ELSEVIER

# Volatility puzzles: a simple framework for gauging return-volatility regressions

## Tim Bollerslev[a], Hao Zhou[b],*

[a]*Department of Economics, Duke University, Post Office Box 90097, Durham, NC 27708, USA*
[b]*Division of Research and Statistics, Federal Reserve Board, Mail Stop 91, Washington, DC 20551, USA*

## Abstract

This paper provides a simple theoretical 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 asymmetric response of current volatility to lagged negative and positive returns, typically referred to as 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 S&P500 aggregate market index.
© 2005 Elsevier B.V. All rights reserved.

---

*Corresponding author. Tel.: 202 452 3360; fax: 202 728 5887.
  *E-mail addresses:* boller@econ.duke.edu (T. Bollerslev), hao.zhou@frb.gov (H. Zhou).

## 1. Introduction

Following the realization in the late 1980s 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., 2004, 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.

Building on the popular Heston (1993) one-factor stochastic volatility model, the present paper provides a simple theoretical framework for reconciling these conflicting empirical findings. Specifically, by postulating a 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; Bali and Peng, 2003; Guo and Whitelaw, 2003; Ghysels et al., 2004) 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; Lettau and Ludvigson, 2004) 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, 2002), 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, 2004). 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 so-called 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, 1990; 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, 2004).

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 Lewis, 1992; Fleming et al., 1995; Fleming, 1998; Lamoureux and Lastrapes, 1993, along with the recent survey in Poon and Granger, 2003). As formally shown in the recent studies by Chernov (2002), Pan (2002), and Bates (2003), 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), which explicitly assumes that the stochastic volatility premium is linear. This in turn allows us to utilize various closed form expressions for the conditional moments previously derived by Andersen et al. (2004) and Bollerslev and Zhou (2002). Although the exact relationship and

---

[1]More general multi-factor models further complicate the risk-return tradeoff relationship, as the projection of the returns on the volatility must control for the influence of other state variables. Hence, in this situation the equilibrium correlation between risk and return may be non-linear and even negative (see e.g., Abel, 1988; Tauchen and Hussey, 1991; Backus and Gregory, 1993; Scruggs, 1998; Lettau and Ludvigson, 2004; Christoffersen and Diebold, 2003; Guo and Whitelaw, 2003). Tauchen (2004), within a general equilibrium context, has also recently called into question the interpretation of the standard ARCH-M specification as a model for the volatility risk premium, arguing instead that the effect represents a time-varying risk premium on consumption (endowment) risk.

[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 a leverage effect.

implication derived in this paper may not hold for other more complicated model structures, the basic idea could in principle be generalized to cases of multiple volatility factors and jumps and/or non-linear volatility premia, albeit at the expense of considerable notational and computational complexity (see e.g., Andersen et al., 2002; Eraker et al., 2003; Chernov et al., 2003). Interestingly however, the theoretical results for the relatively simple one-factor affine Heston model turn out to be rich enough to explain the apparent conflicting empirical findings in regards to the monthly return-volatility regressions for the S&P500 aggregate market index and corresponding realized and implied volatilities.

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 thirteen-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,

$$\mathrm{d}p_t = (\mu + \lambda_s V_t)\,\mathrm{d}t + \sqrt{V_t}\,\mathrm{d}B_t,$$
$$\mathrm{d}V_t = \kappa(\theta - V_t)\,\mathrm{d}t + \sigma\sqrt{V_t}\,\mathrm{d}W_t,$$
$$\mathrm{corr}(\mathrm{d}B_t, \mathrm{d}W_t) = \rho, \tag{1}$$

where the latent stochastic volatility, $V_t$, is assumed to follow a square-root process. Empirical model estimates generally point to a negative instantaneous correlation between the two separate Brownian motions driving the price and volatility processes, or $\rho < 0$. This feature is sometimes referred to as a (model-based) "continuous-time" leverage effect. Similarly, the underlying "continuous-time" volatility feedback effect is captured directly by the risk-return trade-off parameter, $\lambda_s > 0$.[3]

Given this dynamic specification for the underlying price-volatility process, standard pricing arguments imply the existence of the following equivalent

---

[3]It is important to stress the difference between the "continuous-time" volatility feedback effect, or risk-return trade-off, $\lambda_s > 0$, and the "empirical" volatility feedback effect, or risk-return trade-off, as defined by the slope coefficient in a regression of the discrete-time returns on an observable volatility proxy. Similarly, the notion of a "continuous-time" leverage asymmetry effect, or $\rho < 0$, formally differs from the "empirical" leverage asymmetry effect defined by the sample correlation between lagged discrete-time returns and a current volatility proxy. It is the implications of the continuous-time model in (1) for these latter "empirical" measures that is the focus of the present analysis.

Martingale measure, or "risk-neutralized" distribution,

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

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.[4] 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 Eq. (1) by the functional relationships, $\kappa^* = \kappa + \lambda_v$ and $\theta^* = \kappa\theta/(\kappa + \lambda_v)$. The $\lambda_v$ 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 $\kappa^*$, is therefore slower (possibly even explosive) than the mean reversion for the actual volatility, as determined by $\kappa$ (for a more detailed discussion of the connection between the objective and the risk-neutral distributions, see also Benzoni, 2001; Wu, 2001; Chernov, 2002; 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 Eq. (1). Thus, as discussed in the introduction, several studies have resorted to more complicated multi-factor representations as a way to resolve this apparent empirical puzzle (see, e.g., Scruggs, 1998; Guo and Whitelaw, 2003, and the discussion therein).

Meanwhile, consider the continuously compounded returns from time $t$ to $t + \Delta$ 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 $E(\int_t^{t+\Delta} \sqrt{V_u} \, dB_u \times \int_t^{t+\Delta} V_u \, du) = 0$ holds true. However, with a non-zero instantaneous "leverage" coefficient, or $\rho < 0$, the residual and the integrated volatility will be correlated, resulting in a biased estimate for $\lambda_s$.

---

[4]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.

Specifically, consider the population regression,

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

Then as formally shown below, unless $\rho = 0$, the population feedback coefficient $\beta$ will differ from the true feedback coefficient $\lambda_s$. Of course, the integrated volatility is not directly observable, so the sample counterpart to the population regression in (4) is not 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., 2004). 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,

$$R_{t,t+\Delta} = \alpha^* + \beta^* \mathrm{E}_t^* \left( \int_t^{t+\Delta} V_u \, \mathrm{d}u \right) + e_{t,t+\Delta}^*, \tag{5}$$

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 Eq. (3), that is $\beta^* \neq \lambda_s$. Hence, to correctly uncover the volatility feedback parameter from a contemporaneous return-volatility type regression, either the underlying leverage coefficient 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 these biases.[5]

**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 Eq.(4) is then given by*

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

*Let* $a_\Delta = (1 - \mathrm{e}^{-\kappa \Delta})/\kappa$ *and* $a_\Delta^* = (1 - \mathrm{e}^{-\kappa^* \Delta})/\kappa^*$. *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* $0 < \lambda_s < -\rho \kappa / \sigma$, *the two slope parameters are related by,*

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

*while for* $0 < -\rho \kappa/\sigma < \lambda_s < a_\Delta^*/(a_\Delta - a_\Delta^*)\rho \kappa/\sigma$, *we have*

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

---

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

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 $\beta$. Only if the underlying leverage coefficient is zero ($\rho = 0$) will the regression be unbiased for estimating $\lambda_s$. Second, regression estimates based on implied volatility will generally show less of a downward bias and remain positive under all reasonable parameter settings. 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 3.1 below.

