Mathematical Finance, Vol. 11, No. 4 (October 2001), 415–445

RETURN DYNAMICS WHEN PERSISTENCE IS UNOBSERVABLE

Timothy C. Johnson

London Business School

This paper proposes a new theory of the sources of time-varying second (and higher) moments in financial time series. The key idea is that fully rational agents must infer the stochastic degree of persistence of fundamental shocks. Endogenous changes in their uncertainty determine the evolution of conditional moments of returns. The model accounts for the principal observed features of volatility dynamics and implies some new ones. Most strikingly, it implies a relationship between ex post trends, or momentum, and changes in volatility.

Key Words: stochastic volatility, filtering, variable persistence

1. INTRODUCTION

Volatility models in finance have grown increasingly sophisticated in recent years. Spurred by the detailed findings of econometricians and the needs of derivatives practitioners, researchers have proposed specifications that can account for the many complex patterns observed in realized and expected volatility. With complex specifications, however, come the dangers of overfitting and parameter instability. Moreover, the empirical progress has not been matched by a commensurate advance in our understanding of the underlying causes of heteroscedasticity.[1]

Motivated by these considerations, this paper seeks to explain volatility changes in terms of a structural model of economics fundamentals. I derive a theory of heteroscedasticity from a single, simple hypothesis about nature: that some innovations are persistent and some are transient. I show that this basic property of exogenous fundamentals is sufficient to generate the observed features of volatility dynamics, and to imply some new ones.

The outline of the argument is straightforward. Most, if not all, asset pricing models may be viewed as being driven by expectations of a discounted flow of some exogenous fundamental quantity. This may be an earnings flow (for pricing a firm's stock), a money supply flow (for the nominal price level), or a consumption flow (for an endowment claim). If shocks to the flow rate differ over time in their degree of stationarity (or persistence) then investors will vary their reactions to them accordingly: a shock thought

My thanks to George Constantinides, John Cochrane, Nick Polson, Angel Serrat, and Pietro Veronesi for guidance and encouragement. This draft has greatly benefited from the comments of the editor, associate editor, and the two anonymous referees.

Manuscript received June 1999; final revision received July 2000.

Address correspondence to the author at London Business School, Sussex Place, Regent's Park, London NW1 4SA, U.K.; e-mail: tjohnson@lbs.ac.uk.

[1] Although many of the empirical regularities in volatility were noted long ago by Mandelbrot (1963), as recently as 1994 Bollerslev and colleagues, in their comprehensive overview of this area, still contrasted the explosion of descriptive volatility research with "the apparent lack of *any* structural dynamic economic theory explaining the variation of higher order moments" (Bollerslev, Engle, and Nelson 1994), emphasis added.

to have long-term effects will move prices more than one known to be transient. Investors, however, are not likely to know ex ante the exact temporal duration of an innovation. That will only become clear over time. Hence, in determining their immediate reaction to a piece of news, agents must simultaneously solve an inference problem. It is the properties of that inference, then, as much as those of the exogenous series, which determine the properties of price variability.

This point was originally made by Barsky and deLong (1993) in the context of the debate in the finance literature about "excess" stock market volatility. Those authors observed that a fully rational investor who placed *any* prior weight of the possibility that shocks to dividend growth had a permanent component would still hold such a belief with 100 years of data. As a consequence, rationality would dictate much bigger price responses to news—a higher level of volatility—than would be deemed correct by an observer after-the-fact who knew the news to be transient.

The present work extends that logic to the case in which the degree of persistence of shocks need not be fixed. Instead of having a century to estimate a constant parameter, agents must contemporaneously estimate the likely long-term impact of current growth rate changes. Unsurprisingly, their inference problem is never resolved. The parameter uncertainty is not a transient feature of the economy. Instead it affects return dynamics at all times.

Endogenous changes in inferences about persistence thus drive changes in return volatility. GARCH effects—autocorrelation in that volatility—arise for two reasons. First, uncertainty about persistence of *past* shocks means that the current growth rate of fundamentals is unknown, as it depends on their residual influence. This contributes a slowly varying component to current volatility. Second, the direct effect mentioned above is that today's persistence estimate dictates the price response to today's news. But recent persistence carries information about current persistence. Permanent changes are more likely to follow other permanent changes. So this component of volatility is also positively autocorrelated.

Describing the optimal inference process becomes involved mathematically. I have deliberately tried to streamline the economic environment to highlight the role of the expectational dynamics in asset returns. The simplifications involved are pointed out when they arise, and generalizations are considered. The basic properties of the solution are largely robust to the technical specification. The content of the argument is no more than that outlined above.

It is worthwhile to contrast the modeling approach undertaken here with that of simply positing a stochastic volatility process for the exogenous fundamentals. If GARCH effects and other patterns in returns are merely inherited from the underlying information flows driving them, then there is no need for structural theories.

Certainly there is no reason to rule out this possibility. The difficulty is that it seems to fail empirically. Little, if any, of the volatility of returns is traceable to identifiable news (Roll 1984, Cutler, Poterba, and Summers 1989, Berry and Howe 1994). And, when price moves *can* be tied to specific events, the impact on volatility appears to be transient (Jones, Lamont, and Lumsdaine 1995, Andersen and Bollerslev 1997). There are indeed hints of time-varying information flow from the intriguing contemporaneous co-movement between volatility and transactions volume (Tauchen and Pitts 1983, Gallant, Rossi, and Tauchen 1992), but without identification of the news driving both there is no reason to regard volume as any more exogenous than returns.

In contrast it is easy to verify, with hindsight, that fundamental innovations differ in their degree of persistence.[2] Changes to interest rates are clearly persistent when due to policy changes, but otherwise resemble noise. Shocks to a firm's cash flows may be one-time "extraordinary items" or may represent enduring changes to its business conditions. This example illustrates, too, the inference problem facing investors: the distinction between the two types of shocks may be far from obvious. On a macroeconomic level, the question of the permanence or transience of recent productivity increases in the United States (whose answer has enormous implications for both markets and policymakers) is the subject of intense current debate.

In addition to heteroscedastic fundamentals, volatility dynamics may also be influenced by time-varying investor preferences, microstructural effects, heterogeneous information, or investor irrationality. Recent research has begun to examine the implication of each of these. Each still has important shortcomings.

Time-varying marginal utility arguments,[3] since they apply to the pricing of *all* assets, imply that volatility changes should be fully accounted for by systematic components, which they are not. It is also difficult to imagine changes in marginal utilities having much bearing on high frequency patterns.

Microstructural explanations,[4] by contrast, seem ill-suited to explaining low frequency dependencies. Moreover the implication that volatility autocorrelation is a by-product of the trading process itself would seem to be belied by the similarity of the phenomenon across a multitude of instruments and exchange mechanisms.

Models based on heterogeneous information or beliefs[5] would suggest less volatility predictability in markets without significant private information (like currencies and government bonds) which is not the case.

Finally, behavioral models[6] equate volatility persistence with underreaction to news. The implied linkage between return anomalies and heteroscedasticity has yet to receive empirical support.

This is not to doubt that all of the elements above apply to one degree or another. Yet even collectively they are not necessary. The model proposed here gives rise to conditional heteroscedasticity in an environment of perfect markets, rational expectations, homoscedastic fundamentals, symmetric information, and risk neutrality.

The approach of this paper is closely related to that of Detemple (1986), Feldman (1989), Wang (1993), David (1997), and Veronesi (1999). These authors also introduce inferential uncertainty into asset pricing problems, in the form of an unobservable growth rate of fundamentals, and thence derive endogenous volatility processes.[7] The model here complements these works in delineating a new type of parameter uncertainty, which corresponds to a realistic and economically important problem. The dynamics I deduce include features that are potentially significant and distinct from those previously studied. These implications are readily testable.

[2] Evans (1998) reports support for a generalized switching model of dividend growth rate that includes significant shifts in persistence between regimes. Evans and Lewis (1995) find evidence of variable persistence of U.S. interest rate shocks.

[3] See Campbell and Cochrane (1999) or Wang (1996).

[4] Examples include Andersen (1996), Easly and O'Hara (1992), and Foster and Viswanathan (1990).

[5] See Campbell and Hentschel (1992), Wang (1993), and Detemple and Murthy (1994).

[6] Barberis, Shleifer, and Vishny (1998) present a model not unlike the one in this paper of changing *but mistaken* beliefs about persistence of fundamentals. Models of evolving belief populations that explicitly focus on volatility include those of Brock and LeBaron (1996) and Lux (1997).

[7] A parallel literature treats the consumption/portfolio problem of an investor under parameter uncertainty. See Honda (1997) and the references therein.

The outline of the paper is as follows. The next section formalizes the information problem and derives two characterizations of the solution. Section 3 introduces a stylized asset pricing relation, which makes the role of conditional expectations particularly transparent. Using this and the solution techniques of Section 2, I derive the system of equations governing the evolution of return moments and prove a result crucial for empirical work. Section 4 explores the moment system and derives a theorem on the relationship between volatility and trends. Numerical examples investigate the robustness of the assumptions of the theorem and some of its extensions. Section 5 contains some concluding remarks.

2. THE FILTERING PROBLEM

The goal of this section is to formulate and solve the problem of inference about the degree of persistence of shocks to growth rates. For present purposes, the interpretation given to the observed series is irrelevant. The immediate task is mathematical.

The problem is easy to frame in discrete time. If D_t is the fundamental economic series, then the situation is

$$\Delta D_t = \mu_t + \epsilon_{t+1} \tag{2.1}$$

$$\mu_t = \mu_{t-1} + S_t \epsilon_t, \tag{2.2}$$

where ϵ_t is white noise and μ_t, S_t, and ϵ_t are all unobserved. Here S_t determines the degree of persistence: When $S_t = 0$ the growth rate ΔD_t is completely stationary but when $S_t = 1$ it switches to a random walk. The problem is to describe the minimum mean-square estimators (posterior expectations) of μ and S given observations of D.

In fact, it will be necessary to derive the full conditional joint distribution of the unknown variables. To do so it is easier to work in continuous time. Here the natural analogy is the system of stochastic differentials

$$dD_t = \mu_t dt + \sigma_0 \, dW_t \tag{2.3}$$

$$d\mu_t = \sigma_0 S_t \, dW_t \tag{2.4}$$

with W_t a Wiener process and S_t as before. The analogy may be justified formally by appealing to the Wold representation for the observation process. Using the Fubini theorem for stochastic integrals (cf. Protter 1990, VI.45), we may write the continuous-time version as

$$D_t - D_0 = \mu_0 t + \sigma_0 \int_0^t [1 + S_u(t-u)] \, dW_u,$$

which corresponds exactly to that of the discrete-time system (2.1)–(2.2) with integration replaced by summation.

REMARK 2.1. A significant difference between the two systems concerns the magnitude of the long-run effect of a differential shock dW_t. In the discrete system, when $S_t = 1$ the full impact of the shock (of order $(\Delta t)^{1/2}$) is incorporated in all future changes in D. In the continuous system, such a permanent innovation only alters dD_s at times

$s > t$ via the $\mu_s ds$ term, an effect of order only $(\Delta t)^{3/2}$. This suggests employing a different scaling factor in (2.4) to allow stronger effects; that is,

$$d\mu_t = \sigma_1 S_t \, dW_t. \tag{2.5}$$

Mathematically, this would be identical to having S take values in $\{0, \kappa\}$, where $\kappa \equiv \sigma_1/\sigma_0$, instead of $\{0, 1\}$. This generalization could well be important empirically. The results that follow go through with minor changes to some formulas, which I will note in the proofs. For notational simplicity, however, I use the original formulation in the main development.

Proceeding formally, the setting that will be used throughout the paper is as follows. Fix a probability space $(\Omega, \mathcal{F}, P)$ and a filtration $(\mathcal{F}_t)_{t=0}^T$, $\mathcal{F} = \sigma(\cup \mathcal{F}_t)$ satisfying the usual conditions. On this space define two processes: W_t, a standard Wiener process, and S_t a two-state Markov process, independent of W_t, taking values 0 and 1 characterized by transition parameters λ_0 and λ_1. That is, the probability of a switch from 0 to 1 in time Δt is $\lambda_0 \Delta t$, and that from 1 to 0 $\lambda_1 \Delta t$. S_t has the semimartingale representation

$$dS_t = (\lambda_0 - (\lambda_0 + \lambda_1) S_t) \, dt + (1 - 2S_t) \, dN_t^\lambda, \tag{2.6}$$

where N_t^λ is a compensated Poisson process with intensity $\lambda = \lambda(S) = \lambda_1 S + \lambda_0 (1 - S)$.

REMARK 2.2. The discrete nature of the S process is in no sense fundamental to the results, as will be shown bellow. The independence of dN_t^λ and dW is also easy to relax, though correlation effects can result in nontrivial differences in the dynamics. The assumption that is significant, mathematically and economically, is that S is a first-order Markov process. This is tantamount to the assertion that current persistence conveys information about future persistence. Without this, S will only be identified (to observers of D) by the imposition of some other structural assumption. This will fundamentally change the nature of the system.

