Macroeconomic Dynamics, **4**, 2000, 139–169. Printed in the United States of America.

ARTICLES

INFORMATION DYNAMICS IN FINANCIAL MARKETS

PATRICK DE FONTNOUVELLE
Iowa State University

A noisy rational expectations model of asset trading is extended to incorporate costs of information acquisition and expectation formation. Because of the information costs, how much information to acquire becomes an important decision. Agents make this decision by choosing an *expectations strategy* about the future value of information. Because expectation formation is costly, agents often choose strategies that are simpler (and thus cheaper) than rational expectations. The model's dynamics can be expressed in terms of the *market precision*, which represents the amount of information acquired by the average agent. Under certain conditions, market precision follows an unstable and highly irregular time path. This irregularity directly affects observable market quantities. In particular, simulated time series for return volatility and trading volume display a copersistence similar to that found in actual financial data.

Keywords: ARCH, Asymmetric Information, Trading Volume, Noisy Rational Expectations

1. INTRODUCTION

If information about an asset's value is costly, can this value be revealed fully in a competitive equilibrium? Grossman and Stiglitz (1980) show that when the asset in question is a security, the answer is clearly no; in equilibrium, information gatherers must be rewarded for their costly private activities. Suppose that the asset in question is information itself. If computing the value of information is costly, can this value be reflected in an observable market quantity? The question is new to this paper. We show that the answer is also no; if computation is costly, then those who compute the true value of information must be rewarded for their efforts.

The above chain of logic leads to a startling conclusion: Not all agents can have rational expectations about the value of information. If they did, the average quantity of information acquired (which is observable) would reflect its true value.

I would like to thank the Iowa State University Department of Economics for supporting this research while I was an assistant professor there. I also thank William Brock, Cars Hommes, Harvey Lapan, Blake LeBaron, and an anonymous referee for their helpful comments. Address correspondence to: Patrick de Fontnouvelle, U.S. Securities and Exchange Commission, Office of Economic Analysis, Mail Stop 11-5, 450 5th St. NW, Washington, DC 20549, USA; e-mail: defontnouvellep@sec.gov. The Securities and Exchange Commission, as a matter of policy, disclaims responsibility for any private publication or statement by any of its employees. The views expressed herein are those of the author and do not necessarily reflect the views of the Commission or of the author's colleagues upon the staff of the Commission.

 1365-1005/00 $9.50

Agents would then prefer to infer the value of information directly from this quantity, rather than making costly rational expectations computations. This cannot be an equilibrium.

If some agents do not use rational expectations to infer the value of information, what expectations do they use? Recent work on bounded rationality and evolution in economics [e.g., Sargent (1993), Marimon (1997), Evans and Honkapohja (1999)] inspires our answer to this question. Of particular relevance is the complex adaptive systems modeling approach laid out by Holland (1995). Applications of this approach to financial markets include those by Arthur (1995), Arthur et al. (1997), and LeBaron et al. (1999).

Besides forming rational expectations, agents in our model experiment with various other *expectations strategies* that are computationally cheaper than rational expectations. If a strategy generates profits, many agents will use it in the future. If not, few agents will use it. Instead of being governed by a single costless rational expectations strategy, the market's information structure depends upon a population of competing and coevolving expectations strategies.

To formalize the above discussion, we extend a noisy rational expectations model of asset trading [Hellwig (1980); Lang et al. (1992)] to explicitly incorporate costs of information acquisition and expectation formation. Acquiring information is costly because it requires time-consuming research activities such as fundamental analysis. Forming expectations is costly because it involves lengthy computations; rational expectations is the most computationally intensive strategy, and thus the most costly as well.[1]

The agents' decisions of how much private information to acquire (their *information strategies*) depend on the value of this information. They estimate this value by choosing an expectations strategy. The model's dynamics can be expressed in terms of a single variable, *market precision*, which represents the amount of information acquired by the average agent. Market precision is driven by two opposing forces. Agents' desire to reduce computation costs makes them use inexpensive and destabilizing expectations strategies, which push market precision away from equilibrium. Agents' desire to increase trading profits makes them use rational expectations, which pull market precision toward equilibrium. These opposing forces generate a local instability around the equilibrium, which can result in highly irregular time paths for market precision.

The metric for assessing the economic significance of these mathematical effects is empirically inspired. We examine how well the model explains certain empirical features that have proved persistently puzzling to economic theorists: serial correlation in squared returns, serial correlation in trading volume, and correlation between squared returns and volume. These features are referred to jointly as *copersistence* between volatility and trading volume.

Under a rational expectations equilibrium, the degree of information asymmetry in the model does not vary, and thus volatility and trading volume are both constant. Allowing agents to change expectations strategies over time generates variation in market precision, which in turn has a direct effect on trading activity. When market precision is low, trading volume is low because most agents have similar

information. Volatility is also low because there is little private information driving asset prices. Conversely, a high level of market precision will result in high volatility and high trading volume. Under certain configurations of the model, the simulated time series for volatility and trading volume display copersistence similar to that found in actual market data.

The remainder of the paper is organized as follows. Section 2 extends the noisy rational expectations model to include costs of information acquisition and expectation formation. Section 3 summarizes the empirical behavior of volatility and trading volume, and derives the model's implications for these quantities. Section 4 shows that, under certain conditions, the model displays a local instability around the steady state. Section 5 explores the model's dynamics when there are only two expectations strategies. An enlarged strategy space is explored in Section 6. Section 7 concludes.

2. MODEL

The model's construction can be divided into two distinct steps. The first consists of specifying the mechanism through which information asymmetry affects trading dynamics; we use for this step the noisy rational expectations model of Hellwig (1980) and Lang et al. (1992). In this model, agents have private information about the future value of a security. Because the supply of shares is stochastic, the market price does not fully reveal this private information; information is thus a valuable commodity.

We make our main theoretical contribution in the second step, which consists of specifying how the degree of information asymmetry evolves over time. Brock and Hommes (1997; BH) consider a linear cobweb model of commodity supply and demand in which agents can base their views about the commodity's future value on either rational or naive expectations. These authors demonstrate how the interaction of the two types of expectations can lead to highly irregular equilibrium paths. Suppose one thinks of information as a commodity, and of agents as producers of this commodity. Applying a similar framework (as BH) to the process of information production yields a rich yet relatively tractable model of information dynamics.

Goeree and Hommes (2000; GH) extend the results of BH to include general nonlinear supply and demand functions. Their results, in fact, apply to the rational versus naive example presented in Section 5.[2] Whereas GH limit their analysis to two competing expectations strategies, Section 6 of this paper allows for a large number of expectations strateies, which are crucial in explaining the observed behavior of volatility and volume.

2.1. Asset Trading

There is a continuum of agents of measure 1. Asset trading takes place at discrete time periods indexed by $t \in \{1, 2, \ldots, \infty\}$. In each period, there is available one riskless asset and one risky asset. The riskless asset guarantees a rate of return R.

The risky asset has price p_t and pays a single terminal dividend d_t in period $t+1$. Risky assets thus are short-lived,[3] and there is a unique risky asset available in each period t. Agent i's *trading information set* contains the market price p_t and private information $y_{i,t}$:

$$I_{i,t} = \{p_t, y_{i,t}, p_{t-1}, y_{i,t-1}, \ldots\}.$$

This private information consists of the terminal dividend d_t plus an observation error $\epsilon_{i,t}$:

$$y_{i,t} = d_t + \epsilon_{i,t}.$$

The supply of risky assets is a random variable z_t. We assume that d_t, z_t, $\epsilon_{i,t}$, and $\epsilon_{j,t}$ are independent for all i, j, and t. Furthermore, each of these variables is normally distributed:

$$d_t \sim N(\bar{d}, \sigma_d),$$
$$z_t \sim N(\bar{z}, \sigma_z),$$
$$\epsilon_{i,t} \sim N(0, 1/\rho_i).$$

The variance of the private information $y_{i,t}$ is a choice variable for each agent. Choosing a high value for ρ_i will give agent i an informational advantage, and hence higher expected trading profits. Because acquiring information is time-consuming, there is a cost associated with choosing a high ρ_i. The optimal choice of ρ_i thus will entail weighing increased trading profits against increased information costs.

Suppose for now that ρ_i is fixed. Agent i's problem is to choose an amount $x_{i,t}$ to invest in the risky asset. The resulting portfolio generates a trading profit $\pi_{i,t} = x_{i,t}(d_t - Rp_t)$. The agent chooses $x_{i,t}$ in order to maximize expected risk-adjusted trading profits,[4] denoted by $\Pi_{i,t}$:

$$\Pi_{i,t} = E[\pi_{i,t} \mid I_{i,t}] - \frac{r}{2}\operatorname{Var}[\pi_{i,t} \mid I_{i,t}], \tag{1}$$

where r is the risk-aversion parameter. Agent i's demand for the risky asset is then given by

$$x_{i,t} = \frac{E(d_t - Rp_t \mid I_{i,t})}{r \operatorname{Var}(d_t \mid I_{i,t})}. \tag{2}$$

Substituting this demand into (1) yields the expression for agent i's expected risk-adjusted profits as a function of her time t information:

$$\Pi_{i,t} = \frac{E^2(d_t - Rp_t \mid I_{i,t})}{2r \operatorname{Var}(d_t \mid I_{i,t})}$$

2.2. Information Strategy

We refer to ρ_i, the reciprocal of the variance of agent i's private information, as agent i's *information strategy*. We model information costs as an increasing

function $c_\rho(\rho_i)$, and refer to the process of choosing an information strategy (or an expectations strategy) as *strategy revision*.

Asset trading and strategy revision take place on two different time scales.[5] One might imagine that asset trading takes place on a tick-by-tick frequency, and strategy revision on a daily frequency. While agents trade on a tick-by-tick frequency to generate trading profits, they have time to reevaluate their information sources only by night when the markets are closed. We model these two timescales by assuming that all agents revise their strategies at the same times $\{T, 2T, 3T, \ldots, NT, \ldots\}$, $T \geq 1$. This is called synchronous adjustment.[6] Denote asset trading time by t, and strategy revision time by τ, so that one unit of τ-denoted time spans T units of t-denoted time. Quantities relating to asset trading thus are labeled with a t subscript (e.g., $x_{i,t}$), and quantities relating to strategy revision with a τ subscript (e.g., $\rho_{i,\tau}$).

There is a lag of κ periods between the time agent i revises her information strategy and the time at which the variance of $y_{i,t}$ changes to reflect this revision. This lag represents the time required to make changes such as hiring more analysts or subscribing to new commercial information services.

Let $\Pi_{i,\tau}$ denote expected profits averaged over the entire strategy revision period τ:

$$\Pi_{i,\tau} = E[\Pi_{i,t}], \qquad \tau T \leq t < (\tau + 1)T,$$

where $E[\cdot]$ denotes the unconditional expectation. Appendix B shows how to write $\Pi_{i,\tau}$ as a function of time τ variables:

$$\Pi_{i,\tau} = \Pi(B_\tau, \rho_{i,\tau}) = \left[\left(\rho_{i,\tau} + 1/\sigma_d + B_\tau^2/\sigma_z\right)\Lambda(B_\tau) - 1\right]/2r, \tag{3}$$

$$\Lambda(B) \equiv \frac{\sigma_d\sigma_z^2/r^2 + \sigma_z(\sigma_d\sigma_z + \sigma_d B/r)^2 + \overline{z}^2\sigma_d^2\sigma_z^2}{\left(\sigma_z/r + \sigma_d\sigma_z B + \sigma_d B^2/r\right)^2},$$

$$B_\tau \equiv E[\rho_{i,\tau}]/r, \tag{4}$$

where the above expectation is taken across all agents i.