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 $\beta$ 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, and results in unbiased coefficients in population; i.e., $\tilde{\beta} = \lambda_s$ and $\tilde{\alpha} = \mu\Delta$.[6] As such, this regression corresponds more closely to the *implied* return-volatility trade-off regression in (5) that obtains by replacing the expected integrated volatility, $E_t(\int_t^{t+\Delta} V_u \, du)$, with its risk neutral equivalent, $E_t^*(\int_t^{t+\Delta} V_u \, du)$. Of course, the expected integrated volatility will generally depend upon the underlying structural model, but it 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 choice of instruments employed in the estimation, indirectly highlighting the difficulties in accurately approximating $E_t(\int_t^{t+\Delta} V_u \, du)$ (see, e.g., Harvey, 2001; Brandt and Kang, 2004). 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.

---

[6]A formal proof for this result is given in Appendix A along with the proof of Proposition 1.

## 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,[7]

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

and the option implied volatility,

$$\mathrm{E}_t^* \left( \int_t^{t+\Delta} V_u \, \mathrm{d}u \right) = \gamma^* + \delta^* R_{t-\Delta,t} + e_{t-\Delta,t}^*, \tag{11}$$

where the expectation in Eq. (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 ''instantaneous'' leverage, or asymmetry, parameter $\rho$ 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* $a_\Delta = (1 - \mathrm{e}^{-\kappa\Delta})/\kappa$, $a_\Delta^* = (1 - \mathrm{e}^{-\kappa^*\Delta})/\kappa^*$, *and* $c_\Delta = (\mathrm{e}^{-\kappa\Delta} + \kappa\Delta - 1)/\kappa^2$. *The population slope parameters in* (10) *and* (11) *may then be expressed as,*

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

---

[7]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.

*and*

$$\delta^* = \frac{\lambda_s(\theta\sigma^2/2\kappa)a_\Delta^* a_\Delta + \rho\sigma\theta a_\Delta^* a_\Delta}{\theta\Delta + \lambda_s\sigma\theta\big((\lambda_s\sigma/\kappa) + 2\rho\big)c_\Delta}. \tag{13}$$

*Moreover, assuming* $0 < \lambda_s < -(2\rho\kappa/\sigma)$ *it follows that*

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

*while for* $0 < -2\rho\kappa/\sigma < \lambda_s$,

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

The proof of the proposition is given in Appendix B.

It is noteworthy that the "empirical" leverage coefficient, $\delta$ or $\delta^*$, depends critically on both the volatility feedback parameter $\lambda_s$ (positively), as well as the "instantaneous" leverage coefficient $\rho$ (negatively). Thus, while most empirical studies report a negative volatility asymmetry for the aggregate market portfolio ($\delta < 0$), this may help explain the lack of statistical significance and/or the sometimes small economic magnitude of the estimated effect. For the implied volatility regression, the "empirical" leverage coefficient $\delta^*$ also depends directly on the stochastic volatility risk premium $\lambda_v$ through $a_\Delta^*$ (magnitude). Moreover, provided that 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 $\delta^* < \delta < 0$.[8]

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, 2003). In population,

$$\int_t^{t+\Delta} V_u \, \mathrm{d}u = \phi_0 + \phi_1 \mathrm{E}_t^* \left( \int_t^{t+\Delta} V_u \, \mathrm{d}u \right) + e_{t,t+\Delta}. \tag{16}$$

Obviously, for the implied volatility to provide unbiased forecasts, the two projection coefficients should equal $\phi_0 = 0$ and $\phi_1 = 1$, respectively. Meanwhile,

---

[8] There is little empirical evidence in support of the reverse asymmetric effect, or $0 < \delta < \delta^*$, covered by the second case of Proposition 2.

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 a finite-sample equivalent to the population regression in (16) (Christensen and Prabhala, 1998), along with very persistent volatility dynamics rendering standard statistical inference unreliable (Bandi and Perron, 2004). These statistical considerations aside, if the stochastic volatility risk premium, $\lambda_v$, 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,*

$$\phi_0 = b_\Delta - \frac{a_\Delta}{a_\Delta^*} b_\Delta^* \quad and \quad \phi_1 = \frac{a_\Delta}{a_\Delta^*} < 1, \tag{17}$$

*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 $a_\Delta/a_\Delta^*$ becoming less than unity.[9] Thus, any estimate of $\phi_1$ should be gauged against this population bias. Of course, the true structural model parameters, $\kappa$ and $\kappa^*$, are generally unknown and would have to be estimated.[10]

The integrated variance measure entering the population regressions underlying the results in Propositions 1–3 is, of course, not directly observable. However, as noted above, the summation of sufficiently finely sampled high-frequency intraday squared returns may in theory be used in approximating $\int_t^{t+\Delta} V_u \, du$ to any desired degree of accuracy. Meanwhile, a host of market microstructure frictions invariably invalidate the underlying semi-martingale assumption for the returns at the ultra highest frequencies. Hence, following much of the recent literature, the results for the monthly return-volatility regressions reported in Section 3 below are based on so-called realized volatilities constructed from the summation of the five-minute squared returns within each month. The next sub-section demonstrates that this approximation does not materially affect any of the regression estimates.

---

[9]Closely related results, along with a more detailed analysis of the impact of jumps, have recently been derived in concurrent work by Chernov (2002) and Bates (2003). Alternatively, as shown by Bandi and Perron (2001, 2004) and Comte et al. (2003), the bias in the slope coefficient in the (approximate) relationship between the Black-Scholes implied volatility and the expected integrated volatility depends directly on the leverage coefficient, $\rho$.

[10]Conversely, the population return-volatility regressions in Propositions 1–3 (coupled with additional moment restrictions along the lines of Bollerslev and Zhou, 2002) could be employed as a system of equations in estimating the structural model parameters underlying the actual and risk-neutral dynamics in (1) and (2).

Table 1
Simulated benchmark model

| Parameter | Value |
| --- | --- |
| $\kappa$ | 0.20 |
| $\theta$ | 10.00 |
| $\sigma$ | 0.80 |
| $\lambda_v$ | −2.00 |
| $\rho$ | −0.60 |
| $\mu$ | 0.60 |
| $\lambda_s$ | 0.04 |

### 2.4. Realized volatility measurement error

To gauge the impact of the approximation error from using realized volatilities in place of the integrated volatilities in the return-volatility regressions, we report the results from a small scale Monte Carlo study. For comparison purposes with the empirical results reported in the next section, each "day" is divided into 78 "5-minute" intervals, corresponding to a six-and-a-half "hour" trading day. The parameters in the simulated benchmark model (listed in Table 1) implies an annualized return of 12%, an annualized realized volatility of 11%, along with a 19% annualized implied volatility. These particular parameter values were adapted from the estimates for the monthly S&P500 returns and implied volatilities reported in Bollerslev et al. (2005), and correspond fairly closely to other values reported in the literature.