Next, let $\sigma_0 > 0$ be a fixed constant, and define the process μ_t as the unique (strong) solution to (2.4) (in conjunction with (2.6)) having initial value μ_0 independent of S_t and W_t with all moments finite. Last, let D_0 be another such independent random variable and define the process

$$D_t = D_0 + \int_0^t \mu_u \, du + \sigma_0 W_t. \tag{2.7}$$

The information available to an observer of D will be denoted $\mathcal{F}_t^D \equiv \sigma(D_s, 0 \le s \le t)$. Conditional expectations with respect to $\mathcal{F}_t^D$ describe the observer's minimal mean-square inferences. This information set is assumed to include full knowledge of the static parameters describing the transition probabilities of S. (Here that entails knowing λ_0, λ_1, and that the two possible values of S are 0 to 1. Some more general processes for S will be mentioned below.)

It is worth emphasizing, at this point, that the assumption throughout of a known variance parameter σ_0 is without loss of generality from the point of view of filtering. The reason is that in continuous time the instantaneous volatility of any continuous process X_t is $\mathcal{F}_t^X$ measurable. There is no inference to be made. Problems of variance inference are simply not well posed in this setting. But this is not a drawback here, nor is the assumed constancy of σ_0 (which is not a technical necessity). The volatility we wish to study is not that of the fundamental series D, but that of functionals of it (i.e., returns). One of the major goals of the paper is to illustrate that the complicated dynamics of the latter are endogenous to the filtering problem and not inherited from the former.

Readers familiar with filtering problems will recognize some nonstandard features of the above system. The state space is mixed, having a discrete and a continuous variable. Problems with this feature have been studied in the systems literature under the heading hybrid (or jump-linear) systems.[8] More peculiarly, there is no distinction between the "system noise" and the "observation noise": the same Wiener process drives μ_t and D_t. While important results have been obtained for problems containing correlated noises (for nonjumping systems), some degree of independence is typically required or the problem becomes degenerate.[9]

Interestingly, in the current problem this latter degeneracy enables one to identify explicitly the Radon–Nikodym derivative defining the observer's conditional expectations.

PROPOSITION 2.1. *Denote by Ω' the space $(C[0,T] \times \{0,1\}^T \times R)$ with generic element $\omega' = (W', S', \mu_0')$ and endow this space with the Borel σ-algebra. Let $f(\omega)$ be in $L^1(P)$. Put*

$$\zeta_t(\omega';\omega) = \delta(\overline{W}) \cdot e^{-1/2\int_0^t \frac{m_u^2(\omega';\omega)}{\sigma_0^2}\,du + \int_0^t \frac{m_u(\omega';\omega)}{\sigma_0^2}\,dD_u(\omega)}, \tag{2.8}$$

where

$$m_t(\omega';\omega) = e^{-\int_0^t S_u'\,du}\Big[\mu_0' + \int_0^t S_u' e^{\int_0^u S_v'\,dv}\,dD_u(\omega)\Big] \tag{2.9}$$

and $\delta(\cdot)$ places unit mass at the path

$$\overline{W}_s(\omega';\omega) = \sigma_0^{-1}\Big[D_s(\omega) - D_0(\omega) - \int_0^s m_u(\omega';\omega)\,du\Big], \qquad 0 \le s \le t. \tag{2.10}$$

Then

$$E[f|\mathcal{F}_t^D](\omega) \equiv \widetilde{E}_t f = \int_{\Omega'} f(\omega')\zeta_t(\omega';\omega)\,dP(\omega') \int_{\Omega'} \zeta_t(\omega';\omega)\,dP(\omega'). \tag{2.11}$$

Proof. All proofs are in the Appendix.

REMARK 2.3. I use the tilde notation throughout (including $\widetilde{\mathrm{var}}$, $\widetilde{\mathrm{cov}}$, $\widetilde{\mathrm{skew}}$, etc.) to denote posterior expectations with respect to the σ-algebra generated by the observations. If the time subscript is suppressed, the integrand and the information set should be assumed contemporaneous. Thus $\widetilde{\mathrm{var}}(\mu) \equiv \mathrm{var}(\mu_t|\mathcal{F}_t^D)$ and $\widetilde{X} \equiv E[X_t|\mathcal{F}_t^D]$ for any process X_t.

[8] Originally these seem to have arisen in connection with target tracking problems. See Loparo, Roth, and Eckert (1986), Dufour and Bertrand (1994), and Björk (1980, 1982); or in discrete time Krishnamurthy and Evans (1998), Elliott, Dufour, and Sworder (1996), or Mariton (1990).

[9] See Davis (1981), Florchinger (1993), or Florchinger and Gland (1991).

The proposition provides an implementable method of computing posterior expectations for this system. Examination of the result shows that, for an observer of the D process, the set of unknowns is just the history of the switching process S (plus the scalar random variable μ_0). There is no need to try to estimate the Brownian motion path W. This makes the task much more practicable. While both S and W are uncountable collections of random variables, the former is completely characterized by the collection of (exponential) random variables defining its jumps. And in any finite time interval the chance that there are more than K of these declines exponentially in K. So in effect the parameter space over which inferences need to be made is of small dimension. This makes it feasible to perform the function space integration in equation (2.11) by Monte Carlo means.

Such calculations will be used extensively below to study properties of the system. The limitation of this characterization of the filtering solution is that it does not directly give the dynamic evolution of estimators in terms of the data. The next proposition affords a way to update conditional expectations of the quantities of interest recursively. The formulas are given for specific conditional moments, but the technique allows a similar description of arbitrary posterior expectations.

PROPOSITION 2.2. *Let $\widetilde{\mu^k S_t}$ and $\widetilde{\mu_t^k}$ denote $E(\mu_t^k S_t|\mathcal{F}_t^D)$ and $E(\mu_t^k|\mathcal{F}_t^D)$. Define $\widetilde{W}_t \equiv (D_t - D_0 - \int_0^t \tilde{\mu}_s ds)/\sigma_0$. Then*

(I) *$(\widetilde{W}_t, \mathcal{F}_t^D)$ is a Wiener process.*

(II) *$\widetilde{\mu^k S_t}, \widetilde{\mu_t^k}$ exist for all t and all (integer) $k \geq 0$ and satisfy the stochastic differential system*

$$d\widetilde{\mu_t^k} = \left(k(k-1)\widetilde{\mu^{k-2} S_t}\frac{\sigma_0^2}{2}\right)dt + \left(\widetilde{\mu_t^{k+1}} - \widetilde{\mu_t^k}\tilde{\mu}_t + k\widetilde{\mu^{k-1}S_t}\sigma_0^2\right)d\widetilde{W}_t/\sigma_0 \tag{2.12}$$

$$\begin{aligned} d\widetilde{\mu^k S_t} = &\left(\widetilde{\mu_t^k}\lambda_0 - \widetilde{\mu^k S_t}(\lambda_0+\lambda_1) + k(k-1)\widetilde{\mu^{k-2}S_t}\frac{\sigma_0^2}{2}\right)dt \\ &+ \left(\widetilde{\mu^{k+1}S_t} - \widetilde{\mu^k S_t}\tilde{\mu}_t + k\widetilde{\mu^{k-1}S_t}\sigma_0^2\right)d\widetilde{W}_t/\sigma_0 \end{aligned} \tag{2.13}$$

with initial conditions in accordance with the prior.

The first part of the proposition tells us what the D process looks like marginally—that is, with respect to the information in $\mathcal{F}_t^D$. Rewriting the definition of $d\widetilde{W}_t$, we see,

$$dD_t = \widetilde{\mu}_t dt + \sigma_0 d\widetilde{W}_t. \tag{2.14}$$

This representation is a generic filtering result, which does not depend on specific features of the current model. Since $d\widetilde{W}_t$ is $\mathcal{F}_t^D$ measurable, the remarkable implication is that, given his own best estimate of the true drift $\widetilde{\mu}_t \equiv E[\mu_t|\mathcal{F}_t^D]$, an observer of D still views it as an Itô process. D does not appear Markovian, however, because the drift estimate depends on the entire observation history.

The second part of the proposition provides an explicit algorithm for updating that estimate, and all other integer moments of the joint distribution, in real-time. The algorithm may be viewed as consisting of two steps. First, the new datum dD is stripped of its expected component $\tilde{\mu}dt$ to yield the innovation series $\sigma_0 d\widetilde{W}$. Then each old estimate is updated by some multiple of this "shock" plus a trend term, which is just the expected trend of the parameter. The coefficients for the updating, in the case of the integer moments in the proposition, are themselves combinations of integer moments. However the system is not closed: the kth equation contains terms involving $(k+1)$th moments.[10] Nevertheless it can still provide important insights into the model's dynamics. That is the subject of the next section. The remainder of this section considers the effect on the filter in part (II) of some generalizations to the system specification.

The proposition is an application of a classical filtering result due to Fujisaki, Kallianpur, and Kunita (1972) which requires much less structure than has been imposed. Several natural extensions, in fact, lead to only minor changes in the coefficients given above.

If, for example, the innovations to μ and D are not perfectly correlated—perhaps due to an extraneous source of observation noise—then (2.12) becomes

$$d\widetilde{\mu_t^k} = \left(k(k-1)\widetilde{\mu^{k-2} S_t} \frac{\sigma_1^2}{2}\right) dt + (\widetilde{\mu_t^{k+1}} - \widetilde{\mu_t^k}\tilde{\mu}_t + k\widetilde{\mu^{k-1} S_t}\rho\sigma_0\sigma_1) d\widetilde{W}_t/\sigma_0, \tag{2.15}$$

where now σ_1 and σ_0 are the diffusion coefficients of $d\mu$ and dD respectively, and ρ is their correlation. The modification to (2.13) is analogous.

If the persistence process is allowed to take multiple values or even vary continuously, then a full description of the filter will include all cross moments of the form $\widetilde{\mu^k S^j}$. However for any S that has the representation

$$dS_t = b(\overline{S} - S_t)\,dt + dM_t$$

(where M is a martingale and b and $\overline{S}$ are constants), the equations in the proposition require only one change. With the identifications $(\lambda_0 + \lambda_1) = b$ and $\lambda_0 = b\overline{S}$, the optimal filter is as above but with $k(k-1)\widetilde{\mu^{k-2} S_t^2}\sigma_0^2/2$ replacing $k(k-1)\widetilde{\mu^{k-2} S_t}\sigma_0^2/2$ in each drift term.

A third interesting generalization would incorporate more sources of information about the unobserved parameters. Additional information about μ—direct news about growth rates—could be modeled by a multidimensional observation D and a general correlation structure:

$$\begin{aligned} dD_t &= \mu_t \iota\, dt + \Sigma_0\, dW \\ d\mu_t &= S_t \Sigma_1\, dW \end{aligned}$$

with ι a unit vector of appropriate dimension and S_t now possibly a matrix-valued persistence process. The notation becomes more involved. But the expressions for the drift estimator $\tilde{\mu}$, for example, illustrate the fact that the essential structure of the filter is preserved. With scalar observations and $k = 1$ in (2.12), we have

$$d\widetilde{\mu}_t = (\widetilde{S}_t\sigma_0^2 + \widetilde{\mathrm{cov}}_t(\mu))\, d\widetilde{W}_t/\sigma_0.$$

The vector equivalent is just

$$d\widetilde{\mu}_t = (\widetilde{S}_t\Sigma_1\Sigma_0' + \widetilde{\mathrm{cov}}_t(\mu)\iota)(\Sigma_0\Sigma_0')^{-1/2}\, d\widetilde{W}_t,$$

where now the innovation series is $d\widetilde{W}_t = (\Sigma_0\Sigma_0')^{-1/2}(dD - \widetilde{\mu}_t\iota\, dt)$. The two significant points about the comparision are (1) the expected degree of persistence $\widetilde{S}_t$ still just enters

[10] Using the results of Björk (1980), it can be shown that no finite-dimensional filter exists for this system; that is, any other system of statistics describing the joint posterior distribution will likewise be infinite dimensional.

linearly in the diffusion coefficient, and (2) the process $d\widetilde{W}_t$ drives both the observation series and the estimator updates, just as before.

The foregoing discussion demonstrates that the specialized system introduced at the beginning of the section is not required to obtain the type of dynamic equations for the estimators derived in Proposition 2.2. However, as mentioned above, these equations (as well as the generalizations) are difficult to implement empirically due to their recursivity. The earlier Proposition 2.1 did fully exploit the structural restrictions in order to obtain an implementable characterization of inferences. Below I will use both results to study general and specific properties of the system.

3. RETURN DYNAMICS

The interpretation to be given to the filtering problem of the preceding section is of an economic series (dD_t) whose shocks have an unobservable dichotomous trait: persistence or transience. To study the implications this could have for financial time series a mapping is needed between the observer's conditional expectations and prices. This section introduces a generic model that makes the role of expectations particularly transparent. Some interpretations of the pricing equation are considered, and a proposition clarifying the relationship between the univariate information in price histories and agents' full information set, $\mathcal{F}^D$, is established. Next I derive a system of stochastic differential equations describing return dynamics from Proposition 2.2. Then I use Proposition 2.1 to illustrate the dynamics for a specific case with parameters motivated by an exchange rate interpretation of the model. Volatility patterns are shown to be broadly similar to those observed in practice.

3.1. A Generic Asset Pricing Equation

To begin, consider a one-good world with a single risky asset paying dividends continuously at rate D_t. Suppose there is also an elastically supplied risk-free storage technology with fixed return r. Identify the stock market level with the shadow price of the endowment for a single representative agent, risk-neutral and infinitely lived. A standard result is given in Lemma 3.1.