The quantity B_τ, which we refer to as *market precision*, has a straightforward intuitive interpretation. To choose her information strategy $\rho_{i,\tau}$, agent i would like to know the information strategies chosen by the other agents. If most agents choose a precise information strategy (high ρ_τ), then the market price p_t will provide a good signal about the terminal dividend d_t. Because the market price is observed costlessly, the added value of private information will be low. Agent i thus would like to choose an imprecise information strategy (low $\rho_{i,\tau}$). Conversely, if most agents choose an imprecise information strategy, agent i would like to choose a precise strategy. Equation (3) shows that all information relevant to choosing $\rho_{i,\tau}$ can in fact be summarized in one quantity—market precision—which simply represents the information strategy of the average agent.

Because of the κ period lag in information strategy revision, agent i must choose in period τ her information strategy for period $\tau + \kappa$. She wishes to choose $\rho_{\tau+\kappa}$

in order to maximize her profit minus her information costs:

$$\max_{\rho_{i,\tau+\kappa}} \Pi(B_{\tau+\kappa}, \rho_{i,\tau+\kappa}) - c_\rho(\rho_{i,\tau+\kappa}). \tag{5}$$

At time τ, however, agent i does not know the value of $B_{\tau+\kappa}$, and thus cannot evaluate (5) directly. She instead forms a prediction of $B_{\tau+\kappa}$, which we denote $B^e_{i,\tau+\kappa}$, and uses this prediction to approximate the maximization problem (5) as follows[7]:

$$\max_{\rho_{i,\tau+\kappa}} \Pi\left(B^e_{i,\tau+\kappa}, \rho_{i,\tau+\kappa}\right) - c_\rho(\rho_{i,\tau+\kappa}). \tag{6}$$

The first-order condition for (6) is

$$c'_\rho(\rho_{i,\tau+\kappa}) = \Lambda\left(B^e_{i,\tau+\kappa}\right)/(2r), \tag{7}$$

which implicitly defines the information strategy $\rho_{i,\tau+\kappa}$ as a function of expected market precision:

$$\rho_{i,\tau+\kappa} = \rho\left(B^e_{i,\tau+\kappa}\right), \tag{8}$$

where $\rho(B) = (c'_\rho)^{-1}[\Lambda(B)/(2r)]$.

2.3. Expectations Strategy

To use (8) to choose an information strategy, agent i must form $B^e_{i,\tau+\kappa}$, her time τ prediction of time $\tau+\kappa$ market precision. Agents form this prediction by choosing an *expectations strategy* H_j. To be more precise, define the *strategy information set* I_τ to contain common (public) information that all agents observe at the end of period τ; it is straightforward to show that I_τ contains a record of realized market precisions[8]:

$$\{B_\tau, B_{\tau-1}, \ldots\} \subset I_\tau.$$

An expectations strategy H_j is a function mapping I_τ into a predicted market precision $B^e_{j,\tau+\kappa}$. The space of all expectations strategies is denoted $\mathcal{H}$:

$$\mathcal{H} = \{H_1, H_2, \ldots, H_K\}. \tag{9}$$

Each H_j in $\mathcal{H}$ has an associated cost function $c_H(H_j)$; rational expectations is the most complicated and thus the most costly expectations strategy.

Given an expectations strategy H_j, we can write the information strategy and the expected profit measure associated with H_j as functions of lagged market precision:

$$\begin{aligned} \rho_{j,\tau} &= \rho[H_j(I_{\tau-\kappa})], \\ \Pi_{j,\tau} &= \Pi\{B_\tau, \rho[H_j(I_{\tau-\kappa})]\}. \end{aligned} \tag{10}$$

Examples of possible expectations strategies include naive expectations and rational expectations. If agent i has naive expectations, then

$$B^e_{i,\tau+\kappa} = H_{\text{nai}}(I_\tau) = B_\tau.$$

If agent i has rational expectations, then

$$B^e_{i,\tau+\kappa} = H_{\text{rat}}(I_\tau) = B_{\tau+\kappa}.$$

Rational expectations entails perfect foresight because the dynamics of B_τ are deterministic [equation (13)].

The model's dynamics depend crucially upon the specification of the expectations strategy space $\mathcal{H}$. In two special cases, the environment just described reduces to a familiar model, but neither is economically plausible. First suppose that all agents base their information acquisition upon naive expectations, so that $\mathcal{H} = \{H_{\text{nai}}\}$. Whenever current market precision is low, agents will expect future market precision also to be low. They will all choose precise information strategies, which implies that future market precision actually will be high. Likewise, a high current market precision will lead to a low future market precision. The market precision thus follows a periodic cycle, and the model displays classic cobweb behavior. The early rational expectations literature provided a compelling argument against cobweb models: In the presence of large predictable cycles, the benefit to any one agent of choosing a "smarter" expectations strategy would be enormous.[9]

Suppose, on the other hand, that all agents base their information acquisition upon rational expectations, so that $\mathcal{H} = \{H_{\text{rat}}\}$. This model has a unique equilibrium: All agents acquire the same amount of information in every period (market precision is constant). Naive expectations thus would predict future market precision just as accurately as rational expectations. Because forming rational expectations is costly, individual agents have a strong incentive to switch to a simpler and cheaper expectations strategy.

It becomes clear that, in the presence of computation costs, the expectations strategy space $\mathcal{H}$ must contain more than one element in order for the model to be economically sensible. We thus need to specify how agents choose among the various expectations strategies in $\mathcal{H}$. Before choosing an expectations strategy, agents are incapable of forming forward-looking expectations. This is because any forward-looking expectations necessarily are defined by an underlying expectations strategy. Agents thus choose an H_j using backward-looking expectations.

Because risk-adjusted profits $\Pi_\tau(H_j) = \Pi(B_\tau, \rho[H_j(I_{\tau-\kappa})]$ are a function of I_τ, agents know $\Pi_\tau(H_j)$ at the end of period τ for each H_j in $\mathcal{H}$. In other words, the time τ strategy information set contains the profit histories of all strategies $H_j \epsilon \mathcal{H}$:[10]

$$\{\Pi_\tau(\mathcal{H}), \Pi_{\tau-1}(\mathcal{H}), \ldots\} \subset I_\tau.$$

For each strategy H_j, agent i's associated utility (at the beginning of period τ) is composed of a deterministic and a stochastic component[11]:

$$\begin{aligned} U_{i,\tau}(H_j) &= u_j(I_\tau) + \varepsilon_{i,j,\tau}/\beta \\ u_j(I_\tau) &= \Pi_\tau(H_j) - c_\rho\{\rho[H_j(I_{\tau-\kappa})]\} - c_H(H_j). \end{aligned} \tag{11}$$

The deterministic component $u_j(I_\tau)$ is simply the profit measure minus the information and expectation costs. The stochastic components $\varepsilon_{i,j,\tau}$ have an extreme value distribution, and are i.i.d. across agents, expectations strategies, and time.[12]

We are using discrete-choice theory [Manski and McFadden (1981); Anderson et al. (1993)] to model agents' choice of expectations strategy. From a technical perspective, the motivation for the stochastic component $\varepsilon_{i,j,\tau}$ is to reconcile two conflicting objectives of the model. The first objective is to capture, with some degree of realism, a trading environment in which agents have heterogeneous expectations. This objective suggests a complex model with many degrees of freedom. The second objective is tractability, which generally requires a simple model with few degrees of freedom. The random variables $\varepsilon_{i,j,\tau}$ are a way of introducing more degrees of freedom into the simple model without losing tractability.

From an economic perspective, the $\varepsilon_{i,j,\tau}$ represent unmodeled heterogeneity. The motivation here is similar to that behind more conventional applications of discrete-choice theory, such as commuters' choice between modes of transport. Although commuters from the same suburb face similar time and comfort differentials between various modes, they do not all choose the same mode. Similarly, our agents choose different expectations strategies even though they all observe the same profit histories $\{\Pi_\tau(\mathcal{H}), \Pi_{\tau-1}(\mathcal{H}), \ldots\}$. We provide four examples of the unmodeled heterogeneity that the $\varepsilon_{i,j,\tau}$ are meant to capture:

1. *Unobservable characteristics*: An agent may have unmodeled reasons for preferring one particular expectations strategy H_j.
2. *Misspecification*: Certain agents might maximize criteria other than the profit measure $\Pi_\tau(H_j)$.
3. *Limited information*: Certain agents might not know all the available strategies from which to choose in $\mathcal{H}$.
4. *Variation in costs:* There may be variation across agents in information and expectation costs.

From an evolutionary perspective, one may think of discrete choice as a tractable way of modeling the information dynamics as an evolving ecology of expectations strategies, in which each strategy competes against the others for survival. Survival consists of a strategy's being used by agents. In our model, the population of strategies is represented by $\mathcal{H}$, each strategy's fitness by $u_j(\cdot)$, and natural selection by equation (12).[13]

Because the number of agents is infinite and the $\varepsilon_{i,j,\tau}$ are independent across agents, the law of large numbers implies that the fraction of agents choosing expectations strategy H_j is equal to the probability that any individual agent chooses H_j. Because the $\varepsilon_{i,j,\tau}$ have an extreme value distribution, this probability is given by

$$n_{j,\tau} = n_j(\mathcal{H}, I_\tau) = z(\mathcal{H}, I_\tau)e^{\beta u_j(I_\tau)},$$
$$z(\mathcal{H}, I_\tau) = \left[\sum_{k=1}^{K} e^{\beta u_k(I_\tau)}\right]^{-1}. \tag{12}$$

Because (12) specifies how strategies are distributed across agents, we can use equations (4), (9), and (10) to show that the market precision B_τ evolves according to

$$B_\tau = \frac{1}{r}\sum_{j=1}^{K} n_j(\mathcal{H}, B_{\tau-\kappa}, \ldots)\rho[H_j(B_{\tau-\kappa}, \ldots)]. \quad \textbf{(13)}$$

Equation (13) is the main equation of the model; all other quantities of interest can de derived from B_τ.

Because one of the expectations strategies H_j is rational expectations [and $H_{\text{rat}}(I_{\tau-\kappa}) = B_\tau$], B_τ appears in both sides of (13). Solving for B_τ thus is not analytically possible because doing so would require finding the roots of a complicated nonlinear expression. We use numerical methods (the POLYROOT procedure in GAUSS) to solve (13), and rely upon simulation techniques to explore the model's dynamic behavior.

3. VOLATILITY AND TRADING VOLUME

To assess the model's empirical relevance, we investigate whether it can explain three key empirical features found in financial market data:

1. *Persistence in volatility:* Volatility in stock returns displays a high degree of persistence. See Bollerslev et al. (1992) for a survey of the relevant literature.
2. *Persistence in trading volume:* Time series for trading volume display statistically significant serial autocorrelation. This autocorrelation has been documented for indices and individual stocks, as documented by Antoniewicz (1992) and Lamoureux and Lastrapes (1990).
3. *Cross correlation between volatility and trading volume:* Volatility and trading volume are highly contemporaneously correlated, as shown by the numerous studies surveyed by Karpoff (1987).

There has been some recent progress on providing a structural explanation for the dynamics of trading volume; refer to Wang (1994) and references therein. However, there has been little progress on explaining volatility persistence, and virtually none on explaining the joint dynamics of volatility and trading volume. Gallant et al. (1992, p. 202) write, "there seems to be no model with dynamically optimizing, heterogeneous agents that can jointly account for major stylized facts."