The first row in each of the three panels in Table 2, labeled "Integrated", summarizes the mean and median biases along with the root mean-square-error for the finite sample distributions of the regression slope coefficients from the infeasible "monthly" volatility feedback, leverage effect, and implied volatility regressions in Eqs. (4), (10), and (16), respectively, when using the (latent) integrated volatility.[11] The second row in each panel reports the same statistics for the corresponding feasible regressions based on the realized volatilities constructed from the sum of the squared "5-minute" returns. To illustrate the impact of increasing the sample span, we report the results with both 150 and 600 "monthly" observations, corresponding to twelve-and-a-half and fifty "years" of data, respectively. Comparing the two sets of results, it is immediately evident that the measurement errors in the "five-minute" based realized volatilities have only negligible effects on the "monthly" return-volatility regressions.[12]

---

[11]The "monthly" integrated volatility is calculated by the summation of the simulated squared returns over "30-second" intervals. The summary statistics are based on a total of 500 replications.

[12]The asymptotic theory developed by Barndorff-Nielsen and Shephard (2002, 2004) and Andersen et al. (2005) provides a formal framework for incorporating the influence of the measurement errors more generally.

Table 2
Simulated monthly return-volatility regressions

|  | True value | Mean bias | | Median bias | | RMSE | |
|---|---|---|---|---|---|---|---|
|  |  | $T = 150$ | $T = 600$ | $T = 150$ | $T = 600$ | $T = 150$ | $T = 600$ |
| *Volatility feedback effect* |  |  |  |  |  |  |  |
| Integrated | $\beta = -0.1100$ | $-0.0038$ | $-0.0019$ | $-0.0050$ | $-0.0016$ | $0.0652$ | $0.0287$ |
| Realized, 5-min | $\beta = -0.1100$ | $-0.0015$ | $-0.0002$ | $-0.0030$ | $0.0000$ | $0.0649$ | $0.0289$ |
| Realized, 1-day | $\beta = -0.1100$ | $0.0809$ | $0.0778$ | $0.0822$ | $0.0767$ | $0.0997$ | $0.0824$ |
| *Leverage effect* |  |  |  |  |  |  |  |
| Integrated | $\delta = -0.3471$ | $0.0118$ | $0.0077$ | $0.0147$ | $0.0083$ | $0.0960$ | $0.0456$ |
| Realized, 5-min | $\delta = -0.3471$ | $0.0111$ | $0.0080$ | $0.0163$ | $0.0083$ | $0.0960$ | $0.0459$ |
| Realized, 1-day | $\delta = -0.3471$ | $0.0185$ | $0.0124$ | $0.0190$ | $0.0127$ | $0.1341$ | $0.0648$ |
| *Implied volatility forecasting bias* |  |  |  |  |  |  |  |
| Integrated | $\phi_1 = 0.3231$ | $-0.0056$ | $-0.0024$ | $-0.0047$ | $-0.0024$ | $0.0131$ | $0.0058$ |
| Realized, 5-min | $\phi_1 = 0.3231$ | $-0.0055$ | $-0.0023$ | $-0.0053$ | $-0.0022$ | $0.0134$ | $0.0060$ |
| Realized, 1-day | $\phi_1 = 0.3231$ | $-0.0034$ | $-0.0009$ | $-0.0049$ | $-0.0010$ | $0.0343$ | $0.0183$ |

In contrast, consider the results based on "monthly" realized volatilities constructed from the summation of the 22 "daily" squared returns within each "month", as utilized in some of the existing literature. The estimates for the volatility feedback coefficients, in particular, now exhibit large biases (toward zero for the specific parameter values in the simulated model). Interestingly, these biases are not alleviated much with the longer sample. The RMSE's are, of course, also much larger when only "daily" returns are used in approximating the "monthly" integrated volatility.[13]

## 3. Empirical illustration

Our empirical analysis is based on monthly returns and volatilities for the S&P500 composite index spanning the period from January 1990 through February 2002.[14] The monthly continuously compounded percentage returns are constructed from the daily S&P500 closing prices supplied by Standard & Poor. Normalizing the monthly time interval to unity, we will refer to the return over the $(t+1)$th month as $R_{t,t+1}$.

---

[13]Of course, if the same regressions were analyzed at the "daily" level with the realized volatilities constructed from the "five-minute" intraday returns, some of these same biases would likely hold true.

[14]The starting date of January 1990 reflects the availability of risk neutral implied volatilities for the S&P500 (new VIX) available from the Chicago Board of Options Exchange (CBOE) since September 22, 2003. Earlier versions of this paper relied on implied volatilities for the S&P100 (old VIX) and a starting date of January 1986. The results are qualitatively similar and available upon request.

As previously noted, 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 $t + 1$,

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

where the 78 five-minute subintervals represents the normal trading hours from 9:30 am to 4:00 pm, including the close-to-open five-minute interval. As discussed in the previous section, the realized volatility, $\text{RV}_{t,t+1}$, is readily interpreted as a consistent (for increasing sampling frequency), and in the present situation with roughly $22 \times 78 = 1,716$ five-minute returns per month, a highly accurate estimate of the corresponding integrated volatility, $\text{IV}_{t,t+1} \equiv \int_t^{t+1} V_s \, ds$.

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

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

where $\text{E}^*$ refers to the risk-adjusted expectation of the one-month ahead integrated volatility, $\text{IV}_{t,t+1}$. Our measure for $\text{IV}_{t,t+1}^*$ is based on the (new) VIX index for the S&P500 volatility provided by the CBOE. Importantly, these are model-free implied volatilities calculated on the basis of the approach in Britten-Jones and Neuberger (2000); see also Jiang and Tian (2004) for an alternative theoretical derivation of this model-free options implied volatility measure along with a discussion of issues related to its practical implementation.[15]

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 $\text{RV}_{t,t+1}$ and $\text{IV}_{t,t+1}^*$ 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 $\text{RV}_{t,t+1}^{1/2}$ and $\text{IV}_{t,t+1}^{*1/2}$.

Summary statistics for all of the variables are reported in Table 3. 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 fourteen percent. The returns are negatively skewed with fatter tails than the normal distribution. The realized volatilities are systematically lower than the implied volatilities, and their unconditional distributions also deviate less 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,

---

[15]In contrast, the old VIX index (now labeled VXO) provided by the CBOE until September 21, 2003, was based on index options for the S&P100, and relied on the Black-Scholes pricing formula for inverting the options prices.

