satisfying foundation for homogeneous expectations. In the long run, the expectations will follow either the fundamental solution or the martingale solution, but the outcome depends on the starting fractions and a period of stochastic movement among strategies. Given the algebraic symmetry in (18) and (17), there is no meaningful difference between fundamentalism and mysticism from the point of view of the game.

Case 2: Core Belief in Fundamentals and Averages

By augmenting the model with two assumptions about agents' core beliefs, we can make a strong case for convergence to fundamentals.

Condition 1 *The fraction* β_t *is bounded from below by the minimum* $\beta_{\min} > 0$.

Two arguments support the assumption that a core fraction $\beta_{\min}$ of the agents are willing to stick with the reflective strategy regardless of particular realized payoffs. First, equation (20) guarantees that the payoff to the reflective forecast is at least equal to the average payoff. Second, if the process does converge to a steady state with a single dominant strategy, then the reflective forecast will automatically match the winning forecast numerically, guaranteeing its followers the maximum possible payoff. Case 1 is an example of this. While we do not have explicit costs of evaluating or switching strategies, the reflective forecast avoids both while maintaining an above average payoff at all times and matching whatever forecast is eventually dominant.

Condition 2 *The fraction* γ_t *is bounded from below by the minimum* $\gamma_{\min} > 0$.

We attribute the core following for fundamentalism to published research in economics. A large literature in economics attempts to justify the assumption that agents will unanimously agree on the fundamental solution to models with forward expectations. Countless papers simply impose this assumption. We do not assume universal belief in the fundamentalist forecast, but we do assume that some fraction $\gamma_{\min}$ of the agents choose their strategy based on published economic research and follow the fundamentalist forecast regardless of what other agents do.

Our goal with Case 2 is to consider the possibility of convergence to the fundamentalist expectation without simply imposing the assumption that $\gamma_{\min} = 1.00$. In our simulation results, we

that point in Cases 2 and 3, which are the focus of this paper.

take both $\beta_{\min}$ and $\gamma_{\min}$ to equal 0.05. We pick 0.05 to be small relative to 1.00, but large relative to the precentage of the agents who consider the mystical forecast in Case 3 below.

For even these small figures, mysticism will eventually lose its appeal and the fractions will move to the steady state where $\lambda_t = 0$. While mysticism can gain popularity when the martingale accidentally forecasts fundamentals, any period of relatively small squared forecast errors will, according to (27), drain followers from both mysticism and fundamentalism, increasing the following of reflectivism. The limit of this process leaves a remaining core $\gamma_{\min} = 0.05$ of unyielding fundamentalists that outnumbers the remaining $\lambda_{\min} = 0$ followers of mysticism. It is thus the natural advantage of reflectivism and the core of unyielding fundamentalists that lead to the collapse of mysticism in the long run.

The mechanism that causes the eventual failure of mysticism is an important characteristic of our model. In some models of rational bubbles, Evans (1992) and Hall, Psaradakis and Sola (1999), for example, the bubble collapses with some exogenously given probability.[15] The collapse in our model occurs endogenously, given the natural tendency for the fraction β_t following the relective forecast to increase and given that we take $\gamma_{\min}$ to be substantially larger than $\lambda_{\min}$ because the fundamental solution is prominently featured in published economics research. Furthermore, the collapse of mysticism requires only that agents study squared forecast errors, not that they develop some deeper theory about minimum state variables or expectational stability.

Case 3: Evaluating New Information

The central question we address in this paper is whether the convergence to the fundamentalist forecast in Case 2 is robust to a very small fraction of the agents evaluating new information. In particular, suppose the agents consider the mystical forecast, which is based on a martingale solution. We assume that the agents cannot know with certainty that the martingale is extraneous.[16] In our simulation results, we study the effects if $\lambda_{\min} = 0.0001$ of the agents consider the mystical forecast. In practice, we choose $\lambda_{\min}$ to be much lower than the other minima so that, if mysticism is near its minimum, it has very little effect on the asset price and y_t essentially follows the fundamental

[15] Charemza and Deadman (1995) present an alternative model of speculative bubbles analogous to Evans (1992). The explosive term enters additively in Evans (1992) and multiplicatively in Charemza and Deadman (1995).

[16] Even if the agents intend to rule out nonstationary solutions, nonstationarity tests do not yield certainty in finite samples.

solution. Our goal is to find out whether, even after a very minimal beginning, mysticism can gain sufficient following to impact the system, causing bubble-like behavior and inducing ARCH effects in the time series data. Our formal statement of this challenge to stability is

Condition 3 *The fraction λ_t is bounded from below by the minimum $\lambda_{\min} > 0$. If this bound is reached, the martingale restarts at $m_t = 0$.*[17]

Analyzing whether equilibria are robust to the introduction of small fractions of the population using alternative strategies is a common topic in evolutionary game theory. Binmore, Gale and Samuelson (1995) and Binmore and Samuelson (1999), in particular, consider the possibillity that arbitrarily small "drift" in the population fractions can have large impacts on the outcome. Here, we simulate the model while imposing $\lambda_{\min} = 0.0001$ to examine the stability of adherence to the fundamental solution when a very small fraction of the agents consider extraneous information.

7 Simulation Methodology

The remainder of this paper examines the empirical properties of the per share excess return

$$Z_t = d_t + y_t - \alpha^{-1} y_{t-1}, \tag{29}$$

which is a natural measure of investment performance. Expressing y_t as (14) and noting that y_t^* satsifies (7) yields $Z_t = U_t + C$, where $U_t = F_t + G_t$ is the reflectivist forecast error given by (15) and (16) and C is the risk premium that appears in (6). Because C is constant in this paper, the per share excess returns are (up to a constant) equal to the reflectivist forecast errors.

One advantage of studying per share excess returns rather than percentage returns is that our results are invariant to the dynamic structure of the dividend process. We can simulate the serially independent innovations to fundamentals F_t without making assumptions about the dividend process and without actually calculating d_t, y_t, and y_{t-1}.[18]

[17] An alternative would be have a new martingale start every period, but that would lead to a large and variable number of strategies in a given period.