LEMMA 3.1. *Under the above assumptions, and with the system and observation dynamics given by equations* (2.4), (2.6), *and* (2.7), *the price of the stock is*

$$P_t = \frac{D_t}{r} + \frac{\tilde{\mu}_t}{r^2}. \tag{3.1}$$

Equation (3.1) leaves much to be desired as a model of stock prices. It ignores risk-aversion, uses constant discount rates, and permits negative prices. Nevertheless it motivates consideration of the general class of models that may be written

$$P_t = \theta_0 + \theta_1 D_t + \theta_2 \tilde{\mu}_t, \tag{3.2}$$

where the θs are constant, D is some fundamental series, and μ is its growth rate. The notion is that, at least to first order and/or for limited time spans, a variety of pricing problems reduce to adding a capitalized growth rate to a current intrinsic value.

A well-known example in the finance literature is the "dividend-ratio" model of Campbell and Shiller (1988). Using a log-linear approximation to a present-value identity, those authors express the log dividend-price ratio, δ_t, of a stock as

$$\delta_t = \text{constant} - \widetilde{E}_t\left(\sum_j \rho^j \Delta d_{t+j}\right),$$

where ρ is a constant and Δd_t is the dividend growth rate. If the expected value of that growth rate is the same for all future periods (as it is under the specification in the last section), then the term on the right is just a constant times my $\tilde{\mu}_t$. So again one recovers (3.2) (with $\theta_1 = 1$) by solving for the log stock price. The Campbell–Shiller derivation is valid in any economy with constant expected returns, subject to the conditions mentioned. Moreover it shows that (3.2) can be interpreted in log terms, so that changes in P are returns.

As another example, one can interpret P as the log exchange rate in a flexible-price monetary model (as in Frenkel and Mussa 1985). Now D_t is the difference in money supply growth rates between two countries; θ_0 is zero, θ_1 is one, and θ_2 is the semi-elasticity of money demand with respect to interest rates. The argument (outlined in Johnson 1999) essentially only uses purchasing power parity and a log linearization of real money demand.

Under this latter interpretation, (3.2) is just a standard macroeconomic model of exchange rate determination. Much of the empirical work on conditional heteroscedasticity has also focused on currencies. So this case may be particularly relevant. Moreover, as it is expressed in terms of log prices, like the Campbell–Shiller model, negative values pose no conceptual difficulty. For these reasons, I frame most of the discussion of the model below in terms of this interpretation.

Equation (3.2) describes prices as a function of conditional expectations, where the conditioning is with respect to a second series, D. In practice, however, the volatility dynamics of interest are the *univariate* ones. That is, we ultimately want to know what the return series looks like when conditioning only on its own history. So a description of the dynamics in terms of $\mathcal{F}^D$ measurable quantities will have to be projected onto the smaller information set $\mathcal{F}^P$.

This projection is a second filtering problem. Solving it (e.g., obtaining a characterization of $E_t(\tilde{\mu}_t|\mathcal{F}_t^P)$) is, in general, a much harder task than solving the one in Section 2. This is because the system being projected is now infinite dimensional and none of its components is independent of the others. So it appears that the *marginal* price dynamics may be vastly more complicated than the bivariate ones. The following proposition shows that this is not the case.

PROPOSITION 3.1. *If prices are given by a model of the form* (3.2), *where the system and observation dynamics are given by* (2.4), (2.6), *and* (2.7), *then, assuming* $P(\mu_0 \leq \alpha|\mathcal{F}_0^P) = P(\mu_0 \leq \alpha|\mathcal{F}_0^D)$, $\forall\alpha$, *we have*

$$\mathcal{F}_t^P = \mathcal{F}_t^D \, \forall t.$$

The happy result, is that—for filtering purposes—prices are fully revealing whenever (3.2) applies. Conditional expectations with respect to price histories alone not only are determined by but actually coincide with those with respect to the joint history of prices and fundamentals. Thus the models already derived *are* univariate ones. The driving series of (normalized) fundamental innovations $\widetilde{W}_t$ is also the (normalized) price innovation series.

The proposition holds in this model because there is only one exogenous variable, D, whose changes can be recovered from those of P due to the monotonicity of their effect on it. But the proof is not dependent on the restricted nature of the true parameter evolution. In particular, neither the dichotomous property of the persistence variable, nor the absence of outside noise sources in the equations for dD and $d\mu$ is invoked.

Besides greatly simplifying the development, the proposition also makes the analysis more robust in the sense that the results do not depend on a particular identification of D. Stock prices may be driven by consumption or perhaps by dividends. Both money supplies and fiscal deficits are reasonable candidates for the determinant of exchange rates. We need not decide. As long as the dynamics of the exogenous series are described by a system of the form introduced in Section 2, and as long as prices are approximately linear in its growth rate, the details of the economy do not matter. The results below will still be correct.

3.2. Volatility

If asset prices (or log prices) are given by equation (3.2), then changes are just a weighted sum of the change in the exogenous fundamental D and in the endogenous estimate of its drift.

In understanding the behavior of returns, it is instructive to start by examining the dynamics that would obtain without unobservability. If the growth rate μ could be inferred without error, then $\tilde{\mu} = \mu$ and both components of returns are exogenous. The dynamics in this case reveal the content of the assumptions made about the stochastic environment. Combining (2.3), (2.4), and (3.2) yields

$$dP_t = \theta_1 \mu_t \, dt + \sigma_0(\theta_1 + \theta_2 S_t) \, dW_t.$$

This illustrates the intuition given in the introduction. In the presence of *any* persistence, returns will display "excess volatility"—that is, strictly more than contributed by the variation in fundamentals. Changes in persistence translate directly into changes in volatility. In particular the volatility process will inherit the autocorrelation structure of S_t. Note that none of these observations would change if innovations to D and μ were not perfectly correlated. If equation (2.4) were replaced by

$$d\mu_t = \sigma_1 S_t \, dV_t,$$

where dV is a separate Wiener process whose correlation with dW is ρ, then the volatility expression would just be

$$(\theta_1^2\sigma_0^2 + 2\theta_1\theta_2\rho\sigma_0\sigma_1 S_t + \theta_2^2\sigma_1^2 S_t^2)^{1/2},$$

which has the same properties just described.

By restricting agents' information to the fundamental series D, we replace this description of volatility in terms of an unexplained second state-variable with one governed by expectations about that variable. We will see below that the distinction between the two models is, in fact, testable. However the full model retains the conclusions of the preceding paragraph, with expected persistence substituted for true persistence.

Deriving the dynamics of the full model is a straightforward application of Proposition 2.2. Equation (3.2) now gives the system

$$dP_t = \theta_1 \tilde{\mu}_t \, dt + \tilde{h}_t \, d\widetilde{W}_t \tag{3.3}$$

$$d\widetilde{W}_t \equiv (dD_t - \tilde{\mu}_t \, dt)/\sigma_0 \tag{3.4}$$

$$\tilde{h}_t = \theta_1 \sigma_0 + \theta_2 \hbar_t \tag{3.5}$$

$$\hbar_t \equiv \sigma_0 \widetilde{S}_t + \widetilde{\text{var}}_t(\mu)/\sigma_0 \tag{3.6}$$

$$d\tilde{\mu}_t = \hbar_t \, d\widetilde{W}_t \tag{3.7}$$

$$d\hbar_t = \sigma_0[(\lambda_0 + \lambda_1)(\overline{S} - \widetilde{S}_t) + (\widetilde{S}_t - \hbar_t^2/\sigma_0^2)] \, dt + g_t \, d\widetilde{W}_t \tag{3.8}$$

$$\overline{S} \equiv \lambda_0/(\lambda_0 + \lambda_1) \tag{3.9}$$

$$g_t \equiv 3\widetilde{\text{cov}}_t(\mu, S) + \widetilde{\text{skew}}_t(\mu)/\sigma_0^2 \tag{3.10}$$

$$\tilde{g}_t \equiv \theta_2 g_t. \tag{3.11}$$

Look first at equations (3.5) and (3.6). They show that, in addition to replacing S by its expectation, there is another difference from the volatility of the simplified, observable case. This is the contribution of uncertainty about μ. The inferred growth rate is more volatile than the parameter itself, since it is driven by the inferred innovation $d\widetilde{W}$, not dW. In consequence, returns are more volatile when the posterior variance of μ is high. This is the indirect influence of unobservability on the system.

The term $\hbar$ is necessarily positive. If the pricing model is normalized to have $\theta_1 > 0$ and $\theta_2 > 0$ (which just corresponds to placing positive value on the flow D), then $\tilde{h}$, the volatility of P, is positive as well. In fact, it is bounded away from zero.

Equation (3.7) shows that the growth rate estimator inherits the martingale property of the true growth rate. Had the specification included a mean-reverting drift for μ (perhaps to ensure its long-run stationarity) the effect would be to add the expected value of that drift to this equation. The important point is that the diffusion coefficient of $d\tilde{\mu}$ would be unaltered. Since this is the term that generates return heteroscedasticity, our volatility analysis would be robust to this generalization.

Now consider (3.8) which describes the evolution of volatility. The first thing to notice is that this defines a qualitatively new class of continuous time volatility models. Unlike discrete time GARCH models or their degenerate diffusion limits (Kind, Liptser, and Runggaldier 1991, Corradi 1997), it is the price innovations *themselves*, not their square or absolute value, that enter here. This implies that volatility will be path-dependent, and that return shocks will have asymmetric effects on volatility. The sign of the impact is governed by the term g. Since this term is a combination of two posterior moments whose sign may vary, the correlation between returns and volatility may switch sign. This correlation, in turn, determines the skewness of the finite-horizon return distribution. The model thus allows time-varying return skewness, which would imply changes over time in the implied-volatility "smiles" in options prices.

Another distinctive feature of the model is the implication that, although volatility is stochastic, it is completely determined by the path of return innovations, $\widetilde{W}_t$. Unlike ordinary stochastic volatility formulations, or the continuous-time GARCH limits of Nelson (1990), there is no outside source of noise driving the variation in $\hbar$ (or $\tilde{h}$). This, too, has important implications for options pricing as it implies that, despite the randomness in volatility, markets remain complete. (Hobson and Rogers 1998 also recently introduced a class of volatility models with this property.) In fact, in the present model, the completeness is even stronger in that it is not only volatility that is stochastic, but the volatility of volatility, and *its* volatility, and so on.[11]

The drift term in (3.8) is a quadratic form in $\widetilde{S}_t$ and $\widetilde{\mathrm{var}}_t(\mu)/\sigma_0^2$. It is always negative (positive) for high (low) enough value of either argument. Hence the equation implies that volatility will be more strongly mean reverting than in the linear specification most commonly used in stochastic volatility models. We cannot infer, however, that $\hbar$ will be a stationary process. But the equation does tell us that its drift will be zero when $\widetilde{S}_t$ is at its unconditional mean $\overline{S} \equiv \lambda_0/(\lambda_0+\lambda_1)$ and $\hbar = \sigma_0\sqrt{\overline{S}}$. For this reason, it is appropriate to view $\sigma_0(\theta_1 + \theta_2\sqrt{\overline{S}})$ as the steady-state volatility of the system.[12]

Much harder to analyze is the diffusion term in (3.8), denoted g, which involves higher posterior moments of agents' beliefs. As noted above, the system of posterior moment equations is not closed. So while one could use Proposition 2.2 to write down the dynamics of the components $\widetilde{\mathrm{cov}}_t(\mu, S)$ and $\widetilde{\mathrm{skew}}_t(\mu)$, these would only involve higher moments still. Section 4 will make some analytical progress by introducing auxiliary assumptions. For now, though, the best way to gain insight into the system's properties is to exploit the alternative characterization of the agent's beliefs given by the function space integral in Proposition 2.1. I use this next to compute all the posterior moments for a specific example. This produces a representative time-series of volatility and the volatility of volatility, and permits analysis of the individual components.

3.3. Time-Series Characteristics of Volatility

To investigate the volatility patterns described by equation (3.8), I study its realization under a particular parameter configuration chosen to represent plausibly the structural restrictions of an exchange rate model, At issue, specifically, is whether or not the model reproduces the type of volatility autocorrelation found so consistently in practice.

The technique is straightforward. Once the structural parameters are chosen, the true values of the state variables (μ_0, D_0, and the processes S_t and W_t) are drawn from their distributions at daily frequency for 25 years. These values imply values for the true growth rate process μ_t and the observed series D_t, yielding a full history of the economy. Then the history of agents' contemporaneous expectations is calculated from their observation data D, using the formula derived in Proposition 2.1.[13] These expectations are then converted to returns using equation (3.2) with coefficients chosen as described below. Besides the return series, histories of the instantaneous volatility $\tilde{h}_t$ and its volatility $\tilde{g}_t$ are likewise constructed.

To select parameter values relevant in an exchange rate context, I employ the simple monetary model mentioned in Section 3.1. The D series is then identified as the difference in log real money balances between two countries. In the United States the standard

[11] It is worth remarking that the perfect correlation between volatility and returns is unrelated to the perfect correlation assumed between the observable series D and its true drift μ (cf. equation (2.15) above). Rather, it is simply a consequence of the fact that both quantities are now endogenous, and hence must be driven by the same external news.

