Volatility Puzzles: A Unified Framework for Gauging Return-Volatility Regressions*

Tim Bollerslev†
and
Hao Zhou‡

August 20, 2003

Abstract

This paper provides a simple unified framework for assessing the empirical linkages between returns and realized and implied volatilities. First, we show that whereas the volatility feedback effect as measured by the sign of the correlation between contemporaneous return and realized volatility depends importantly on the underlying structural model parameters, the correlation between return and implied volatility is unambiguously positive for all reasonable parameter configurations. Second, the lagged return-volatility asymmetry, or the leverage effect, is always stronger for implied than realized volatility. Third, implied volatilities generally provide downward biased forecasts of subsequent realized volatilities. Our results help explain previous findings reported in the extant empirical literature, and is further corroborated by new estimation results for a sample of monthly returns and implied and realized volatilities for the aggregate S&P market index.

JEL Classification: G12, C51, C22.

Keywords: Leverage Asymmetry, Volatility Feedback, Implied Volatility Forecast, Realized Volatility, Stochastic Volatility Model, Model Misspecification, Estimation Bias.

---

*The work of Bollerslev was supported by a grant from the NSF to the NBER. We have benefited from discussions with Jim Clouse and Mike Gibson. Matthew Chesnes provided excellent research assistance. We would also like to thank George Jiang, Neil Shephard, Rossen Valkanov, along with seminar participants in the University of Arizona, and the Symposium of New Frontiers in Financial Volatility Modeling (Florence, Italy) for their helpful comments and suggestions. The views presented here are solely those of the authors and do not necessarily represent those of the Federal Reserve Board or its staff.

†Department of Economics, Duke University, Post Office Box 90097, Durham NC 27708 USA, Email boller@econ.duke.edu, Phone 919-660-1846, Fax 919-684-8974.

‡Division of Research and Statistics, Federal Reserve Board, Mail Stop 91, Washington DC 20551 USA, E-mail hao.zhou@frb.gov, Phone 202-452-3360, Fax 202-728-5889.

1 Introduction

Following the realization in the late eighties that financial market volatility is both time-varying and predictable, empirical investigations into the temporal linkages between aggregate stock market volatility and returns have figured very prominently in the literature. Of course, volatility per se is not directly observable, and several different volatility proxies have been employed in empirically assessing the linkages, including (i) model-based procedures that explicitly parameterize the volatility process as an ARCH or stochastic volatility model, (ii) direct market-based realized volatilities constructed by the summation of intra-period higher-frequency squared returns, and (iii) forward looking market-based implied volatilities inferred from options prices (see Andersen et al., 2003, for further discussion of the various volatility concepts and procedures). Meanwhile, a cursory read of the burgeoning volatility literature reveals a perplexing set of results, with the sign and the size of the reported volatility-return relationships differing significantly across competing studies and procedures.

The present paper provides a unified theoretical framework for reconciling these conflicting empirical findings. Specifically, by postulating a relatively simple parametric volatility model for the dynamic dependencies in the underlying returns, we show how the sign and the magnitude of the linear relationships between (i) the contemporaneous returns and the market-based volatilities, which we refer to as the volatility feedback effect, (ii) the lagged returns and the current market-based volatilities, which we refer to as the leverage effect, and (iii) the two different market-based volatilities, which we refer to as the implied volatility forecasting bias, all depend importantly on the parameters of the underlying structural model and the stochastic volatility risk premium.

The classical Intertemporal CAPM (ICAPM) model of Merton (1980) implies that the excess return on the aggregate market portfolio should be positively and directly proportionally related to the volatility of the market (see also Pindyck, 1984). This volatility feedback effect also underlies the ARCH-M model originally developed by Engle et al. (1987). However, empirical applications of the ARCH-M, and related stochastic volatility models, have met with mixed success. Some studies (see, e.g., French et al., 1987; Chou, 1988; Campbell and Hentschel, 1992; Ghysels et al., 2002) have reported consistently positive and significant estimates of the risk premium, while others (see, e.g., Campbell, 1987; Turner et al., 1989; Breen et al., 1989; Chou et al., 1992; Glosten et al., 1993) document negative values, unstable signs, or otherwise insignificant estimates. Moreover, the contemporaneous risk-return tradeoff appears sensitive to the use of ARCH as opposed to stochastic volatility formulations (Koopman and Uspensky, 1999), the length of the return horizon (Harrison and Zhang, 1999), along with the instruments and conditioning information used in empirically estimating the relationship (Harvey, 2001; Brandt and Kang, 2002). As we show below, these conflicting results are not necessarily inconsistent with the basic ICAPM model, in that the risk-return tradeoff relationship depends importantly on the particular volatility measure employed in the empirical investigations.[1]

The leverage effect, which predicts a negative correlation between current returns and future volatilities, was first discussed by Black (1976) and Christie (1982). The effect (and the name) may (in part) be attributed to a chain of events according to which a negative return causes an increase in the debt-to-equity ratio, in turn resulting in an increase in the future volatility of the return to equity.[2] Empirical evidence along these lines generally confirms that aggregate market volatility responds asymmetrically to negative and positive returns, but the economic magnitude is often small and not always statistically significant (e.g., Schwert, 1989; Nelson, 1991; Gallant et al., 1992; Glosten et al., 1993; Engle and Ng, 1993; Duffee, 1995; Bekaert and Wu, 2000). Moreover, the evidence tends to be weaker for individual stocks (e.g., Tauchen et al., 1996; Andersen et al., 2001). Importantly, the magnitude also depends on the volatility proxy employed in the estimation, with options implied volatilities generally exhibiting much more pronounced asymmetry (e.g., Bates, 2000; Wu and Xiao, 2002; Eraker, 2003)

A closely related issue concerns the bias in options implied volatilities as forecasts of the corresponding future realized volatilities. An extensive literature has documented that the market-based expectations embedded in options prices generally exceed the realized volatilities resulting in positive intercepts and slope coefficients less than unity in regression-based unbiasedness tests (see, e.g., Canina and Figlewski, 1993; Christensen and Prabhala, 1998; Day and M.Lewis, 1992; Fleming et al., 1995; Fleming, 1998; Lamoureux and Lastrapes, 1993, along with the recent survey in Poon and Granger, 2002). As formally shown in the

[1]More general multi-factor models also complicate the risk-return tradeoff relationship, as the projection of the returns on the volatility must now control for other state variables (see, e.g., Abel, 1988; Tauchen and Hussey, 1991; Backus and Gregory, 1993; Scruggs, 1998).

[2]Note, the volatility feedback effect, along with the well-documented persistent volatility dynamics, also implies an observationally equivalent negative correlation between current returns and future volatility, as a shock to the volatility will require an immediate return adjustment to compensate for the increased future risk. We follow the convention in the literature of referring to the negative correlation between future volatility and current returns as the leverage effect.

recent studies by Bates (2002), Chernov (2002), and Pan (2002), this bias is intimately related to the market price of volatility risk, and some of our theoretical results in regards to the implied volatility forecasting bias parallel the developments in these concurrent studies.

Our theoretical results are based on the one-factor continuous-time stochastic volatility model popularized by Heston (1993). This allows us to utilize various closed form expressions for the conditional moments previously derived by Andersen et al. (2002b) and Bollerslev and Zhou (2002). However, the same basic idea could in principle be generalized to other more complicated model structures, including multiple volatility factors and jumps, at the expense of notational and computational complexity (see, e.g. Andersen et al., 2002a; Eraker et al., 2003; Chernov et al., 2003). Nonetheless, the relatively simple one-factor affine Heston model is rich enough to explain our empirical findings in regards to the monthly return-volatility regressions for the Standard & Poor's aggregate market index.

The plan for the rest of the paper is as follows. Section 2 starts out by a discussion of the basic model structure, followed by the theoretical predictions related to the volatility feedback effect, the leverage effect, and the implied volatility forecasting bias, respectively. Section 3 provides confirmatory empirical evidence based on a fifteen-year sample of monthly returns, and high-frequency-based realized and implied volatilities for the Standard & Poor's composite index. Section 4 concludes. All of the derivations are given in a technical Appendix.

2 Theoretical Model Structure

Let p_t denote the time-t logarithmic price of the risky asset, or portfolio. The one-factor continuous-time affine stochastic volatility model of Heston (1993) then postulates the following dynamics for the instantaneous returns,

$$\begin{aligned} dp_t &= (\mu + \lambda_s V_t)dt + \sqrt{V_t}dB_t, \\ dV_t &= \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}dW_t, \\ \mathrm{corr}(dB_t, dW_t) &= \rho, \end{aligned} \tag{1}$$

where the latent stochastic volatility, V_t, is assumed to follow a square-root process. A negative instantaneous correlation between the two separate Brownian motions driving the price and volatility processes, or $\rho < 0$, is directly associated with the leverage effect in the raw returns; i.e., the tendency for contemporaneous returns and volatility to be negatively correlated. Similarly, the volatility feedback effect is captured directly by the risk-return trade-off parameter, $\lambda_s > 0$. Other more complicated model structures, including multiple latent volatility factors along with jumps in the price and/or volatility, could in principle be analyzed by similar means. However, for expositional purposes we restrict our analysis to the relatively simple model in equation (1).

Given this dynamics for the underlying price process, standard pricing arguments imply the existence of the following equivalent Martingale measure, or "risk-neutralized" distribution,

$$\begin{aligned} dp_t &= (r_t^* - d_t)dt + \sqrt{V_t}dB_t^*, \\ dV_t &= \kappa^*(\theta^* - V_t)dt + \sigma\sqrt{V_t}dW_t^*, \\ \text{corr}(dB_t^*, dW_t^*) &= \rho, \end{aligned} \tag{2}$$

where d_t refers to the dividend payout rate and r_t^* denotes the risk-neutral interest rate. The value of any contingent claim written on the underlying asset is now readily evaluated by calculating the expected payoff in this risk-neutral distribution.[3] We will refer to this expectation by the superscript *, as in $E^*(\cdot)$. The values of the risk-neutral parameters in (2) are directly related to the parameters of the actual price process in equation (1) by the functional relationships, $\kappa^* = \kappa + \lambda_v$ and $\theta^* = \kappa\theta/(\kappa + \lambda_v)$. The λ_v parameter refers to the stochastic volatility risk premium, which is generally estimated to be negative. Hence, the degree of mean reversion for the risk-neutralized volatility process, as determined by κ^*, is therefore slower (possibly even explosive) than the mean reversion for the actual volatility, as determined by κ (for a more detailed discussion of the connection between the objective and the risk-neutral distributions, see also Benzoni, 2001; Chernov, 2002; Wu, 2001; Pan, 2002).