[18] In calculations not reported here we have checked the differences that might result from assuming a specific dividend process in order to make possible direct calculation of d_t, y_t, and gross percentage returns $(d_t + y_t - y_{t-1})/y_{t-1}$. Results for gross percentage returns very similar to those in Tables 1 and 2 for excess returns can be obtained if dividends are assumed to follow an AR(1) process with an autoregression parameter equal to 0.95.

The conditional volatility of excess returns Z_t depends on three main parameters: agent aggressiveness θ, martingale volatility σ_η, and the volatility of fundamentals σ_F. We adopt the normalization $\sigma_F = 1.0$ so that the volatility of the fundamentals is fixed. Some normalization is necessary because multiplying θ by a constant ϕ while dividing F_t, G_t, and A_t by ϕ would leave the weighted payoffs (26) unchanged. The normalization $\sigma_F = 1.0$ has the intuitive appeal of holding the properties of fundamentals fixed while focusing on agent behavior and on the nature of the extraneous martingale.

Other parameters are fixed at reasonable values that are constant across simulations. These parameter values include the agents' risk aversion parameter $a = 0.25$, the discount factor $\alpha = 0.99$, and the total supply of shares per investor $z^s = 1.0$.

All simulations begin at the potentially stable point where reflectivism has its maximum number of followers given Conditions 2 and 3. The initial fraction $\lambda_{\min} = 0.0001$ of the population using the mystical forecast is much lower than the other minima, $\gamma_{\min} = 0.05$ and $\beta_{\min} = 0.05$. While we assume that one agent in 20 is an unyielding believer in the true fundamentals, we assume that only one agent in 10,000 is attracted to a mystical forecast under first consideration. The initial minimum fraction following fundamentalism thus dominates the initial fraction following mysticism, making the reflectivist forecast nearly identical to the fundamentalist forecast. These initial values are intended to ensure that the ARCH effects in the simulated data do not arise spuriously from the size of the initial fraction following mysticism.

If, after some initial increase, the following λ_t for the mystical forecast again falls to the minimum $\lambda_{\min}$, we implement Condition 3 by allowing λ_t to remain at $\lambda_{\min}$ and resetting the martingale to $m_{t-1} = 0$. This effectively represents a new mystical forecast. The model in this paper could be extended to include many different mystical forecasts operating simultaneously, but we focus on a single mystical forecast for clarity. An alternative approach, which is common in the computational finance literature (LeBaron (2000, 2006)) would be to regularly introduce new strategies into the population.

8 Volatility Clustering and Excess Kurtosis

Volatility clustering and fat tails are two of the most striking properties of excess returns in financial market data.[19] In this section, we explore the relation between these phenomena on the one hand and agent aggressiveness and the variance of the martingale innovation on the other.

Figures 1, 2, and 3 show typical realized simulations for the model. The top graph in each figure is the realized share price.[20] The second graph shows the per share excess return. The bottom three graphs show the shares of the population using each strategy across time.

Figures 1 and 2 ilustrate the dramatic differences between cases where agents basically agree on a single forecast and cases where agents hold heterogeneous, evolving expectations. Figure 1 ($\sigma_\eta = 1.0$, $\theta = 0.5$) shows a case where mysticism plays no role. The population share for mysticism remains at its minimum, and the asset price remains close to the fundamental solution.

Figure 2 ($\sigma_\eta = 1.0$, $\theta = 1.0$) shows how mysticism becomes a factor when agents are more aggressive. There are periods when mysticism succeeds in attracting adherents, becoming the dominant strategy at times. Stretches of time when mysticism dominates often show bubble behavior in the asset price and large deviations from zero in the excess returns. Such bubbles do not, however, last indefinitely. As we note in Section 6, the unyielding fraction $\gamma_{\min}$ following fundamentals is much larger than the minimum fraction $\lambda_{\min}$ for mysticism, and reflectivism has a natural advantage. These two factors lead to an eventual collapse of mysticism.

Figure 3 ($\sigma_\eta = 2.0$, $\theta = 2.0$) shows a further increase in martingale volatility sufficient to cause briefer, but more decisive episodes of mysticism. Volatility clustering in the returns is quite evident and occurs around short outbursts of mysticism. Engle (2001) suggests that clusters of large shocks must be the result of news. We can interpret our simulations as agents temporarily responding to a new variable, but quickly discarding it as irrelevant. Mysticism can gain a large following but it usually lasts for less than 10 periods before it is rejected. Outbreaks of mysticism that last longer, as shown in Figure 2, tend to occur for more moderate values of θ and σ_η.

To formally confirm the volatility clustering apparent in these figures, we also calculate Engle's

[19] Pagan (1996).

[20] As we note in the previous section, we need the variance of the innovation to fundamentals, but not a specific dividend process to calculate excess returns and the fractions β_t, γ_t, and λ_t. To calculate realized share price for the graph, we use the auxiliary assumption that dividends follow an AR(1) process with autoregression parameter 0.95.

(1982) test for ARCH in the simulated returns (29), setting the lag parameter to 5.[21] The critical value for Engle's test for ARCH is 11.07. The test statistics for Figures 1, 2, and 3 are 2.68, 39.01, and 62.83, respectively.

Table 1 gives a more comprehensive view of some features of the simulated data. The four panels show degrees of agent aggressiveness, and the rows within the panels show degrees of martingale volatility. Each entry in the first column shows the percentage of the sample runs, out of 10,000 trials, for which one would reject the null hypothesis of homoskedasticity using the 0.05 critical value for Engle's test with the lag parameter set to 5. Column 2 shows the average estimate of the long memory parameter d for the Geweke and Porter-Hudak (1983) estimation procedure for squared excess returns. Column 3 shows the average kurtosis over the trials. The three columns on the right show the average percentage (over all trials and all periods) of the agents following each of the three forecasting strategies.

The basic finding in Table 1 is that agent aggressiveness and martingale volatility both contribute to the likelihood of diagnosing ARCH and excess kurtosis in the simulated data. Agent aggressiveness, parameterized by the curvature of the payoff function, is clearly important. If the parameter θ determining agent aggressiveness is less than 1.00, there is no evidence of ARCH for any level of martingale volatility. For $\theta \geq 1.00$, however, ARCH is a dominant feature of the data as agents aggressively switch forecasts searching for the optimal strategy.