3.1. Model's Implications for Volatility and Volume

Hellwig (1980) shows that, in each asset trading period t, the risky asset's price is determined according to

$$p_t = \phi_0(B_\tau) + \phi(B_\tau)d_t - \gamma(B_\tau)z_t, \quad \textbf{(14)}$$

where the functions ϕ_0, ϕ, and γ are defined in Appendix A. If we define volatility in trading period t to be $v_t = (p_t - p_{t-1})^2$, it follows that

$$v_t = [\phi(B_\tau)(d_t - d_{t-1}) - \gamma(B_\tau)(z_t - z_{t-1})]^2. \tag{15}$$

The average volatility over strategy revision period τ is given by the expectation of (15) conditional on B_τ:

$$v_\tau = 2\phi^2(B_\tau)\sigma_d + 2\gamma^2(B_\tau)\sigma_z. \tag{16}$$

As in Lang et al. (1992), we define *trading volume* in period t to be the expected value (taken across all agents) of the absolute change in asset demand between $t-1$ and t:

$$V_t = \frac{1}{2}E|x_{i,t} - x_{i,t-1}|.$$

Define $x_{i,t}^j$ to be the demand for the risky asset of an agent i who has expectations strategy H_j; then, V_t can be rewritten as

$$V_t = \frac{1}{2}\sum_{j=1}^{K} n_{j,\tau-\kappa} E\left|x_{i,t}^j - x_{i,t-1}^j\right|,$$

where again the expectation is taken across all agents. One then can take the average of V_t over the strategy revision period τ to obtain V_τ as a function of the market precision B_τ. These calculations are contained in Appendix C.

4. LOCAL STABILITY OF THE STEADY STATE

Suppose that the space $\mathcal{H}$ of expectations strategies always contains at least two elements, rational expectations (H_{rat}) and naive expectations [$H_{\text{nai}(0)}$]:[14]

$$H_{\text{rat}}(I_{\tau-\kappa}) = B_\tau,$$
$$H_{\text{nai}(0)}(I_{\tau-\kappa}) = B_{\tau-\kappa}.$$

We parameterize the information cost function to be quadratic in ρ_τ (the information strategy):

$$c_\rho(\rho_\tau) = \varphi\rho_\tau^2. \tag{17}$$

The first-order condition (7) for maximizing risk-adjusted profits with respect to ρ implies that

$$\rho_{\tau,\text{rat}} = \Lambda(B_\tau)/(4r\varphi),$$
$$\rho_{\tau,\text{nai}(0)} = \Lambda(B_{\tau-\kappa})/(4r\varphi).$$

Let B^* denote the model's rational expectations equilibrium, which is defined by $B^* = \rho(B^*)$.

Assumption 1. For each expectations strategy H_j, $H_j(B^*, \ldots, B^*) = B^*$.

Assumption 1 requires all expectations strategies in $\mathcal{H}$ to be unbiased at B^*: If market precision has been B^* for a long time, then each strategy will expect it to remain B^* in the future.

PROPOSITION 1. *Suppose that $H_{\text{nai}(0)}$ has the lowest computation cost of any expectations strategy, and that the information cost is quadratic as in* (17). *If Assumption* 1 *holds, then there exist constants β^c and $\{\varphi^c, r^c, \sigma_z^c, \bar{z}^c\}$ such that the steady-state B^* is locally unstable whenever*

(*a*) *intensity of choice is high*: $\beta > \beta^c$,
and one of the following conditions holds:
(*b*) *The cost of information is low*: $\varphi < \varphi^c$;
(*c*) *Risk aversion is low*: $r < r^c$;
(*d*) *There is a high degree of share supply uncertainty*: $\sigma_z > \sigma_z^c$;
(*e*) *The supply of shares is large*: $\bar{z} > \bar{z}^c$.

Proof. See Appendix D. ■

On a heuristic level, the proof proceeds as follows: First, we show that conditions (b)–(e) lead to instability at B^* whenever all agents are restricted to having naive expectations $H_{\text{nai}(0)}$.

Then suppose that B^* is in fact a stable steady state (when agents' expectations are unrestricted), so that B_τ converges to B^* once it falls within a certain basin of attraction. As B_τ approaches B^*, $H_{\text{nai}(0)}$ becomes an increasingly accurate predictor of market precision. Because it is the cheapest strategy, $H_{\text{nai}(0)}$ will eventually dominate all other expectations strategies. If β is large, agents are highly sensitive to profit differences between expectations strategies; almost all agents will thus choose $H_{\text{nai}(0)}$. A continuity argument yields the desired instability result.

To understand the economic content of Proposition 1, it is worth briefly reviewing conditions (a)–(e). In addition to establishing the conditions' plausibility, the discussion highlights what sorts of issues may be addressed (in the future) within our modeling framework:

(a) *High intensity of choice.* The intensity of choice β determines how important the stochastic component is in determining each agent's choice of expectations strategy. A large β implies that the stochastic components play a relatively small role, and that profit is the main motivator behind this choice. Conversely, a small value of β implies that profit plays only a minor role in the choice of expectations strategy. Condition (a) implies that the more agents' decisions are driven by profit differences between strategies, the more likely information-driven instability becomes.

(b) *Low information cost.* Low information costs can be interpreted as representing the increased prevalence of cheap, high-quality electronic information. They also can be interpreted as a proxy for the transparency of finanial market activity, and for rules requiring corporate disclosure of financially significant events. Both interpretations seem plausible, and represent actual conditions in many major financial markets. That

increased transparency and availability of information leads to instability may seem counterintuitive. The result is thus worth reexplaining.

Suppose that in the current time period (period 1), most agents have naive expectations as well as small amounts of private information. Information is scarce and thus will turn out to be quite valuable. Because it is so cheap, the naive agents then will purchase a lot of it in the next period (period 2). Information then loses its value, and so, naive agents purchase even less in period 2 than they did in period 1. This oscillation in market precision becomes wider and wider until more agents acquire other expectations strategies.

(c) *Low risk aversion.* In addition to describing agents' attitude toward risk, the parameter r also can provide a measure of market completeness. Derivatives, for example, allow risks to be unbundled and placed where they are most easily born, thus lowering the market's effective level of risk aversion. The existence of derivatives should, in the context of this model, increase the likelihood of information-driven instability.

(d) *High share supply uncertainty.* Share supply uncerainty can be interpreted as emanating from either noise trading [e.g., De Long et al. (1990)], changes in agents' risk aversion [e.g., Campbell et al. (1993)], or changes in agents' private investment opportunities [e.g., Wang (1994)]. In all three cases, the supply uncertainty dilutes the information content of market prices, and forces traders wishing to remain informed to purchase more private information.

(e) *High share supply.* Because a high share supply requires agents to hold more shares, it also increases their appetitie for private information. By increasing agents' appetite for information, conditions (d) and (e) both lead to the cyclic pattern outlined under condition (b), and thus to instability.

5. RATIONAL VERSUS NAIVE EXPECTATIONS

Proposition 1 suggests that the model might display interesting dynamic behavior when β is large. The following two sections explore this possibility using numerical simulation. First, suppose that the space $\mathcal{H}$ of expectations strategies has only the two elements: rational and naive expectations. Let the lag in information strategy revision be one period ($\kappa = 1$), so that

$$H_{\text{rat}}(I_{\tau-1}) = B_\tau,$$
$$H_{\text{nai}(0)}(I_{\tau-1}) = B_{\tau-1}.$$

Denote by $\alpha = c_H(H_{\text{rat}}) - c_H[H_{\text{nai}(0)}]$ the difference in cost between rational and naive expectations. We set the riskless interest rate R and the risk-aversion coefficient r both equal to 1. The mean terminal dividend $\bar{d}$ for the risky asset is also 1, and its variance σ_d is 0.015. The mean supply $\bar{z}$ of shares is 1, and the variance σ_z is 0.1. The fixed cost α of forming rational expectations is 0.03; the cost coefficient φ for the information strategy is 0.00025.

Figure 1 presents a bifurcation diagram that shows how the model's long-run behavior changes as β is increased from 10 to 700.[15] When β is small, the equilibrium relation (13) has a stable steady state. As β increases, the model begins to display periodic behavior. Figure 1 shows periods of 2, 4, 8, and 16 as β increases

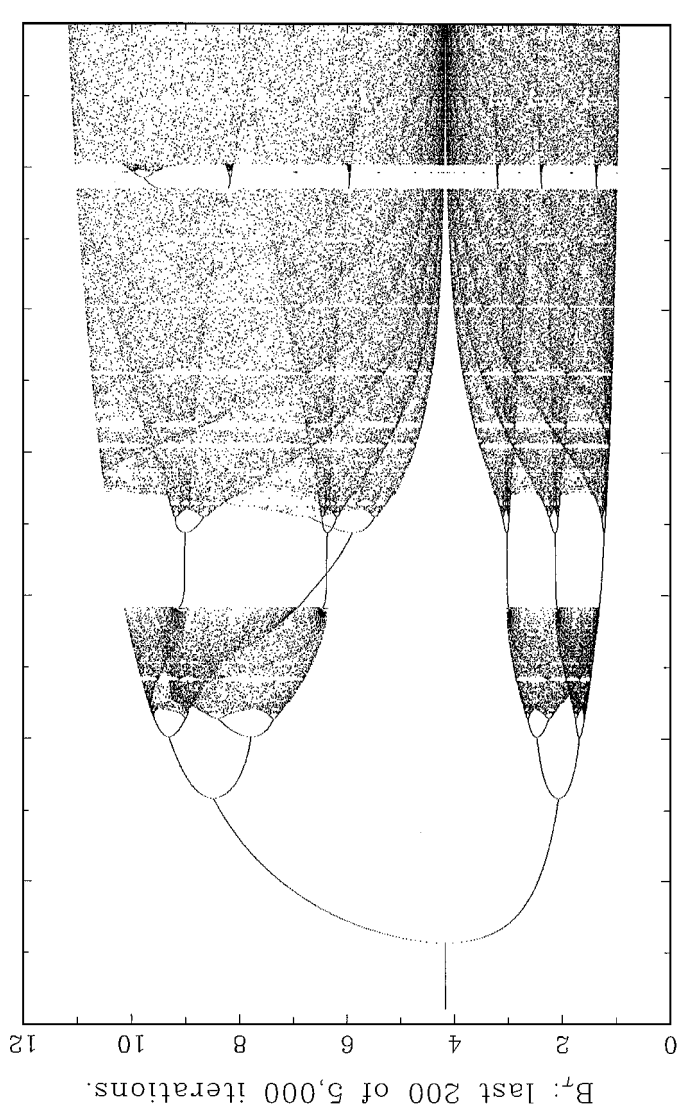

FIGURE 1. Bifurcated diagram.

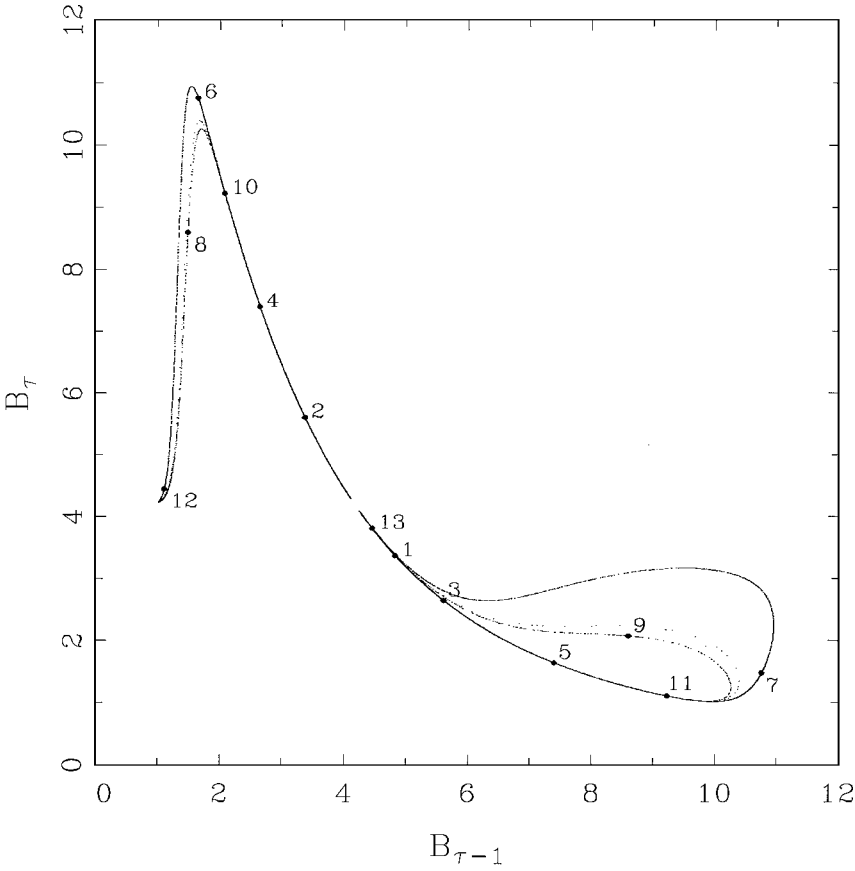

FIGURE 2. Final 25,000 points in $\{B_{\tau-1}, B_\tau\}$ space.

from 10 to about 220. At approximately $\beta = 225$, the model begins to display extremely complicated behavior, visiting many points B_τ on the vertical axis.