We next turn to our discussion of the corresponding model-based implications for the different return-volatility regressions, starting with the volatility feedback effect.

2.1 Volatility Feedback Effect

Empirical assessments of the relationship between returns and contemporaneous volatility have typically found the volatility feedback effect to be statistically insignificant, and sometimes even negative. These results may appear at odds with the ICAPM and the corresponding one-factor model in equation (1). Thus, as discussed in the introduction, several studies have resorted to more complicated multi-factor representations as a way to resolve

[3]Notice, that in the presence of stochastic volatility it is generally not possible to perfectly hedge contingent claims payoff, and options are therefore no longer redundant assets.

this apparent empirical puzzle (see, e.g., Scruggs, 1998, and the discussion therein). Meanwhile, consider the continuously compounded returns from time t to t+Δ implied by the simple model in (1),

$$R_{t,t+\Delta} = p_{t+\Delta} - p_t = \mu\Delta + \lambda_s \int_t^{t+\Delta} V_u du + \int_t^{t+\Delta} \sqrt{V_u} dB_u. \tag{3}$$

Although the "residual" defined by $\int_t^{t+\Delta} \sqrt{V_u} dB_u$ is heteroskedastic, the population regression of the returns on a constant and the integrated volatility would correctly uncover the volatility feedback effect ($\lambda_s > 0$), provided that the orthogonality condition $E\left(\int_t^{t+\Delta} \sqrt{V_u} dB_u \times \int_t^{t+\Delta} V_u du\right) = 0$ holds true. However, with a non-zero "leverage" effect, or $\rho < 0$, the residual and the integrated volatility will be correlated, resulting in a biased estimate for λ_s.

Specifically, consider the population regression,

$$R_{t,t+\Delta} = \alpha + \beta \int_t^{t+\Delta} V_u du + e_{t,t+\Delta}. \tag{4}$$

Then as formally shown below, unless $\rho = 0$, the population feedback coefficient β will differ from the true feedback coefficient λ_s. Of course, the integrated volatility is not directly observable, so the sample counterpart to the population regression in (4) isn't actually feasible. However, the integrated volatility may in theory be approximated arbitrarily well by the corresponding *realized* volatility constructed by the summation of sufficiently finely sampled high-frequency squared returns (see, e.g., Andersen et al., 2003). This approach, which is now routinely employed in the literature, also underlies our empirical analysis in Section 3 below.

Alternatively, consider the corresponding *implied* volatility-return regression,

$$R_{t,t+\Delta} = \alpha^* + \beta^* E_t^* \left(\int_t^{t+\Delta} V_u du\right) + e^*_{t,t+\Delta}, \tag{5}$$

where the risk neutral expectation is taken under the distribution in (2). In this situation, unless the stochastic volatility risk-premium equals zero, or $\lambda_v = 0$, the population feedback coefficient will again differ from the true feedback coefficient in equation (3), that is $\beta^* \neq \lambda_s$. Hence, to correctly uncover the volatility feedback parameter from a contemporaneous return-volatility type regression, either the leverage effect must be zero if the regression is based on a realized volatility proxy, or the stochastic volatility risk premium must be zero when using options implied volatilities. Of course neither case is likely to hold empirically. Proposition 1 characterizes the exact form of the resulting biases.[4]

Proposition 1 *Assume that the parameters in (1) and (2) adhere to the standard sign restrictions,* $\kappa > 0$, $\theta > 0$, $\sigma > 0$, $\rho < 0$, $\lambda_v < 0$, $\lambda_s > 0$, *and that* $\mu \neq 0$. *The population feedback coefficient in the integrated volatility regression in equation (4) is then given by,*

$$\beta = \lambda_s + \frac{\rho\kappa}{\sigma} < \lambda_s. \tag{6}$$

Let $a_\Delta = (1 - e^{-\kappa\Delta})/\kappa$ *and* $a^*_\Delta = (1 - e^{-\kappa^*\Delta})/\kappa^*$. *The population feedback coefficient in the implied volatility regression in (5) may then be expressed as,*

$$\beta^* = \lambda_s \frac{a_\Delta}{a^*_\Delta} < \lambda_s. \tag{7}$$

Moreover, assuming that $0 < \lambda_s < -\frac{\rho\kappa}{\sigma}$, *the two slope parameters are related by,*

$$\beta < 0 < \beta^* < \lambda_s, \tag{8}$$

while for $0 < -\frac{\rho\kappa}{\sigma} < \lambda_s < \frac{a^*_\Delta}{a_\Delta - a^*_\Delta}\frac{\rho\kappa}{\sigma}$, *we have*

$$0 < \beta < \beta^* < \lambda_s. \tag{9}$$

The proof of the proposition is given in the technical Appendix A.

The implications of the proposition for empirical studies designed to uncover the volatility feedback effect are immediate. First, regression-based procedures utilizing realized volatility proxies will invariably result in a downward biased slope estimate, with the sign and magnitude depending on the underlying structural parameters. This, of course, is entirely consistent with the extant literature discussed above reporting inconclusive and sometimes even negative estimates for β. Only if the leverage, or asymmetry, effect is zero ($\rho = 0$) will the regression be unbiased for estimating λ_s. Second, regression estimates based on implied volatility will generally show less of a downward bias and remain positive under the most general parameter setting. However, only if the stochastic volatility risk premium equals zero ($\lambda_v = 0$) will the bias completely disappear. Again, this is directly in line with the existing literature discussed above, as well as the new empirical results reported in Section 3 below.

[4] As shown in the appendix, the population intercepts α and α^* will also generally differ from the true drift in equation (3), that is $\mu\Delta$. However, we will focus our discussion on the slope coefficients which are typically associated with the volatility feedback effect.

This result also helps explain why various versions of filtered volatility (obtained by projecting on lagged historical squared and/or absolute returns) may produce less biased or even positive β estimates. Specifically, instead of the *realized* return - *realized* volatility trade-off regression in (4), consider the *realized* return - *expected* volatility trade-off regression

$$R_{t,t+\Delta} = \tilde{\alpha} + \tilde{\beta} E_t \left(\int_t^{t+\Delta} V_u du \right) + e_{t,t+\Delta}.$$

This regression explicitly purges the simultaneous correlation between the return and volatility innovations. As such, this regression corresponds more closely to the *implied* return-volatility trade-off regression in (5) that obtain by replacing the expected integrated volatility, $E_t \left(\int_t^{t+\Delta} V_u du \right)$, with its risk neutral equivalent, $E_t^* \left(\int_t^{t+\Delta} V_u du \right)$. Of course, the expected integrated volatility will generally depend upon the underlying structural model, but may be approximated empirically through the use of instrumental variables procedures. However, as previously noted, the resulting estimates for the risk-return trade-off relationship are often very sensitive to the particular *ad hoc* choice of instruments employed in the estimation (Harvey, 2001; Brandt and Kang, 2002). We shall return to this issue in the empirical Section 3.1 below.

At a more general level Proposition 1 clearly highlights the importance of the volatility proxy used in the estimation of the risk-return trade-off relationship, and as such indirectly explains the instability in the estimates reported in the extant literature in regards to the model choice, instrument control, and return horizon. Similar issues arise in the empirical estimation of the leverage effect, to which we turn next.

2.2 Leverage Effect

Several different parametric volatility models and volatility-return regressions have been employed in the literature for empirically assessing the leverage effect (see e.g., the discussion in Bekaert and Wu (2000), along with the surveys of the ARCH literature in Bollerslev et al. (1992) and Bollerslev et al. (1994)). Although most estimates support the hypothesis that aggregate stock market volatility responds asymmetrically to past negative and positive returns, as discussed in the introduction, the magnitude and the statistical significance of the estimated effect is quite sensitive to the return horizon and the particular volatility proxy employed in the estimation.

At the most basic level the leverage effect is generally associated with a negative correlation between current volatility and lagged returns. To formally quantify this correlation, consider the corresponding population regressions for the integrated volatility,[5]

$$\int_t^{t+\Delta} V_u du = \gamma + \delta R_{t-\Delta,t} + e_{t,t+\Delta}, \tag{10}$$

and the option implied volatility,

$$E_t^* \left(\int_t^{t+\Delta} V_u du \right) = \gamma^* + \delta^* R_{t-\Delta,t} + e^*_{t-\Delta,t}, \tag{11}$$

where the expectation in equation (11) is again taken with respect to the risk-neutral distribution. Of course, the slope parameters in the simplified asymmetry regressions in (10) and (11) do not correspond directly to the leverage, or asymmetry, parameter ρ determining the correlation between the two Brownian motions in (1). However, as the following proposition makes clear, the population regression parameters may be expressed as explicit nonlinear functions of the underlying structural parameters in (1) and (2). These functional relationships in turn explain the stronger asymmetry observed empirically between implied volatility and lagged-returns.

Proposition 2 *Assume that the parameters in (1) and (2) adhere to the standard sign restrictions,* $\kappa > 0$, $\theta > 0$, $\sigma > 0$, $\rho < 0$, $\lambda_v < 0$, *and* $\lambda_s > 0$. *Let* $a_\Delta = (1 - e^{-\kappa\Delta})/\kappa$, $a^*_\Delta = (1 - e^{-\kappa^*\Delta})/\kappa^*$, *and* $c_\Delta = (e^{-\kappa\Delta} + \kappa\Delta - 1)/\kappa$. *The population slope parameters in (10) and (11) may then be expressed as,*

$$\delta = \frac{\lambda_s \frac{\theta\sigma^2}{2\kappa} a_\Delta^2 + \rho\sigma\theta a_\Delta^2}{\theta\Delta + \frac{\lambda_s\sigma\theta}{\kappa}\left(\frac{\lambda_s\sigma}{\kappa} + \rho\right) c_\Delta}, \tag{12}$$

and

$$\delta^* = \frac{\rho\sigma\theta a^*_\Delta a_\Delta}{\theta\Delta + \frac{\lambda_s\sigma\theta}{\kappa}\left(\frac{\lambda_s\sigma}{\kappa} + \rho\right) c_\Delta}. \tag{13}$$

Moreover, assuming $0 < \lambda_s < -\frac{2\rho\kappa}{\sigma}$ *it follows that,*

$$\delta^* < \delta < 0, \tag{14}$$

while for $0 < -\frac{2\rho\kappa}{\sigma} < \lambda_s$,

$$\delta^* < 0 < \delta. \tag{15}$$

[5] In the empirical section we also report the results from a longer regression in which we include the lagged volatility along with different response coefficients for positive and negative returns.