For a fixed level of agent aggressiveness θ, Table 1 shows that mysticism is more likely to be an important factor for a large σ_η. Small martingale innovations lead to small differences in the payoffs ((17), (18) and (19)), making the likely gains in popularity for mysticism small even when A_{t-1} does match the sign of U_t. For a larger martingale innovation variance, fortuitious values for A_{t-1} have a greater chance of causing large differences in the payoffs and gains for mysticism.

The estimates of the long memory parameter for squared excess returns, on the other hand, show no indication that long memory is an important feature of the simulations.[22] Several sets of results that are not reported in Table 1 also show no evidence of long memory. The serial correlation

[21] Engle's test is a joint test of the coefficients of the first k lags of the squared residuals in a regression on the current squared residual. Calculations for a range of alternative values for k show that the results are not particularly sensitive to this parameter for our simulations.

[22] Pagan (1996, p.30) notes the common finding of long memory in empirical squared excess returns.

coefficients for squared and absolute excess returns have the pattern that, in cases where the first-order serial correlation is noticeably greater than zero, the higher order serial correlations drop toward zero faster than is consistent with long memory. The long memory parameter estimates for the absolute value of excess returns are similar to those for squared excess returns, but smaller. For excess returns , the estimates of d are all very near zero.

The most striking feature of the population fractions in Table 1 is that the average percentage following mysticism does not have to be very large to induce ARCH. For $\theta = 1.0$, the average fraction following mysticism is less than 0.015 in all cases even though the probability of rejecting homoskedasticity is as high as 0.700. In no case in Table 1 does the average fraction following mysticism exceed 0.12.

9 GARCH

A natural next step is to examine the simulated data that showed ARCH effects to see whether it is well represented by a GARCH (Generalized ARCH) model. GARCH models, introduced by Bollerslev (1986), are commonly used to examine financial market data and offer a useful extension of the ARCH approach. We examine the simulated data with the GARCH(1,1) model that Engle (2004) terms "the workhorse of financial applications."[23] This approach models the conditional variance of the errors $E_{t-1}\left(\varepsilon_t^2\right) = \sigma_{R,t}^2$ as

$$\sigma_{R,t}^2 = \kappa + \varphi\sigma_{R,t-1}^2 + \psi\varepsilon_{t-1}^2, \tag{30}$$

where κ, φ, and ψ are constants. The conditional variance depends on the previous period's conditional variance and on the previous period's squared error. The advantage of this specification is its parsimony. It does not require multiple lags of ε_t^2, and it separates the effects of the long term conditional variance $\sigma_{R,t-1}^2$ and the short term squared errors ε_{t-1}^2. Of course it is possible to include further lags of either variable, but, as Engle (2001) notes, the GARCH(1,1) model has proved sufficient for most financial market data.

Table 2 reports estimation results for GARCH(1,1) models for the data described for Table 1.

[23] Bollerslev, Chou, and Kroner (1992) survey GARCH modeling in finance.

Following the organization of Table 1, each section of the table gives results for a given level of agent aggressiveness, and each row summarizes the results for a given standard deviation of the martingale innovation. Each row is calculated using 1,000 sample runs of 1,000 periods. All the reported statistics are calculated for just those simulations for which Engle's test rejects the null of no ARCH, making the GARCH results conditional on a preliminary diagnosis of ARCH.

The estimate of φ is the best initial indication of whether the data is well represented by (30). For those cases where Engle's test rejects the null of no ARCH, Column 1 gives the conditional probability that a significance test for the estimate of φ will reject the null of zero, implying that the conditional variance is serially dependent. Columns 2 gives the mean of the estimates of φ for the cases where $H_0 : \varphi \leq 0$ is rejected.

We also report rejection probabilities for $H_0 : \psi \leq 0$ in Column 3. The probability in Column 3 is conditional on the Engle test rejecting the null of no ARCH and the test of φ rejecting $H_0 : \varphi \leq 0$. Column 4 shows the mean of the estimates of ψ where the null is rejected. Finally, we conduct a diagnostic test (Enders (2004, p. 136)) for serial correlation in the squared residuals, using a Ljung-Box Q test on the runs for which φ is significantly greater than zero. The final column reports the percentage of the runs for which the Q test rejects the null of no serial correlation in the squared error terms.

We draw two main conclusions from Table 2. First, as in Table 1, agent aggressiveness θ and martingale volatility σ_η are both important in determining the relevance of a GARCH(1,1) empirical model. For small values of these parameters, few trials are diagnosed as having ARCH by Engle's test and, for those few trials for which the null of no ARCH is rejected, the frequency of rejecting $H_0 : \varphi \leq 0$ is modest. For larger values of θ and σ_η, the probability of rejecting $H_0 : \varphi \leq 0$ is over 0.5 in several cases. For those cases, the probability of rejecting $H_0 : \psi \leq 0$ is near one and the probability that the diagnostic test finds serial correlation in the squared residuals is near the nominal 5% size of the test. For several rows in Table 2, therefore, there is a good chance that standard econometric tests would support the conclusion that the data is GARCH(1,1).

Second, Table 2 hints at the empirical appeal of generalizing the basic ARCH/GARCH models. At the highest level of agents aggressiveness $\theta = 4.0$ and for the larger values of martingale volatility σ_η, the results for Engle's test in Table 1 are decisively in favor of ARCH. For those same parameter values, Table 2 does not lend much support to fitting a GARCH(1,1) model. The probability of rejecting $H_0 : \varphi \leq 0$ is relatively low, as is the mean of the significant estimates of the ARCH autoregression parameter φ in the two rows for the larger values of σ_η. While the rejection probabilities for the diagnostic test are not large for any rows of the table, the larger rejection probabilities occur for the higher values of agent aggressiveness θ. These results point in the direction of a more extensive search over the generalized family of ARCH models.