[12] If the persistence parameter κ differs from unity (cf. Remark 2.1), then $\hbar$ becomes $\sigma_0\kappa\widetilde{S}_t + \widetilde{\mathrm{var}}_t(\mu)/\sigma_0$ and $d\hbar$ becomes $[\sigma_0\kappa(\lambda_0 - (\lambda_0+\lambda_1)\widetilde{S}_t) + (\sigma_0\kappa^2\widetilde{S}_t - \hbar_t^2/\sigma_0)]\,dt + [3\kappa\,\widetilde{\mathrm{cov}}_t(\mu, S) + \widetilde{\mathrm{skew}}_t(\mu)/\sigma_0^2]\,d\widetilde{W}_t$. The steady-state volatility is thus $\sigma_0(\theta_1 + \theta_2\kappa\sqrt{\overline{S}})$.

[13] The function space integration involved in this step is carried out by Monte Carlo sampling until the numerical error for all quantities is less than $\pm 2.5\%$. Euler approximations are used for the stochastic integrals. Adding higher order corrections of the Milstein type produces very small changes in the expectations and does not affect the conclusions. Notice also that the proposition gives the exact continuous-time Radon–Nikodym derivative for each day, not that of a discretized approximation. So the choice of the time interval does not affect the expectations.

deviation of changes in M1 has been between 1% and 3% per annum over the last 20 years, with similar values for other large economies. Across countries money growth covaries as business cycles do, suggesting a volatility, σ_0, of D of perhaps 3%.

The parameters λ_0 and λ_1 may be interpreted as the inverse of the expected duration (in years) of regimes of transitory and permanent shocks to D, respectively. Episodes of permanent shocks, which may be thought of as periods of monetary policy adjustment, are likely to be both infrequent and short-lived. I take $\lambda_0 = 1/3$ and $\lambda_1 = 3$, implying $\overline{S} = 0.10$, which corresponds to assuming that, on average, the economy is in its stable phase 90% of the time, being interrupted every 3 years by a structural change taking, on average, 4 months to resolve. It is difficult to gauge the validity of these exact choices as the current specification has not been considered in the empirical monetary literature.[14]

In choosing the preference parameters θ_0, θ_1, and θ_2, there is, in fact, just one further degree of freedom. The monetary model implies that $\theta_0 = 0$ (a consequence of purchasing power parity) and $\theta_1 = 1$ (via the Fisher hypothesis). The theory also imposes constraints on θ_2. This coefficient corresponds to the semielasticity of money demand with respect to nominal interest rates. Empirical estimates of this quantity typically find values between 5 and 10.[15] I choose a figure within this range, the precise value of which (9.55) is taken to set the unconditional volatility of the simulated returns to 12% per annum, a typical number for a major exchange rate in the last decade. The fact that the economic theory supports this calibration reinforces the plausibility of the parameter configuration as a whole.

The key properties of the simulated returns can be seen in Figure 3.1. The top panel shows the log exchange rate history, along with its two components $\theta_1 D$ and $\theta_2 \tilde{\mu}$. The components are perfectly correlated and about the same size. But the growth rate term contributes most of the volatility. The second panel shows the history of that volatility, $\tilde{h}_t$, and of the two components that comprise it. Of these, the posterior variance term (middle line) is both larger and more volatile than the expected persistence term (lowest line). However these components too are nearly perfectly positively correlated. This observation, not obvious from the analysis in the last section, indicates that *both* components may be viewed as being driven by changing beliefs about persistence.

The composite volatility term, $\tilde{h}_t$, clearly exhibits the positive autocorrelation typical of many financial series. In addition, the nonlinear mean reversion is apparent, with large spikes in volatility being damped more strongly than small changes about the long-run mean. This implies a behavior that would seem nearly integrated locally (or in small samples) but stationary globally.

Neither stationarity nor ergodicity of the $\tilde{h}_t$ process has been shown analytically. So properties of its autocovariance function remain conjectural, and may not resemble simulated values. Nevertheless, the top panel of Figure 3.2, showing the sample autocorrelation, strongly supports the visual impression. This particular parameterization exhibits significantly positive volatility autocorrelations to lags of at least two years. In this sense, the model can be said to imply the existence of GARCH effects.

Properties of the true volatility process are not, however, directly comparable to actual data, since, in practice, volatility cannot be estimated perfectly. It is not clear whether

[14] Kaminsky and Lewis (1996) find evidence of regime shifts in U.S. monetary fundamentals in the second half of the 1980s. The regimes (which are restricted to have the same switching probabilites) have expected duration of between 77 and 280 weeks, depending on the specification.

[15] See Baba, Hendry, and Starr (1992) and Ball (1998) for evidence regarding this parameter in the United States. The model implicitly assumes the same value for both countries. Also note that the semi-elasticity is a negative number. It becomes positive, as in the text, when the exchange rate is in units of the foreign currency, which is consistent with $\theta_1 = +1$.

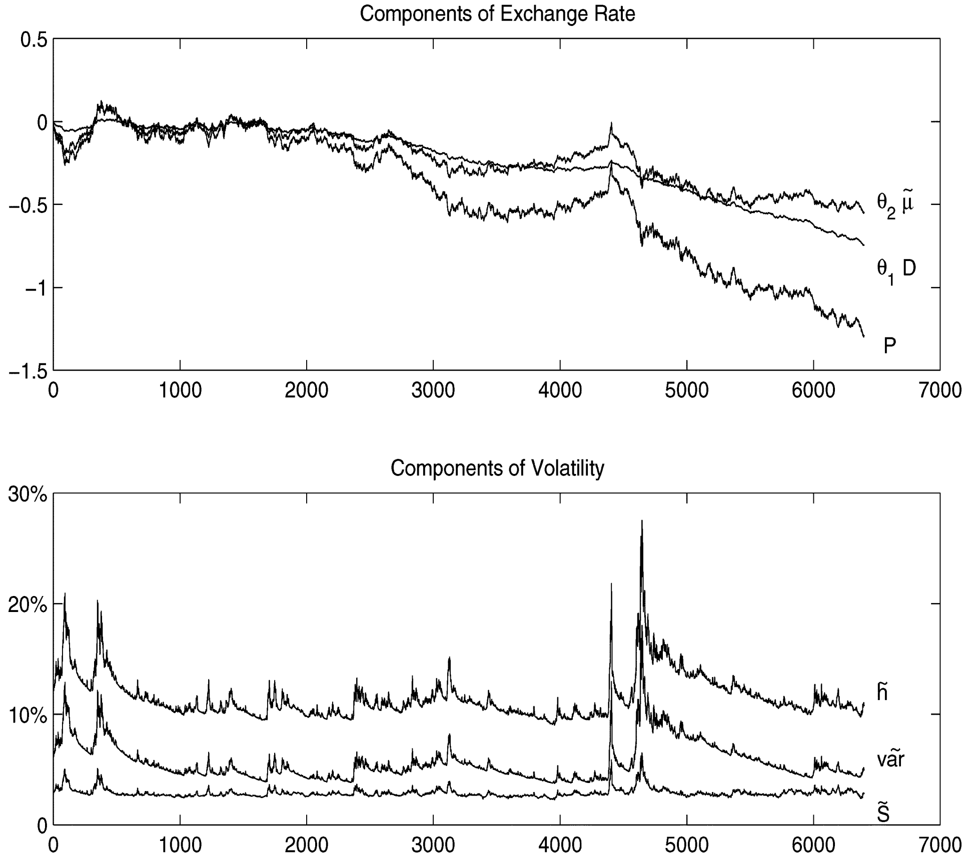

FIGURE 3.1. Simulated data and volatility. The top panel shows the simulated log exchange rate series (P) computed using the formula of Proposition 2.1 and the parameters of Section 3.3. The two components of the series are also shown (cf. equation (3.3)). The bottom panel plots the instantaneous volatility, $\tilde{h}$, for the simulation and its two time-varying components (cf. equations (3.5), (3.6)). Note: the series labeled $\widetilde{\text{var}}$ is $\theta_2 \widetilde{\text{var}}_t(\mu)/\sigma_0$ and that labeled $\tilde{S}$ is $\theta_2 \sigma_0 \widetilde{S}_t$.

the *degree* of autocorrelation found is consistent with observed behavior of exchange rates. Empirical researchers typically measure GARCH phenomena by the sample autocorrelation of squared returns. The second panel of the figure plots this function for the simulated return series here.

The square returns exhibit far less autocorrelation than the true volatility process, a consequence of the former being the latter multiplied by an i.i.d. innovation. The figure, in fact, closely resembles the curves reported in studies of daily returns (Baillie and Bollerslev 1989, Ding, Granger, and Engle 1993), with rapid decay to less than 0.10 after the first few days and then very slow decay, with statistically positive values at lags of well over a year. Hence, in addition to delivering GARCH effects of the required magnitude, the model also offers a potential explanation for the "long memory" in volatility phenomenon (Baillie, Bollerslev, and Mikkelsen 1996, Andersen and Bollerslev 1997).

Table 3.1 compares the simulated returns with those of some representative currency series. Here the benchmark for assessing GARCH effects is the Ljung and Box (1978)

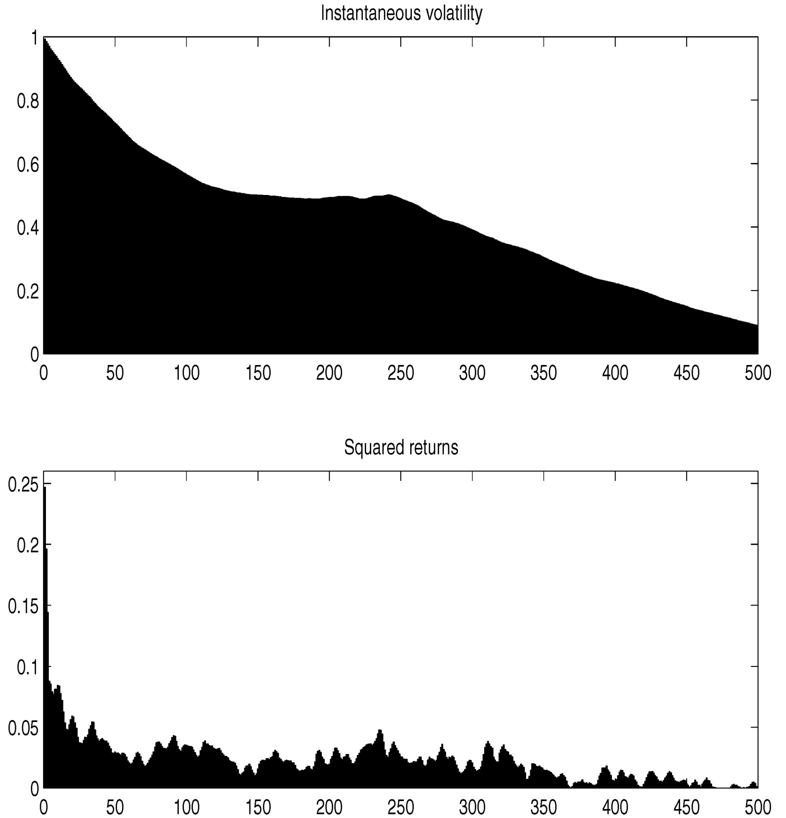

FIGURE 3.2. Volatility autocorrelation functions. The top panel shows the sample autocorrelation function of the true volatility series $\tilde{h}_t$ for the simulation of Section 3.3. The bottom panel shows the autocorrelation for the simulated squared returns. This function has been smoothed using a Bartlett kernel of width 3. The horizontal axis is in units of days.

TABLE 3.1
Daily Return Statistics

Sample	$\rho_1(r)$	$\rho_2(r)$	$Q(10)$	$\rho_1(r^2)$	$\rho_2(r^2)$	$Q^2(10)$	κ
Model Simulation	.0097	−.0118	9.95	.0706	.0567	456.10	3.87
USD/GBP 3/80-2/85	.0004	.0013	7.64	.1474	.0551	74.79	4.81
DEM/USD 3/80-2/85	.0529	−.0030	9.16	.0934	.0462	133.08	4.21

The table compares a 25-year simulated return series of the unobservable persistence model with two samples of log exchange rate changes, taken from Bollerslev (1987). Shown are the first two autocorrelation coefficients (ρ), and the Ljung and Box (1978) statistic (Q) for both returns and square returns. The Ljung–Box statistic is asymptotically chi-square distributed with 10 degrees of freedom. The last column gives the sample kurtosis of the returns.

portmanteau statistic. This statistic is a weighted average of autocorrelation coefficients, and scales like the sample size. The comparison numbers are from a data set of approximately 1/5th the length of the simulation. Hence the values reported for the squared returns, being five times greater, are consistent with true autocorrelation coefficients of the same magnitude.

It is also important to note that the fidelity of the model in terms of second moments does not come at the cost of inducing spurious effects in other moments. From the table, we can also see that autocorrelations of the returns themselves are insignificant, as in the data. Also interestingly, the sample fourth moment, though not as large as that typical of daily currency returns, is far larger than normally reproducible by diffusion models, or daily models with normal innovations. The current model does not fully explain the kurtosis in the data. However its explanation of second moment behavior may contribute part of the answer.