The proof of the proposition is given in Appendix B.

It is noteworthy that for the integrated volatility regression, the "leverage" coefficient, δ, depends critically on both the volatility feedback parameter λ_s (positively), as well as the "structural" leverage coefficient ρ (negatively). Thus, although most empirical studies report strong realized volatility asymmetry for the aggregate market portfolio ($\delta < 0$), this may help explain the lack of statistical significance, and sometimes even reverse asymmetry reported occasionally. For the implied volatility regression, the "leverage" coefficient δ^* depends directly on ρ (negatively), and the stochastic volatility risk premium λ_v through a^*_Δ (magnitude). In contrast to the integrated volatility regression, the coefficient in the implied volatility regression is unambiguously negative provided that $\rho < 0$. Moreover, provided that the volatility feedback effect is positive ($\lambda_s > 0$), and the stochastic volatility risk premium is negative ($\lambda_v < 0$), as it is commonly assumed in the literature, the magnitude of the implied volatility asymmetry always exceeds that of the integrated volatility, that is $\delta^* < \delta$.

Similar considerations help to explain the downward bias in the implied-realized volatility forecasting regressions, to which we now turn.

2.3 Implied Volatility Forecasting Bias

The two previous subsections demonstrate how the use of realized or implied volatility proxies can result in quite different population parameters in the contemporaneous and lagged return-volatility regressions. A closely related question, concerns the extent to which implied volatilities provide unbiased forecasts of the corresponding future realized volatilities. The most common approach employed in the literature for assessing the forecasting bias is based on regressing the ex-post realized volatility over some time period, say $[t, t+\Delta]$, on a constant and the time t implied volatility for an option maturing at $t+\Delta$ (for a recent survey of this extensive empirical literature see Poon and Granger, 2002). In population,

$$\int_t^{t+\Delta} V_u du = \phi_0 + \phi_1 E_t^* \left(\int_t^{t+\Delta} V_u du \right) + e_{t,t+\Delta}. \tag{16}$$

Obviously, for the implied volatility to provide unbiased forecasts, the two projection coefficients should equal $\phi_0 = 0$ and $\phi_1 = 1$, respectively. Meanwhile, most empirical studies report statistically significant biases in the direction of $\phi_0 > 0$ and $\phi_1 < 1$. These empirical biases have in part been explained by a standard errors-in-variables type problem arising from the use of finite-sample equivalents to the population regression in (16) (Christensen and

Prabhala, 1998), along with very persistent volatility dynamics rendering standard statistical inference unreliable (Bandi and Perron, 2003). However, these statistical considerations aside, it follows that if the stochastic volatility risk premium, λ_v, differs from zero, the two population regression coefficients in (16) implied by the structural model in (1) and (2) will not equal zero and unity, respectively.

Proposition 3 *Assume that the parameters in (1) and (2) adhere to the standard sign restrictions, $\kappa > 0$, $\theta > 0$, $\sigma > 0$, $\rho < 0$, $\lambda_v < 0$, and $\lambda_s > 0$. The population parameters in the regression in (16) are then given by,*

$$\phi_0 = b_\Delta - \frac{a_\Delta}{a^*_\Delta} b^*_\Delta \qquad \textit{and} \qquad \phi_1 = \frac{a_\Delta}{a^*_\Delta} < 1, \tag{17}$$

where $a_\Delta = (1 - e^{-\kappa\Delta})/\kappa$, $a^*_\Delta = (1 - e^{-\kappa^*\Delta})/\kappa^*$, $b_\Delta = \theta(\Delta - a_\Delta)$, and $b^*_\Delta = \theta^*(\Delta - a^*_\Delta)$.

The proof of the proposition is given in Appendix C.

The proposition immediately explains the typical finding of a downward bias in the estimated slope coefficient. Intuitively, for $\lambda_v < 0$, the stochastic volatility risk premium reduces the degree of mean reversion in the risk-neutral volatility process relative to that of the actual volatility process ($\kappa^* < \kappa$), in turn resulting in the ratio a_Δ/a^*_Δ becoming less than one.[6] Thus, any estimate of ϕ_1 should be gauged against this population bias. Of course, the true structural model parameters, κ and κ^*, are generally unknown and would have to be estimated.

Conversely, the population return-volatility regressions in Propositions 1-3 (coupled with additional moment restrictions along the lines of Bollerslev and Zhou, 2002) could in principle be employed as a system of equations in estimating the structural model parameters underlying the actual and risk-neutral distributions in (1) and (2). We shall not pursue this approach any further here. Instead, we next turn to a concrete empirical implementation of the five regression equations based on aggregate U.S. stock market volatility and returns for directly illustrating the various biases implied by Propositions 1-3.

[6] Closely related results, along with a more detailed analysis of the impact of jumps, have recently been derived in concurrent work by Bates (2002) and Chernov (2002).

3 Empirical Illustration

Our empirical analysis is based on monthly returns and volatilities for the S&P composite index spanning the period from January 1986 through February 2002.[7] The monthly continuously compounded percentage returns are constructed from the daily S&P500 closing prices supplied by the Wall Street Journal. Normalizing the monthly time interval to unity, we will refer to the return over the $t+1$'th month as $R_{t,t+1}$.

The corresponding realized volatilities are based on the summation of the five-minute squared returns within the month. The high-frequency data for the S&P500 index is provided by the Institute of Financial Markets. Specifically, with n_{t+1} trading days in month $t+1$,

$$\mathrm{RV}_{t,t+1} \equiv \sum_{i=1}^{78 \cdot n_{t+1}} (\log P_{t+i/78 \cdot n_{t+1}} - \log P_{t+(i-1)/78 \cdot n_{t+1}})^2, \tag{18}$$

where the 78 five-minute subintervals represents the normal trading hours from 9:30am to 4:00pm, including the close-to-open five-minute interval. As discussed in the previous section, the realized volatility, $\mathrm{RV}_{t,t+1}$, is readily interpreted as a consistent (for increasing sampling frequency) estimate of the corresponding integrated volatility, $\mathrm{IV}_{t,t+1} \equiv \int_t^{t+1} V_s ds$.

The monthly implied volatility (variance) is formally defined by,

$$\mathrm{IV}^*_{t,t+1} \equiv E^* \left[\int_t^{t+1} V_s ds | \mathcal{F}_t \right], \tag{19}$$

where E^* refers to the risk-adjusted expectation of the one-month ahead integrated volatility, $\mathrm{IV}_{t,t+1}$.[8] The most actively traded equity index options are written on the S&P100 index. Hence for liquidity reasons, we rely on the corresponding option implied volatilities (VIX) provide by the CBOE in empirically quantifying $\mathrm{IV}^*_{t,t+1}$. The VIX index (reported in annualized percentage standard deviation form) is based on a weighted average of the one-month ahead volatilities inverted from eight near-the-money puts and calls, and is now widely regarded among market participants as *the* "implied volatility index" (see Fleming

[7] The start date of January 1986 reflects the availability of both the S&P high-frequency return data and VIX implied volatility index. However, all of the regression results reported below are materially unaffected by excluding the October 1987 stock market crash and starting the sample in January 1988. These results are available upon request.

[8] This measure is also sometimes referred to as the "risk-neutral integrated volatility." In empirical studies the implied volatility concept is used almost exclusively in reference to the volatility that equates the Black-Scholes price with the actual price of an option. This generally provides an accurate approximation for short-lived at-the-money options; see, e.g., Ledoit et al. (2002).

et al., 1995; Fleming, 1998, for a precise definition and further discussion of the VIX index).[9]

To facilitate the theoretical derivations, all of the volatility regressions analyzed in the previous section were cast in the form of variances corresponding to the empirical $\mathrm{RV}_{t,t+1}$ and $\mathrm{IV}^*_{t,t+1}$ measures defined above. However, for robustness reasons previous empirical studies have often been implemented in the form of standard deviations. Hence, we augment the variance regressions with the analogous regressions based on $\mathrm{RV}^{1/2}_{t,t+1}$ and $\mathrm{IV}^{*1/2}_{t,t+1}$.

Summary statistics for all of the variables are reported in Table 1. For comparison purposes the standard deviations and the variances are converted to percentage and squared percentage points, respectively. From the first column, the average annualized return on the market was about ten percent, with a sample standard deviation of around sixteen percent. The returns are negatively skewed with much fatter tails than the normal distribution. The implied volatilities are systematically higher than the realized volatilities, and their unconditional distributions also deviate more from the normal. The returns are approximately serially uncorrelated, while the volatility series (both in standard deviation and variance forms) exhibit pronounced own temporal dependencies. In fact, the first ten autocorrelations reported in the bottom part of the table are all highly significant with the gradual, but very slow, decay suggestive of long-memory type features. This is also evident from the time series plots for each of the five series given in Figure 1.