Figure 2 plots the evolution of the model in $\{B_{\tau-1}, B_\tau\}$ space when the intensity of choice $\beta = 550$. The model was simulated for 30,000 periods, of which the final 25,000 are displayed. One can see easily that the model is unstable by noting that the slope of the curve is less than -1 at the intersection with the 45-deg line $B_\tau = B_{\tau-1}$. A small oscillation will be mapped farther and farther from B^* until B_τ reaches the upward-sloping portion on the left side of Figure 2. This upward-sloping portion then sends B_τ near B^* again.

More concretely, begin at point 1, where B_1 is near its steady-state value B^*. Because B_τ does not fluctuate much in this region, the naive expectations strategy performs well compared to rational expectations, and so, most agents choose naive expectations and $n_{\text{rat},1}$ is small. Because almost all of the agents are using naive expectations, the model has a cobweb-like instability. As this instability grows over time, the fluctuations in B_τ become larger. This stage can be seen at points

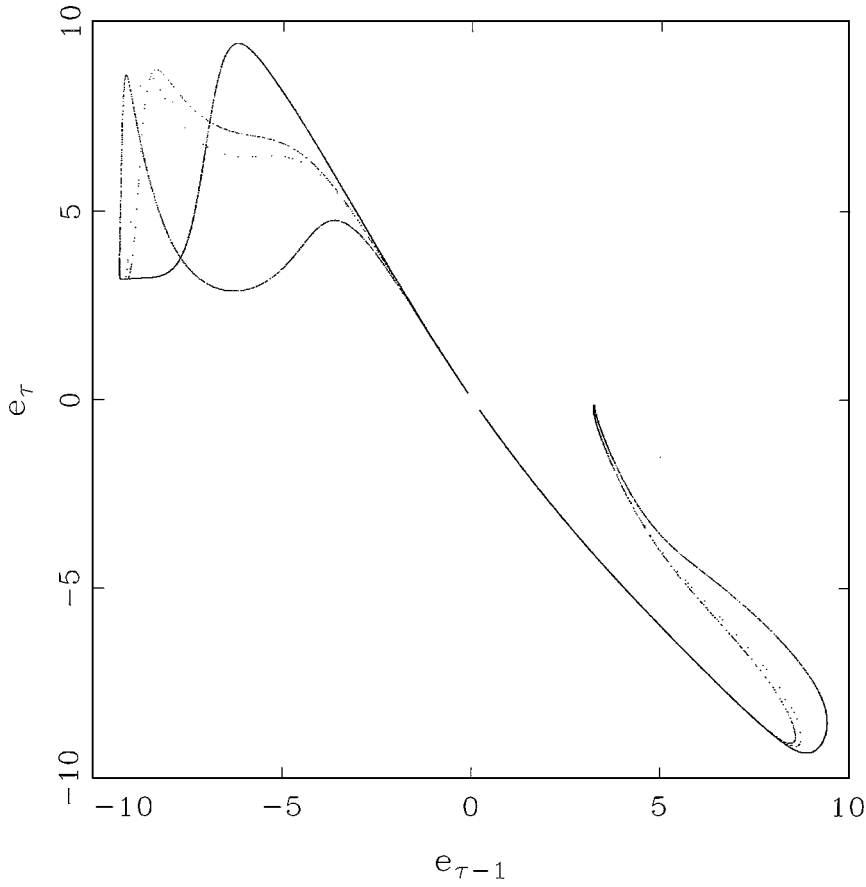

FIGURE 3. Predictable forecast errors.

2 through 12. Because the fluctuations are growing, the performance of the naive expectations strategy is deteriorating, so that n_{rat} is increasing. As the fluctuations become extremely large, it becomes worthwhile to pay the fixed cost of forming rational expectations, so $n_{\text{rat},13}$ is near 1 at point 13. Because many agents are now rational, B_{13} is very near the steady-state B^*, and the amplification of the small deviation from equilibrium begins again.

Define $e_\tau = B_\tau - B_{\tau-1}$ to be the forecast error associated with naive expectations. Figure 3, which plots e_τ against $e_{\tau-1}$, suggests that the limited strategy space $\mathcal{H} = \{H_{\text{rat}}, H_{\text{nai}(0)}\}$ does not make complete economic sense, for it seems unreasonable that agents would not revise an expectations strategy that makes such consistent and easily predictable forecast errors. These errors persist because agents can eliminate them only by switching to the rational expectations strategy. Unless the errors are large, paying the associated computation cost simply is not worthwhile.

Figure 4 displays the simulated time series for volatility and trading volume for 40 out of the 30,000 simulated strategy revision periods. As in Figure 2, small

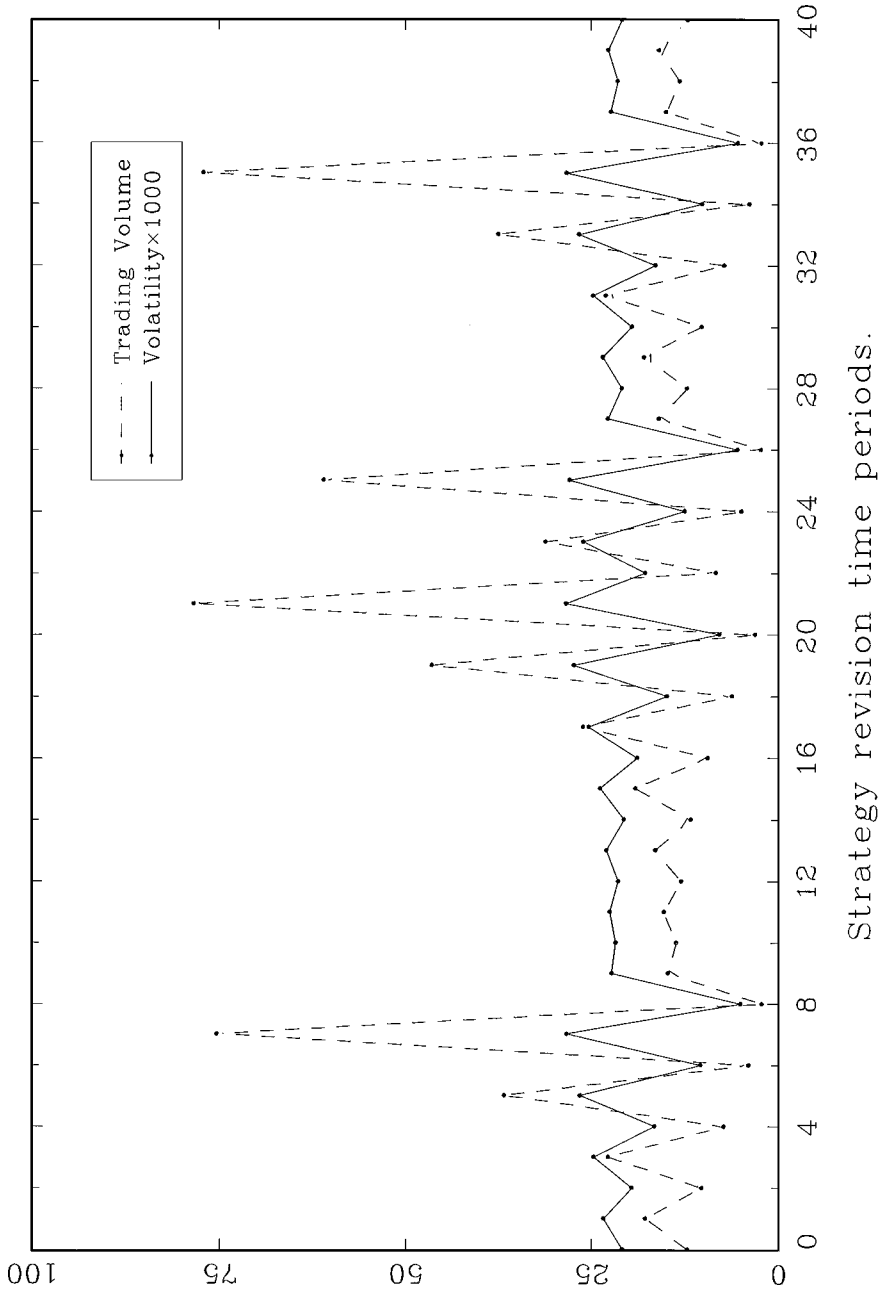

FIGURE 4. Volatility and trading volume.

deviations are amplified repeatedly until volatility and volume return near their steady-state values. There is thus considerable time variation in both series, which is a typical feature of market data. Both series also display a marked sawtooth pattern: If volatility is high in one period, it always will be low in the next period. Like the predictable forecast errors, this pattern is a result of the limited expectations strategy space $\mathcal{H}$. Agents using naive expectations generate the cobweb-like pattern, which is not arbitraged away because rational expectations are expensive and there are no other expectations strategies in $\mathcal{H}$. The sawtooth pattern is of course inconsistent with the empirical findings reviewed in Section 3, and thus poses another problem for this simple version of the model.

6. ENLARGING THE STRATEGY SPACE

In the simple model of the preceding section, agents make repeated systematic forecast errors because the cost of forming rational expectations is usually greater than the benefits of eliminating the errors. Agents do not experiment with other expectations strategies because, by definition, $\mathcal{H}$ has only two elements. Enlarging the expectations strategy space to include a wider range of backward-looking expectations allows agents to reduce forecast errors without incurring the full cost of rational expectations.

Let H_{rat} denote rational expectations, and let $H_{\text{nai}(l)}$ denote l-lagged naive expectations:

$$\begin{aligned} H_{\text{rat}}(I_{\tau-\kappa}) &= B_\tau, \\ H_{\text{nai}(l)}(I_{\tau-\kappa}) &= B_{\tau-\kappa-l}, \quad 0 \le l \le 30. \end{aligned} \tag{18}$$

The expectations strategies $H_{\text{nai}(l)}$ allow agents to detect and take advantage of any cyclical patterns that might appear in the data, and thus eliminate the most easily predictable forecast errors. The strategy $H_{\text{nai}(1)}$ would, for example, allow agents to detect and eliminate sawtooth pattern (of period 2) seen in the preceding section. The information cost function is still quadratic in the signal precision $[c_\rho(\rho_\tau) = \varphi\rho_\tau^2]$, and the fixed cost for naive expectations increases in l:

$$\begin{aligned} c_H(H_{\text{rat}}) &= \alpha_{\text{rat}}, \\ c_H[H_{\text{nai}(l)}] &= \alpha_{\text{nai}} l. \end{aligned} \tag{19}$$

Specification (19) represents increasing memory costs associated with storing past values of B_τ.