Overall, the results in Table 2 demonstrate that letting agents choose among competing forecasts can produce ARCH/GARCH effects that are very similar to those found in data for financial markets. We can decompose the parameter combinations of martingale volatility and agent aggressiveness into three regions. For low martingale volatility and/or low agent aggressiveness, the process is nearly always at the fundamentalist solution. For somewhat higher martingale volatility and/or agent aggressiveness, mysticism becomes an important factor, the data frequently exhibit symptoms of ARCH, and a GARCH(1,1) model often fits the data very well. For very high martingale volatility and very aggressive agents, Engle's (1982) test for ARCH rejects the null hypothesis decisively, but the volatility in the system cannot always be adequately modeled by a GARCH(1,1) model. In this situation, econometricians committed to an ARCH approach would likely construct a model more complex than GARCH(1,1).

10 Conclusion

While ARCH/GARCH models have proved to be extremely successful empirical econometric techniques, explaining the underlying causes of conditional volatility in financial markets has been a difficult challenge. This paper presents a formal model explaining how such effects arise endogenously when forward-looking agents choose among forecasting strategies.

The leading candidate for the source of conditional volatility has long been news of some kind (Engle (2001, 2004)). Our results are driven by the arrival of new information, but they do not require any assumption that information arrives at nonuniform rates. We show instead that the process of experimenting with and rejecting sources of information can be a key factor explaining conditional volatility. In fact, our parameter specifying how aggressively agents search for the optimal forecasting strategy is a primary factor in accounting for conditional volatility.

Our results do not require any radical departure from rationality. We consider only forward-looking mean-variance optimizing agents who differ only because they choose among three forecasting strategies that are all consistent with the notion of rational expectations. The fundamentalists use the rational expectations solution dominant in that literature. Mysticism follows the same principles, but mistakenly experiments with the idea that the martingale innovations are fundamental. Both fundamentalism and mysticism are, of course, only fully rational if they are adopted by all agents. The reflectivists focus on this point and adopt a forecast taking full account of the behavior of the other two groups of agents.

The characteristics of the realized asset price differ dramatically across parameter values. If the martingale variance is small or agents are not very aggressive in pursuing the optimal forecasting strategy, agents tend to agree numerically on the fundamentalist forecast and there is little evidence of volatility clustering. For larger martingale variances and/or more aggressive agents, ARCH/GARCH effects appear in significant fractions of the sample runs.

Econometric tests of the simulated data in the latter cases detect ARCH and GARCH effects similar to those found in financial market data. We test for ARCH using Engle's (1982) test and then estimate the GARCH(1,1) model often used in practice. For a range of martingale volatility and a range of agent aggressiveness, the sample statistics indicate that the simulated data is well represented by a GARCH(1,1) model. For the combination of a large martingale innovation variance and very aggressive agents, the evidence points in the direction of ARCH, but in a form more complicated than a GARCH(1,1) model. As we propose in the introduction to this paper, these results confirm that empirical ARCH/GARCH effects can be an artifact of viewing data generated by heterogeneous expectations from the perspective of a model that assumes a single expectation.

One implication of these results is that a test for ARCH can be viewed as a specification test for the assumption that agents agree on a single expectation. ARCH will be observed if the levels of martingale volatility and agent aggressiveness are high enough to make divergent expectations a common feature of the data. From this perspective, the widespread econometric evidence in favor of ARCH/GARCH for variables such as inflation, interest rates, exchange rates, and returns on financial assets presents a challenge to the assumption that agents in models explaining these variables agree on a single expectation.

Appendix

To derive the approximate equations of motion (27) and (28), it proves convenient to turn (24) upside-down to produce

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

For β_t this becomes

$$\frac{\beta_t}{\beta_{t+1}} = \frac{\beta_t exp(\theta^2 \pi_{\beta,t}) + \gamma_t exp(\theta^2 \pi_{\gamma,t}) + \lambda_t exp(\theta^2 \pi_{\lambda,t})}{exp(\theta^2 \pi_{\beta,t})},$$

which can be written as

$$\frac{\beta_t}{\beta_{t+1}} = \beta_t + \gamma_t exp(\theta^2(\pi_{\gamma,t} - \pi_{\beta,t})) + \lambda_t exp(\theta^2(\pi_{\lambda,t} - \pi_{\beta,t})).$$

A two-term Taylor series approximation to each exponential function (about $A_t = 0$) yields

$$\frac{\beta_t}{\beta_{t+1}} \cong 1 + \gamma_t(\theta^2(\pi_{\gamma,t} - \pi_{\beta,t})) + \lambda_t(\theta^2(\pi_{\lambda,t} - \pi_{\beta,t})) + \frac{\gamma_t}{2}(\theta^2(\pi_{\gamma,t} - \pi_{\beta,t}))^2 + \frac{\lambda_t}{2}(\theta^2(\pi_{\lambda,t} - \pi_{\beta,t}))^2.$$

Substituting

$$\pi_{\gamma,t} - \pi_{\beta,t} = -2n_{t-1}A_{t-1}U_t - n_{t-1}^2 A_{t-1}^2$$

and

$$\pi_{\lambda,t} - \pi_{\beta,t} = 2\left(1 - n_{t-1}\right) A_{t-1}U_t - \left(1 - n_{t-1}\right)^2 A_{t-1}^2$$

from (17), (18), and (19) introduces a variety of terms involving A_{t-1} and U_t. We retain the terms involving A_{t-1}^2 and $A_{t-1}^2 U_t^2$. The terms involving $A_{t-1}U_t$ cancel. We leave out the terms involving $A_{t-1}^3 U_t$ on the grounds that their expectation is near zero.[24] We leave out the terms involving A_{t-1}^4 on the grounds that they will be dominated by the terms involving A_{t-1}^2 for small variances of the martingale innovation, especially early in the history of the martingale.[25] Our approximate

[24] The weights in the martingale innovation in G_t change between period $t-1$ and period t, leaving open a tenuous possibility of some correlation with the lagged martingale.