Table 3
Summary statistics for monthly returns and volatilities

| Statistics | $R_{t,t+1}$ | $RV_{t,t+1}^{1/2}$ | $IV_{t,t+1}^{*1/2}$ | $RV_{t,t+1}$ | $IV_{t,t+1}^{*}$ |
|---|---|---|---|---|---|
| Mean | 0.8365 | 3.4289 | 5.6402 | 14.0982 | 35.0116 |
| Std. Dev. | 4.1449 | 1.5353 | 1.7949 | 13.3458 | 23.6188 |
| Skewness | −0.6214 | 1.0590 | 0.8805 | 1.9579 | 2.0850 |
| Kurtosis | 4.1696 | 3.6945 | 4.2386 | 7.0473 | 10.2722 |
| | | | | | |
| Minimum | −15.7586 | 1.3658 | 3.0686 | 1.8655 | 9.4164 |
| 5% Qntl. | −6.0935 | 1.6828 | 3.3602 | 2.8319 | 11.2909 |
| 25% Qntl. | −1.9522 | 2.2147 | 3.9881 | 4.9050 | 15.9046 |
| 50% Qntl. | 1.0429 | 3.0027 | 5.6205 | 9.0161 | 31.5901 |
| 75% Qntl. | 3.6992 | 4.2845 | 6.8423 | 18.3569 | 46.8177 |
| 95% Qntl. | 7.1228 | 6.4501 | 8.5772 | 41.6279 | 73.5703 |
| Maximum | 10.5790 | 8.3521 | 12.7825 | 69.7571 | 163.3932 |
| | | | | | |
| $\rho_1$ | −0.0622 | 0.8019 | 0.8197 | 0.7254 | 0.7470 |
| $\rho_2$ | −0.0273 | 0.6641 | 0.6767 | 0.5111 | 0.5280 |
| $\rho_3$ | −0.0208 | 0.5904 | 0.5866 | 0.4208 | 0.4231 |
| $\rho_4$ | −0.0544 | 0.5579 | 0.5639 | 0.3872 | 0.4036 |
| $\rho_5$ | 0.0296 | 0.5938 | 0.5770 | 0.4511 | 0.4311 |
| $\rho_6$ | 0.0029 | 0.6120 | 0.5547 | 0.4990 | 0.4051 |
| $\rho_7$ | 0.0899 | 0.5804 | 0.5457 | 0.4452 | 0.3927 |
| $\rho_8$ | 0.0124 | 0.5774 | 0.5502 | 0.4381 | 0.4109 |
| $\rho_9$ | 0.0956 | 0.5638 | 0.5618 | 0.4149 | 0.4323 |
| $\rho_{10}$ | 0.1203 | 0.5810 | 0.5990 | 0.4505 | 0.5108 |

decay suggestive of long-memory type features. This is also evident from the time series plots for each of the five series given in Fig. 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 Eqs. (4) and (5),

$$R_{t,t+1} = \alpha + \beta RV_{t,t+1} + u_{t+1}, \tag{20}$$

$$R_{t,t+1} = \alpha^* + \beta^* IV_{t,t+1}^* + u_{t+1}^*, \tag{21}$$

along with the corresponding robust regressions in standard deviation form,

$$R_{t,t+1} = \alpha + \beta RV_{t,t+1}^{1/2} + u_{t+1}, \tag{22}$$

$$R_{t,t+1} = \alpha^* + \beta^* IV_{t,t+1}^{*1/2} + u_{t+1}^*, \tag{23}$$

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
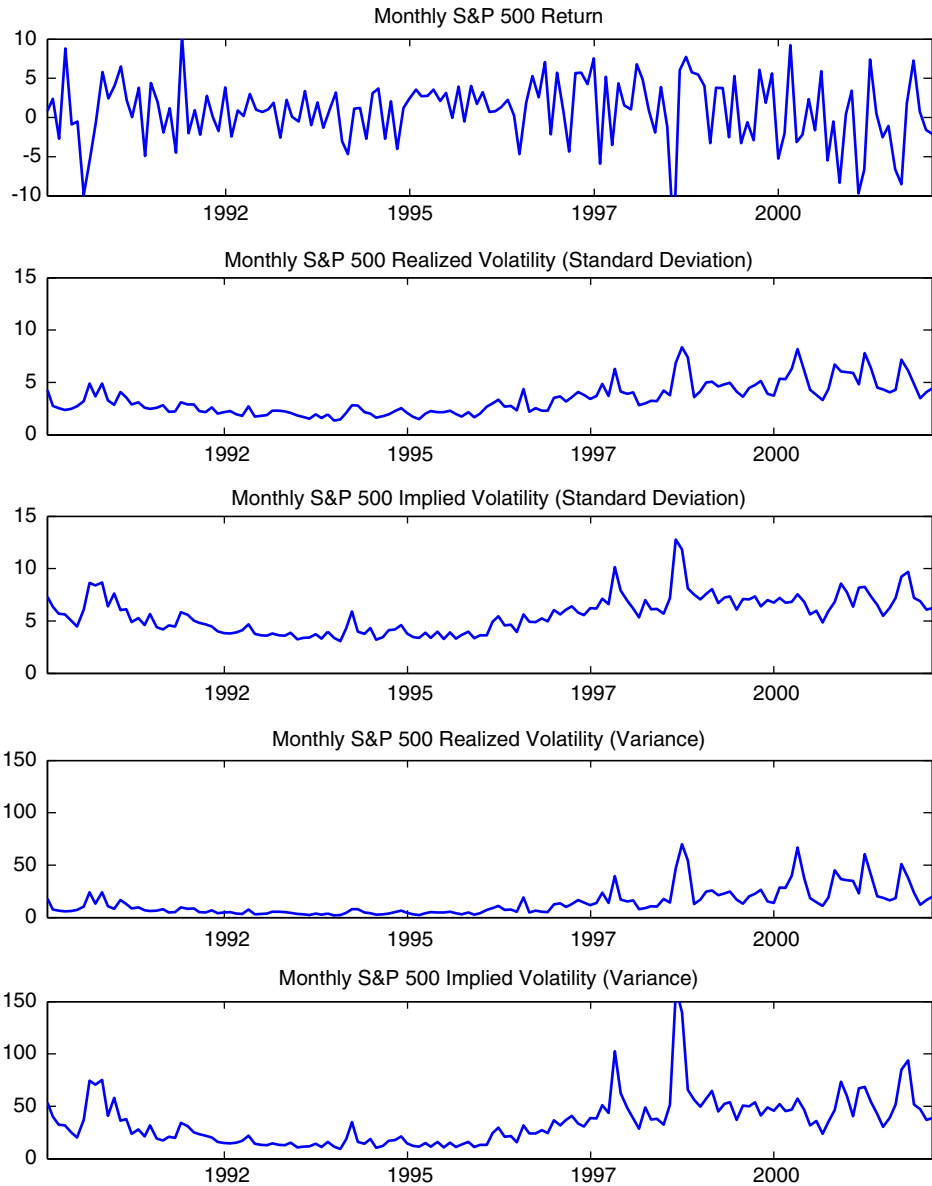
Fig. 1. Time series plot of returns and volatilities.

based on a Newey–West covariance matrix estimator allowing for a two-month lag, are reported in Table 4.[16] Interestingly, the two realized volatility regressions both

---

[16]The residuals under the Heston (1993) specification is clearly heteroskedastic, and any misspecification is likely to result in serial correlation.