Besides showing the model's ability to generate observed heteroscedasticity, the simulation affords a means to distinguish the model from standard econometric specifications. As remarked in Section 3.2, the predicted perfect correlation between volatility and returns implies that large price moves can potentially coincide with reductions in instantaneous variance. This can happen in stochastic volatility models too, where the processes are driven by separate (possibly correlated) innovations. In that case, plotting these changes against each other will produce a bivariate normal cloud. In a GARCH(1,1) specification, by contrast, the same plot will yield a U-shaped curve. In the usual notation of that literature,

$$h_{t+1}^2 = \omega + \alpha r_t^2 + \beta h_t^2 \Rightarrow \Delta h_t^2 \approx \omega + \alpha r_t^2,$$

since, empirically, $\beta \approx 1$ in high-frequency samples.

Figure 3.3 plots the variance impact curves for the simulated data. Unlike either alternative, the quasi-instantaneous (1 day) relationship is characterized by a distinctive butterfly shape. As with a GARCH model, small returns are associated with slightly lower variance. But, as in a correlated bivariate model, large returns coincide with large changes. Here, though, the sign of the correlation is itself random. The symmetry of the figure suggests that it is as likely to be positive as negative. At longer horizons the expected U-shape does emerge. So the basic intuition of the GARCH model—that realized volatility drives future volatility—is preserved.

Like volatility autocorrelation functions, these impact functions are not directly computable from data. Sample counterparts may not be sufficiently informative to discriminate between models.[16] However the graphs clearly point to the sign-switching behavior of the volatility-return correlation as a defining characteristic of the unobservable persistence model. To decide if this behavior is counterfactual, what is needed is a prediction of precisely when such switching should occur. That is the subject of the next section.

The conclusions of this section, being simulation based, are necessarily contingent. I have not shown that the unobservable persistence model generates return-like dynamics in general. I have only demonstrated that this can happen over long time intervals for a particular parameter configuration. These parameters were motivated by the underlying economic story, as it applies to exchange rates. In this context, the model does appear to offer a structural explanation for conditional heteroscedasticity, while, at the same time, implying sufficiently novel features to enable it to be falsified.

[16] To my knowledge, there are no empirical studies of the volatility response function of exchange rates.

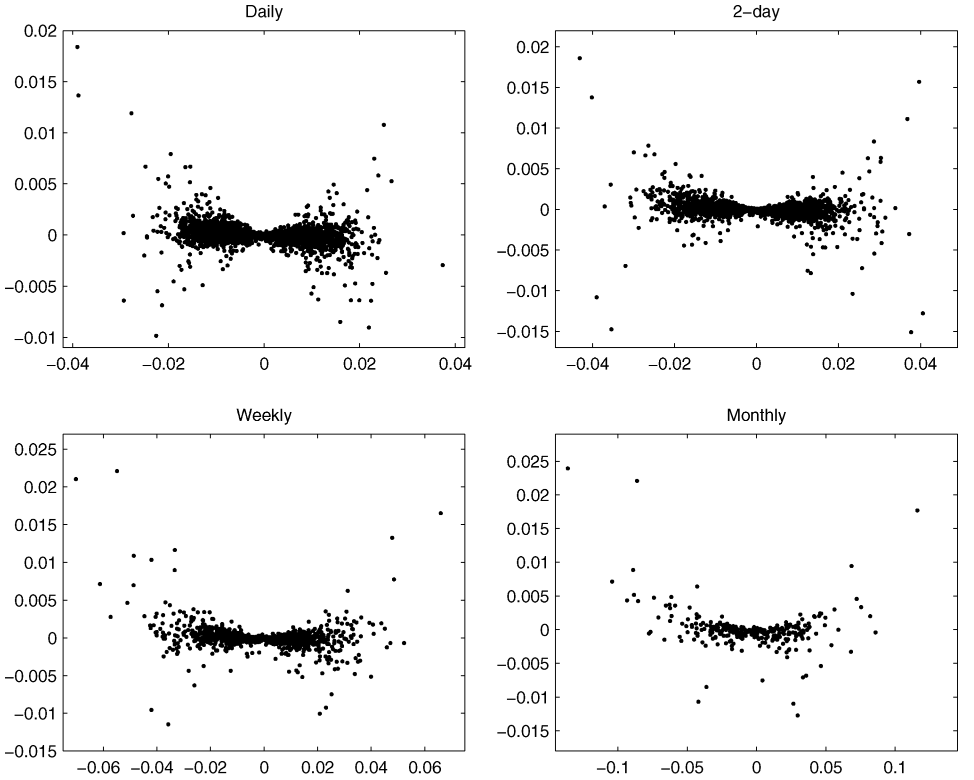

FIGURE 3.3. Volatility impact functions. The figure plots changes in instantaneous variance $\tilde{h}_t^2$ against return shocks over different time intervals using the simulated series of Section 3.3.

4. A TESTABLE IMPLICATION

The preceding section established that, for reasonable parameter values, the unobservable persistence model can reproduce empirical patterns in second moments of returns. Yet the model also implies complex pathwise dynamics, including counterintuitive decreases in volatility under some circumstances. The goal of this section is to extract a general implication from this complexity which will enable its predicted behavior to be uniquely distinguished.

To do this, it is necessary to understand the evolution of the volatility of volatility. We have seen that this variable, $\tilde{g}_t$, is determined by two moments $\widetilde{\mathrm{cov}}_t$ and $\widetilde{\mathrm{skew}}_t$ of agents' posterior beliefs, and that the evolution of these are governed by still higher moments. However under some conditions it is possible to infer the signs of these terms, and, thence, the sign of $\tilde{g}_t$. That quantity is the instantaneous correlation of volatility and returns. So the result is a prediction about the slope of the volatility impact function at each point in time.

After first defining some conditions and discussing their plausibility, the next subsection presents the main proposition. A second subsection then explores the applicability of the conditions, and the extent to which the result may be expected to be stronger than that formally established.

4.1. Volatility and Trends

To begin, consider the occurrence (or the set in the sample space)

$$A_\tau = \{\omega : d\langle \widetilde{S}_t, \tilde{h}_t\rangle/dt > 0,\ t = \tau\}. \tag{4.1}$$

When this obtains, $\widetilde{S}_t$ experiences an unexpected shock if and only if $\tilde{h}_t$ experiences one in the same direction. Since the two series are driven by the same $\widetilde{W}_t$, this means at t they are perfectly correlated. An equivalent condition (from Proposition 2.2) is

$$\{\mathrm{sign}(\widetilde{\mathrm{cov}}_\tau(\mu, S)) = \mathrm{sign}(3\,\widetilde{\mathrm{cov}}_\tau(\mu, S) + \widetilde{\mathrm{skew}}_\tau(\mu)/\sigma_0^2)\},$$

which effectively requires the skew to be small relative to the covariance when their signs disagree. Intuitively, most of the time the signs *should* agree. This is because the agent's distribution for μ is a mixture of the distribution $\mu|S = 1$ and $\mu|S = 0$. If each of these is roughly symmetric, and if the first has less mass ($S = 1$ is less likely), the superposition will be skewed in the direction of the $S = 1$ mode, which is also the direction of the covariance of μ and S. Below I will present numerical evidence that A_τ is an event of high probability for many parameter combinations.

Next, I introduce a restriction enabling an approximation to the cross-moment $\widetilde{\mathrm{cov}}_t(\mu, S)$. By conditioning on the current value of the state S, this may always be written $\widetilde{E}_t((\mu - \tilde{\mu})S) = \widetilde{S}_t(1 - \widetilde{S}_t)[\widetilde{E}_t(\mu|S_t = 1) - \widetilde{E}_t(\mu|S_t = 0)]$. The sign is determined only by the term in brackets, which indicates the direction in which the μ estimate changes as $\widetilde{S}_t$ goes up.

In the proof of Proposition 2.1, I establish the representation

$$\widetilde{\mathrm{E}}_t(\mu_t) = \widetilde{\mathrm{E}}_t\left(\mu_0 e^{-\int_0^t S_u\,du} + \int_0^t S_v e^{-\int_v^t S_u\,du}\,dD_v\right),$$

whence

$$\widetilde{\mathrm{E}}_t(\mu_t|S_t = i) = \widetilde{\mathrm{E}}_t\left(\mu_0 e^{-\int_0^t S_u\,du} + \int_0^t S_v e^{-\int_v^t S_u\,du}\,dD_v|S_t = i\right). \tag{4.2}$$

Ignoring the term involving μ_0 and taking the posterior expectation inside the dD integral,[17] we see that one forms expectations about μ_t by weighting the observation series by the kernel $\widetilde{\mathrm{E}}_t(S_v e^{-\int_v^t S_u\,du})$ and summing.

Now if the filtering problem is hard enough, then the data will rarely provide enough information to cause the observer to greatly alter his prior expectations of these weights. This seems especially likely when conditioning on the current value of S: given the state, posterior and prior expectations of that discount factor must coincide for v close to t, and their difference for earlier values will be ameliorated by the exponential term. So, in this case, the weights will not vary much with current information. That is, the expectations involved will be close to their unconditional values.

By restricting consideration to successively less informative histories, we may make this approximation as close as desired. The exact condition I will use to restrict the informativeness is

$$B_\tau = \left\{\omega : \mathrm{sign}\Big(\int_\Omega \mu_\tau(S_\tau - \widetilde{S}_\tau)\zeta(\omega)\,dP\Big) = \mathrm{sign}\Big(\int_\Omega \mu_\tau(S_\tau - \overline{S})\delta(\overline{W}(\omega))\,dP\Big)\right\}, \tag{4.3}$$

where ζ is the Radon–Nikodym derivative given in Proposition 2.1, δ is given by equation (2.10), and $\overline{S} \equiv \lambda_0/(\lambda_0 + \lambda_1)$. The two integrands only differ by the exponential

[17] Recall that since the integrand is of bounded variation, this integral may be treated as a Lebesgue–Steiltjes pathwise limit.

component of the Radon–Nikodym derivative (cf. equation (2.8)), and the change from $\widetilde{S}_t$ to its mean $\overline{S}$. So the condition will hold when the exponential is close to constant under the prior. That is, the D history does not discriminate strongly between different S histories. The probability of B_τ will also be shown below to be close to 1 for many relevant cases.

A third condition to be used below restricts consideration to a set on which the path of the volatility $\hbar$ is close to its steady-state value $\hat{h} \equiv \sigma_0 \overline{S}^{1/2}$. Let $\hat{\mu}_t$ be the drift estimator obtained by approximating the former by the latter:

$$\hat{\mu}_t \equiv \tilde{\mu}_0 + \int_0^t \hat{h}\, d\widetilde{W}_u. \tag{4.4}$$

The condition weights the history of errors $(\hat{\mu}_t - \tilde{\mu}_t)$ by a kernel φ, to be defined later. This weighted sum is compared to the similarly weighted observation history:

$$C_\tau = \left\{\omega : \left| \int_0^\tau \varphi(\tau - u)(\hat{\mu}_u - \tilde{\mu}_u)\, du \right| \le \left| \int_0^\tau \varphi(\tau - u)\, dD_u \right| \right\}. \tag{4.5}$$

This set, too, is not unlikely if (as will be the case) the kernel places most of its mass on recent history. Then, even if $\hbar$ strays far from its steady-state value, the observation changes dD, being of order $(dt)^{1/2}$, will make the right-hand integral much larger than the dt integral on the left.

Last, I need to formalize the conditions under which the initial conditions do not significantly affect the posterior moments. For this it suffices to imagine that the initial time was long ago. Notationally, it is convenient to do this by taking the current time τ to infinity.

Under the preceding assumptions, the volatility of volatility can be signed.

PROPOSITION 4.1. *Let $g_t \equiv 3\,\widetilde{\text{cov}}_t(\mu, S) + \widetilde{\text{skew}}_t(\mu)/\sigma_0^2$. There exists a weighting function $\varphi(\cdot)$ (given in the Appendix) such that for all large τ,*

(I) *on $A_\tau \cap B_\tau$:*

$$\text{sign}(g_\tau) = \text{sign}\Big(\int_0^\tau \varphi(\tau - u)\, dD_u \Big). \tag{4.6}$$

(II) *Moreover, on $A_\tau \cap B_\tau \cap C_\tau$:*

$$\text{sign}(g_\tau) = \text{sign}\Big(\int_0^\tau \psi(\tau - u)\, d\widetilde{W}_u \Big) \tag{4.7}$$

where $\psi(x) \equiv \sigma_0[\varphi(x) + \overline{S}^{1/2}\phi(x)]$ and $\phi(x) \equiv \int_0^x \varphi(y)dy$.

The proposition gives the sign of the volatility-return correlation in terms of two related data weighting function. Both these are close to an ordinary exponential smoother. (They are graphed in Figures 4.1 for the parameter values used in Section 3.3.) Hence the right-hand sides in parts (I) and (II) are essentially measures of the local trend in the observations (I), or the normalized innovations (II).