[25] We consider values of σ_η as large as 4.0, and this approximation does not apply to such large martingale innovation variances. This point has no effect on our results because the simulations use the complete nonlinear equations of motion, not the approximations in this appendix.

equation describing the motion of β_t is thus

$$\frac{\beta_t}{\beta_{t+1}} \cong 1 - (1-\beta_t)(\theta^2 n_{t-1}\,(1-n_{t-1})\,A_{t-1}^2 + 2\theta^4 n_{t-1}\,(1-n_t)\,A_{t-1}^2 U_t^2).$$

The equation of motion (28) for n_t follows from

$$n_{t+1} = \frac{\lambda_t w(\pi_{\lambda,t})}{\lambda_t w(\pi_{\lambda,t}) + \gamma_t w(\pi_{\gamma,t})}.$$

If $n_{t+1} > 0$ and $w(\pi_{\lambda,t}) > 0$, some rearrangement yields

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

(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. A one-term Taylor series approximation about $\pi_{\beta,t}$ for each exponential function yields

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

The factor w'/w measures curvature of the payoff weighting function, which we identify as agent aggressiveness. The potential for large changes in n_t depends on this measure of agent aggressiveness.

References

Binmore, K., Samuelson, L., 1999. Evolutionary drift and equilibrium selection. Review of Economic Studies 66, 363-393.

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.

Blume, L.E., Easley, D., 1992. Evolution and market behavior. Journal of Economic Theory 58, 9-40.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31(3), 307-327.

Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance: a review of the theory and empirical evidence. Journal of Econometrics 52, 5-59.

Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. Econometrica (65), 1059-1095.

Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics and Control 22, 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.W., Deadman, D.F., 1995. Speculative bubbles with stochastic explosive roots: the failure of unit root testing. Journal of Empirical Finance 2, 153-163.

Chiarella, C., He, X., 2002. Heterogeneous beliefs, risk and learning in a simple asset pricing model. Computational Economics 19, 95-132.

Cochrane, J.H., 2001. Asset Pricing. Princeton University Press, Princeton, NJ.

DeGrauwe, P., 1993, Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets. Blackwell, Cambridge, MA.

DeLong, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. Journal of Political Economy 98, 703-738.

den Haan, W.J., Spear, S., 1998. Volatility clustering in real interest rates: theory and evidence. Journal of Monetary Economics 41, 431-453.

Diebold, F.X., Nerlove, M., 1989. The dynamics of exchange rate volatility: a multivariate latent-factor ARCH model. Journal of Applied Econometrics 4, 1-22.

Enders, W., 2004. Applied Economic Time Series, Wiley, Hoboken, NJ.

Engle, R.F., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U.K inflation. Econometrica 50, 987-1008.

Engle, R.F., 2001. GARCH 101: the use of ARCH/GARCH models in applied econometrics. Journal of Economic Perspectives 15(4), 157-168.

Engle, R.F., 2004. Risk and volatility: econometric models and financial practice. American Economic Review 94(3), 405-420.

Engle, R.F., Lilien, D.M., Robbins, R.P., 1987. Estimating time varying risk premia in the term structure: the ARCH-M model. Econometrica 55, 391-407.

Evans, G.W., 1992. Pitfalls in testing for explosive bubbles in asset prices. American Economic Review 81, 922-30.

Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton, NJ.

Evans, G.W., Guesnerie, R., 2003. Coordination on saddle path solutions: the eductive viewpoint – linear multivariate models. Macroeconomic Dynamics 7(1), 42-62.

Farmer, R.E.A., 1999. The macroeconomics of sulf-fulfilling prophecies, second edition. The MIT Press, Cambridge, Massachusetts

Föllmer, H, Horst, U., Kirman, A., 2005. Equilibria in financial markets with heterogeneous agents: a probabilistic perspective. Journal of Mathematical Economics 41, 123-155.

Gaunersdorfer, A., Hommes, C.H., 2006. A nonlinear structural model for volatility clustering. Forthcoming in G. Teyseeiere and A. Kirman, eds., Long Memory in Economics, Springer Verlag.

Gaunersdorfer, A., Hommes, C.H., Wagener, F.O.O., 2003. Bifurcation routes to volatility clustering under evolutionary learning, working paper.

Geweke, J. and Porter-Hudak, S., 1983, The Estimation and Application of Long Memory Time Series Models. Journal of Time Series Analysis 4, 221-238.

Grandmont, J.-M., 1998. Expectations formation and stability of large socioeconomic systems. Econometrica 66, 741-781.

Hall, S.G., Psaradakis, Z., Sola, M., 1999. Detecting periodically collapsing bubbles: a markov-switching unit root test. Journal of Applied Econometrics 14(2), 143-154.

Hofbauer, J., Weibull, J.W., 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. H., 2006. Heterogeneous agent models in economics and finance. Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics, Edited by L. Tesfatsion and K.L. Judd, Elsevier Science B.V., in press.

Howitt, P., McAfee, R.P., 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, NY.

LeBaron, B., 2006. Agents-based computational finance. Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics, Edited by L. Tesfatsion and K.L. Judd, Elsevier Science B.V., in press.

LeBaron, B., 2000. Agents-based computational finance: suggested readings and early research. Journal of Economic Dynamics and Control 24, 679-702.

LeBaron, B., Arthur, W.B., Palmer, R., 1999. Time series properties of an artificial stock market. Journal of Economic Dynamics and Control 23, 1487-1516.

Lux, T., Marchesi, M., 2000. Volatility clustering in financial markets: a microsimulation of interacting agents. International Journal of Theoretical and Applied Finance 3, 675-702.

Marcet, A., Sargent, T.J., 1989. 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.

McCallum, B.T., 1997. The role of the minimum state variables criterion in rational expectations models. International Journal of Tax and Finance 6(4), 621-639.

Pagan, A., 1996. The econometrics of financial markets. Journal of Empirical Finance 3, 15-102.

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

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

Woodford, M., 1990. Learning to Believe in Sunspots. Econometrica 58, 277-307.