Table 4
Volatility feedback effect

| Standard deviation | | | Variance | | |
|---|---|---|---|---|---|
| Realized | Implied | Expected | Realized | Implied | Expected |
| $\alpha = 2.661$ | $\alpha^* = -0.489$ | $\tilde{\alpha} = 0.961$ | $\alpha = 1.700$ | $\alpha^* = -0.001$ | $\tilde{\alpha} = 0.889$ |
| (0.712) | (0.914) | (1.048) | (0.401) | (0.458) | (0.395) |
| $\beta = -0.532$ | $\beta^* = 0.235$ | $\tilde{\beta} = -0.036$ | $\beta = -0.061$ | $\beta^* = 0.023$ | $\tilde{\beta} = -0.003$ |
| (0.239) | (0.175) | (0.210) | (0.030) | (0.012) | (0.024) |
| | | | | | |
| $R^2 = 0.038$ | $R^2 = 0.010$ | $R^2 = 0.005$ | $R^2 = 0.038$ | $R^2 = 0.018$ | $R^2 = 0.004$ |

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

result in significant estimates for the intercepts (positive) and slopes (negative). In contrast, the implied volatility regressions produce insignificant (negative) intercepts and marginally significant (positive) slopes. 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, $\beta < 0 < \beta^*$, provided that the condition $0 < \lambda_s < -\rho\kappa/\sigma$ is satisfied by the underlying structural model parameters.[17]

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 6 are also consistent with the recent empirical findings by Brandt and Kang (2004), 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. (2004) 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 4 report the results from a standard instrumental variables procedure in which we rely on the lagged squared 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 four percent to virtually zero. Moreover, the corresponding standard errors are also much larger, and the

---

[17]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 $0 < -\rho\kappa/\sigma < \lambda_s < a_{\Delta}^*/(a_{\Delta} - a_{\Delta}^*)\rho\kappa/\sigma$.

parameter estimates are no longer statistically significant.[18] As such, these results further highlight the sensitivity to the particular volatility proxy and instrument choice employed in the reduced form volatility feedback regressions.

## 3.2. Leverage effect

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

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

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

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,

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

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

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,

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

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

In these longer regressions, weak asymmetry would again be implied by negative $\delta$'s, while strong asymmetry would have the $\alpha$'s be negative as well. Similarly, if the implied volatility responds more asymmetrically to past returns than do the realized volatility, we would expect to find that the estimates for the $\delta$'s satisfy the relation $\delta^* < \delta < 0$. These same considerations apply to the pair of robust standard deviation regressions,

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

$$\mathrm{IV}^{*1/2}_{t,t+1} = \gamma^* + \beta^* \mathrm{IV}^{*1/2}_{t-1,t} + \alpha^*|R_{t-1,t}| - \delta^*|R_{t-1,t}|I_{(R_{t-1,t} \leqslant 0)} + u^*_t. \tag{31}$$

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

---

[18]We also experimented with a number of other instrumental variables, including the lagged raw and absolute returns and the lagged volatilities in standard deviation and variance forms. The results (available upon request) were qualitatively similar, but only the lagged squared returns produced positive $R$-squares.

Table 5
Leverage effect

**Short regression**

| Standard deviation | | Variance | |
|---|---|---|---|
| Realized | Implied | Realized | Implied |
| $\gamma = 3.509$ | $\gamma^* = 5.731$ | $\gamma = 14.916$ | $\gamma^* = 36.373$ |
| (0.206) | (0.235) | (1.737) | (3.076) |
| $\delta = -0.088$ | $\delta^* = -0.118$ | $\delta = -0.901$ | $\delta^* = -1.749$ |
| (0.039) | (0.040) | (0.408) | (0.651) |
| | | | |
| $R^2 = 0.056$ | $R^2 = 0.075$ | $R^2 = 0.078$ | $R^2 = 0.094$ |

**Long regression**

| Standard deviation | | Variance | |
|---|---|---|---|
| Realized | Implied | Realized | Implied |
| $\gamma = 0.693$ | $\gamma^* = 0.864$ | $\gamma = 3.829$ | $\gamma^* = 7.009$ |
| (0.192) | (0.164) | (1.067) | (1.475) |
| $\beta = 0.718$ | $\beta^* = 0.782$ | $\beta = 0.604$ | $\beta^* = 0.723$ |
| (0.058) | (0.029) | (0.066) | (0.042) |
| $\alpha = 0.055$ | $\alpha^* = -0.009$ | $\alpha = 0.059$ | $\alpha^* = -0.085$ |
| (0.040) | (0.027) | (0.063) | (0.053) |
| $\delta = -0.074$ | $\delta^* = -0.308$ | $\delta = -0.102$ | $\delta^* = -0.551$ |
| (0.050) | (0.034) | (0.071) | (0.054) |
| | | | |
| $R^2 = 0.672$ | $R^2 = 0.843$ | $R^2 = 0.606$ | $R^2 = 0.829$ |

The actual S&P estimation results for the leverage regressions are reported in Table 5. 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 $\delta$'s from the variance regression, $\delta^* = -1.75$ and $\delta = -0.90$, also adhere 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 $\delta^* = -0.12$ and $\delta = -0.09$, 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 $\delta^*$ 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 stronger asymmetry for the implied as opposed to the realized volatility regressions. The well-documented strong own temporal dependencies in the

Table 6
Implied volatility forecasting bias

| Standard deviation | Variance |
|---|---|
| $\phi_0 = -0.381$ | $\phi_0 = -0.356$ |
| (0.248) | (1.019) |
| $\phi_1 = 0.675$ | $\phi_1 = 0.412$ |
| (0.053) | (0.040) |
| | |
| $R^2 = 0.623$ | $R^2 = 0.533$ |

volatility also result in large and highly significant estimates for the $\beta$'s, along with much higher $R$-squares for the long volatility regressions (0.83–0.84 for the implied and 0.61–0.67 for the realized volatilities).

All-in-all, the at first somewhat puzzling empirical findings for the different regressions reported in Table 5 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 5 are also consistent with previous empirical evidence in the literature related to the significance, or the lack thereof, of the volatility asymmetry effect for other markets and time periods (see, e.g., Schwert, 1990; 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,

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

or in terms of standard deviations,

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

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 stochastic volatility risk is priced.[19]

The actual estimation results reported in Table 6 also do not support the unbiasedness hypothesis. Both of the estimates for $\phi_1$ are significantly less than

---

[19]The existence of a non-zero stochastic volatility risk premium for explaining estimates of $\phi_1 < 1$, along the lines of Proposition 3, has previously been discussed by Benzoni (2001), Chernov (2002), Pan (2002), and Bates (2003), among others.

unity, and the estimates for $\phi_0$ are less than zero, albeit not significantly so. The regression in standard deviation form results in a fairly high $R$-square of 0.62, while the $R$-square from the less robust variance regression equals 0.53. 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., Neely, 2003; Martens and Zeins, 2004 and the survey by Poon and Granger, 2003). Moreover, the direction of the estimated biases are exactly as expected from Proposition 3, and as such do not necessarily suggest any inefficiencies.

## 4. Conclusion

The continuous-time framework developed 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 biased forecasts of subsequent realized volatilities.

It would be interesting to extend the empirical analysis for the aggregate S&P500 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 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.