3.1 Volatility Feedback Effect

Our estimates of the volatility feedback effect are based on the empirical equivalents to the two population regressions in equations (4) and (5),

$$\mathrm{R}_{t,t+1} = \alpha + \beta \mathrm{RV}_{t,t+1} + u_{t+1}, \tag{20}$$
$$\mathrm{R}_{t,t+1} = \alpha^* + \beta^* \mathrm{IV}^*_{t,t+1} + u^*_{t+1}, \tag{21}$$

along with the corresponding robust regressions in standard deviation form,

$$\mathrm{R}_{t,t+1} = \alpha + \beta \mathrm{RV}^{1/2}_{t,t+1} + u_{t+1}, \tag{22}$$
$$\mathrm{R}_{t,t+1} = \alpha^* + \beta^* \mathrm{IV}^{*1/2}_{t,t+1} + u^*_{t+1}, \tag{23}$$

[9] Ideally, we would want realized and implied volatilities and returns for the same index. However, reliable high-frequency data, required for the construction of the realized volatilities, are only available for the S&P500 index, while liquidity considerations in the options market dictates the use of the S&P100 VIX index. Of course, the returns on the two indexes are very close, with a monthly sample correlation of 0.988.

where the residuals from the regressions are generically denoted by u_{t+1} and u^*_{t+1}, respectively. The coefficient estimates, along with their asymptotic standard errors based on a Newey-West covariance matrix estimator allowing for a two-month lag, are reported in Table 2. Interestingly, for the regressions involving the implied volatility, the intercept and slope coefficients are all insignificant, and both of the R-squares are very close to zero. In contrast, the two realized volatility regressions both result in highly significant estimates for the intercepts (positive) and slopes (negative). These contemporaneous return-volatility regressions also explain about ten and thirteen percent of the variability in the ex-post monthly returns, respectively. Although the empirical finding of a significant negative relationship between aggregated stock market returns and realized volatility may appear counter intuitive, the result is, of course, entirely consistent with Proposition 1 and the ranking of the corresponding population parameters, $\beta < 0 < \beta^*$, given that the condition $0 < \lambda_s < -\frac{\rho\kappa}{\sigma}$ is satisfied by the underlying structural model parameters.[10]

It is worth noting, that while all of the regressions above involve a trade-off between monthly returns and volatilities over the identical time horizon, the one-month implied volatility is determined at time t, whereas the realized volatility is not observable until $t+1$. As such, the results in Table 2 are also consistent with the recent empirical findings by Brandt and Kang (2002), who report a (puzzling) negative contemporaneous relation between the conditional mean and the conditional variance of the market returns, along with a more conventional positive tradeoff for the one-month lagged volatility. Similarly, Ghysels et al. (2002) report a significant positive trade-off relationship when the squared returns 20-50 days in the past are weighted most heavily in their realized volatility constructs.

In order to further illustrate this point, the last columns in each of the panels in Table 2 report the results from a standard instrumental variables procedure in which we rely on the absolute lagged returns as instruments for the realized volatilities in the two regressions in (20) and (22). Although the slope coefficient estimates for both the standard deviation and the variance formulation remain negative, they are clearly much closer to zero, and the regression R-squares drop from 10-13% to about 5-6%. Moreover, the corresponding standard errors are also much larger so that the estimates are no longer statistically significant. It is also noteworthy that on using the lagged raw returns as instruments, the estimates for the

[10]Recall that the positive return volatility trade-off observed in some empirical studies, $0 < \beta < \beta^*$, may similarly be justified by the second case of Proposition 1, when the underlying structural parameters satisfy the condition $0 < -\frac{\rho\kappa}{\sigma} < \lambda_s < \frac{a^*_\Delta}{a_\Delta - a^*_\Delta}\frac{\rho\kappa}{\sigma}$.

two slope coefficients (available upon request) are both positive, albeit even closer to zero statistically.[11] As such, these results further highlight the sensitivity to the particular volatility proxy and instrument choice employed in the reduced form volatility feedback regressions, despite the fact that some instruments may produce conventional positive feedback effect.

3.2 Leverage Effect

The empirical equivalents to the simple leverage regressions analyzed in Section 2.2 above take the form,

$$\mathrm{RV}_{t,t+1} = \gamma + \delta \mathrm{R}_{t-1,t} + u_{t+1}, \tag{24}$$

$$\mathrm{IV}^*_{t,t+1} = \gamma^* + \delta^* \mathrm{R}_{t-1,t} + u^*_t, \tag{25}$$

with the theoretical prediction from Proposition 2 that in population $\delta^* < \delta < 0$; i.e., the implied volatility is more responsive to the lagged return than the realized volatility. Again, for robustness reasons, the asymmetry implications may alternatively be tested in standard deviation form,

$$\mathrm{RV}^{1/2}_{t,t+1} = \gamma + \delta \mathrm{R}_{t-1,t} + u_{t+1}, \tag{26}$$

$$\mathrm{IV}^{*1/2}_{t,t+1} = \gamma^* + \delta^* \mathrm{R}_{t-1,t} + u^*_t, \tag{27}$$

with the similar predictions in regards to the sign and ordering of the slope coefficients. Moreover, to account for the strong own temporal dependencies in the volatility and to allow for different impacts from past negative and positive returns, a slightly longer asymmetry regression is often estimated empirically,

$$\mathrm{RV}_{t,t+1} = \gamma + \beta \mathrm{RV}_{t-1,t} + \alpha \left(\mathrm{R}_{t-1,t}\right)^2 - \delta \left(\mathrm{R}_{t-1,t}\right)^2 I_{\left(\mathrm{R}_{t-1,t} \leq 0\right)} + u_{t+1}, \tag{28}$$

$$\mathrm{IV}^*_{t,t+1} = \gamma^* + \beta^* \mathrm{IV}^*_{t-1,t} + \alpha^* \left(\mathrm{R}_{t-1,t}\right)^2 - \delta^* \left(\mathrm{R}_{t-1,t}\right)^2 I_{\left(\mathrm{R}_{t-1,t} \leq 0\right)} + u^*_t. \tag{29}$$

In these longer regressions, weak asymmetry would again be implied by negative δ's, while strong asymmetry would have the α's be negative as well. Similarly, if the implied volatility responds more asymmetrically to past returns than do the realized volatility, we would

[11] We also experimented with a number of other instrumental variables, including the lagged squared returns and the lagged volatilities in standard deviation and variance forms. Although not statistically significant, only the results based on the absolute lagged returns were generally consistent in terms of their signs across different subsamples.

expect to find that the estimates for the δ's satisfy the relation $\delta^* < \delta < 0$. These same considerations apply to the pair of robust standard deviation regressions,

$$\mathrm{RV}_{t,t+1}^{1/2} = \gamma + \beta \mathrm{RV}_{t-1,t}^{1/2} + \alpha \left|\mathrm{R}_{t-1,t}\right| - \delta \left|\mathrm{R}_{t-1,t}\right| I_{\left(\mathrm{R}_{t-1,t} \leq 0\right)} + u_{t+1}, \tag{30}$$

$$\mathrm{IV}_{t,t+1}^{*1/2} = \gamma^* + \beta^* \mathrm{IV}_{t-1,t}^{*1/2} + \alpha^* \left|\mathrm{R}_{t-1,t}\right| - \delta^* \left|\mathrm{R}_{t-1,t}\right| I_{\left(\mathrm{R}_{t-1,t} \leq 0\right)} + u_t^*. \tag{31}$$

Note that for $\beta = 0$ and $\delta = 2\alpha$ the long regression in equation (30) collapses to the short regression in equation (26). Likewise, for $\beta^* = 0$ and $\delta^* = 2\alpha^*$ the two risk-neutral regressions in equations (31) and (27) coincide.

The actual S&P estimation results for the leverage regressions are reported in Table 3. The intercepts and slope coefficients for the short regressions in the first panel of the table are all highly significant. The R-squares for the realized volatility regressions are systematically lower than the corresponding R-squares for the implied volatilities. Importantly, the estimates of the δ's from the variance regression, $\delta^* = -3.31$ and $\delta = -0.82$, also adhere exactly to the theoretical predictions from Proposition 2 of negative and more pronounced asymmetry for the implied volatility. Similarly, the short regressions in standard deviation form results in estimates of $\delta^* = -0.17$ and $\delta = -0.08$, both of which are significantly less than zero. Turning to the longer regressions, it is noteworthy that the asymmetry in the realized volatility are no longer statistically significant, while the two estimates of δ^* from the implied volatility regressions are both overwhelmingly significant, and economically large. Again, this is directly in line with the implications from Proposition 2 of strong asymmetry for the implied volatility along with weak, and possibly even reverse, asymmetry for the realized volatility.[12] The well-documented strong own temporal dependencies in the volatility also result in fairly large and highly significant estimates for the β's, along with much higher R-squares for the long volatility regressions (84-88% for the implied and 38-55% for the realized volatilities).

All-in-all, the at first somewhat puzzling empirical findings for the different regressions reported in Table 3 again highlight the importance of properly interpreting the estimated asymmetry in lieu of the theoretical implications for the different volatility proxies detailed in Proposition 2. In this regard, the results in Table 3 are also consistent with previous empirical evidence in the literature related to the significance, or the lack thereof, of the

[12] Indeed, subsample analysis excluding 1986-1987 and 1998-2002, result in reverse asymmetry for the realized volatility, while the implied volatility asymmetry remains very strong. This is therefore consistent with the second case of Proposition 2, i.e., $\delta^* < 0 < \delta$. Further details concerning these subsample results are again available upon request.

volatility asymmetry effect for other markets and time periods (see, e.g., Schwert, 1989; Nelson, 1991; Gallant et al., 1992; Engle and Ng, 1993; Duffee, 1995; Bekaert and Wu, 2000; Wu, 2001, among others). We next turn to discussion of the related empirical evidence concerning the unbiasedness regressions directly linking the implied and realized volatility.

3.3 Implied Volatility Forecasting Bias

The question of whether implied volatilities provide unbiased and informationally efficient forecasts of the corresponding future realized volatilities have been studied extensively in the empirical finance literature. The typical regression employed in the literature takes the form,

$$\mathrm{RV}_{t,t+1} = \phi_0 + \phi_1 \mathrm{IV}^*_{t,t+1} + u_{t+1}, \tag{32}$$

or in terms of standard deviations,

$$\mathrm{RV}^{1/2}_{t,t+1} = \phi_0 + \phi_1 \mathrm{IV}^{*1/2}_{t,t+1} + u_{t+1}, \tag{33}$$

where unbiasedness would be associated with $\phi_0 = 0$ and $\phi_1 = 1$. Of course, the theoretical results in Proposition 3 implies that these are not the values to be expected empirically if the stochastic volatility risk is priced.[13]

The actual estimation results reported in Table 4 also do not support the unbiasedness hypothesis. Both of the estimates for ϕ_1 are significantly less than unity, and the estimates for ϕ_0 are greater than zero, albeit not significantly so. The regression in standard deviation form results in a fairly high R-square of 0.51, while the R-square from the less robust variance regression equals 0.34. These findings of a downward bias in the implied volatility forecasts along with fairly high explanatory power when judged by the high-frequence based realized volatility measures are directly in line with recent empirical results in the literature (see, e.g., Martens and Zein, 2002; Neely, 2003, and the survey by Poon and Granger, 2002). Moreover, the direction of the estimated biases are exactly as expected from Proposition 3, and as such do not necessarily suggest any inefficiencies.