We set the riskless interest rate R and the risk-aversion coefficient r both equal to 1. The mean payoff $\overline{v}$ for the risky asset is also 1, and its variance σ_d is 0.015. The mean supply $\overline{z}$ of shares is 1,000, and the variance σ_z is 0.1. The fixed cost α_{rat} for purchasing rational expectations is 13, the fixed cost α_{nai} for purchasing naive expectations is 0.1, and the cost coefficient φ for signal precision is 2.5. The

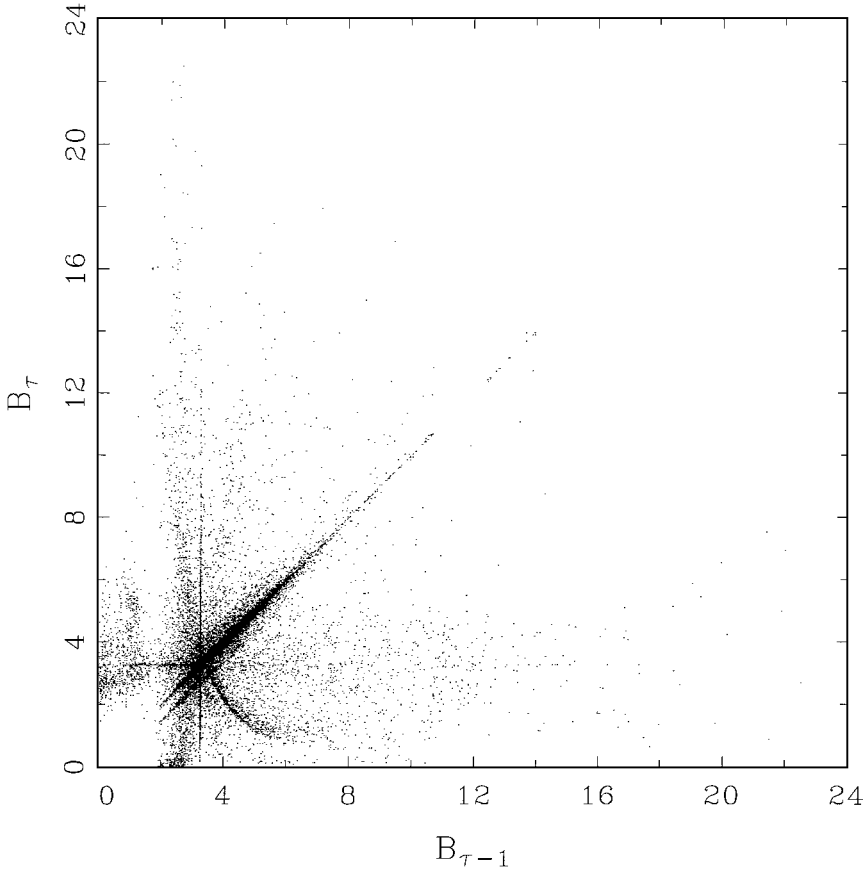

FIGURE 5. Model iterations displayed as 25,000 points in $\{B_{\tau-1}, B_\tau\}$ space.

intensity of choice β is 2.5. The lag in information strategy revision is $\kappa = 10$, so that changes in information strategy take 10 periods to come into effect. Because a strategy revision period τ corresponds to one day, this specification corresponds to a two-week lag in information strategy revision.[16]

Figure 5 displays 25,000 iterations of the model in $\{B_{\tau-1}, B_\tau\}$ space, and shows a clear positive correlation between the value of $B_{\tau-1}$ and that of B_τ. This positive correlation suggests that related quantities such as volatility and volume should display positive autocorrelation. Naive agents cannot exploit (and thus eliminate) this short-run autocorrelation because of the lag in information strategy revision.

The forecast error for the strategy $H_{\text{nai}(0)}$ is $e_\tau = B_\tau - B_{\tau-10}$. Figure 6 shows the same 25,000 iterations in $\{e_{\tau-10}, e_\tau\}$ space.[17] These forecast errors are less predictable than those displayed in Figure 3. Although there may be some predictability remaining, predictability matters much less in the enlarged strategy space

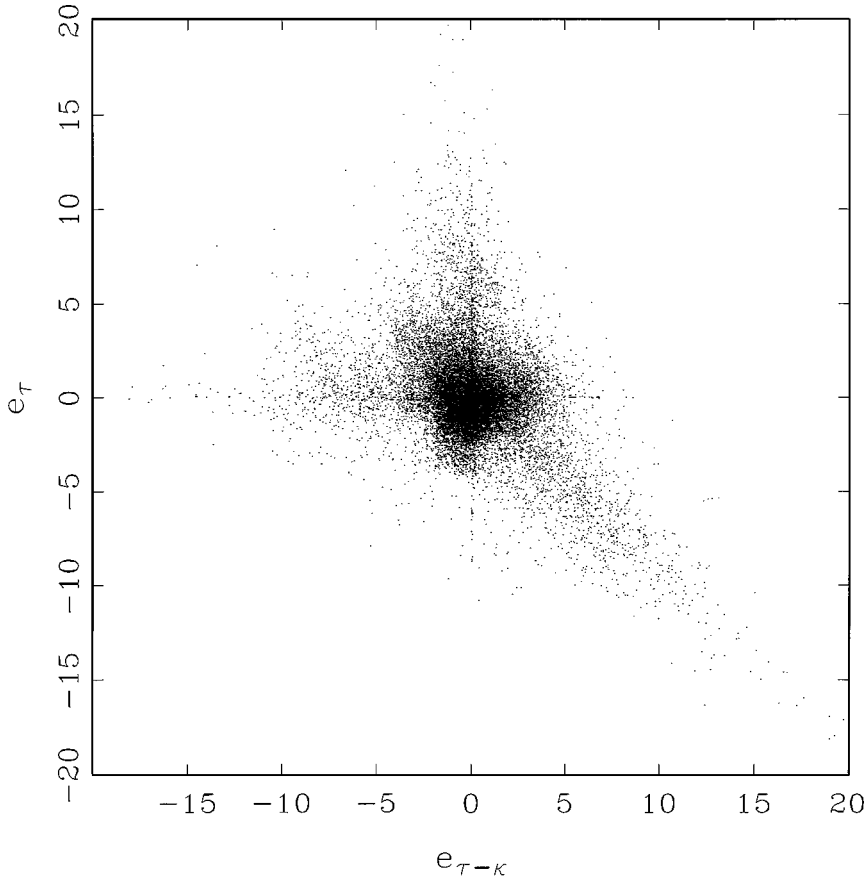

FIGURE 6. Less predictable forecast errors.

model because agents can always switch to some other inexpensive expectations strategy. When strategy $H_{\text{nai}(0)}$ makes a large forecast error, the number of agents using $H_{\text{nai}(0)}$ will be small in the next period; the next period's error is thus of little economic consequence.

Figure 7 displays the simulated time series for volatility and volume for 100 out of the 30,000 simulated strategy revision periods. As in Figure 4, both series covary considerably over time. In addition, Figure 7 does not display the persistent, predictable sawtooth pattern seen in Figure 4. Although there are brief oscillatory periods, these are interspersed with relatively stable periods in which neither volatility nor volume changes significantly from one period to the next. The overall result is that the simulated time series are quite plausible looking. We investigate their plausibility more rigorously in Figures 8 to 10.

Figure 8 compares the autocorrelogram generated from the simulated volatility series with that generated from daily data on IBM stock returns,[18] which provide

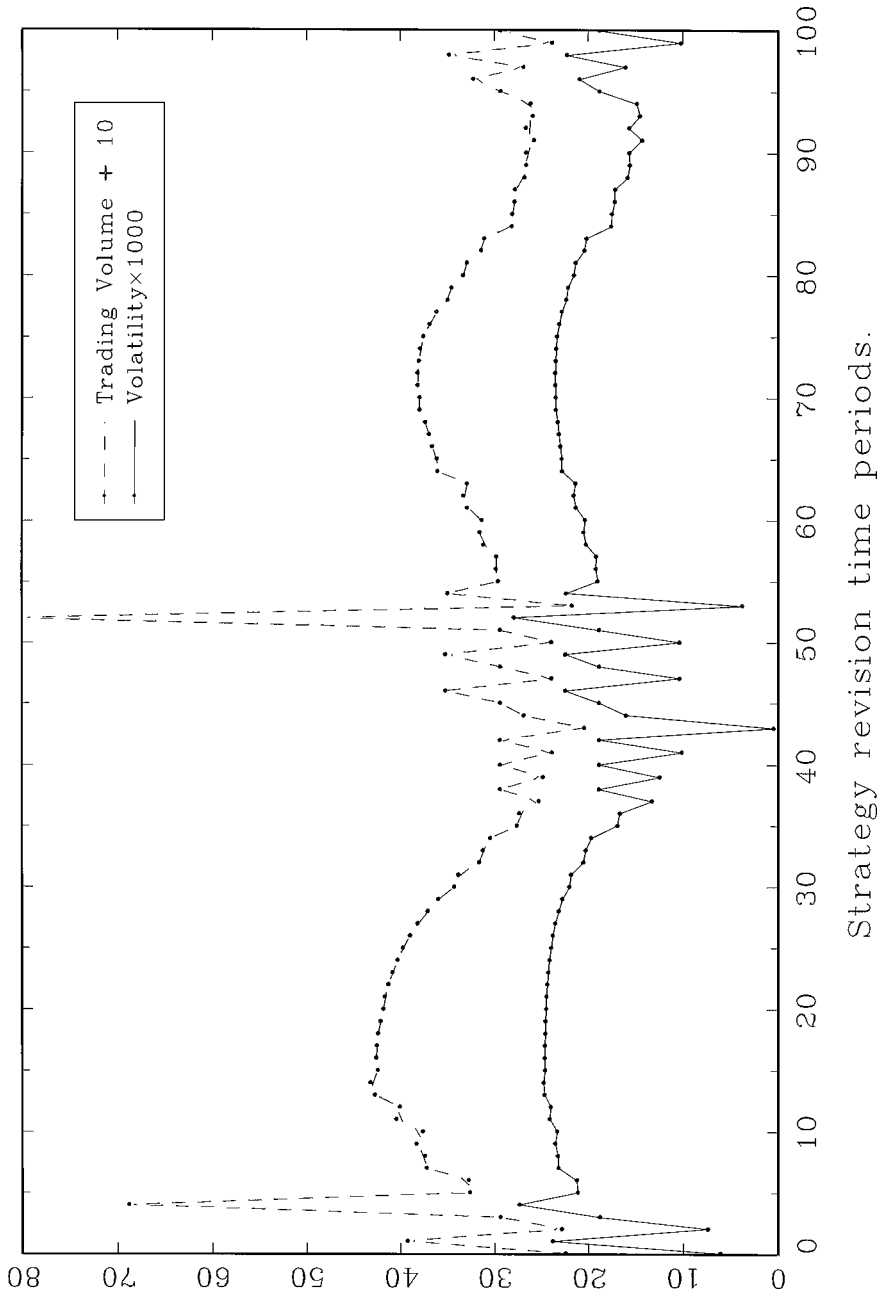

FIGURE 7. Volatility and trading volume.