## Appendix 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 $c_\Delta = (e^{-\kappa\Delta} + \kappa\Delta - 1)/\kappa^2$. 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. (2004), the variance term may be written as,

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

Similarly, the covariance term takes the form,

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

Rearranging and integrating by parts, the second term becomes,

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

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

$$\int_t^{t+\Delta} \mathrm{E}\left(\int_t^u \sqrt{V_s}\,\mathrm{d}B_s V_u\right)\mathrm{d}u = \int_t^{t+\Delta} -\kappa \int_t^u \mathrm{E}\left(\int_t^s \sqrt{V_\tau}\,\mathrm{d}B_\tau V_s\right)\mathrm{d}s\,\mathrm{d}u$$
$$+ \int_t^{t+\Delta}\left(\int_t^u \rho\sigma\theta\,\mathrm{d}s\right)\mathrm{d}u,$$

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

$$\frac{\mathrm{d}E\left(\int_t^u \sqrt{V_s}\,\mathrm{d}B_s V_u\right)}{\mathrm{d}u} = -\kappa E\left(\int_t^u \sqrt{V_s}\,\mathrm{d}B_s V_u\right) + \rho\sigma\theta. \tag{A.1}$$

Solving this equation yields,

$$\mathrm{E}\left(\int_t^u \sqrt{V_s}\,\mathrm{d}B_s V_u\right) = \mathrm{e}^{-\kappa(u-t)}\mathrm{E}\left(\int_t^t \sqrt{V_s}\,\mathrm{d}B_s V_t\right) + \frac{\rho\sigma\theta}{\kappa}(1 - \mathrm{e}^{-\kappa(u-t)}). \tag{A.2}$$

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

$$\int_t^{t+\Delta}\left(\int_t^u \sqrt{V_s}\,\mathrm{d}B_s V_u\right)\mathrm{d}u = \frac{\rho\sigma\theta}{\kappa^2}(\mathrm{e}^{-\kappa\Delta} + \kappa\Delta - 1) = \rho\sigma\theta c_\Delta. \tag{A.3}$$

Therefore,

$$\beta = \frac{\mathrm{COV}\left(R_{t,t+\Delta}, \int_t^{t+\Delta} V_u\,\mathrm{d}u\right)}{\mathrm{VAR}\left(\int_t^{t+\Delta} V_u\,\mathrm{d}u\right)} = \lambda_s + \frac{\rho\kappa}{\sigma}, \tag{A.4}$$

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

$$\alpha = \mathrm{E}(R_{t,t+\Delta}) - \beta\mathrm{E}\left(\int_t^{t+\Delta} V_u\,\mathrm{d}u\right) = \mu\Delta + \lambda_s\theta\Delta - \beta\theta\Delta = \mu\Delta - \frac{\rho\kappa\theta}{\sigma}\Delta. \tag{A.5}$$

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

$$\mathrm{VAR}\left[\mathrm{E}_t^*\left(\int_t^{t+\Delta} V_u\,\mathrm{d}u\right)\right] = \mathrm{VAR}(a_\Delta^* V_t + b_\Delta^*) = \frac{\sigma^2\theta}{2\kappa}a_\Delta^* a_\Delta^*.$$

Similarly, for the covariance term

$$\mathrm{COV}\left[R_{t,t+\varDelta}, \mathrm{E}_t^*\left(\int_t^{t+\varDelta} V_u\,\mathrm{d}u\right)\right]$$

$$= \mathrm{COV}\left(\mu\varDelta + \lambda_s\int_t^{t+\varDelta} V_u\,\mathrm{d}u + \int_t^{t+\varDelta}\sqrt{V_u}\,\mathrm{d}B_u, a_\varDelta^* V_t + b_\varDelta^*\right)$$

$$= \lambda_s a_\varDelta^*\,\mathrm{COV}\left(\int_t^{t+\varDelta} V_u\,\mathrm{d}u, V_t\right) + a_\varDelta^*\,\mathrm{COV}\left(\int_t^{t+\varDelta}\sqrt{V_u}\,\mathrm{d}B_u, V_t\right)$$

$$= \lambda_s a_\varDelta^*\,\mathrm{COV}(a_\varDelta V_t + b_\varDelta, V_t) + a_\varDelta^*\,\mathrm{E}\left(\int_t^{t+\varDelta}\sqrt{V_u}\,\mathrm{d}B_u V_t\right)$$

$$= \lambda_s\frac{\sigma^2\theta}{2\kappa}a_\varDelta a_\varDelta^*.$$

Hence,

$$\beta^* = \frac{\mathrm{COV}\left[\int_t^{t+\varDelta} V_u\,\mathrm{d}u, \mathrm{E}_t^*\left(\int_t^{t+\varDelta} V_u\,\mathrm{d}u\right)\right]}{\mathrm{VAR}\left[\mathrm{E}_t^*\left(\int_t^{t+\varDelta} V_u\,\mathrm{d}u\right)\right]} = \lambda_s\frac{(\sigma^2\theta/2\kappa)a_\varDelta a_\varDelta^*}{(\sigma^2\theta/2\kappa)a_\varDelta^* a_\varDelta^*} = \lambda_s\frac{a_\varDelta}{a_\varDelta^*}, \qquad (\mathrm{A.6})$$

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

$$\alpha^* = \mathrm{E}(R_{t,t+\varDelta}) - \beta^*\mathrm{E}\left[\mathrm{E}_t^*\left(\int_t^{t+\varDelta} V_u\,\mathrm{d}u\right)\right] = \mu\varDelta + \lambda_s\theta\varDelta - \beta^*(a_\varDelta^*\theta + b_\varDelta^*). \quad (\mathrm{A.7})$$

Finally, combining the results above, it follows readily that if $0 < \lambda_s < -\rho\kappa/\sigma$,

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

while for $0 < -\rho\kappa/\sigma < \lambda_s < a_\varDelta^*/(a_\varDelta - a_\varDelta^*)\rho\kappa/\sigma$,

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

To establish unbiasedness of the realized return - expected variance regression discussed in the main text in Section 2.1, it suffices to show that

$$\mathrm{COV}\left[R_{t,t+\varDelta}, \mathrm{E}_t\left(\int_t^{t+\varDelta} V_u\,\mathrm{d}u\right)\right]$$

$$= \mathrm{COV}\left(\mu\varDelta + \lambda_s\int_t^{t+\varDelta} V_u\,\mathrm{d}u + \int_t^{t+\varDelta}\sqrt{V_u}\,\mathrm{d}B_u, a_\varDelta V_t + b_\varDelta\right)$$

$$= \lambda_s a_\varDelta\mathrm{COV}\left(\int_t^{t+\varDelta} V_u\,\mathrm{d}u, V_t\right) + a_\varDelta\mathrm{COV}\left(\int_t^{t+\varDelta}\sqrt{V_u}\,\mathrm{d}B_u, V_t\right)$$

$$= \lambda_s a_\varDelta\mathrm{COV}(a_\varDelta V_t + b_\varDelta, V_t) + a_\varDelta\mathrm{E}\left(\int_t^{t+\varDelta}\sqrt{V_u}\,\mathrm{d}B_u V_t\right)$$

$$= \lambda_s\frac{\sigma^2\theta}{2\kappa}a_\varDelta^2.$$

Then, utilizing the results in the first part of the proof for the implied volatility feedback effect to evaluate the variance of the expectation, it follows that $\tilde{\beta} = \lambda_s$ and $\tilde{\alpha} = \mu\Delta$.