[13] The existence of a non-zero stochastic volatility risk premium for explaining estimates of $\phi_0 > 0$ and $\phi_1 < 1$, along the lines of Proposition 3, has previously been discussed by Bates (2002), Benzoni (2001), Chernov (2002), and Pan (2002), among others. Meanwhile, as in Fleming (1998), different subsamples result in estimates of $\phi_0 \leq 0$, which is still consistent with Proposition 3. These results are available upon request.

4 Conclusion

The simple unified continuous-time framework derived in this paper for assessing the linkages between discretely observed returns and realized and implied volatilities help explain a number of puzzling findings in the extant empirical literature. In particular, we show that whereas the sign of the correlation between return and implied volatility is unambiguously positive, the correlation between contemporaneous return and realized volatility is generally undetermined. Similarly, the lagged return-volatility asymmetry is always stronger for implied than realized volatility. Also, implied volatilities generally provide downward biased forecasts of subsequent realized volatilities.

It would be interesting to extend the empirical analysis for the aggregate S&P market index presented here to other markets. In particular, the volatility feedback and asymmetry effects may not be as important for other markets, and consequently result in qualitatively different return-volatility linkages. The theoretical analysis of more complicated model structures allowing for jumps in the volatility and/or returns along with multiple volatility factors may also give rise to additional new insights. The regression-based implications derived here could also be used in directly estimating the underlying objective and risk-neutral dynamics, including the stochastic volatility risk premium, by appropriately matching the sample and population moments for the realized and implied volatilities. We leave further work along these lines for future research.

References

Abel, Andrew B. (1988), "Stock Prices under Time-Varying Dividend Risk: An Exact Solution in an Infinite-Horizon General Equilibrium Model," *Journal of Monetary Economics*, vol. 22, 375–393.

Andersen, Torben G., Luca Benzoni, and Jesper Lund (2002a), "An Empirical Investigation of Continuous-Time Equity Return Models," *Journal of Finance*, vol. 57, 1239–1284.

Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold (2003), *Handbook of Financial Econometrics*, chap. Parametric and Nonparametric Volatility Measurement, Elsevier Science B.V., Amsterdam, forthcoming.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens (2001), "The Distribution of Realized Stock Return Volatility," *Journal of Financial Economics*, vol. 61, 43–76.

Andersen, Torben G., Tim Bollerslev, and Nour Meddahi (2002b), "Analytical Evaluation of Volatility Forecasts," *Working Paper*, Northwestern, Duke, and University of Montreal.

Backus, David K. and Allen W. Gregory (1993), "Theoretical Relations between Risk Premiums and Conditional Variance," *Journal of Business and Economic Statistics*.

Bandi, Federico and Benoit Perron (2003), "Long Memory and the Relation Between Implied and Realized Volatility," *Working Paper*, Department of Economics, University of Montreal.

Bates, David S. (2000), "Post-'87 Crash Fears in the S&P 500 Futures Option Market," *Journal of Econometrics*, vol. 94, 181–238.

Bates, David S. (2002), "Empirical Option Pricing: A Retrospection," *Journal of Econometrics*, forthcoming.

Bekaert, Geert and Guojun Wu (2000), "Asymmetric Volatility and Risk in Equity Markets," *Review of Financial Studies*, vol. 13, 1–42.

Benzoni, Luca (2001), "Pricing Options under Stochastic Volatility: An Empirical Investigation," *Working Paper*, University of Minnesota.

Black, Fischer (1976), "Studies in Stock Price Volatility Changes," *Proceedings of the 1976 Business Meeting of the Business and Economics Section, American Statistical Association*, pages 177–181.

Bollerslev, Tim, Ray Y. Chou, and Kenneth F. Kroner (1992), "ARCH Modeling in Finance," *Journal of Econometrics*, vol. 52, 5–59.

Bollerslev, Tim, Robert F. Engle, and Daniel B. Nelson (1994), "ARCH Models," in "Handbook of Econometrics," (edited by Engle, Robert F. and Daniel L. McFadden), vol. IV, Amsterdam: North-Holland.

Bollerslev, Tim and Hao Zhou (2002), "Estimating Stochastic Volatility Diffusion Using Conditional Moments of Integrated Volatility," *Journal of Econometrics*, vol. 109, 33–65.

Brandt, Michael W. and Qiang Kang (2002), "On the Relationship Between the Conditional Mean and Volatility of Stock Returns: A Latent VAR Approach," *Journal Financial Economics*, forthcoming.

Breen, William, Lawrence R. Glosten, and Ravi Jagannathan (1989), "Economic Significance of Predictable Variations in Stock Index Returns," *Journal of Finance*, vol. 44, 1177–1189.

Campbell, John Y. (1987), "Stock Returns and the Term Structure," *Journal of Financial Economics*, vol. 18, 373–399.

Campbell, John Y. and Ludger Hentschel (1992), "No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns," *Journal of Financial Economics*.

Canina, Linda and Stephen Figlewski (1993), "The Informational Content of Implied Volatility," *Review of Financial Studies*, vol. 6, 659–681.

Chernov, Mikhail (2002), "On the Role of Volatility Risk Premia in Implied Volatilities Based Forecasting Regressions," *Working Paper*, Columbia University.

Chernov, Mikhail, A. Ronald Gallant, Eric Ghysels, and George Tauchen (2003), "Alternative Models for Stock Price Dynamics," *Journal of Econometrics*, forthcoming.

Chou, Ray Y. (1988), "Volatility Persistence and Stock Valuations: Some Empirical Evidence Using GARCH," *Journal of Econometrics*, vol. 3, 279–294.

Chou, Ray Y., Robert F. Engle, and Alex Kane (1992), "Measuring Risk Aversion from Excess Returns on a Stock Index," *Journal of Econometrics*, vol. 52, 201–224.

Christensen, B. J. and N. R. Prabhala (1998), "The Relation between Implied and Realized Volatility," *Journal of Financial Economics*, vol. 50, 125–150.

Christie, Andrew A. (1982), "The Stochastic Behavior of Common Stock Variances—Value, Leverage and Interest Rate Effects," *Journal of Financial Economics*, vol. 10, 407–432.

Day, Theodore E. and Craig M.Lewis (1992), "Stock Market volatility and the Information Content of Stock Index Options," *Journal of Econometrics*, vol. 52, 267–287.

Duffee, Gregory R. (1995), "Stock Returns and Volatility: A Firm-Level Analysis," *Journal of Financial Economics*, vol. 37, 399–420.

Engle, Robert F., David M. Lilien, and Russell P. Robins (1987), "Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model," *Econometrica*, vol. 55, 391–407.

Engle, Robert F. and Victor K. Ng (1993), "Measuring and Testing the Impact of News on Volatility," *Journal of Finance*, vol. 48, 1749–1778.

Eraker, Bjørn (2003), "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices," *Journal of Finance*, forthcoming.

Eraker, Bjørn, Michael Johannes, and Nicholas Polson (2003), "The Impact of Jumps in Equity Index Volatility and Returns," *Journal of Finance*, forthcoming.

Fleming, Jeff (1998), "The Quality of Market Volatility Forecasts Implied by S&P 100 Index Option Prices," *Journal of Empirical Finance*, vol. 5, 317–345.

Fleming, Jeff, Barbara Ostdiek, and Robert E. Whaley (1995), "Predicting Stock Market Volatility: A New Measure," *Journal of Futures Markets*, vol. 15, 265–302.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh (1987), "Expected Stock Returns and Volatility," *Journal of Financial Economics*, vol. 19, 3–30.

Gallant, A. Ronald, Peter E. Rossi, and George Tauchen (1992), "Stock Prices and Volume," *Review of Financial Studies*, vol. 5, 199–242.

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov (2002), "There is a Risk-Return Tradeoff After All!" *Working Paper*, UNC and UCLA.

Glosten, Larry R., Ravi Jagannathan, and David Runkle (1993), "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, vol. 48, 1779–1801.

Harrison, Paul and Harold H. Zhang (1999), "An Investigation of the Risk and Return Relation at Long Horizons," *Review of Economics and Statistics*, vol. 81, 399–408.

Harvey, Campbell R. (2001), "The Specification of Conditional Expectations," *Journal of Empirical Finance*, vol. 8, 573–637.

Heston, Steven (1993), "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, vol. 6, 327–343.

Koopman, S. J. and E. H. Uspensky (1999), "The Stochastic Volatility in Mean Model: Empirical Evidence from International Stock Markets," *Working Paper*, Free University of Amsterdam.

Lamoureux, Christopher G. and William D. Lastrapes (1993), "Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities," *Review of Financial Studies*, vol. 6, 293–326.

Ledoit, Olivier, Pedro Santa-Clara, , and Shu Yan (2002), "Relative Pricing of Options with Stochastic Volatility," *Working Paper*, anderson Graduate School of Management, UCLA.

Martens, Martin and Jason Zein (2002), "Predicting Financial Volatility: High-Frequency Time-Series Forecasts vis-a-vis Implied Volatility," *Working Paper*, Econometric Institute, Erasmus University Rotterdam, The Netherlands.

Merton, Robert C. (1980), "On Estimating the Expected Return on the Market," *Journal of Financial Economics*, vol. 8, 323–361.

Neely, Christopher J. (2003), "Forecasting Foreign Exchange Volatility: Is Implied Volatility the Best We Can Do?" *Working Paper*, Federal Reserve Bank of St. Louis.

Nelson, Daniel B. (1991), "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, vol. 59, 347–370.

Pan, Jun (2002), "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study," *Journal of Financial Economics*, vol. 63, 3–50.

Pindyck, Robert (1984), "Risk, Inflation and the Stock Market," *American Economic Review*, vol. 74, 335–351.

Poon, Ser-Huang and Clive Granger (2002), "Forecasting Volatility in Financial Markets: A Review," *Journal of Economic Literature*, forthcoming.

Schwert, G. William (1989), "Stock Volatility and the Crash of 1987," *Review of Financial Studies*, pages 77–102.

Scruggs, John T. (1998), "Resolving the Puzzling Intertemporal Relation between the Market Risk Premium and Conditional Market Variance: A Two-Factor Approach," *Journal of Finance*, vol. 52, 575–603.

Tauchen, George and Robert Hussey (1991), "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, vol. 59, 371–396.