For this reason, the proposition has a very simple interpretation: volatility goes up when trends continue, and goes down when trends reverse. To see this, suppose that at t the conditions hold and g_t is positive. By (4.7) that implies the trend in returns has been positive prior to t. Then the continuation of that trend over the next instant means that

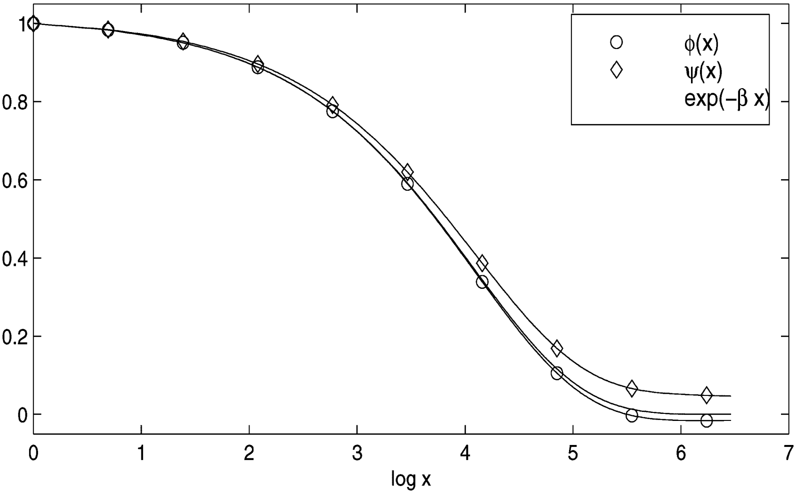

FIGURE 4.1. Data weighting functions. The figure shows the two kernel functions derived in Proposition 4.1 evaluated for the parameter set used in Section 3.3. The horizontal axis is the log of lag time (in units of days). Also plotted is the pure exponential function with decay parameter $\beta = 1 + \lambda_0 + \lambda_1$.

the next innovation $d\widetilde{W}_t$ is positive. But then the volatility shock, $g_t d\widetilde{W}_t$, is also positive because g_t is. Had the trend reversed, $d\widetilde{W}_t < 0$ would imply $g_t d\widetilde{W}_t < 0$, a decrease in volatility.[18]

The intuition behind this surprising realtionship may be described in terms of the effect of a recent trend on the two components of $\hbar_t$: $\widetilde{\text{var}}_t(\mu)$ and $\widetilde{S}_t$. First, a trend is tantamount to systematic forecast errors of the same sign. Even if the errors are small, this raises the likelihood that they resulted from a biased estimate of the drift μ, rather than from chance. Hence posterior uncertainty, $\widetilde{\text{var}}_t(\mu)$, about that estimate rises.[19] At the same time, the rational inference from the recent trend is that μ is likely to have been shocked recently. But this only happens in the persistent state (i.e., when $S = 1$). Since states change infrequently, it is (relatively more) likely that this is still the state. This causes the observer to update his estimate $\widetilde{S}_t$.

It should be emphasized that the proposition only relates instantaneous changes in volatility to instantaneous changes in trends. It does not say that volatility will be high following a trend, nor even that volatility will have gone up over the course of such a trend. Moreover, in only equating signs, the proposition does not assert any relationship between the strength of trends and the magnitude of g.

Nevertheless, the result provides the basis for the primary implication of the unobservable persistence model: changes in momentum (in absolute value) should predict changes in volatility. The theory delivers this implication in terms of a specific definition of momentum—one that is, in fact, $\mathcal{F}^P$-measurable and hence directly computable from returns.

[18] The argument here refers to the unexpected component of return and volatility changes. Over the instant dt, the expected components may be ignored.

[19] This effect might not be present if μ were expected to mean-revert, as, for example in the model of Wang (1993). Then, by the time a trend were detected, the likely change to μ might have dissipated. Also models with regime switching in μ, such as Veronesi (1999), would not share this feature. Then a trend would signal a switch to the new level. But since that level is known, posterior uncertainty could even decrease.

4.2. Applicability of the Trend Result

To be meaningful, the volatility-momentum relationship derived above must hold under a wide range of circumstances. I now check that the conditions imposed in the proposition are plausible, as asserted, in the sense of occurring with high frequency for diverse parameter configurations. Beyond that, I examine the extent to which the magnitude of trends predicts the magnitude of the volatility of volatility. Although not explicitly predicted by the proposition, a reliable linkage between these two quantities would greatly strengthen the result by rendering g easily computable from returns. Finally, the practical relevance of the result depends on the scale of the predicted effects. Some simple calculations illustrate the likely response of volatility to trends.

To address these issues, I return to the numerical technique of Section 3.3. As there, long sample paths are simulated for different parameter values, and the exact conditional moments are calculated along each using the integration result of Proposition 2.1. Now, in addition, I compute the trend indicator functions $\int_0^t \varphi(t-u)\, dD_u$ and $\int_0^t \psi(t-u)\, d\widetilde{W}_u$ using the formulas for the kernels given in the Appendix. These are then compared to $g_t \equiv 3\,\widetilde{\text{cov}}_t(\mu, S) + \widetilde{\text{skew}}_t(\mu)/\sigma_0^2$.

For this exercise, I again use the structural parameters λ_0 and λ_1 chosen in Section 3.3. As these were not particularly restricted by economic theory, I vary each by an order of magnitude, roughly bracketing the initial case. These are the only free parameters at issue: the extra structural variable σ_0 serves only to scale all the quantities, and hence drops out of all the comparisons below. The numbers reported will be numerically identical for any specific choice.

The results, shown in Table 4.1, confirm that Proposition 4.1 should apply quite generally. The first three columns give the percentage of simulated days on which the conditions A_τ, B_τ, and $A_\tau \cap B_\tau$ held. For all cases, the latter frequency is between 81% and 93%. Columns 5 and 6 show that the third condition C_τ is highly likely, with the composite intersection of all three holding only slightly less often—between 81% and 90% of the time. Columns 4 and 7 show the frequencies with which the actual conclusions of parts (I) and (II) of the proposition obtained. The relationships do hold somewhat more generally than proven. But the numbers imply that the conditions imposed were very nearly necessary. The systems simulated have not been shown to be stationary or ergodic. So the reported frequencies may not well approximate either long-run averages or unconditional probabilities. Nonetheless, the table demonstrates that the occurrence of the events in question is not confined to brief periods or specific parameter choices.

The proof of Proposition 4.1 is based on approximation arguments which hold out the possibility that, in addition to having the same sign, the sizes of the quantities studied may be related. To assess this, I regress the actual values of g_t in each simulation on the trend indicator $\int_0^t \psi(t-u)\, d\widetilde{W}_u$. The results are summarized by the rightmost column of Table 4.1. For all cases considered, a strong linear relationship does, in fact, emerge. Over a time-span of 6400 days, a remarkable 69%–84% of the variability of the volatility of volatility is accounted for by the simplest possible specification. It appears, then, that the conclusion of the proposition can be significantly strengthened in a useful sense: the trend indicator function tells us *how much* volatility will change in response to return shocks as well as in which direction.

REMARK 4.1. Having validated this relationship, we can now see the essential form of the volatility dynamics under the unobservable persistence model without having actually

TABLE 4.1
Conditions of Proposition 4.1

λ_1	λ_0	A_τ	B_τ	$A_\tau \cap B_\tau$	(I)	C_τ	$A_\tau \cap B_\tau \cap C_\tau$	(II)	R^2
10.0	1.0	93.5	88.4	85.3	88.8	99.1	85.0	88.7	84.1
10.0	.10	83.1	95.7	81.1	83.5	99.3	81.0	83.6	69.2
3.0	.33	95.5	89.4	87.7	90.4	96.0	85.9	90.3	83.6
1.0	1.0	96.5	95.7	92.3	92.4	95.2	89.9	90.9	78.1
1.0	.10	95.2	92.8	91.7	95.5	87.3	85.7	92.0	74.6

The table shows the frequency of the occurrence of the conditions and conclusions of Proposition 4.1 in 25 year simulations for each of the parameter choices shown, using the method described in Section 3.3. The conditions A_τ, B_τ, and C_τ can be defined respectively by equations (4.1), (4.3), and (4.5) in the text. The columns labeled (I) and (II) show the frequencies of the conclusion obtaining for parts (I) and (II) of the proposition. The last column reports the R^2 from the regression of the true volatility of volatility on the approximating process $\int_0^t \psi(t-u)\,d\widetilde{W}_u$.

closed the filtering system. If $g_t \approx c_1 \int_0^t \psi(t-u) d\widetilde{W}_u + c_2$ where c_1 and c_2 are roughly locally constant, then

$$\tilde{h}_t \approx \int_0^t c_0(u)\,du + c_1 \int_0^t \int_0^s \psi(s-u)\,d\widetilde{W}_u\,d\widetilde{W}_s + c_2 \widetilde{W}_t.$$

The first term is just the predictable component: the initial level plus accumulated mean-reversion. The last term is like a constant-elasticity-of-variance contribution whose sign is determined by the asymmetry of fundamental risk. The surprising component is the middle term. As in the HARCH model of Müller et al. (1997) and the QARCH of Sentana (1995), here the cross moments of return innovations have explanatory power for levels of volatility. Strong empirical evidence in favor of such a contribution is presented in Dacorogna et al. (1998). The model introduced here provides both a theoretical justification for this term and testable restrictions on the weighting of the cross moments.

A last topic the simulations can address is the strength of the volatility response to trends. To illustrate the magnitudes involved, I tabulate the average absolute value of the trend indicator for the scenario of Section 3.3. Multiplying this number by the value of θ_2 used there gives a typical value for $\tilde{g}$ of 0.025. In such a case, a three-standard deviation return the next day (i.e, $\Delta\widetilde{W}_t = 3\sqrt{\Delta t}$) would raise or lower return volatility by approximately 0.005 (half of a volatility point, or about 4% for this example) depending on whether the shock reversed or continued the preceding trend. Changes of this magnitude are economically significant for foreign exchange derivatives, and should be readily detectable in high-frequency data.

This section has shown that the unobservable persistence model predicts a novel effect in volatility dynamics: changes in momentum should lead to changes in volatility. This feature emerges from the structure of investors' inference problem: autocorrelated shocks are more likely to be persistent shocks ex post. For all parameter configurations considered, numerical examples demonstrated that the conditions of the theorem were generally satisfied over long periods. Further, the simulations established that an important extension of the theorem held almost as broadly: the actual magnitude of the trend indicator function predicted the magnitude of the instantaneous volatility-return covariance, not just its sign, consistently and accurately. These findings constitute a remarkable—and testable—suggestion about the pathwise properties of volatility.

5. CONCLUDING REMARKS

This paper shows that the major empirical regularities in the volatility of financial time series can be viewed as arising from the changing inferences of investors about the degree of persistence of fundamental shocks.

In advancing a structural explanation of these regularities, the work contributes to a growing branch (some of which was cited in the introduction) of asset pricing research which seeks an understanding of why and how specific features of the economic environment lead to specific time-series properties. The full story of heteroscedasticity in financial returns undoubtedly involves elements of market frictions, investor irrationality, time-varying discount rates, and time-varying information flows. Theorists have recently begun to map out the precise second moment implications of these differing mechanisms. In similar fashion, this paper has built a theory from a realistic, almost self-evident assumption: that some shocks have persistent effects and some do not. From this I derived a full description of return dynamics resulting from the time-varying inferences of investors.

Most notably, I deduced a strong testable implication about the relationship between volatility and realized trends in the data. Interesting and potentially significant in its own right, the result further demonstrates that the utility of the modeling approach goes beyond the initial goal of explaining known patterns of heteroscedasticity.

APPENDIX

Proof of Proposition 2.1. First note that μ_t has the (pointwise) integral representation $e^{\int_0^t S_u\,du}[\mu_0 + \int_0^t S_u e^{-\int_0^u S_v\,dv} dD_u]$ which may be verified by integration by parts (the integrand is of bounded variation). Thus given μ_0 and the histories of D and S we know μ_t and hence W_t from $\sigma_0 W_t = D_t - D_0 - \int_0^t \mu_u\,du$. Or $E(g(W', S', \mu_0)|\mathcal{F}^{D,S,\mu_0}) = g(\overline{W}(\omega, \omega), S, \mu_0)$ where $\overline{W}$ is defined by (2.10). So

$$E(g(W', S', \mu_0)|\mathcal{F}^D) = E(g(\overline{W}, S, \mu_0)|\mathcal{F}^D). \tag{A.1}$$

Conditionally, only the path $\overline{W}$ is possible. So it suffices to find the marginal distribution over the subspace $\Theta \equiv (\{0, 1\}^T \times R)$ of values for S and μ_0 induced by the realization D.

The rest of the proof consists of verifying sufficient conditions to apply a function space version of Bayes' theorem due to Kallianpur and Striebel (1968), which will give us this marginal.