## Appendix B. Proof of Proposition 2

By definition

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

The denominator may be rewritten as,

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

where the second equality follows from the results in Andersen et al. (2004), Bollerslev and Zhou (2002), and Proposition 1 above. The numerator may be expressed as,

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

where the first term uses the result from Andersen et al. (2004), and the second term uses the result from the proof of Proposition 1. Combining the two equations it follows therefore that

$$\delta = \frac{\lambda_s(\theta\sigma^2/2\kappa)a_\Delta^2 + \rho\sigma\theta a_\Delta^2}{\theta\Delta + \lambda_s\sigma\theta((\lambda_s\sigma/\kappa) + 2\rho)c_\Delta}. \tag{B.1}$$

This readily implies that the intercept in the realized volatility asymmetry regression equals,

$$\gamma = \theta\Delta - \delta(\mu\Delta + \lambda_s\theta\Delta). \tag{B.2}$$

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

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

The numerator may be rewritten as,

$$
\begin{aligned}
\text{COV}&\left[\text{E}_t^*\left(\int_t^{t+\Delta} V_u\, du\right), R_{t-\Delta,t}\right]\\
&= \text{COV}[a_\Delta^* V_t + b_\Delta^*, R_{t-\Delta,t}]\\
&= a_\Delta^* \lambda_s \text{COV}\left(V_t, \int_{t-\Delta}^t V_u\, du\right) + a_\Delta^* \text{COV}\left(V_t, \int_{t-\Delta}^t \sqrt{V_u}\, dB_u\right)\\
&= \lambda_s \frac{\theta\sigma^2}{2\kappa} a_\Delta^* a_\Delta + \rho\sigma\theta a_\Delta^* a_\Delta,
\end{aligned}
$$

where $a_\Delta^* = \frac{1}{\kappa^*}(1 - e^{-\kappa^*\Delta})$ and $b_\Delta^* = \theta^*(\Delta - a_\Delta^*)$, and the last line of the proof utilizes the results from the proof of Proposition 1. Now combing the different equations it follows that

$$\delta^* = \frac{\lambda_s(\theta\sigma^2/2\kappa)a_\Delta^* a_\Delta + \rho\sigma\theta a_\Delta^* a_\Delta}{\theta\Delta + \lambda_s\sigma\theta((\lambda_s\sigma/\kappa) + 2\rho)c_\Delta}, \tag{B.3}$$

while the intercept in the implied volatility asymmetry regression takes the form

$$\gamma^* = (a_\Delta^*\theta + b_\Delta^*) - \delta^*(\mu\Delta + \lambda_s\theta\Delta). \tag{B.4}$$

Lastly, note that the common denominator of $\delta$ and $\delta^*$ (corresponding to a variance) is always positive. If $0 < \lambda_s < -2\rho\kappa/\sigma$ we therefore have,

$$\delta^* < \delta < 0, \tag{B.5}$$

while for $0 < -2\rho\kappa/\sigma < \lambda_s$,

$$0 < \delta < \delta^*. \tag{B.6}$$

The assumption that $\lambda_v < 0$ ensures that $a_\Delta^* > a_\Delta > 0$. Thus the usual parameter restrictions in the Proposition guarantees the ordering of the $\delta$'s.

## Appendix C. Proof of Proposition 3

From the proof of Proposition 1,

$$
\begin{aligned}
\phi_1 &= \frac{\text{COV}[\int_t^{t+\Delta} V_u\, du, \text{E}_t^*(\int_t^{t+\Delta} V_u\, du)]}{\text{VAR}[\text{E}_t^*(\int_t^{t+\Delta} V_u\, du)]}\\
&= (\sigma^2\theta/2\kappa)a_\Delta a_\Delta^* / (\sigma^2\theta/2\kappa)a_\Delta^* a_\Delta^* = \frac{a_\Delta}{a_\Delta^*} = \frac{(1 - e^{-\kappa\Delta})\kappa^*}{(1 - e^{-\kappa^*\Delta})\kappa} < 1,
\end{aligned}
$$

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

$$\phi_0 = \mathrm{E}\left(\int_t^{t+\Delta} V_u \, \mathrm{d}u\right) - \phi_1 \mathrm{E}\left[\mathrm{E}_t^*\left(\int_t^{t+\Delta} V_u \, \mathrm{d}u\right)\right]$$

$$= \theta\Delta - \frac{a_\Delta}{a_\Delta^*}(a_\Delta^*\theta + b_\Delta^*) = b_\Delta - \frac{a_\Delta}{a_\Delta^*}b_\Delta^*,$$

which can generally not be signed.

# References

Abel, A.B., 1988. Stock prices under time-varying dividend risk: an exact solution in an infinite-horizon general equilibrium model. Journal of Monetary Economics 22, 375–393.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2001. The distribution of realized stock return volatility. Journal of Financial Economics 61, 43–76.

Andersen, T.G., Benzoni, L., Lund, J., 2002. An empirical investigation of continuous-time equity return models. Journal of Finance 57, 1239–1284.

Andersen, T.G., Bollerslev, T., Diebold, F.X., 2004. Handbook of Financial Econometrics, Chapter Parametric and Nonparametric Volatility Measurement, Elsevier, Amsterdam, forthcoming.

Andersen, T.G., Bollerslev, T., Meddahi, N., 2005. Correcting the errors: volatility forecast evaluation based on high-frequency data and realized volatilities. Econometrica 73, 279–296.

Andersen, T.G., Bollerslev, T., Meddahi, N., 2004. Analytical evaluation of volatility forecasts. International Economic Review, forthcoming.

Backus, D.K., Gregory, A.W., 1993. Theoretical relations between risk premiums and conditional variance. Journal of Business and Economic Statistics 11, 177–185.

Bali, T.G., Peng, L., 2003. Is there a risk-return tradeoff? Evidence from high-frequency data. Working paper, Baruch College, CUNY.

Bandi, F., Perron, B., 2001. A fundamentally different interpretation of the relation between implied and realized volatility. Working paper, Department of Economics, University of Montreal.

Bandi, F., Perron, B., 2004. Long memory and the relation between implied and realized volatility. Working paper, Department of Economics, University of Montreal.

Barndorff-Nielsen, O., Shephard, N., 2002. Econometric analysis of realised volatility and its use in estimating stochastic volatility models. Journal of Royal Statistical Society, Series B 64.

Barndorff-Nielsen, O., Shephard, N., 2004. A feasible central limit theory for realised volatility under leverage. Working paper, Nuffield College, Oxford University.

Bates, D.S., 2000. Post-'87 crash fears in the S&P futures option market. Journal of Econometrics 94, 181–238.

Bates, D.S., 2003. Empirical option pricing: a retrospection. Journal of Econometrics 116, 387–404.

Bekaert, G., Wu, G., 2000. Asymmetric volatility and risk in equity markets. Review of Financial Studies 13, 1–42.

Benzoni, L., 2001. Pricing options under stochastic volatility: an empirical investigation. Working paper, University of Minnesota.