Tauchen, George, Harold Zhang, and Ming Liu (1996), "Volume, Volatility, and Leverage: A Dynamic Analysis," *Journal of Econometrics*, vol. 74, 177–208.

Turner, Christopher M., Richard Startz, and Charles R. Nelson (1989), "A Markov Model of Heteroskedasticity, Risk, and Learning in the Stock Market," *Journal of Financial Economics*, vol. 25, 3–22.

Wu, Guojun (2001), "The Determinants of Asymmetric Volatility," *Review of Financial Studies*, vol. 14, 837–859.

Wu, Guojun and Zhijie Xiao (2002), "A Generallized Partially Linear Model of Asymmetric Volatility," *Journal of Empirical Finance*, vol. 9, 287–319.

A Proof of Proposition 1

To simplify the exposition, define $a_\Delta = (1 - e^{-\kappa\Delta})/\kappa$, $a^*_\Delta = (1 - e^{-\kappa^*\Delta})/\kappa^*$, $b_\Delta = \theta(\Delta - a_\Delta)$, $b^*_\Delta = \theta^*(\Delta - a^*_\Delta)$, and $c_\Delta = (e^{-\kappa\Delta} + \kappa\Delta - 1)/\kappa$. The proof consists of three steps.

To determine the projection coefficients in the realized volatility-return trade-off relationship, note that from Andersen et al. (2002b), the variance term may be written as,

$$VAR\left(\int_t^{t+\Delta} V_u du\right) = \frac{\theta\sigma^2}{\kappa^3}\left(e^{-\kappa\Delta} + \kappa\Delta - 1\right) = \frac{\theta\sigma^2}{\kappa^2}c_\Delta.$$

Similarly, the covariance term takes the form,

$$\begin{aligned}
& COV\left(R_{t,t+\Delta}, \textstyle\int_t^{t+\Delta} V_u du\right) \\
= \; & COV\left(\mu\Delta + \lambda_s \textstyle\int_t^{t+\Delta} V_u du + \int_t^{t+\Delta}\sqrt{V_u}dB_u, \int_t^{t+\Delta} V_u du\right) \\
= \; & \lambda_s VAR\left(\textstyle\int_t^{t+\Delta} V_u du\right) + COV\left(\int_t^{t+\Delta}\sqrt{V_u}dB_u, \int_t^{t+\Delta} V_u du\right).
\end{aligned}$$

Rearranging and integrating by parts, the second term becomes,

$$\begin{aligned}
& COV\left(\textstyle\int_t^{t+\Delta}\sqrt{V_u}dB_u, \int_t^{t+\Delta} V_u du\right) \\
= \; & E\left(\textstyle\int_t^{t+\Delta}\sqrt{V_u}dB_u \int_t^{t+\Delta} V_u du\right) \\
= \; & E\left[\textstyle\int_t^{t+\Delta}\left(\int_t^u \sqrt{V_s}dB_s\right)V_u du + \int_t^{t+\Delta}\left(\int_t^u V_s ds\right)\sqrt{V_u}dB_u\right] \\
= \; & E\left[\textstyle\int_t^{t+\Delta}\left(\int_t^u \sqrt{V_s}dB_s\right)V_u du\right] \\
= \; & E\left[\textstyle\int_t^{t+\Delta}\left(\int_t^u \sqrt{V_s}dB_s\right)\left(V_t + \int_t^u \kappa(\theta - V_s)ds + \int_t^u \sigma\sqrt{V_s}dW_s\right)du\right] \\
= \; & E\left[\textstyle\int_t^{t+\Delta} -\kappa\left(\int_t^u \sqrt{V_s}dB_s \int_t^u V_s ds\right)\right] + E\left[\int_t^{t+\Delta}\left(\int_t^u \sqrt{V_s}dB_s \int_t^u \sigma\sqrt{V_s}dW_s\right)du\right] \\
= \; & E\textstyle\int_t^{t+\Delta}\left[-\kappa\int_t^u\left(\int_t^s \sqrt{V_\tau}dB_\tau\right)V_s ds\right]du + E\left(\int_t^{t+\Delta}\int_t^u \rho\sigma V_s ds du\right).
\end{aligned}$$

Notice that the recursive structure within the Riemann integral $\int_t^{t+\Delta}(\cdot)du$ must hold over any time interval Δ, so that in particular,

$$\begin{aligned}
\int_t^{t+\Delta} E\left(\int_t^u \sqrt{V_s}dB_s V_u\right)du \; = \; & \int_t^{t+\Delta} -\kappa\int_t^u E\left(\int_t^s \sqrt{V_\tau}dB_\tau V_s\right)ds du \\
& + \int_t^{t+\Delta}\left(\int_t^u \rho\sigma\theta ds\right)du,
\end{aligned}$$

which gives rise to the linear first order ordinary differential equation,

$$\frac{dE\left(\int_t^u \sqrt{V_s}dB_s V_u\right)}{du} = -\kappa E\left(\int_t^u \sqrt{V_s}dB_s V_u\right) + \rho\sigma\theta. \tag{A1}$$

Solving this equation yields,

$$E\left(\int_t^u \sqrt{V_s}dB_s V_u\right) = e^{-\kappa(u-t)}E\left(\int_t^t \sqrt{V_s}dB_s V_t\right) + \frac{\rho\sigma\theta}{\kappa}\left(1 - e^{-\kappa(u-t)}\right). \tag{A2}$$

Realizing that the first term on the right-hand-side equals zero, and completing the outside integration operator now yields,

$$\int_t^{t+\Delta}\left(\int_t^u \sqrt{V_s}dB_sV_u\right)du = \frac{\rho\sigma\theta}{\kappa^2}\left(e^{-\kappa\Delta}+\kappa\Delta-1\right) = \frac{\rho\sigma\theta}{\kappa}c_\Delta. \tag{A3}$$

Therefore,

$$\beta = \frac{COV\left(R_{t,t+\Delta}, \int_t^{t+\Delta}V_udu\right)}{VAR\left(\int_t^{t+\Delta}V_udu\right)} = \lambda_s + \frac{\rho\kappa}{\sigma}, \tag{A4}$$

which is less than zero provided that $0 < \lambda_s < -\frac{\rho\kappa}{\sigma}$ and great than zero if $0 < -\frac{\rho\kappa}{\sigma} < \lambda_s$. In addition, the intercept may be written as,

$$\alpha = E\left(R_{t,t+\Delta}\right) - \beta E\left(\int_t^{t+\Delta}V_udu\right) = \mu\Delta + \lambda_s\theta\Delta - \beta\theta\Delta = \mu\Delta - \frac{\rho\kappa\theta}{\sigma}\Delta, \tag{A5}$$

which is larger than $\mu\Delta$ under the usual parameter restrictions stipulated in the Proposition.

The determination of the projection coefficients in the implied volatility-return regression proceed by analogous arguments. First, the variance term may be written as,

$$VAR\left[E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right] = VAR\left(a_\Delta^*V_t + b_\Delta^*\right) = \frac{\sigma^2\theta}{2\kappa}a_\Delta^*a_\Delta^*.$$

Similarly, for the covariance term

$$\begin{aligned} & COV\left[R_{t,t+\Delta}, E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right] \\ = \; & COV\left[\mu\Delta + \lambda_s\int_t^{t+\Delta}V_udu + \int_t^{t+\Delta}\sqrt{V_u}dB_u, a_\Delta^*V_t + b_\Delta^*\right] \\ = \; & \lambda_s COV\left[\int_t^{t+\Delta}V_udu, a_\Delta^*V_t\right] + a_\Delta^*COV\left(\int_t^{t+\Delta}\sqrt{V_u}dB_u, V_t\right), \end{aligned}$$

where

$$\begin{aligned} & COV\left[\int_t^{t+\Delta}V_udu, E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right] \\ = \; & E\left[\left(\int_t^{t+\Delta}V_udu\right)E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right] - E\left(\int_t^{t+\Delta}V_udu\right)E\left[E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right] \\ = \; & E\left[E_t\left(\int_t^{t+\Delta}V_udu\right)E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right] - E\left[E_t\left(\int_t^{t+\Delta}V_udu\right)\right]E\left[E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right] \\ = \; & COV\left[E_t\left(\int_t^{t+\Delta}V_udu\right), E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right] \\ = \; & \tfrac{\sigma^2\theta}{2\kappa}a_\Delta a_\Delta^*, \end{aligned}$$

while the second term equals zero,

$$COV\left(\int_t^{t+\Delta}\sqrt{V_u}dB_u, V_t\right) = E\left[\left(\int_t^{t+\Delta}\sqrt{V_u}dB_u\right)V_t\right] = E\left[E_t\left(\int_t^{t+\Delta}\sqrt{V_u}dB_u\right)V_t\right] = 0.$$

Hence,

$$\beta^* = \frac{COV\left[\int_t^{t+\Delta}V_udu, E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right]}{VAR\left[E_t^*\left(\int_t^{t+\Delta}V_udu\right)\right]} = \lambda_s\frac{\frac{\sigma^2\theta}{2\kappa}a_\Delta a_\Delta^*}{\frac{\sigma^2\theta}{2\kappa}a_\Delta^*a_\Delta^*} = \lambda_s\frac{a_\Delta}{a_\Delta^*}, \tag{A6}$$

which is less than λ_s but larger than zero, provided that $\lambda_v < 0$ and $\lambda_s > 0$. Also,

$$\alpha^* = E\left(R_{t,t+\Delta}\right) - \beta^* E\left[E_t^*\left(\int_t^{t+\Delta} V_u du\right)\right] = \mu\Delta + \lambda_s\theta\Delta - \lambda_s\frac{a_\Delta}{a_\Delta^*}\theta\Delta, \tag{A7}$$

which is larger than $\mu\Delta$ given $\lambda_v < 0$ and $\lambda_s > 0$.