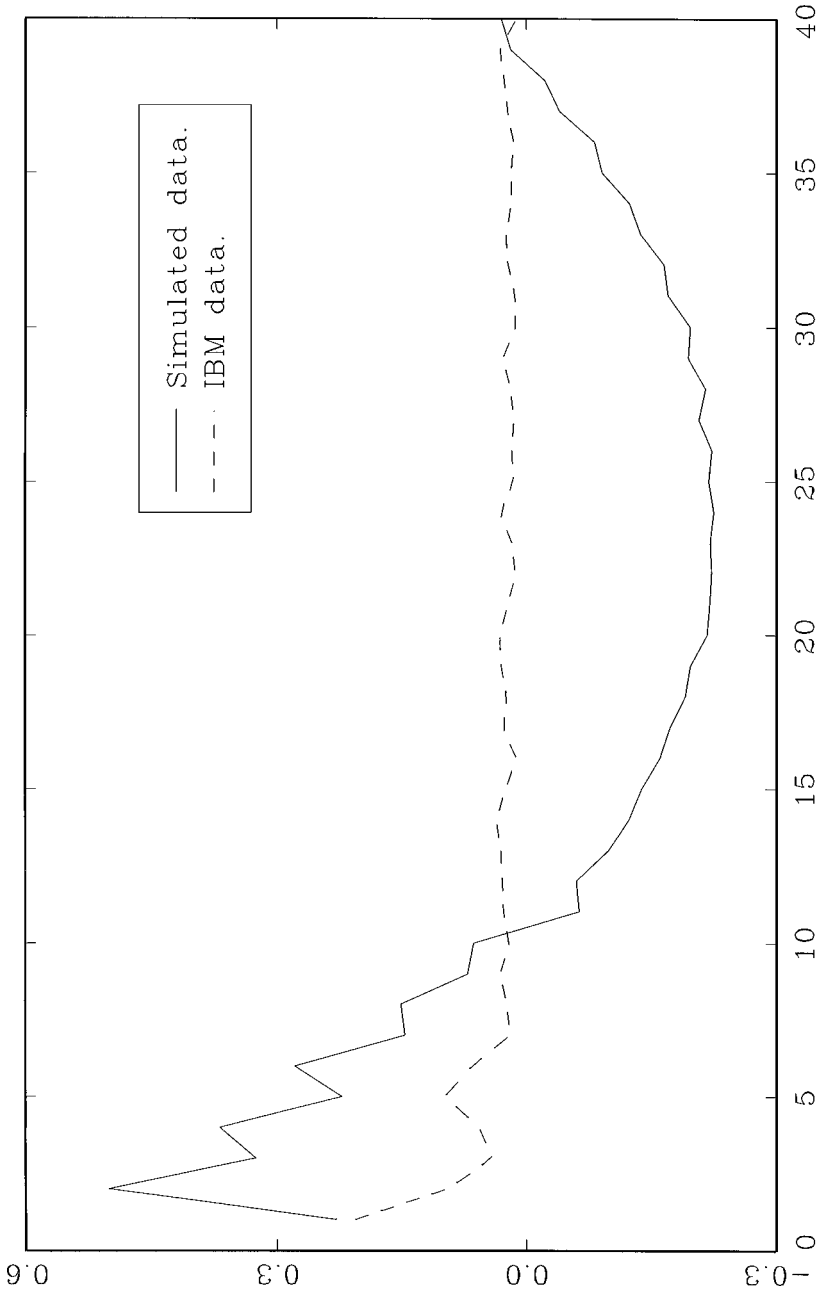

FIGURE 8. Autocorrelogram of volatility.

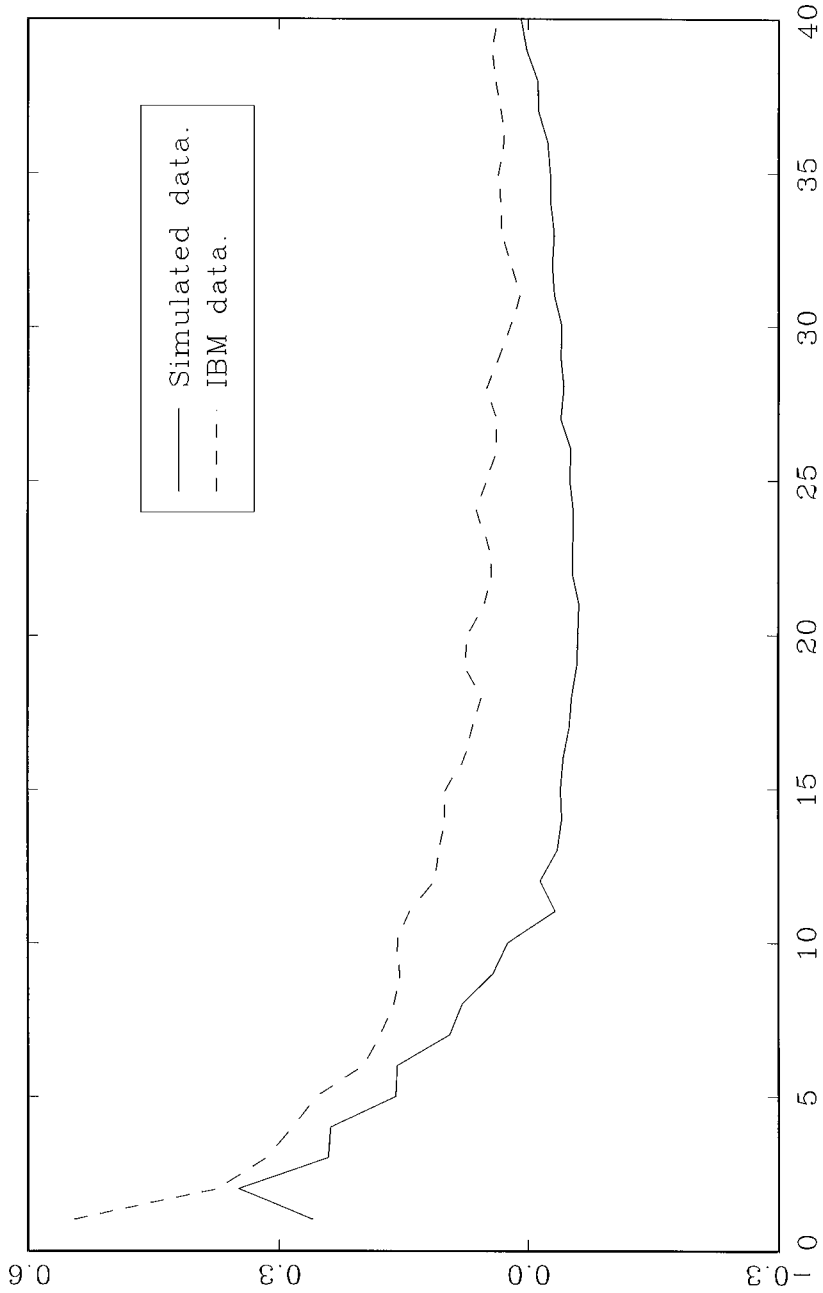

FIGURE 9. Autocorrelogram of volume.

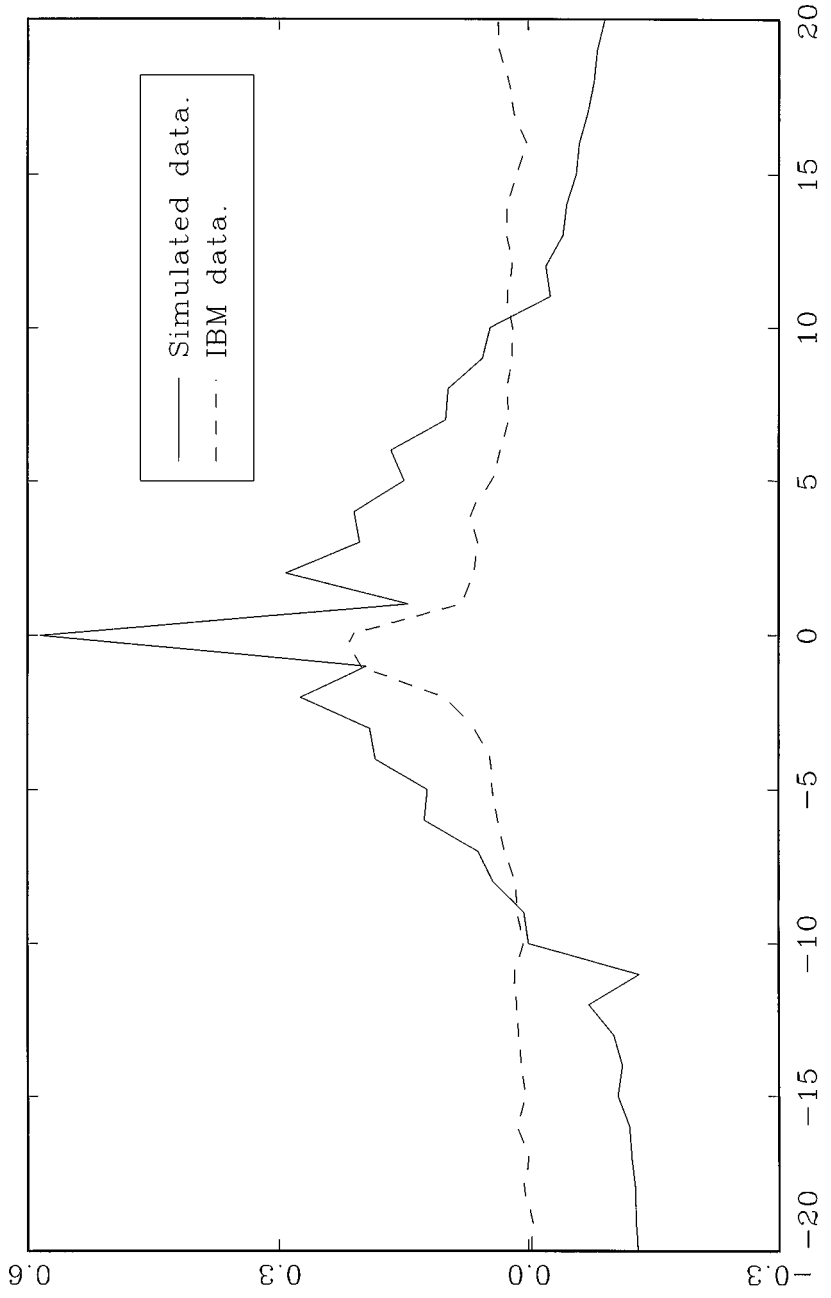

FIGURE 10. Cross correlogram of volatility and volume.

a qualitative benchmark for the model. The first-order autocorrelation coefficient from the simulations is clearly positive, and between 2 and 10 lags, the simulations also display positive autocorrelation coefficients. So, the model succeeds in capturing the volatility persistence seen in the IBM data. The empirical and simulated correlograms also match well near 40 lags, both displaying correlation coefficients near zero.

At intermediate lags, however, the simulations generate negative autocorrelation coefficients, which are not consistent with those from the IBM data. As discussed previously, naive agents expect market precision to be high in the future when it is high today, thus inducing negative autocorrelation in market precision—and hence in volatility and trading volume. Although enlarging the strategy space does allow other strategies to take advantage of—and thus reduce—the negative autocorrelation, it does not completely eliminate the effect.[19]

Figure 9 displays the autocorrelograms of trading volume, and indicates a very good match between the simulated and actual data. Both display significant autocorrelation at low lags, which diminishes rapidly to a near-zero level at longer lags. Figure 10 displays the cross correlograms between volatility and trading volume. As in Figure 8, the simulated data match the IBM data qualitatively but not quantitatively. Although both correlograms display peaks at 0 lags, the peak for the simulated data is much more pronounced than that for the actual data. In addition, the simulated data display negative correlation between volume and volatility at 20 lags, whereas the actual data display no correlation at 20 lags. We conclude that the model can explain some but not all of the features found in market data, and we interpret the results as an encouragement for future work.

7. CONCLUSIONS

Are the forces driving volatility and trading volume in financial markets the same as those driving asset returns? This question is important not only to those seeking a better theoretical understanding of trading dynamics, but also to those with applied goals such as volatility prediction or policy evaluation.

The empirical literature suggests a negative answer to this question. In ARCH models [see Bollerslev et al. (1993)], volatility is specified explicitly to be a deterministic function of past-return innovations. Stochastic volatility models [e.g., Jacquier et al. (1994); Andersen (1996)] decouple volatility and trading volume from returns by allowing them to depend on an exogenous directing process. Geweke (1995) finds that stochastic volatility models empirically outperform ARCH models precisely because they do not constrain volatility to depend on past realized returns.