Figure 1

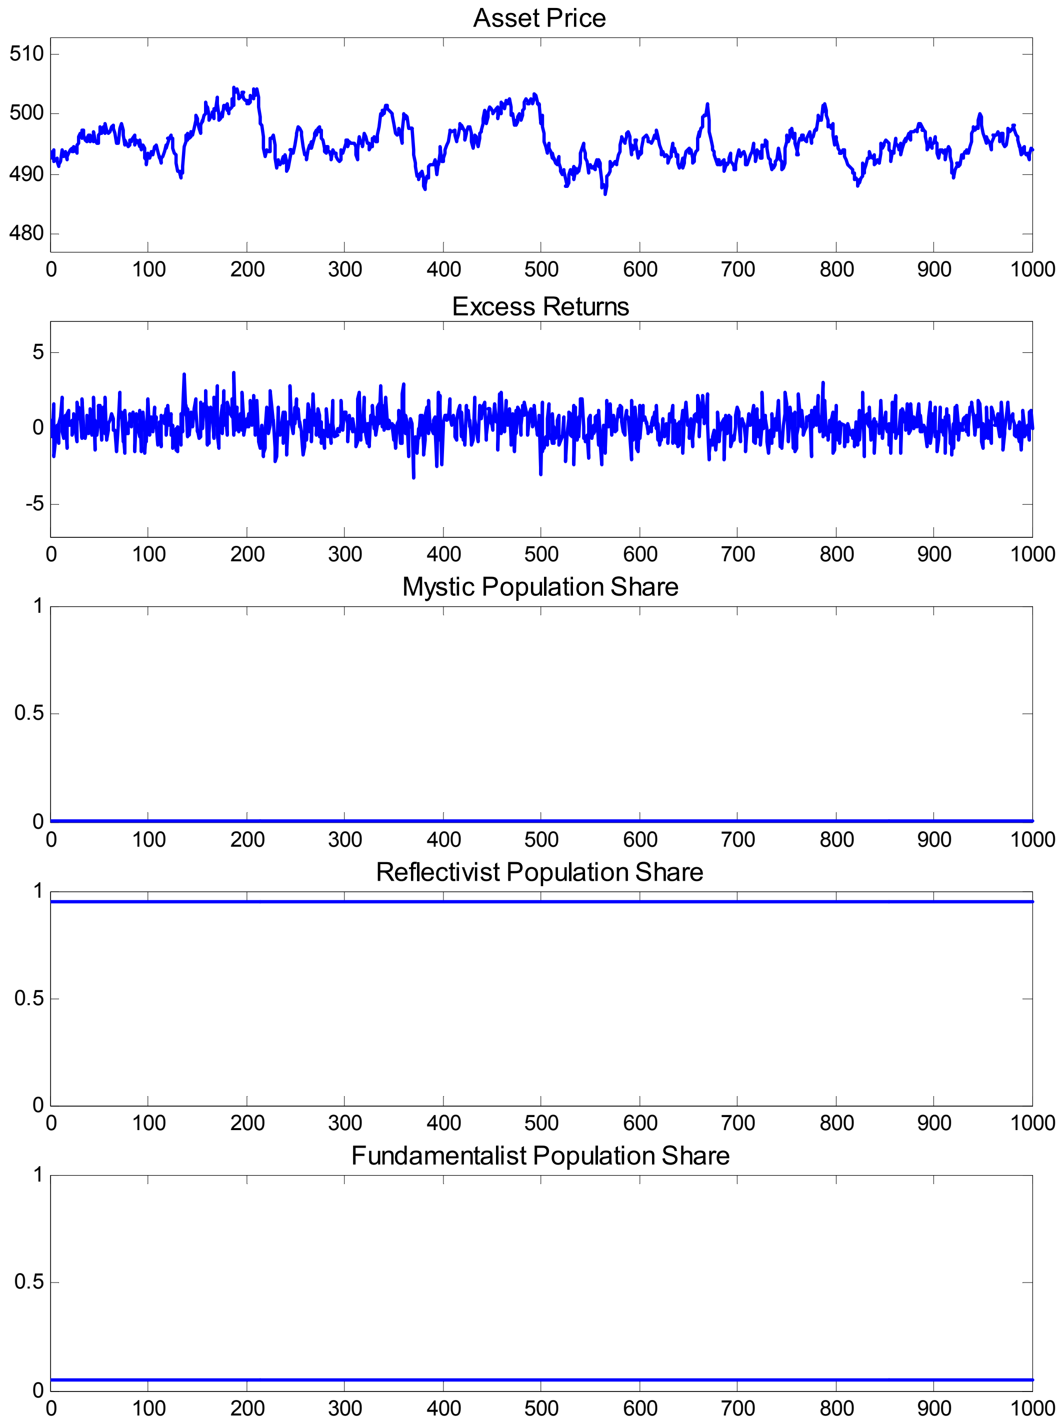

$\sigma_F = 1.0,\ \sigma_\eta = 1.0,\ \theta = 0.5$

Figure 2

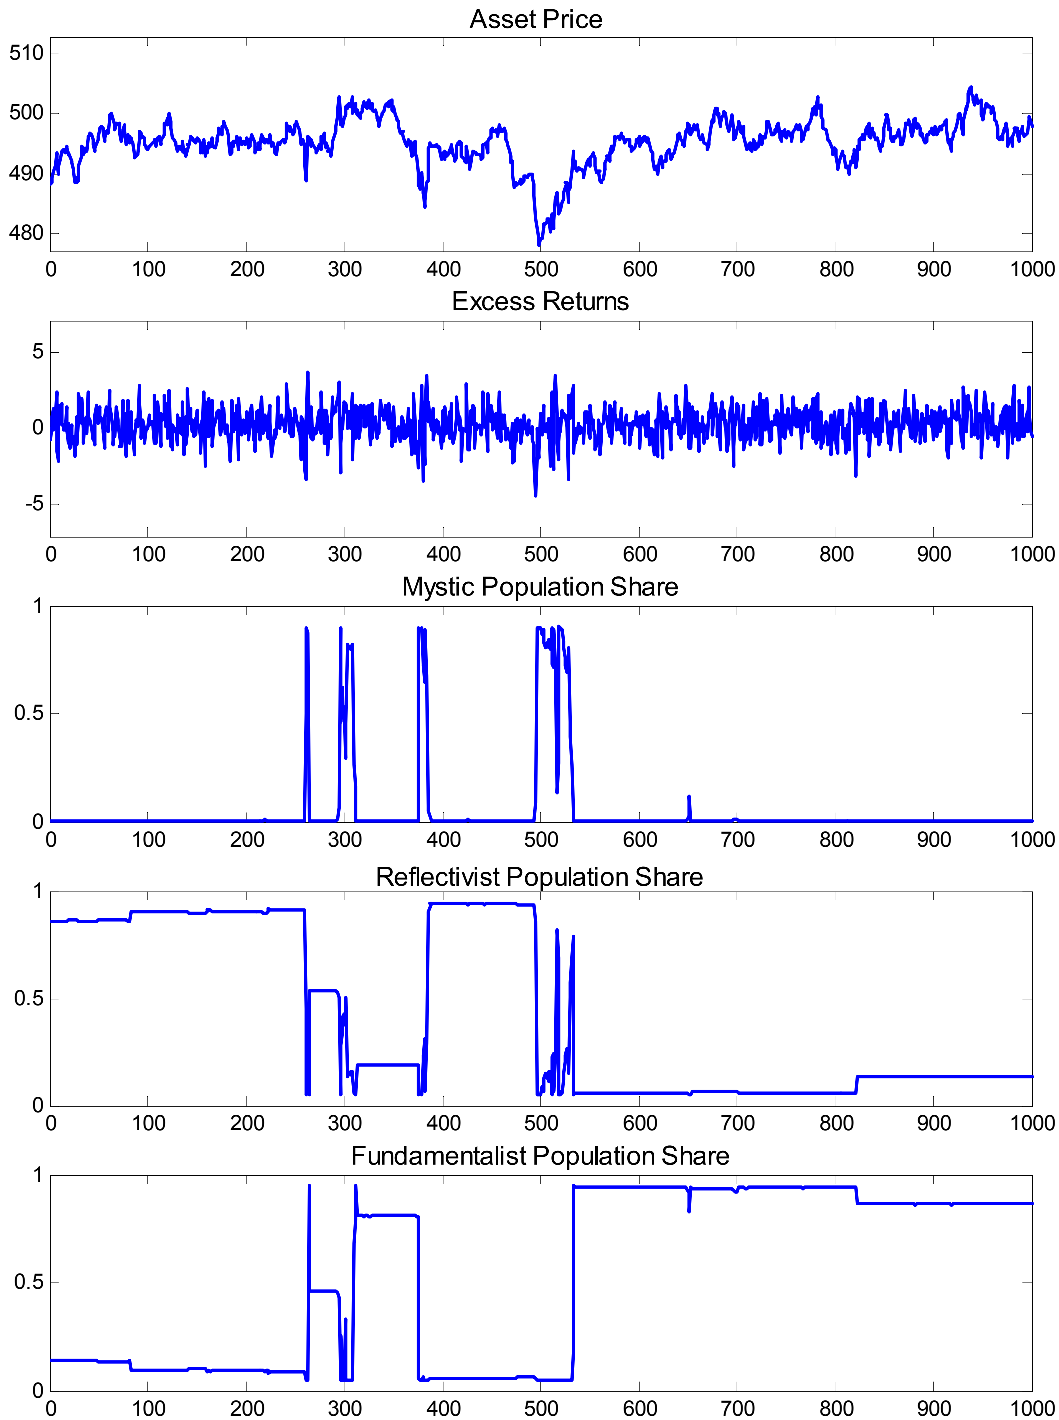

$\sigma_F = 1.0$, $\sigma_\eta = 1.0$, $\theta = 1.0$

Figure 3

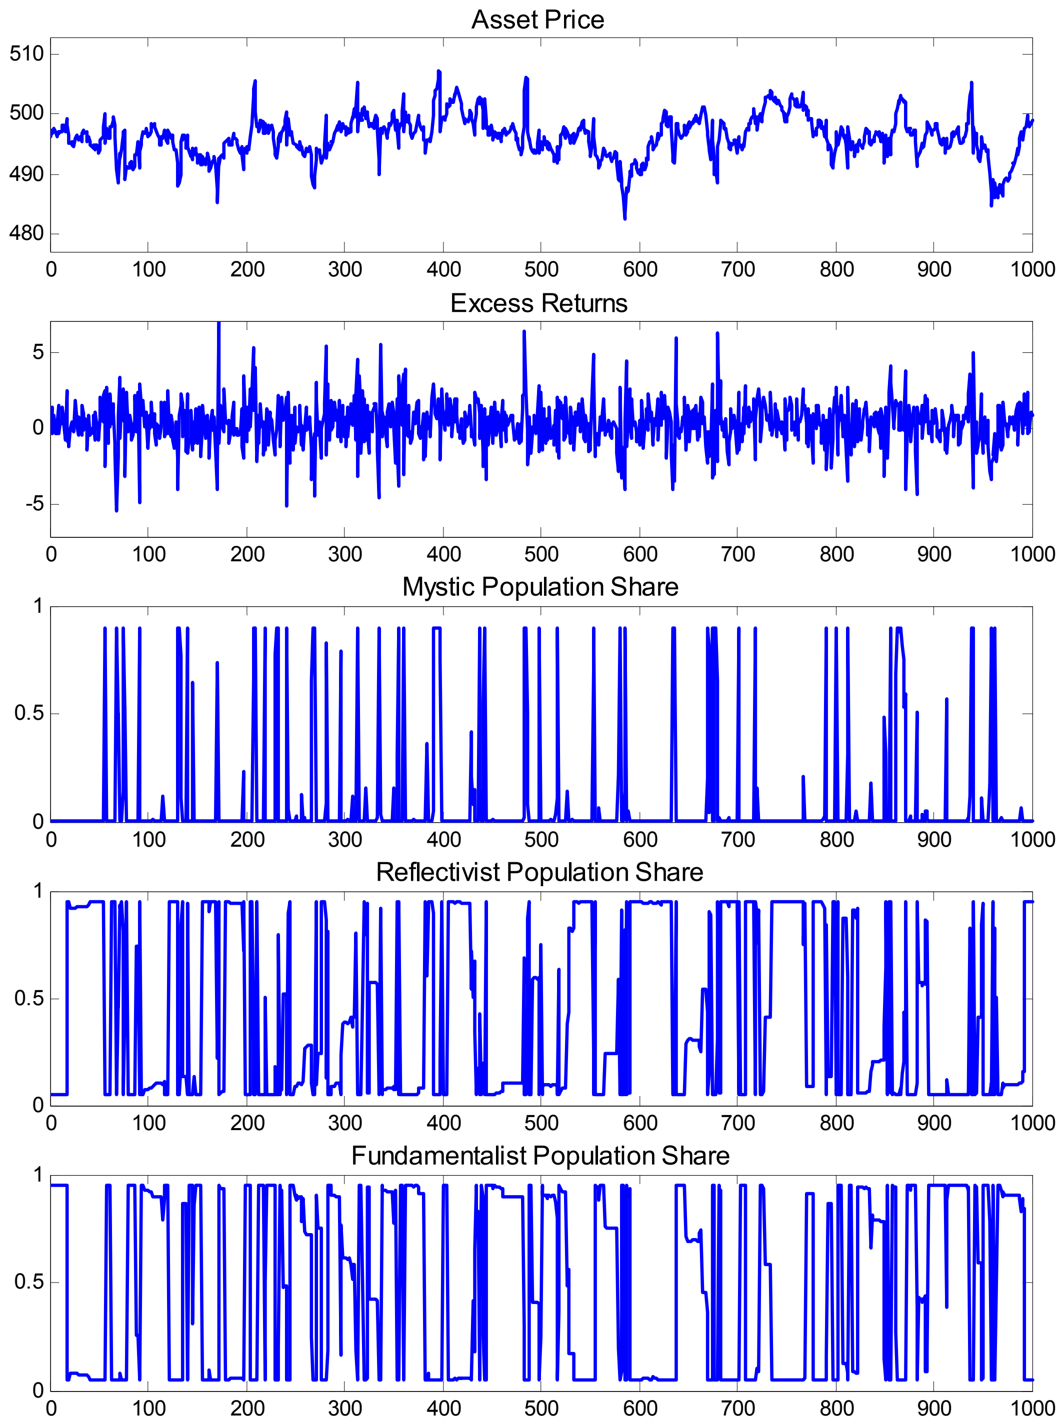

$\sigma_F = 1.0,\ \sigma_\eta = 2.0,\ \theta = 2.0$

Table 1
Volatility Clustering and Excess Kurtosis

	σ_η	Prob of Rejecting H_o: no ARCH	Long Memory Parameter d (sq. exc. ret.)	Average Kurtosis	Average Population Shares Fund. γ_t	Myst. λ_t	Refl. β_t
θ=1/2							
	0.5	0.042	0.000	3.00	0.051	0.000	0.949
	1	0.043	0.000	3.00	0.052	0.000	0.948
	2	0.042	-0.001	3.00	0.053	0.000	0.947
	4	0.046	0.000	3.01	0.055	0.000	0.945
θ=1							
	0.5	0.098	0.023	3.18	0.531	0.015	0.454
	1	0.282	0.055	3.84	0.609	0.013	0.378
	2	0.643	0.029	5.75	0.659	0.007	0.334
	4	0.700	0.007	9.02	0.649	0.003	0.348
θ=2							
	0.5	0.156	0.005	3.31	0.300	0.102	0.599
	1	0.821	0.008	4.68	0.364	0.108	0.529
	2	0.993	-0.002	7.64	0.412	0.070	0.519
	4	0.997	-0.009	16.64	0.374	0.033	0.593
θ=4							
	0.5	0.060	0.002	3.15	0.136	0.103	0.762