Finally, combining the results above, it follows readily that if $0 < \lambda_s < -\frac{\rho\kappa}{\sigma}$ we have

$$\beta < 0 < \beta^*, \tag{A8}$$

while for $0 < -\frac{\rho\kappa}{\sigma} < \lambda_s < \frac{a_\Delta^*}{a_\Delta - a_\Delta^*}\frac{\rho\kappa}{\sigma}$ we have

$$0 < \beta < \beta^*. \tag{A9}$$

B Proof of Proposition 2

By definition

$$\delta = \frac{COV\left(\int_t^{t+\Delta} V_u du, R_{t-\Delta,t}\right)}{VAR\left(R_{t-\Delta,t}\right)}.$$

The denominator may be rewritten as,

$$\begin{aligned}
& VAR\left(R_{t-\Delta,t}\right) \\
= \; & \lambda_s^2 VAR\left(\textstyle\int_{t-\Delta}^t V_u du\right) + E\left(\textstyle\int_{t-\Delta}^t V_u du\right) + \lambda_s COV\left(\textstyle\int_{t-\Delta}^t V_u du, \int_{t-\Delta}^t \sqrt{V_u} dB_u\right) \\
= \; & \lambda_s^2 \tfrac{\theta\sigma^2}{\kappa^3}\left(e^{-\kappa\Delta} + \kappa\Delta - 1\right) + \theta\Delta + \lambda_s\tfrac{\rho\sigma\theta}{\kappa^2}\left(e^{-\kappa\Delta} + \kappa\Delta - 1\right) \\
= \; & \theta\Delta + \tfrac{\lambda_s\sigma\theta}{\kappa}\left(\tfrac{\lambda_s\sigma}{\kappa} + \rho\right) c_\Delta,
\end{aligned}$$

where the second equality follows from the results in Andersen et al. (2002b), Bollerslev and Zhou (2002), and Proposition 1 above. The numerator may be expressed as,

$$\begin{aligned}
& COV\left(\textstyle\int_t^{t+\Delta} V_u du, R_{t-\Delta,t}\right) \\
= \; & \lambda_s COV\left(\textstyle\int_t^{t+\Delta} V_u du, \int_{t-\Delta}^t V_u du\right) + COV\left(\textstyle\int_t^{t+\Delta} V_u du, \int_{t-\Delta}^t \sqrt{V_u} dB_u\right) \\
= \; & \lambda_s \tfrac{\theta\sigma^2}{2\kappa^3}\left(1 - e^{-\kappa\Delta}\right)^2 + \tfrac{\rho\sigma\theta}{\kappa^2}\left(1 - e^{-\kappa\Delta}\right)^2 \\
= \; & \lambda_s \tfrac{\theta\sigma^2}{2\kappa} a_\Delta^2 + \rho\sigma\theta a_\Delta^2,
\end{aligned}$$

where the first term uses the result from Andersen et al. (2002b), and the second term uses the result from the proof of Proposition 1,

$$\begin{aligned}
COV\left(\textstyle\int_t^{t+\Delta} V_u du, \int_{t-\Delta}^t \sqrt{V_u} dB_u\right) &= E\left(\textstyle\int_t^{t+\Delta} V_u du \int_{t-\Delta}^t \sqrt{V_u} dB_u\right) \\
&= E\left[E_t\left(\textstyle\int_t^{t+\Delta} V_u du\right)\int_{t-\Delta}^t \sqrt{V_u} dB_u\right] \\
&= E\left(a_\Delta V_t \textstyle\int_{t-\Delta}^t \sqrt{V_u} dB_u\right) \\
&= \rho\sigma\theta a_\Delta^2.
\end{aligned}$$

Combining the two equations it follows therefore that

$$\delta = \frac{\lambda_s \frac{\theta\sigma^2}{2\kappa} a_\Delta^2 + \rho\sigma\theta a_\Delta^2}{\theta\Delta + \frac{\lambda_s \sigma\theta}{\kappa}\left(\frac{\lambda_s\sigma}{\kappa} + \rho\right) c_\Delta}. \tag{A10}$$

The intercept in the realized volatility asymmetry regression can be easily shown as

$$\gamma = \theta\Delta - \delta(\mu\Delta + \lambda_s\theta\Delta). \tag{A11}$$

The coefficient of the implied volatility asymmetry is similarly defined by,

$$\delta^* = \frac{COV\left[E_t^*\left(\int_t^{t+\Delta} V_u du\right), R_{t-\Delta,t}\right]}{VAR\left(R_{t-\Delta,t}\right)}.$$

The numerator may be rewritten as,

$$\begin{aligned}
& COV\left[E_t^*\left(\int_t^{t+\Delta} V_u du\right), R_{t-\Delta,t}\right] \\
= \; & COV\left[a_\Delta^* V_t + b_\Delta^*, R_{t-\Delta,t}\right] \\
= \; & a_\Delta^* \lambda_s COV\left(V_t, \int_{t-\Delta}^t V_u du\right) + a_\Delta^* COV\left(V_t, \int_{t-\Delta}^t \sqrt{V_u} dB_u\right) \\
= \; & -a_\Delta^* \lambda_s 0 + a_\Delta^* \frac{\rho\sigma\theta}{\kappa}\left(1 - e^{-\kappa\Delta}\right) \\
= \; & \rho\sigma\theta a_\Delta^* a_\Delta,
\end{aligned}$$

where $a_\Delta^* = \frac{1}{\kappa^*}\left(1 - e^{-\kappa^*\Delta}\right)$ and $b_\Delta^* = \theta^*(\Delta - a_\Delta^*)$, and the last line of the proof utilizes the results from the proof of Proposition 1 that

$$\begin{aligned}
& COV\left(V_t, \int_{t-\Delta}^t V_u du\right) \\
= \; & E\left(V_t \int_{t-\Delta}^t V_u du\right) - \theta^2\Delta \\
= \; & E\left(V_{t-\Delta} \int_{t-\Delta}^t V_u du\right) + E\left(\kappa\theta \int_{t-\Delta}^t du \int_{t-\Delta}^t V_u du\right) \\
& + E\left(-\kappa \int_{t-\Delta}^t V_u du \int_{t-\Delta}^t V_u du\right) + E\left(\sigma \int_{t-\Delta}^t \sqrt{V_u} dW_u \int_{t-\Delta}^t V_u du\right) - \theta^2\Delta \\
= \; & E\left[V_{t-\Delta}(a_\Delta V_{t-\Delta} + b_\Delta)\right] + \kappa\theta^2\Delta^2 - \kappa\left[VAR\left(\int_{t-\Delta}^t V_u du\right) + \theta^2\Delta^2\right] \\
& + \frac{\sigma^2\theta}{\kappa^2}\left(e^{-\kappa\Delta} + \kappa\Delta - 1\right) - \theta^2\Delta \\
= \; & \theta^2\Delta - \frac{\theta\sigma^2}{\kappa^2}\left(e^{-\kappa\Delta} + \kappa\Delta - 1\right) + \frac{\sigma^2\theta}{\kappa^2}\left(e^{-\kappa\Delta} + \kappa\Delta - 1\right) - \theta^2\Delta \\
= \; & 0,
\end{aligned}$$

and,

$$\begin{aligned}
& COV\left(V_t, \int_{t-\Delta}^t \sqrt{V_u} dB_u\right) \\
= \; & E\left(V_t \int_{t-\Delta}^t \sqrt{V_u} dB_u\right) \\
= \; & \frac{\rho\sigma\theta}{\kappa}\left(1 - e^{-\kappa\Delta}\right) \\
= \; & \rho\sigma\theta a_\Delta.
\end{aligned}$$

Now combing the different equations it follows that

$$\delta^* = \frac{\rho\sigma\theta a^*_\Delta a_\Delta}{\theta\Delta + \frac{\lambda_s\sigma\theta}{\kappa}\left(\frac{\lambda_s\sigma}{\kappa} + \rho\right) c_\Delta}. \tag{A12}$$

The intercept in the implied volatility asymmetry regression can be easily shown as

$$\gamma^* = (a^*_\Delta\theta + b^*_\Delta) - \delta^*(\mu\Delta + \lambda_s\theta\Delta). \tag{A13}$$

Lastly, note that the common denominator of δ and δ^* (corresponding to a variance) is always positive. The requirement that $\delta^* < \delta$, or

$$\rho\sigma\theta a^*_\Delta a_\Delta < \lambda_s \frac{\theta\sigma^2}{2\kappa} a^2_\Delta + \rho\sigma\theta a^2_\Delta,$$

is equivalent to

$$\lambda_s \frac{\theta\sigma^2}{2\kappa} a^2_\Delta > 0 > \rho\sigma\theta a_\Delta (a^*_\Delta - a_\Delta).$$

The assumption that $\lambda_v < 0$ ensures that $a^*_\Delta - a_\Delta > 0$. Thus, given the usual parameter restrictions in the Proposition, the assumptions that $\lambda_s > 0$ and $\rho < 0$ are sufficient to guarantee the ordering of the δ's.

C Proof of Proposition 3

From the proof of Proposition 1,

$$\phi_1 = \frac{COV\left[\int_t^{t+\Delta} V_u du, E^*_t\left(\int_t^{t+\Delta} V_u du\right)\right]}{VAR\left[E^*_t\left(\int_t^{t+\Delta} V_u du\right)\right]} = \frac{\frac{\sigma^2\theta}{2\kappa} a_\Delta a^*_\Delta}{\frac{\sigma^2\theta}{2\kappa} a^*_\Delta a^*_\Delta} = \frac{a_\Delta}{a^*_\Delta} = \frac{\left(1 - e^{-\kappa\Delta}\right)\kappa^*}{\left(1 - e^{-\kappa^*\Delta}\right)\kappa} < 1,$$

where the last inequality follows directly by the assumption that $\kappa^* = \kappa + \lambda_v < \kappa$. Similarly, the intercept may be evaluated as,

$$\phi_0 = E\left(\int_t^{t+\Delta} V_u du\right) - \phi_1 E\left[E^*_t\left(\int_t^{t+\Delta} V_u du\right)\right] = \theta\Delta - \frac{a_\Delta}{a^*_\Delta}\left(a^*_\Delta\theta + b^*_\Delta\right) = b_\Delta - \frac{a_\Delta}{a^*_\Delta} b^*_\Delta,$$

which can generally not be signed.