Since μ depends only on $\theta \in \Theta$ and the path of D, we may write

$$dD_t = A_t(\theta, D)dt + \sigma_0 dW_t. \tag{A.2}$$

Then we note that θ is independent of W. In addition we need the following:

(A) D is by definition the strong solution to (A.2).
(B) $A_t(\theta, D) = \mu_t(\omega)$ is continuous on [0, 1], hence $\int_0^T |A_t| dt < \infty$ for all ω.
(C) $\sigma_0 > 0$ is assumed.
(D) By continuity again $\int_0^T A_t^2\, dt < \infty$. Moreover $A_t(\theta, W) = e^{-\int_0^t S_u\, du} \cdot [\mu_0 + \int_0^t S_u e^{\int_0^u S_v\, dv} \sigma_0\, dW_u]$ also has continuous paths so the same applies.
(E) $E|A_t(\theta, D)| \le (E\mu_t^2)^{1/2} = (E\mu_0^2 + E(\int_0^t S_u \sigma_0\, dW_u)^2)^{1/2} \le (E\mu_0^2 + \sigma_0 t)^{1/2}$. So $\int_0^T E|A_t|\, dt < \infty$. For any subsequence $\mathcal{G}_t$ of sub σ-algebras of $\mathcal{F}_t$;$\infty > \int_0^T EA_t^2\, dt = E\int_0^T E(A_t^2|\mathcal{G}_t)\, dt \ge E\int_0^T E(A_t|\mathcal{G}_t)^2 dt$. Thus $\int_0^T E(A_t|\mathcal{G}_t)^2 dt < \infty$ a.s.

Under these conditions, Theorem 7.23 of Liptser and Shiryayev (1977) applies, and we may conclude that the process $(D_u)_{u=0}^t$ induces on Θ a measure having unnormalized density with respect to the (marginal) prior $P_\theta = P|_{\mathcal{F}\cap\Theta}$ given by

$$\rho(\omega'; \omega, t) = e^{-\frac{1}{2}\int_0^t \frac{m_u^2(\omega';\omega)}{\sigma_0^2}\, du + \int_0^t \frac{m_u(\omega';\omega)}{\sigma_0^2}\, dD_u(\omega)}. \tag{A.3}$$

Hence

$$E(g(\overline{W}, S, \mu_0)|\mathcal{F}_t^D) = \int_\Theta g(\overline{W}, S, \mu_0)\rho(\omega'; \omega, t)\, dP_\theta \Big/ \int_\Theta \rho(\omega'; \omega, t)\, dP_\theta,$$

which combined with (A.1) implies (2.11). □

Proof of Proposition 2.2. The proposition is an application of Theorem 8.1 in Liptser and Shiryayev (1977) once appropriate regularity conditions have been checked. For all of these it will suffice to verify that $\mu_t - \mu_0 \equiv X_t$ has moments of all orders, bounded in t. But this is clear. $E|X_t|^p = E[E(|X_t|^p|S_u, 0 \le u \le t)]$. Given the S path, X_t is normal with 0 mean and variance $\int_0^t S_u\, du$. So the conditional pth moment is less than that of $\mathcal{N}(0, t)$, hence the same is true of the unconditional moment. Also $E(|\mu^k S_t|^p) \le (ES_t)^{1/2}(E\mu_t^{2kp})^{1/2} \le (E\mu_t^{2kp})^{1/2}$. Thus these moments have the same property.

If α_t denotes generically the drift term in the differential representation of either $\mu^k S_t$ or μ_t^k, then by Ito's Lemma, α_t will be a linear combination of terms of the same form (cf. equations (2.6), (2.4)). These considerations imply that $\int_0^T E\alpha_t^2 dt$, $\int_0^T E|\mu^k S_t|^2 dt$, and $\int_0^T E|\mu_t^k|^2 dt$ are all finite, which suffices to apply the desired result. □

Proof of Lemma 3.1. With the stated assumptions the value of the equity claim is just $\widetilde{E}_t \int_t^\infty e^{-r(u-t)} D_u\, du = \widetilde{E}_t \int_t^\infty e^{-r(u-t)}[\sigma_0(W_u - W_t) + D_t + \int_t^u \mu_v\, dv]\, du$. Examining the integrands on the right, we can interchange the integrations of the first term (since $\int_t^\infty e^{-r(u-t)}(u-t)\, du < \infty$) to get rid of it. The second term is D_t/r. The third is

$$\begin{aligned}
&\widetilde{E}_t \int_t^\infty \int_t^u \left(\mu_t + \int_t^v \sigma_0 S_r\, dW_r\right) dv\ e^{-r(u-t)}\, du \\
&= \frac{\tilde{\mu}}{r^2} + \sigma_0 \widetilde{E}_t \int_t^\infty \int_t^u \int_t^u S_r\, dW_r\, dv\, e^{-r(u-t)}\, du.
\end{aligned}$$

To kill the quadruple integral we only need to be careful in applying Fubini's theorem. If $Y_v \equiv \int_t^v S_r\, dW_r$ then $\widetilde{E}_t Y_v^2 = \widetilde{E}_t(\widetilde{E}_t(Y_v^2|S_r, t \le r \le v)) = \widetilde{E}_t \int_t^v S_r\, dr \le (v-t)$.

So $\widetilde{E}_t|Y_v| \le (v-t)^{1/2}$, and $\int_t^u (v-t)^{1/2}\,dv < \infty \Rightarrow \widetilde{E}_t \int_t^u |Y_v|\,dv = \int_t^u \widetilde{E}_t|Y_v|\,dv$. Since this last integral is less than $\int_t^u (v-t)^{1/2}\,dv = \frac{2}{3}(u-t)^{3/2}$ and $\int_t^\infty e^{-r(u-t)}(u-t)^{3/2}\,du < \infty$, this permits $\widetilde{E}_t \int_t^\infty e^{-r(u-t)} \int_t^u Y_v\,dv\,du = \int_t^\infty e^{-r(u-t)}\widetilde{E}_t \int_t^u Y_v\,dv\,du$. We have already seen that $\widetilde{E}_t \int_t^u Y_v\,dv = \int_t^u \widetilde{E}_t Y_v\,dv$. Conditioning again on the outcome of S, we see that $\widetilde{E}_t Y_v = 0, \forall v$. □

Proof of Proposition 3.1. $\mathcal{F}_t^P \subset \mathcal{F}_t^D$ is obvious. Next, I claim that $\tilde{\mu}_t$ is $\mathcal{F}_t^P$ measurable for all t. This will imply the desired result since for the models under consideration P_t is an invertible (in fact, linear here) function of D_t and $\tilde{\mu}_t$.

To establish the claim, write the model generically as $P_t = \theta_0 + \theta_1 D_t + \theta_2 \tilde{\mu}_t$, and assume $\theta_1 \neq 0, \theta_2 \neq 0$. (If $\theta_1 = 0$ the claim is obvious. If $\theta_2 = 0$ the Proposition is obvious. If both are zero the result is superfluous.) Then we have

$$dP_t = \theta_1 dD_t + \theta_2 d\tilde{\mu}_t = \theta_1 \tilde{\mu}_t dt + \tilde{h}_t d\widetilde{W}_t$$

and

$$d\tilde{\mu}_t = \hbar_t d\widetilde{W}_t,$$

where

$$\tilde{h}_t \equiv \theta_1\sigma_0 + \theta_2\hbar_t \qquad \hbar_t \equiv \sigma_0\widetilde{S}_t + \widetilde{\mathrm{var}}_t(\mu_t)/\sigma_0.$$

But then ($\tilde{h}_t \neq 0, \forall t$)

$$\gamma_t dP_t = \gamma_t\theta_1\tilde{\mu}_t dt + \hbar_t d\widetilde{W}_t \Rightarrow$$

$$d\tilde{\mu}_t = -\gamma_t\theta_1\tilde{\mu}_t dt + \gamma_t dP_t,$$

where $\gamma_t \equiv \hbar_t/\tilde{h}_t$. This last equation has solution

$$\tilde{\mu}_t = e^{-\theta_1\int_0^t \gamma_s\,ds}\left(\tilde{\mu}_0 + \theta_1\int_0^t \gamma_u e^{\theta_1\int_0^u \gamma_s\,ds} dP_u\right).$$

Moreover, as γ_t is bounded, this solution is unique (Protter 1990, Thm. V.7, suffices).

Now the hypothesis of the proposition implies that $\tilde{\mu}_0$ is $\mathcal{F}_0^P$ (hence $\mathcal{F}_t^P$) measurable. But also $\hbar_t$, being the rate of change of the quadratic variation of P_t, is $\mathcal{F}_t^P$ measurable. Hence $\hbar_t$ is. Hence γ_t is. So the entire right side of the last equation is. □

Proof of Proposition 4.1. On the set A_τ it suffices to show the conclusion (I) with $\widetilde{\mathrm{cov}}_\tau(\mu, S)$ on the left-hand side of (4.6). On B_τ that covariance has the same sign as

$$\int_\Theta\int_C \mu_\tau(S_\tau - \overline{S})\delta(\overline{W})\,dP_{W|\theta}\,dP_\theta$$

where $\theta \equiv \{\mu_0, S\} \in \Theta \equiv R \times \{0,1\}^T$, $P_{W|\theta}$ is the marginal prior over Θ, and $P_{W|\theta}$ is Weiner measure (by the assumed independence). Integrating over the function δ just has the effect of fixing the Weiner path as a function of the data and the other parameters, hence identifying μ. This leaves

$$\int_\Theta (S_\tau - \overline{S})\left[\mu_0 e^{-\int_0^\tau S_u\,du} + \int_0^\tau S_v e^{-\int_v^\tau S_u\,du}\,dD_v(\omega)\right]dP_\theta.$$

Since $\overline{S} = \int_\Theta S_\tau dP_\theta$, we can rewrite that integral as

$$\overline{S}(1-\overline{S})(E_\theta[\mu|S_\tau = 1] - E_\theta[\mu|S_\tau = 0]) \tag{A.4}$$

where,

$$E_\theta[\mu_\tau|S_\tau = 1](\omega) \equiv \int_\Theta S_\tau\left[\mu_0 e^{-\int_0^\tau S_u\,du} + \int_0^\tau S_v e^{-\int_v^\tau S_u\,du} dD_v(\omega)\right]dP_\theta \Big/ \int_\Theta S_\tau\,dP_\theta.$$

The factor $\overline{S}(1-\overline{S})$ has no effect on the sign. So it suffices to characterize the difference in conditional expectations in (A.4).

Now

$$E_\theta\Big[\int_0^\tau S_v e^{-\int_v^\tau S_u\,du}\,dD_v(\omega)\big|S_\tau=i\Big]=\int_0^\tau \frac{d}{dv}E_\theta\big[e^{-\int_v^\tau S_u\,du}|S_\tau=i\big]\,dD_v(\omega)$$

since the integrand is nonnegative and bounded (uniformly in v). And then

$$\begin{aligned}&E_\theta\big[e^{-\int_v^\tau S_u\,du}|S_\tau=i\big]\\&\quad=\overline{S}E_\theta\big[e^{-\int_v^\tau S_u\,du}|S_\tau=i,S_0=1\big]+(1-\overline{S})E_\theta\big[e^{-\int_v^\tau S_u\,du}|S_\tau=i,S_0=0\big].\end{aligned}$$

The E_θ terms on the right, which are just unconditional (prior) expectations, can be computed as the solution to a two-point boundary problem. For large τ, however, the solutions will be the same as those of the time-reversed expectations $E_\theta[e^{-\int_0^{\tau-v} S_u\,du}|S_0=i]$, which will also equal $E_\theta[e^{-\int_v^\tau S_u\,du}|S_\tau=i]$. Likewise, using the time-reversed measure,

$$E_\theta[\mu_0 e^{-\int_0^\tau S_u\,du}|S_\tau=i,S_0=j]\to E_\theta(\mu_0)E_\theta[e^{-\int_0^\infty S_u\,du}|S_0=i]=0$$

since $P_\theta(\int_0^\infty S_u\,du<\infty)=0$.

Standard Markov chain limiting arguments now imply that the vector $\eta(\tau-v)\equiv[E_\theta(e^{-\int_0^{\tau-v} S_u\,du}|S_0=0),E_\theta(e^{-\int_0^{\tau-v} S_u\,du}|S_0=1)]'$ must solve the initial value problem

$$\dot\eta=F_\eta,\qquad \eta(0)=\begin{pmatrix}1\\1\end{pmatrix},\qquad F=\begin{pmatrix}-\lambda_0 & \lambda_0\\ \lambda_1 & -\lambda_1-1\end{pmatrix}.$$

So $\eta(\tau-v)=\exp\{(\tau-v)F\}\cdot\eta(0)$. Since $[E_\theta(S_v e^{-\int_v^\tau S_u\,du}|S_\tau=0),E_\theta(S_v e^{-\int_v^\tau S_u\,du}|S_\tau=1)]'=-\dot\eta(\tau-v)$, we have

$$E_\theta(\mu|S_\tau=1)-E_\theta(\mu|S_\tau=0)=-\int_0^\tau[-1,1]\dot\eta(\tau-v)\,dD_v.$$

This is the representation that was claimed in (I). The integrand on the right (which is $\varphi(\tau-v)$ in the statement of the theorem) can be written in closed form as follows:

$$\begin{aligned}\varphi(u)&=e^{-\frac{\beta}{2}u}\Big(\cosh(\alpha u/2)-\frac{\beta}{\alpha}\sinh(\alpha u/2)\Big)\\ \beta&\equiv 1+\lambda_0+\lambda_1\\ \alpha&\equiv\sqrt{\beta^2-4\lambda_0}.\end{aligned}$$

This is always real since $\beta\geq 1+\lambda_0\geq 2\sqrt{\lambda_0}$. The function starts at one and flips sign once. In addition, we may directly compute that

$$\phi(u)\equiv\int_0^u\varphi(v)dv=\frac{2}{\alpha}e^{-\frac{\beta}{2}u}\sinh(\alpha u/2)\tag{A.5}$$

which arises below.