	1	0.513	0.001	4.03	0.168	0.118	0.715
	2	0.992	0.001	6.76	0.168	0.083	0.749
	4	1.000	0.000	14.88	0.136	0.044	0.820

The first column is the fraction of the runs where the null of no ARCH is rejected at the 5% significance level. The second and third columns are the average long memory parameter d for squared excess returns and the average kurtosis. The 3 columns on the right show the average followings for each strategy, where the average is over both time periods and simulation trials. Each row is based on 10,000 simulations.

Table 2

Estimates of GARCH(1,1) Models

σ_η	Conditional Prob of Rejecting $H_o: \varphi=0$	Mean of Estimates of φ	Conditional Prob of Rejecting $H_o: \psi=0$	Mean of Estimates of ψ	Conditional Rejection Prob Q test
$\theta = 1/2$					
0.5	0.150	0.771	0.333	0.044	0.000
1	0.087	0.731	0.250	0.048	0.000
2	0.139	0.742	0.200	0.041	0.200
4	0.182	0.622	0.300	0.067	0.000
$\theta = 1$					
0.5	0.258	0.747	0.708	0.053	0.042
1	0.483	0.760	0.884	0.060	0.055
2	0.672	0.642	0.973	0.085	0.055
4	0.601	0.509	0.988	0.138	0.059
$\theta = 2$					
0.5	0.348	0.696	0.704	0.064	0.111
1	0.350	0.442	0.997	0.178	0.080
2	0.702	0.198	1.000	0.493	0.039
4	0.586	0.104	1.000	0.838	0.048
$\theta = 4$					
0.5	0.258	0.822	0.375	0.034	0.000
1	0.193	0.621	0.790	0.087	0.084
2	0.222	0.223	0.996	0.390	0.086
4	0.281	0.110	1.000	0.859	0.075

The first column is the fraction of the runs where the GARCH parameter φ is significantly positive, conditional on the presence of ARCH effects. The following column shows the mean of the estimates of φ for the runs where ARCH effects are present and φ is significantly positive. The next column shows the fraction of the runs where the ARCH parameter ψ is significantly positive, conditional on ARCH effects being present and φ being significantly positive. The following column shows the mean of the estimates of ψ where ARCH effects are present and ψ and φ are significantly positive. The last column shows the fraction of the runs rejected by the diagnostic Q test among those that had ARCH effects and a significantly positive φ. Each row is based on 1,000 simulations.