D Tables and Figures

Table 1: Summary Statistics for Monthly Returns and Volatilities

	Return	Standard Deviation		Variance	
		Implied	Realized	Implied	Realized
Mean	0.857	6.074	3.363	41.190	13.558
Std. Dev.	4.595	2.078	1.504	33.673	13.907
Skewness	-1.213	1.634	1.476	3.967	2.916
Kurtosis	7.611	8.602	5.637	27.516	14.883
Minimum	-24.543	2.835	1.366	8.036	1.866
5% Qntl.	-6.634	3.465	1.733	12.003	3.004
25% Qntl.	-1.914	4.615	2.261	21.300	5.113
50% Qntl.	1.232	5.828	2.881	33.970	8.301
75% Qntl.	3.863	7.137	4.093	50.934	16.751
95% Qntl.	7.074	9.724	6.287	94.565	39.523
Maximum	12.378	17.728	10.284	314.266	105.767
ρ_1	-0.011	0.773	0.738	0.642	0.616
ρ_2	-0.077	0.627	0.588	0.438	0.413
ρ_3	-0.034	0.523	0.519	0.321	0.343
ρ_4	-0.117	0.485	0.458	0.275	0.273
ρ_5	0.065	0.483	0.487	0.272	0.319
ρ_6	0.012	0.485	0.512	0.289	0.373
ρ_7	0.066	0.460	0.472	0.251	0.307
ρ_8	-0.032	0.447	0.462	0.244	0.300
ρ_9	-0.028	0.441	0.457	0.243	0.293
ρ_{10}	0.122	0.422	0.441	0.233	0.291

Table 2: Volatility Feedback Effect

Standard Deviation			Variance		
Implied	Realized	Expected	Implied	Realized	Expected
$\alpha^* = 0.637$	$\alpha = 4.156$	$\tilde{\alpha} = 1.717$	$\alpha^* = 0.828$	$\alpha = 2.486$	$\tilde{\alpha} = 1.231$
(0.932)	(1.171)	(1.838)	(0.489)	(0.552)	(12.91)
$\beta^* = 0.036$	β= -0.981	$\tilde{\beta}$= -0.264	$\beta^* = 0.001$	β = -0.120	$\tilde{\beta}$ = -0.029
(0.170)	(0.390)	(0.594)	(0.013)	(0.047)	(0.965)
$R^2 = 0.000$	$R^2 = 0.103$	$R^2 = 0.047$	R^2 =0.000	$R^2 = 0.132$	$R^2 = 0.057$

Note: The "Expected" volatility regressions refer to the instrumental variables regressions using the absolute lagged returns as instruments for the realized volatilities.

Table 3: Leverage Effect

Short Regression			
Standard Deviation		Variance	
Implied	Realized	Implied	Realized
$\gamma^* = 6.225$	$\gamma = 3.438$	$\gamma^* = 44.087$	$\gamma = 14.317$
(0.238)	(0.168)	(3.891)	(1.471)
δ^* = -0.170	δ = -0.081	δ^* = -3.305	δ = -0.824
(0.059)	(0.028)	(1.384)	(0.283)
$R^2 = 0.141$	$R^2 = 0.062$	$R^2 = 0.204$	$R^2 = 0.074$

Long Regression			
Standard Deviation		Variance	
Implied	Realized	Implied	Realized
$\gamma^* = 1.218$	$\gamma = 0.900$	$\gamma^* = 12.110$	$\gamma = 4.967$
(0.279)	(0.195)	(2.494)	(1.183)
$\beta^* = 0.689$	$\beta = 0.723$	$\beta^* = 0.580$	$\beta = 0.654$
(0.059)	(0.054)	(0.067)	(0.068)
$\alpha^* = 0.072$	$\alpha = 0.007$	$\alpha^* = 0.032$	α = -0.003
(0.037)	(0.036)	(0.052)	(0.044)
δ^* = -0.309	δ = -0.009	δ^* = -0.444	$\delta = 0.015$
(0.033)	(0.050)	(0.062)	(0.058)
$R^2 = 0.844$	$R^2 = 0.547$	$R^2 = 0.884$	$R^2 = 0.383$

Table 4: Implied Volatility Forecasting Bias

Standard Deviation	Variance
$\phi_0 = 0.231$	$\phi_0 = 3.705$
(0.392)	(2.643)
$\phi_1 = 0.516$	$\phi_1 = 0.239$
(0.069)	(0.066)
$R^2 = 0.507$	$R^2 = 0.336$

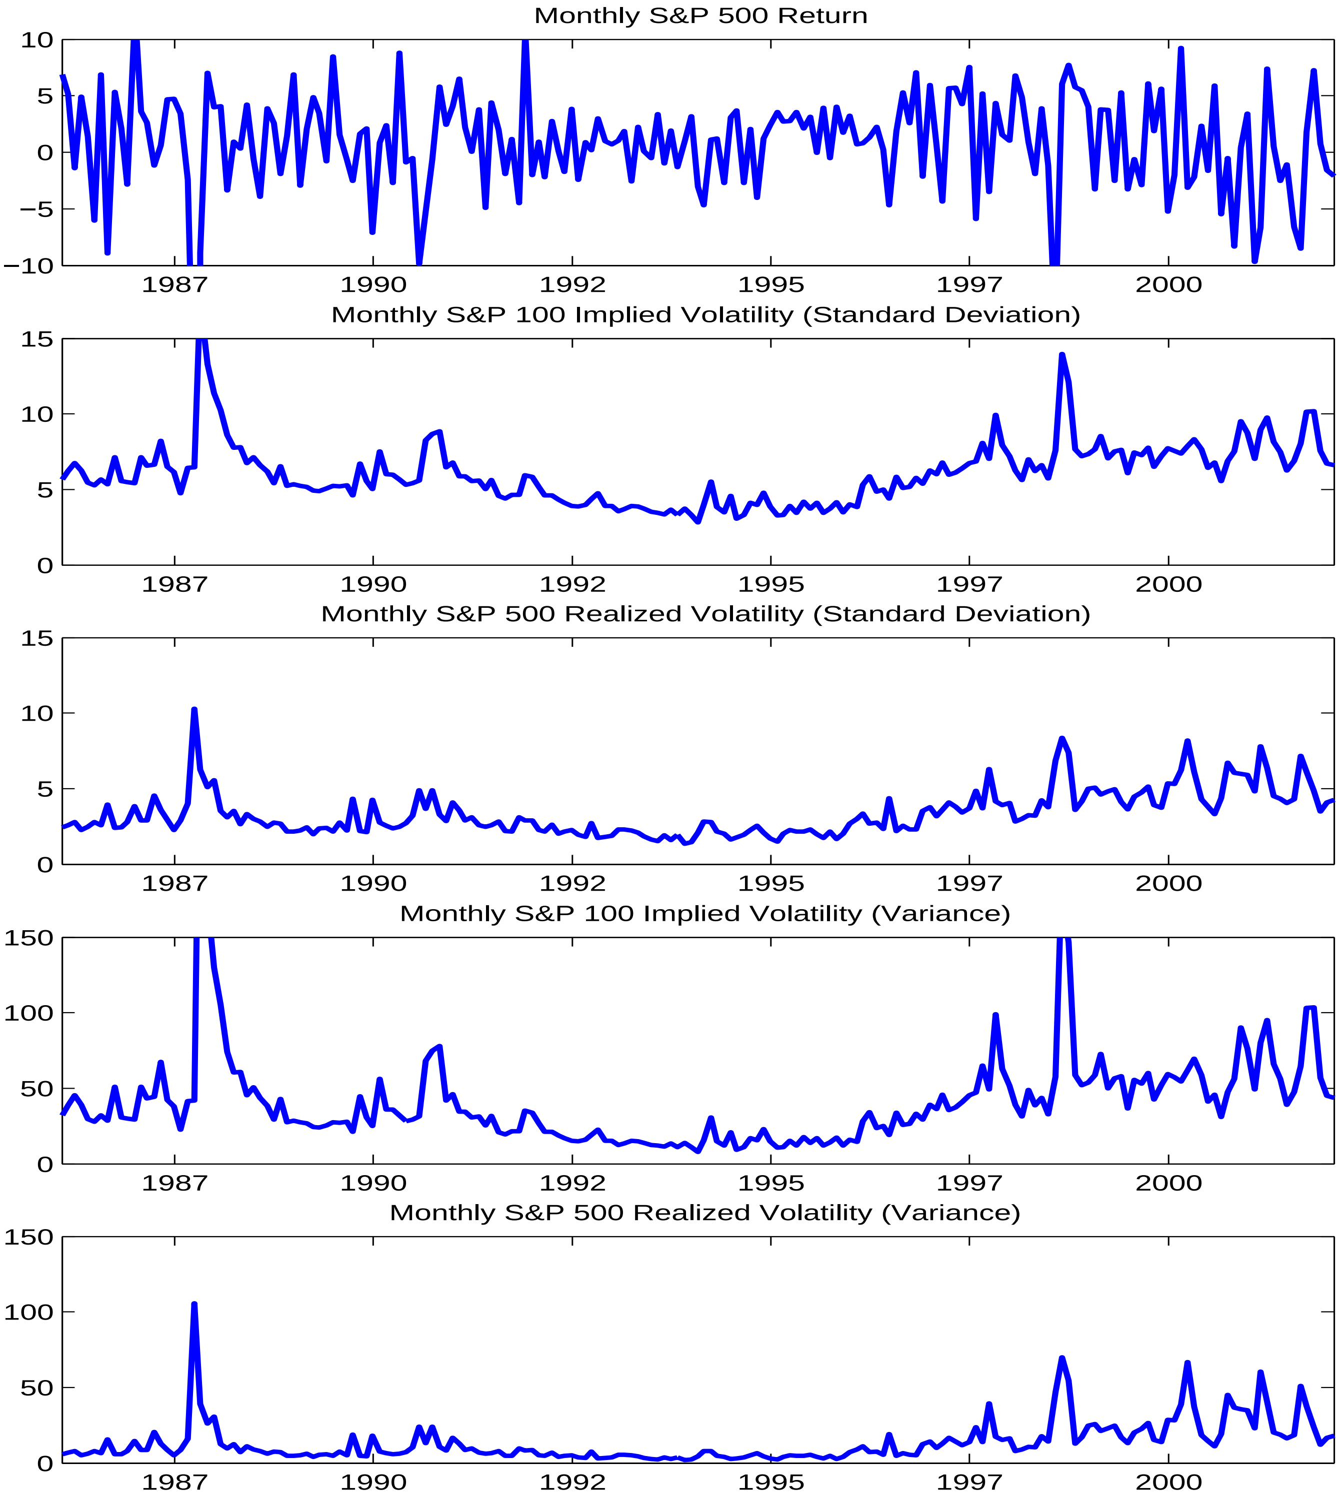

Figure 1: Time Series Plot of Returns and Volatilities.