Our model provides a theoretical explanation of how the forces driving volatility and trading volume could differ from those driving returns. In fact, one may view it as a structural[20] counterpart to stochastic volatility models, which are purely statistical. Returns are driven on a short timescale by dividend innovations. Our "directing process," market precision, is decoupled from returns because it evolves on a long timescale.

The decoupling of market precision from returns ensures consistency with the large literature on efficient markets [see, e.g., Fama (1991) or Lo (1996)], which asserts that returns are very difficult to predict from publicly available information.[21] This is an important point, for one would expect a priori that a model implying predictability in two major market quantities (volatility and volume) also would imply predictability in the third (returns).

In addition to providing insight into the market's information dynamics, a structural model can address policy questions that purely statistical models cannot answer. How might changes in market structure affect trading dynamics? Can the government stabilize financial markets? Is such stabilization desirable? The main hurdle to be cleared before the model can credibly address such questions is to construct a strategy space that is general yet parsimoniously parameterized. Generality is necessary to avoid restricting agents to narrow or implausible behavioral patterns (as in Section 5). Parsimony is necessary for the model's behavior to depend on meaningful economic assumptions, rather than on unobserved behavioral parameters (as in Section 6). Our encouraging results suggest that this is an important area for future research.[22]

NOTES

1. One might wonder why computation costs should be important for this model, in which all computations are relatively straightforward. To analyze the effects of these costs, however, we need a model that does not overwhelm our own computational abilities. We thus regard the model as a simplification of a more complex trading environment, in which forming rational expectations requires intensive computation. See Evans and Ramey (1992, p. 211) for related discussion.

2. Taken together, Assumption 1 and Proposition 1 of Goeree and Hommes (2000) correspond closely to this paper's Proposition 1.

3. A related model with long-lived assets is explored in de Fontnouvelle (1996).

4. Because returns on the risky asset are normally distributed, the mean-variance criterion (1) implies the same demand for risky assets as the expected utility approach using a negative exponential utility function [e.g., Campbell et al. (1993), p. 927]. The mean-variance approach, however, allows us to calculate (in Section 2.3) an exact expression for the fraction of agents using each expectations strategy at each time t.

5. The idea of two different timescales has a long history in economics, and often relates (as it does in this paper) to explaining heteroskedastic effects in returns. In mixture-of-distributions models [e.g., Clark (1973); Andersen (1996)], for example, the number of trades per day is governed by an exogenous directing process, which evolves slowly relative to asset trading time.

6. The assumption of synchronous adjustment is made for tractablity rather than realism. A tractable way of modeling nonsynchronous adjustment may be to allow a fraction q of agents to revise their strategies during each asset trading period; we do not believe that such a modification would substantively alter the paper's conclusions. A full examination of whether such an approach is indeed tractable, and of its effects on the model's dynamics, is left to future research.

7. The evolution of B_τ is deterministic, so that rational expectations implies perfect foresight. There is thus no approximation error between (5) and (6) for agents with rational expectations. An agent using another expectations strategy simply accepts the approximation error as one of the drawbacks of that strategy.

8. To see why I_τ contains a record of realized market precisions, recall that all agents know p_t for each t in the strategy revision period τ, so that they know the sample average $\hat{p}_\tau = (\sum_{t\in\tau} p_t)/T$. If T is large, then the law of large numbers implies that $\hat{p}_\tau$ is a precise approximation of $p_\tau = E_\tau[p_t]$, the expected market price for strategy revision period τ. Equation (14) shows that B_τ is a root of a

quadratic equation of the form $B_\tau^2(\theta_1 - \theta_2 p_\tau) + B_\tau(\theta_3 - \theta_4 p_\tau) + (\theta_5 - \theta_6 p_\tau) = 0$. This equation has one positive root, and thus B_τ is a well-defined function of the observed information p_τ.

9. Bray (1982) shows that if agents periodically update their backward-looking expectations, cobweb models can in fact converge to the rational expectations equilibrium. Models that generate cobwebs by prohibiting agents from updating or experimenting with their expectations thus are not economically plausible. See also Townsend (1983) and Bray and Savin (1986).

10. Alternatively, one can specify that at the end of each period τ, the profitability of each expectations strategy over τ is announced. This specification is very similar to aggregate statistic models in game theory, in which relevant population aggregates are announced at the end of each playing round [see, e.g., Fudenberg and Levine (1998)]. The main difference here is that there is an explicit mathematical mechanism through which agents could deduce the population aggregates (profitabilities) from the observed price history.

11. The meaning of the parameter β is discussed immediately following Proposition 1.

12. Different assumptions about the serial dependence in the $\varepsilon_{i,j,\tau}$ series lead to different formulas for trading volume (but not for volatility or market precision). To check that the results do not hinge too strongly on the assumption of serial independence, we recalculated the volume series under the assumption that $\varepsilon_{i,j,\tau} = \varepsilon_{i,j,\tau+\omega}$ for all i, j, τ, and ω. The simulation results do not change appreciably under this alternative assumption.

13. For other applications of discrete choice to dynamic structural modeling, see Follmer (1974), Durlauf (1989), Blume (1993), Aoki (1995), and Brock and LeBaron (1996).

14. See equation (18) to understand this notation.

15. The parameter β was increased from 10 to 700 by increments of 1. For each value of β on the horizontal axis, the model was iterated 5,000 times, and the values of B_τ for the final 200 iterations were plotted on the vertical axis.

16. Clearly, the value chosen for κ will affect the amount of persistence in the model: The longer κ is, the more difficult it will be for agents to exploit short-run predictability in market precision. Experimentation confirms that the shape of the autocorrelograms for volatility and trading volume, and thus the quantitative behavior of the model, does depend on the value of κ. The model's qualitative behavior (the fact that the model displays copersistence), however, is invariant to reasonable changes in κ.

17. Because $\kappa = 10$, agents can exploit dependence between $e_{\tau-j}$ and e_τ only for $j \geq 10$.

18. The returns data are daily returns on IBM stock from July 6, 1962, through December 31, 1993. The volume series is generated from the turnover ratio, the number of traded shares divided by the number of outstanding shares. This ratio was detrended using a 100-day moving average, and then log transformed.

19. Some alternate specification of the strategy space $\mathcal{H}$ probably would reduce the negative autocorrelation displayed by the model. However, there is no systematic method of constructing and exploring strategy spaces, and it is unclear how one might interpret the results of an attempted search over all possible strategy spaces. Developing such a systematic method is an important goal for future work.

20. Our model is structural because the directing process (market precision) is endogenous.

21. To ensure consistency with the efficient-markets literature, one also must assume that the traders' signal $y_{i,t}$ is not available to the academic econometrician. Because these signals are intended to represent private information, this seems an innocuous assumption.

22. Work along these lines is developed by Brock and Hommes (1998).

REFERENCES

Andersen, T. (1996) Return volatility and trading volume: An information flow interpretation of stochastic volatility. *Journal of Finance* 51, 169–204.

Anderson, S., A. de Palma & J. Thisse (1993) *Discrete Choice Theory of Product Differentiation*. Cambridge, MA: MIT Press.

Antoniewicz, R. (1992) *An Empirical Analysis of Stock Returns and Volume*. Ph.D. Dissertation, University of Wisconsin at Madison.

Aoki, M. (1995) Economic fluctuations with interactive agents: Dynamic and stochastic externalities. *Japanese Economic Review* 46.

Arthur, B. (1995) Complexity in economic and financial markets. *Complexity* 1, 20–25.

Arthur, B., Holland, J., LeBaron, B., R. Palmer & P. Taylor (1997) Asset pricing under endogenous expectations, in an artificial stock market. In B. Arthur, S. Durlauf & D. Lane (eds.), *The Economy as an Evolving Complex System: II*, pp. 15–44. Reading, MA: Addison-Wesley.

Blume, L. (1993) The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 387–426.

Bollerslev, T., R. Chou & K. Kroner (1992) ARCH modeling in finance. *Journal of Econometrics* 52, 5–59.

Bollerslev, T., R. Engle & D. Nelson (1993) ARCH models. In R. Engle & D. McFadden (eds.), *Handbook of Econometrics*, vol. IV, pp. 2959–3038. Amsterdam: North-Holland.

Bray, M. (1982) Learning, estimation and the stability of rational expectations equilibria. *Journal of Economic Theory* 26, 318–339.

Bray, M. & N. Savin (1986) Rational expectations equilibria, learning, and model specification. *Econometrica* 54, 1129–1160.

Brock, W. & C. Hommes (1997) A rational route to randomness. *Econometrica* 65, 1059–1095.

Brock, W. & B. LeBaron (1996) A dynamic structural model for stock return volatility and trading volume. *Review of Economics and Statistics* 78, 94–110.

Campbell, J., S. Grossman & J. Wang (1993) Trading volume and serial correlation in stock returns. *Quarterly Journal of Economics* 108, 905–939.

Clark, P. (1973) A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* 41, 135–155.

de Fontnouvelle, P. (1996) Stock Returns and Trading Volume: A Structural Explanation of the Empirical Regularities. Working paper, Iowa State University.

De Long, J., A. Shleifer, L. Summers & R. Waldmann (1990) Noise trader risk in financial markets. *Journal of Political Economy* 98, 703–738.

Durlauf, S. (1989) Locally Interacting Systems, Coordination Failure, and the Behavior of Aggregate Activity. Working paper, Stanford University.

Evans, G. & S. Honkapohja (1999) Learning dynamics. In J. Taylor & M. Woodford (eds.), *Handbook of Macroeconomics*, vol. 1, Ch. 7, pp. 499–542. Amsterdam: North-Holland.

Evans, G. & G. Ramey (1992) Expectation calculation and macroeconomic dynamics. *American Economic Review* 82, 207–224.

Fama, E. (1991) Efficient capital markets: II. *Journal of Finance* 46, 1575–1618.

Follmer, H. (1974) Random economies with many interacting agents. *Journal of Mathematical Economics* 1, 51–62.

Fudenberg, D. & D. Levine (1998) *The Theory of Learning in Games*. Cambridge, MA: MIT Press.

Gallant, A., P. Rossi & G. Tauchen (1992) Stock prices and volume. *Review of Financial Studies* 5, 199–242.

Geweke, J. (1995) The Bayesian Comparison of Econometric Models. Working paper, University of Minnesota.

Goeree, J. & C. Hommes (2000) Heterogeneous beliefs and the non-linear cobweb model. *Journal of Economic Dynamics and Control* 24, 761–798.

Grossman, S. & J. Stiglitz (1980) On the impossibility of informationally efficient markets. *American Economic Review* 70, 393–408.

Hellwig, M. (1980) On the aggregation of information in competitive markets. *Journal of Economic Theory* 22, 477–498.

Holland, J. (1995) *Hidden Order: How Adaptation Builds Complexity*. Reading, MA: Addison-Wesley.

Jacquier, E., N. Polson & P. Rossi (1994) Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* 12, 371–388.

Karpoff, J. (1987) The relation between price changes and trading volume: A survey. *Journal of Financial and Quantitative Analysis* 22, 109–126.
Lamoureux, C. & W. Lastrapes (1990) Heteroskedasticity in stock return data: Volume versus GARCH effects. *Journal of Finance* 45, 221–229.
Lang, L., R. Litzenberger & V. Madrigal (1992) Testing financial market equilibrium under asymmetric information. *Journal of Political Economy* 111, 317–348.
LeBaron, B., B. Arthur & R. Palmer (1999) The time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23, 1487–1516.
Lo, A. (ed.) (1996) *Market Efficiency: Stock Market Behaviour in Theory and Practice*. London: Edward Elgar.
Manski, C. & D. McFadden (1981) *Structural Analysis of Discrete Data with Econometric Applications*. Cambridge, MA: MIT Press.
Marimon, R. (1997) Learning from learning in economics. In Kreps, D. & K. Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications*, vol. I, *Seventh World Congress of the Econometric Society*, Cambridge, UK: Cambridge University Press.
Sargent, T. (1993) *Bounded Rationality in Macroeconomics*. New York: Oxford University Press.
Townsend, R. (1983) Forecasting the forecasts of others. *Journal of Political Economy* 91, 546–588.
Wang, J. (1994) A model of competitive stock trading volume. *Journal of Political Economy* 102, 127–168.