If the degree of persistence parameter $\kappa\neq 1$ is present (cf. Remark 2.1 in the text), it multiplies S in all the exponential integral terms. This causes the matrix F above to become $\left(\begin{smallmatrix}-\lambda_0 & \lambda_0\\ \lambda_1 & -\lambda_1-\kappa\end{smallmatrix}\right)$. The solution φ gets multiplied by κ and the constants α and β become $\sqrt{\beta^2-4\kappa\lambda_0}$ and $\kappa+\lambda_0+\lambda_1$ respectively. In addition, the definition of the set

$\underset{\smile}{A_t}$ on which the theorem holds is altered to $\mathrm{sign}(\widetilde{\mathrm{cov}}_t(\mu, S)) = \mathrm{sign}(3\kappa\widetilde{\mathrm{cov}}_t(\mu, S) + \widetilde{\mathrm{skew}}_t(\mu)/\sigma_0^2)$, which will be more or less likely than before depending on whether the impact parameter is over or under one.

To get from (4.6) to (4.7), use equation (2.14) in the text to replace dD as the integrating variable. Then, by the Fubini theorem for stochastic integrals (Protter 1990, VI.45),

$$\begin{aligned}\int_0^\tau \varphi(\tau - u)(\tilde{\mu}_u\, du) &= \tilde{\mu}_0 \int_0^\tau \varphi(\tau - u)\, du + \int_0^\tau \int_0^u \varphi(\tau - u)\hbar_v\, d\widetilde{W}_v\, du \\ &= \int_0^\tau \left[\int_0^{\tau - v} \varphi(u)\, du \right] \hbar_v\, d\widetilde{W}_v \\ &= \int_0^\tau \phi(\tau - v)\hbar_v\, d\widetilde{W}_v.\end{aligned}$$

Here the term involving the initial conditions on the second line is suppressed because $\int_0^\infty \varphi(x)\, dx = 0$ (cf. (A.5) above) and τ is taken to be arbitrarily large. From the last line and part (I) we conclude

$$\mathrm{sign}(g_\tau) = \mathrm{sign}\Big(\int_0^\tau [\varphi(\tau - v) + \phi(\tau - v)\hbar_v]\sigma_0\, d\widetilde{W}_v \Big).$$

Now consider approximating the last integrand with $\hbar_v$ replaced by the constant $\hat{h}$. The result is the right-hand side of (4.7). To conclude that the two have the same sign, it suffices that the absolute value of their difference is less than the maximum of the absolute value of either one.

In fact, the restriction C_τ imposed more than that. To see this, write the difference as

$$\int_0^\tau \phi(\tau - v)(\hat{h} - \hbar_v)\, d\widetilde{W}_v = \int_0^\tau \varphi(\tau - u)(\hat{\mu}_u - \tilde{\mu}_u)\, du$$

(using the definition of $\hat{\mu}$). This is the quantity restricted on C_τ to be not greater than $|\int_0^\tau \varphi(\tau - u) dD_u|$, which is not greater than

$$\max\Big[\Big| \int_0^\tau \varphi(\tau - u)\, dD_u \Big|, \Big| \int_0^\tau \psi(\tau - u)\, d\widetilde{W}_u \Big| \Big],$$

which are the two quantities in question. This establishes (II). □

REFERENCES

ANDERSEN, T. G. (1996): Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility, *J. Finance* 51, 169–204.

ANDERSEN, T. G., and T. BOLLERSLEV (1997): Deutsche Mark Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies, *J. Finance* 52, 219–265.

BABA, Y., D. F. HENDRY, and R. M. STARR (1992): The Demand for M1 in the USA: 1960–1988. *Rev. Economic Studies* 59, 25–62.

BAILLIE, R. T., and T. BOLLERSLEV (1989): The Message in Daily Exchange Rates: A Conditional Variance Tale, *J. Business & Econ. Statistics* 7, 295–307.

BAILLIE, R. T., T. BOLLERSLEV, and H. O. MIKKELSEN (1996): Fractionally Integrated Generalized Autoregressive Conditional Heteroscedasticity, *J. Econometrics* 74, 3–30.

BALL, L. (1998): Another Look at D long-Run Money Demand; NBER Working paper #6597.

BARBERIS, N., A. SHLEIFER, and R. VISHNY (1998): A Model of Investor Sentiment, *J. Financial Econ.* 49, 307–344.

BARSKY, R. B., and J. B. DELONG (1993): Why Does the Stock Market Fluctuate?, *Quart. J. Econ.* 108, 291–311.

BERRY, T. D., and K. M. HOWE (1994): Public Information Arrival, *J. Finance* 49, 1331–1346.

BJÖRK, T. (1980): Finite Dimensional Optimal Filters for a Class of Itô Processes with Jumping Parameters, *Stochastics* 4, 167–183.

BJÖRK, T. (1982): Finite Optimal Filters for a Class of Nonlinear Diffusions with Jumping Parameters, *Stochastics* 6, 121–138.

BOLLERSLEV, T. (1987): A Conditional Heteroskedastic Time Series Model for Speculative Prices and Returns, *Rev. Econ. Stat.* 69, 542–547.

BOLLERSLEV, T., R. F. ENGLE, and D. B. NELSON (1994): ARCH Models, in *Handbook of Econometrics,* vol. IV, ed. R. F. Engle, and D. L. McFadden. Amsterdam: Elsevier Science.

BROCK, W. A., and B. D. LEBARON (1996): A Dynamic Structural Model for Stock Return Volatility and Trading Volume, *Rev. Econ. Stat.* 78, 94–110.

CAMPBELL, J. Y., and J. H. COCHRANE (1999): By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior, *J. Political Econ.* 107, 205–251.

CAMPBELL, J. Y., and L. HENTSCHEL (1992): No News Is Good News: An Asymmetric Model of Changing Volatility in Stock Returns, *J. Financial Econ.* 31, 281–318.

CAMPBELL, J. Y., and R. J. SHILLER (1988): Stock Prices, Earnings and Expected Dividends, *J. Finance* 43, 661–676.

CORRADI, V. (1997): Degenerate Continuous Time Limits of GARCH and GARCH-type Processes, Working paper, University of Pennsylvania.

CUTLER, D. M., J. M. POTERBA, and L. H. SUMMERS (1989): What Moves Stock Prices?, *J. Portfolio Mgt.* 15, 4–12.

DACOROGNA, M. M., U. A. MÜLLER, R. B. OLSEN, and O. V. PICTET (1998): Modelling Short-Term Volatility with GARCH and HARCH Models; in *Nonlinear Modelling of High Frequency Financial Time Series*, eds. C. Dunis and B. Zhou. New York: John Wiley.

DAVID, A. (1997): Fluctuating Confidence in Stock Markets: Implications for Returns and Volatility, *J. Financial Quant. Anal.* 32, 427–462.

DAVIS, M. H. A. (1981): Pathwise Nonlinear Filtering, in *Stochastic Systems. The Mathematics of Filtering and Identification and Applications,* eds. M. Hazewinkel and J. Willems. Dordrecht: Reidel.

DETEMPLE, J. B. (1986): Asset Pricing in a Production Economy with Incomplete Information, *J. Finance* 41, 383–391.

DETEMPLE, J. B., and S. MURTHY (1994): Intertemporal Asset Pricing with Heterogeneous Beliefs, *J. Econ. Theory* 62, 294–320.

DING, Z., C. W. J. GRANGER, and R. F. ENGLE (1993): A Long Memory Property of Stock Returns and a New Model, *J. Empirical Finance* 1, 83–106.

DUFOUR, F., and P. BERTRAND (1994): The Filtering Problem for Continuous-Time Linear Systems with Markovian Switching Coefficients, *Systems Control Letters,* 23, 453–461.

EASLY, D., and M. O'HARA (1992): Time and the Process of Security Price Adjustment, *J. Finance* 47, 577–605.

ELLIOTT, R. T., F. DUFOUR, and D. D. SWORDER (1996): Exact Hybrid Filters in Discrete Time, *IEEE Transactions on Automatic Control* 41, 1807–1810.

EVANS, M. D. D. (1998): Dividend Variability and Stock Market Swings, *Rev. Econ. Stud.* 65, 711–740.

EVANS, M. D. D., and K. K. LEWIS (1995): Do Expected Shifts in Inflation Affect Estimates of the Long-Run Fisher Relation?, *J. Finance* 50, 225–253.

FELDMAN, D. (1989): The Term Structure of Interest Rates in a Partially Observable Economy, *J. Finance* 44, 789–812.

FLORCHINGER, P. (1993): Zakai Equation of Nonlinear Filtering with Unbounded Coefficients. The Case of Dependent Noises, *Sys. Control Lett.* 21, 413–422.

FLORCHINGER, P., and F. L. GLAND (1991): Time-Discretization of the Zakai Equation for Diffusion Processes Observed in Correlated Noise, *Stoch. Stoch. Rep.* 35, 233–256.

FOSTER, F. D., and S. VISWANATHAN (1990): A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Securities Markets, *Rev. Financial Stud.* 3, 593–624.

FRENKEL, J. A., and M. L. MUSSA (1985): Asset Markets, Exchange Rates and the Balance of Payments, in *Handbook of International Economics,* vol. III, eds. R. W. Jones and P. B. Kenen. Amsterdam: North-Holland.

FUJISAKI, M., G. KALLIANPUR, and H. KUNITA (1972): Stochastic Differential Equations for the Nonlinear Filtering Problem, *Osaka J. Math.* 9, 19–40.

GALLANT, A. R., P. E. ROSSI, and G. TAUCHEN (1992): Stock Prices and Volume, *Rev. Financial Stud.* 5, 199–242.

HOBSON, D. G., and L. C. G. ROGERS (1998): Complete Models with Stochastic Volatility, *Math. Finance* 8, 27–48.

HONDA, T. (1997): Optimal Portfolio Choice when the Mean Return is an Unobservable Regime Switching Process; Working paper, Stanford University.

JOHNSON, T. C. (1999): Volatility, Momentum, and Time-Varying Skewness in Foreign Exchange Returns; Working paper, University of Chicago.

JONES, C., O. LAMONT, and R. LUMSDAINE (1995): Public Information and the Persistence of Bond Market Volatility, *J. Financial Econ.* 47, 315–337.

KALLIANPUR, G., and C. STRIEBEL (1968): Estimation of Stochastic Systems: Arbitrary System Process with Additive White Noise Observation Errors, *Ann. Math. Statist.* 39, 785–801.

KAMINSKY, G., and K. K. LEWIS (1996): Does Exchange Rate Intervention Signal Future Monetary Policy?, *J. Monetary Econ.* 37, 285–312.

KIND, P., R. LIPTSER, and W. RUNGGALDIER (1991): Diffusion Approximation in Past-Dependent Models and Applications to Options Pricing, *Ann. Appl. Prob.* 1, 379–405.

KRISHNAMURTHY, V., and J. EVANS (1998): Finite Dimensional Filters for Passive Tracking of Markov Jump Linear Systems; Working paper, University of Melbourne.

LIPTSER, R. S., and A. N. SHIRYAYEV (1977): *Statistics of Random Processes I, General Theory.* Berlin: Springer-Verlag.

LJUNG, G. M., and G. E. P. BOX (1978): On a Measure of Lack of Fit in Time-Series Models, *Biometrika* 65, 297–303.

LOPARO, K., A. Z. ROTH, and S. J. ECKERT (1986): Nonlinear Filtering for Systems with Random Structure, *IEEE Trans. Automatic Cont.* 31, 1064–1068.

LUX, T. (1997): Time Variation of Second Moments from a Noise Trader/Infection Model, *J. Econ. Dynam. Control,* 22, 1–38.

MANDELBROT, B. (1963): The Variation of Certain Speculative Prices, *J. Business,* 36, 349–419.

MARITON, M. (1990): *Jump Linear Systems in Automatic Control.* New York: Marcel Dekker.

MÜLLER, U. A., M. M. DACOROGNA, R. D. DAVÉ, R. B. OLSEN, O. V. PICTET, and J. E. VON WEIZSÄCKER (1997): Volatilities of Different Time Resolutions—Analyzing the Dynamics of Market Components, *J. Empirical Finance*, 4, 213–240.

NELSON, D. B. (1990): ARCH Models as Diffusion Approximations, *J. Econometrics* 45, 7–38.

PROTTER, P. (1990): *Stochastic Integration and Differential Equations.* New York: Springer-Verlag.

ROLL, R. (1984): Orange Juice and Weather, *Am. Econ. Rev.* 74, 861–880.

SENTANA, E. (1995): Quadratic ARCH Models, *Rev. Econ. Stud.* 62, 639–661.

TAUCHEN, G., and M. PITTS (1983): The Price Variability-Volume Relationship on Speculative Markets, *Econometrica* 51, 485–505.

VERONESI, P. (1999): Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model, *Rev. Financial Stud.* 12, 975–1007.

WANG, J. (1993): A Model of Intertemporal Asset Prices Under Asymmetric Information, *Rev. Econ. Stud.* 60, 249–282.

WANG, J. (1996): The Term Structure of Interest Rates in a Pure Exchange Economy with Heterogeneous Investors, *J. Financial Econ.* 41, 75–110.