Black, F., 1976. Studies in stock price volatility changes. Proceedings of the 1976 Business Meeting of the Business and Economics Section, American Statistical Association, pp. 177–181.

Bollerslev, T., Zhou, H., 2002. Estimating stochastic volatility diffusion using conditional moments of integrated volatility. Journal of Econometrics 109, 33–65.

Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance. Journal of Econometrics 52, 5–59.

Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH models. In: Engle, Robert F., McFadden D.L. (Eds.), Handbook of Econometrics, vol. IV, North-Holland, Amsterdam.

Bollerslev, T., Gibson, M., Zhou, H., 2005. Estimating volatility risk premium and investor risk aversion from model-free option-implied and realized volatilities. Working paper, Division of Research and Statistics, Federal Reserve Board.

Brandt, M.W., Kang, Q., 2004. On the relationship between the conditional mean and volatility of stock returns: a latent VAR approach. Journal Financial Economics, forthcoming.

Breen, W., Glosten, L.R., Jagannathan, R., 1989. Economic significance of predictable variations in stock index returns. Journal of Finance 44, 1177–1189.

Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes, and stochastic volatility. Journal of Finance 55, 839–866.

Campbell, J.Y., 1987. Stock returns and the term structure. Journal of Financial Economics 18, 373–399.

Campbell, J.Y., Hentschel, L., 1992. No news is good news: an asymmetric model of changing volatility in stock returns. Journal of Financial Economics 31, 281–318.

Canina, L., Figlewski, S., 1993. The informational content of implied volatility. Review of Financial Studies 6, 659–681.

Chernov, M., 2002. On the role of volatility risk premia in implied volatilities based forecasting regressions. Working paper, Columbia University.

Chernov, M.A., Gallant, R., Ghysels, E., Tauchen, G., 2003. Alternative models for stock price dynamics. Journal of Econometrics 116, 225–257.

Chou, R.Y., 1988. Volatility persistence and stock valuations: some empirical evidence using GARCH. Journal of Econometrics 3, 279–294.

Chou, R.Y., Engle, R.F., Kane, A., 1992. Measuring risk aversion from excess returns on a stock index. Journal of Econometrics 52, 201–224.

Christensen, B.J., Prabhala, N.R., 1998. The relation between implied and realized volatility. Journal of Financial Economics 50, 125–150.

Christie, A.A., 1982. The stochastic behavior of common stock variances—value, leverage and interest rate effects. Journal of Financial Economics 10, 407–432.

Christoffersen, P.F., Diebold, F.X., 2003. Financial asset returns, direction-of-change forecasting, and volatility dynamics. Working paper, McGill University and University of Pennsylvania.

Comte, F., Coutin, L., Renault, E., 2003. Affine fractional stochastic volatility models with application to option pricing. Working paper, University of Montreal.

Day, T.E., Lewis, C.M., 1992. Stock market volatility and the information content of stock index options. Journal of Econometrics 52, 267–287.

Duffee, G.R., 1995. Stock returns and volatility: a firm-level analysis. Journal of Financial Economics 37, 399–420.

Engle, R.F., Ng, V.K., 1993. Measuring and testing the impact of news on volatility. Journal of Finance 48, 1749–1778.

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

Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. Journal of Finance 59, 1367–1403.

Eraker, B., Johannes, M., Polsan, N., 2003. The impact of jumps in equity index volatility and returns. Journal of Finance 58, 1269–1300.

Fleming, J., 1998. The quality of market volatility forecasts implied by S&P 100 index option prices. Journal of Empirical Finance 5, 317–345.

Fleming, J., Ostdiek, B., Whaley, R.E., 1995. Predicting stock market volatility: a new measure. Journal of Futures Markets 15, 265–302.

French, K.R., Schwert, G.W., Stambaugh, R.F., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–30.

Gallant, A.R., Rossi, P.E., Tauchen, G., 1992. Stock prices and volume. Review of Financial Studies 5, 199–242.

Ghysels, E.P., Santa-Clara, P., Valkanov, R., 2004. There is a risk-return tradeoff after all!. Journal of Financial Economics, forthcoming.

Glosten, L.R., Jagannathan, R., Runkle, D., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. Journal of Finance 48, 1779–1801.

Guo, H., Whitelaw, R.F., 2003. Uncovering the risk-return relation in the stock market. NBER Working paper series, No. 9927.

Harrison, P., Zhang, H.H., 1999. An investigation of the risk and return relation at long horizons. Review of Economics and Statistics 81, 399–408.

Harvey, C.R., 2001. The specification of conditional expectations. Journal of Empirical Finance 8, 573–637.

Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies 6, 327–343.

Jiang, G., Tian, Y., 2004. Model-free implied volatility and its information content. Review of Financial Studies, forthcoming.

Koopman, S.J., Uspensky, E.H., 2002. The stochastic volatility in mean model: empirical evidence from international stock markets. Journal of Applied Econometrics 17, 667–689.

Lamoureux, C.G., Lastrapes, W.D., 1993. Forecasting stock-return variance: toward an understanding of stochastic implied volatilities. Review of Financial Studies 6, 293–326.

Lettau, M., Ludvigson, S., 2004. Handbook of Financial Econometrics, Chapter Measuring and Modeling Variation in the Risk-Return Tradeoff. Elsevier, Amsterdam, forthcoming.

Martens, M., Zein, J., 2004. Predicting financial volatility: high-frequency time-series forecasts vis-a-vis implied volatility. Journal of Futures Markets 24, 1005–1028.

Merton, R.C., 1980. On estimating the expected return on the market. Journal of Financial Economics 8, 323–361.

Neely, C.J., 2003. Forecasting foreign exchange volatility: is implied volatility the best we can do? Working paper, Federal Reserve Bank of St. Louis.

Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: a new approach. Econometrica 59, 347–370.

Pan, J., 2002. The jump-risk premia implicit in options: evidence from an integrated time-series study. Journal of Financial Economics 63, 3–50.

Pindyck, R., 1984. Risk, inflation and the stock market. American Economic Review 74, 335–351.

Poon, S.-H., Granger, C., 2003. Forecasting volatility in financial markets: a review. Journal of Economic Literature 41, 478–539.

Schwert, G.W., 1990. Stock volatility and the crash of 1987. Review of Financial Studies 3, 77–102.

Scruggs, J.T., 1998. Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: a two-factor approach. Journal of Finance 52, 575–603.

Tauchen, G., 2004. Stochastic volatility in general equilibrium. Manuscript, Department of Economics, Duke University.

Tauchen, G., Hussey, R., 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. Econometrica 59, 371–396.

Tauchen, G., Zhang, H., Liu, M., 1996. Volume, volatility, and leverage: a dynamic analysis. Journal of Econometrics 74, 177–208.

Turner, C.M., Startz, R., Nelson, R., 1989. A Markov model of heteroskedasticity, risk, and learning in the stock market. Journal of Financial Economics 25, 3–22.

Wu, G., 2001. The determinants of asymmetric volatility. Review of Financial Studies 14, 837–859.

Wu, G., Xiao, Z., 2002. A generalized partially linear model of asymmetric volatility. Journal of Empirical Finance 9, 287–319.