APPENDIX A. HELLWIG'S COEFFICIENTS FOR THE PRICE FUNCTION (14)

$$\phi_0(B) = \frac{\bar{d}\sigma_z/r + \sigma_d \bar{z} B/r}{RM(B)},$$

$$\phi(B) = \frac{\sigma_d B\sigma_z + \sigma_d B^2/r}{RM(B)},$$

$$\gamma(B) = \frac{\sigma_d\sigma_z + \sigma_d B/r}{RM(B)},$$

$$M(B) = \sigma_z/r + \sigma_d B\sigma_z + \sigma_d B^2/r.$$

APPENDIX B. CALCULATING EXPECTED PROFITS

Define $s_{i,\tau} = 1/\rho_{i,\tau}$. Hellwig (1980, p. 493) shows that in each asset trading period t, agent i's conditional expectation and conditional variance are given by the following formulas:

$$E[d_t \mid y_{i,t},\, p_t] = \frac{\sigma_z s_{i,\tau}\bar{d} + \sigma_z\sigma_d y_{i,t}}{D_\tau} + \frac{\sigma_d s_{i,\tau} B_\tau^2[p_t - \phi_0(B_\tau) + \gamma(B_\tau)\bar{z}]}{\phi(B_\tau)D_\tau} \qquad \textbf{(B.1)}$$

$$\begin{aligned} \mathrm{Var}[d_t \mid y_{i,t},\, p_t] &= \sigma_z\sigma_d s_{i,\tau}/D_\tau, \\ D_\tau &= \sigma_z\sigma_d + \sigma_z s_{i,\tau} + \sigma_d s_{i,\tau} B_\tau^2. \end{aligned} \qquad \textbf{(B.2)}$$

Rearranging into statistically independent terms, and noting that $\phi(B) = B\gamma(B)$, it follows from (14) and (B.1) that

$$\begin{aligned} E(d_t - Rp_t \mid y_{i,t},\, p_t) &= \bar{d}[1 - R\phi(B_\tau)] + \bar{z}R\gamma(B_\tau) - R\phi_0(B_\tau) \\ &+ (d_t - \bar{d})[1 - R\phi(B_\tau) - \sigma_z s_{i,\tau}/D_\tau] + (z_t - \bar{z})[R\gamma(B_\tau) - \sigma_d s_{i,\tau} B_\tau/D_\tau] \\ &+ \epsilon_{i,t}\sigma_z\sigma_d/D_\tau, \end{aligned}$$

so that

$$\begin{aligned} E_\tau[E^2(d_t - Rp_t \mid y_{i,t},\, p_t)] &= \mathrm{Var}_\tau[E(d_t - Rp_t \mid y_{i,t},\, p_t)] + E_\tau^2[E(d_t - Rp_t \mid y_{i,t},\, p_t)] \\ &= \sigma_d[1 - R\phi(B_\tau)]^2 + R^2\gamma(B_\tau)^2\sigma_z - \sigma_d\sigma_z s_{i,\tau}/D_\tau \\ &\quad + \left\{\bar{d}[1 - R\phi(B_\tau)] + \bar{z}R\gamma(B_\tau) - R\phi_0(B_\tau)\right\}^2. \end{aligned}$$

Noting that

$$\frac{1}{\mathrm{Var}(d_t - Rp_t \mid y_{i,t},\, p_t)} = \frac{D_\tau}{\sigma_z\sigma_d s_{i,\tau}} = 1/s_{i,\tau} + 1/\sigma_d + B_\tau^2/\sigma_z$$

leads to the result (3).

APPENDIX C. CALCULATING EXPECTED VOLUME

From equations (2), (B.1), and (B.2), it follows that the change in agent i's holdings of the risky asset between $t-1$ and t is given by

$$x_{i,t} - x_{i,t-1} = \left[(d_t - d_{t-1})B_\tau^2 - (z_t - z_{t-1})B_\tau\right]/r\sigma_z + \rho_{i,\tau}(y_{i,t} - y_{i,t-1})/r$$

The expectation and variance (taken over information trading period τ) of this quantity are

$$\begin{aligned} E_\tau(x_{i,t} - x_{i,t-1}) &= 0 \\ \mathrm{Var}_\tau(x_{i,t} - x_{i,t-1}) &= 2\left[\rho_{i,\tau} + \sigma_d\left(\rho_{i,\tau}^2 + B_\tau^4/\sigma_z^2\right) + B_\tau^2/\sigma_z\right]/r^2 \end{aligned} \qquad \textbf{(C.1)}$$

Because the expectation of the absolute value of a random variable x distributed normally $N(0, b)$ is $E|x| = \sqrt{2b/\pi}$, the formula for expected volume is

$$V_\tau = \sum_{j=1}^{K} n_{j,\tau-\kappa}\sqrt{\frac{\rho_{j,\tau} + \sigma_d\left(\rho_{j,\tau}^2 + B_\tau^4/\sigma_z^2\right) + B_\tau^2/\sigma_z}{\pi r^2}}. \qquad \textbf{(C.2)}$$

APPENDIX D. PROOF OF PROPOSITION 1

From (13), we know that B_τ evolves according to an $(L-\kappa+1)$th-order difference equation:

$$\begin{aligned} B_\tau &= F(\beta, B_{\tau-\kappa}, \ldots, B_{\tau-L}) = F(\beta, \boldsymbol{B}_{\tau-\kappa}) \\ &= \frac{1}{r}\sum_{j=1}^{K} n_j(\mathcal{H}, \boldsymbol{B}_{\tau-\kappa})\rho[H_j(\boldsymbol{B}_{\tau-\kappa})], \end{aligned} \tag{D.1}$$

where L depends on how many lagged values of B_τ enter into the strategies H_j. Using Assumption 1, differentiate (D.1) with respect to $B_{\tau-k}$ and evaluate at B^*:

$$\begin{aligned} \left.\frac{dF(\beta, \boldsymbol{B}_{\tau-\kappa})}{dB_{\tau-k}}\right|_{\boldsymbol{B}_{\tau-\kappa}=B^*} &= \frac{\rho(B^*)}{r}\sum_{j=1}^{K}\left.\frac{dn_j(\mathcal{H}, \boldsymbol{B}_{\tau-\kappa})}{dB_{\tau-k}}\right|_{\boldsymbol{B}_{\tau-\kappa}=B^*} \\ &+ \frac{1}{r}\sum_{j=1}^{K} n_j^*(\mathcal{H})\left.\frac{d\rho[H_j(\boldsymbol{B}_{\tau-\kappa})]}{dB_{\tau-k}}\right|_{\boldsymbol{B}_{\tau-\kappa}=B^*}, \end{aligned}$$

where $n_j^*(\mathcal{H}) = n_j(\mathcal{H}, B^*, \ldots, B^*)$. Because the fractions $n_j(\cdot)$ sum to 1, $\sum_{j=1}^{K} dn_j(\cdot)\, dB_{\tau-k} = 0$ so that the first term on the right-hand side of the above equation is zero. Because $H_{\text{nai}(0)}$ has the lowest cost of any strategy, it follows from (11) and (12) that $\lim_{\beta\to\infty} n^*_{\text{nai}(0)}(\mathcal{H}) = 1$, so that

$$\left.\frac{dF(\infty, \boldsymbol{B}_{\tau-\kappa})}{dB_{\tau-k}}\right|_{\boldsymbol{B}_{\tau-\kappa}=B^*} = \frac{1}{r}\left.\frac{d\rho(B_{\tau-\kappa})}{dB_{\tau-k}}\right|_{\boldsymbol{B}_{\tau-\kappa}=B^*}.$$

We have just shown that at B^*, $F(\infty, \boldsymbol{B}_{\tau-\kappa})$ has the same Jacobian as the following system:

$$B_\tau = f(B_{\tau-\kappa}), \tag{D.2}$$

$$f(B) = \frac{\sigma_d\sigma_z^2/r^2 + \sigma_z(\sigma_d\sigma_z + \sigma_d B/r)^2 + \bar{z}^2\sigma_d^2\sigma_z^2}{4r^2\varphi M(B)^2}. \tag{D.3}$$

The system (D.2)–(D.3) describes the model's behavior when all agents use naive expectations, so that $n_{\text{nai}(0)} = 1$ in every period. Differentiating (D.3) with respect to B yields

$$\begin{aligned} f'(B) &= [\sigma_d\sigma_z(\sigma_d\sigma_z + \sigma_d B/r)]/[2r^3\varphi M(B)^2] \\ &- \left[\sigma_d\sigma_z^2/r^2 + \sigma_z(\sigma_d\sigma_z + \sigma_d B/r)^2 + \bar{z}^2\sigma_d^2\sigma_z^2\right][\sigma_d\sigma_z + 2\sigma_d B/r]/[2r^2\varphi M(B)^3]. \end{aligned} \tag{D.4}$$

Now, for each value of φ, let the function $B^*(\varphi)$ denote the steady-state value of B, the market precision. Substituting $B^*(\varphi)$ into (D.3) yields

$$1 = \frac{\sigma_d\sigma_z^2/r^2 + \sigma_z\left[\sigma_d\sigma_z + \sigma_d B^*(\varphi)/r\right]^2 + \bar{z}^2\sigma_d^2\sigma_z^2}{4r^2\varphi M[B^*(\varphi)]^2 B^*(\varphi)}. \tag{D.5}$$

It is then clear that $\lim_{\varphi\to 0} B^*(\varphi) = \infty$, for if not, then the right-hand side of (D.5) could not remain at 1 as $\varphi \to 0$. So, (D.5) implies that

$$\lim_{\varphi\to 0} \varphi B^*(\varphi)^3 = \sigma_z \big/ 4r^2. \tag{**D.6**}$$

Equation (D.4) then implies that $\lim_{\varphi\to 0} f'[B^*(\varphi)] = -\sigma_z/[2r^2\varphi B^*(\varphi)^3]$, which, when combined with (D.6), yields

$$\lim_{\varphi\to 0} f'[B^*(\varphi)] = -2. \tag{**D.7**}$$

By repeating the argument between equation (D.4) and (D.7) for the parameters r, σ_z, and $\bar{z}$, it is straightforward to derive the following limits:

$$\lim_{r\to 0} f'[B^*(r)] = -2,$$

$$\lim_{\sigma_z\to\infty} f'[B^*(\sigma_z)] = -2,$$

$$\lim_{\bar{z}\to\infty} f'[B^*(\bar{z})] = -4,$$

where $B^*(\cdot)$ is defined similarly for r, σ_z, and $\bar{z}$ as it is for φ. There thus exist constants $\{\varphi^c, r^c, \sigma_z^c, \bar{z}^c\}$ such that the system (D.2)–(D.3) is locally unstable whenever one of the four conditions (a)–(d) given in the statement of the proposition is satisfied.

Because $F(\infty, \boldsymbol{B}_{\tau-\kappa})$ and $f(B_{\tau-\kappa})$ have the same Jacobian at B^*, it follows that $F(\infty, \boldsymbol{B}_{\tau-\kappa})$ is also locally unstable at B^* whenever one of these four conditions is met. At least one of the eigenvalues of the Jacobian $DF(\infty, B^*)$ thus must lie strictly outside the unit circle. Because the eigenvalues are continuous functions of the model's parameters, there must exist some critical value β^c such that an eigenvalue of $DF(\beta, B^*)$ lies outside the unit circle whenever $\beta > \beta^c$. We conclude that the system $F(\beta, \boldsymbol{B}_{\tau-\kappa})$ is locally unstable whenever $\beta > \beta^c$ and one of conditions (a)–(d) is met. □