Risk and Return in Stochastic Volatility Models: Volatility Feedback Matters!

Daniel R. Smith*
Faculty of Business Administration
Simon Fraser University

December 13, 2007

Abstract

We develop a model of stock return volatility that includes a positive risk-return relation, a significant volatility feedback effect and explains a rich set of empirical phenomena. It explains the negative correlation between returns and volatility we observe in stock data. We find that including volatility feedback dramatically strengthens the risk-return relation. Contrary to some previous research we find that volatility feedback is economically significant, explaining around 13 percent of daily, and 28 percent of monthly, stock return volatility. We demonstrate that previous studies have found an economically insignificant feedback effect because of their choice of either empirical methodology or model specification.

JEL: C22, G12.

Keywords: Stochastic volatility, risk-return relation, volatility feedback, excess volatility.

*8888 University Drive, Burnaby, BC, Canada, V5A 1S6, Email: drsmith@sfu.ca, http://www.sfu.ca/~drsmith.
I would like to thank Peter Christoffersen, Adlai Fisher, Avi Kamara, Andy Siegel and seminar participants at Pennsylvania State University, the Queensland University of Technology, the University of Technology Sydney, the Australian National University, the University of Southern Queensland, the 2005 UBC Summer Finance Conference in Toffino, the 2005 Northern Finance Association in Vancouver, the 2005 Pacific Northwest Finance Conference in Seattle, and the Likelihood Methods in Finance conference at Princeton University for their comments on earlier drafts. Any remaining errors are my responsibility.

1 Introduction

The risk-return relation is fundamental to much of finance. Many asset pricing models predict a positive relation between the conditional variance and the expected return on the market portfolio in excess of a conditionally risk-free return. However, the evidence for the risk-return relation is at best mixed. Financial economists have also found it difficult to reconcile the high volatility of stock returns with the relative stability of dividends and earnings. Finally, it has been difficult to explain the negative correlation between returns and volatility. This asymmetric return-volatility relationship is typically referred to as a "leverage" effect, yet changes in financial leverage cannot explain the magnitude of the asymmetry. In this paper we develop a very simple model of stock return volatility and include a volatility feedback effect that allows changes in volatility to affect stock returns through changes in future expected returns. This model generates a strong risk-return relation and explains all three of these challenging empirical regularities.

The volatility feedback effect (French, Schwert, and Stambaugh, 1987) is very intuitive. When volatility is persistent and positively related to future expected returns, a positive shock to volatility will result in a negative realized return. The reason for this feedback is simple: an increase in the current level of volatility causes agents to increase their forecasts of future volatility and therefore raise their future required returns. All things being equal, this unexpected increase in volatility will cause stock prices to drop so investors can earn the higher future required return. We develop a stochastic volatility model in which log-volatility evolves as an AR(1) process and incorporate volatility feedback effects. This model is a discrete-time analog of many popular continuous-time volatility process and has been widely applied in finance to model many different financial time series. We demonstrate that this simple volatility model with feedback effects generates rich dynamics in stock returns. Volatility feedback is able to generate a negative correlation between returns and future volatility that matches the empirical correlation remarkably well. We find that accounting for feedback effects results in more precise estimates of the risk-return relation with much lower standard errors. Interestingly the sign of the risk-return relation changes after allowing for volatility feedback. We also find that volatility feedback is very economically significant, explaining around 13 percent of the total variance of daily returns and about 28 percent of the variance of monthly returns.

There is an extensive literature studying the relationship between returns and volatility. French, Schwert, and Stambaugh (1987), Baillie and DeGennaro (1990), Campbell and Hentschel (1992) and more recently Ghysels, Santa-Clara, and Valkanov (2005), Guo and Whitelaw (2006), and Pastor, Sinha, and Swaminathan (2006) find a positive risk-return tradeoff. However, several other studies, including Cambpell (1987), Nelson (1991), and Whitelaw (1994) find a negative relation. Finally Glosten, Jagannathan, and Runkle (1993), Turner, Starz, and Nelson (1989), and Harvey (2001) find both positive and negative relations depending on the setup used. Able (1988), Backus and Gregory (1993), and Whitelaw (2000) respond to this weak evidence by pointing out that not all asset pricing models predict a positive risk-return relation. Alternatively, Scruggs (1998) and Guo and Whitelaw (2006) demonstrate that including hedging demands induced by time-varying expected returns strengthens the risk-return relation. Finally, Harvey (2001) and Ghysels, Santa-Clara, and Valkanov (2005) show that evidence for the risk-return relation depends on the way volatility is estimated and that more sophisticated estimates perform better.

In an interesting recent paper Lundblad (2005) argues that one would typically need extremely long samples to uncover a positive risk-return relation. Lundblad (2005) demonstrates using a Monte Carlo experiment that negative relationships arise spuriously when estimated using sample sizes that are routinely used in the literature. He argues that standard volatility-in-mean tests would require several centuries of return data to uncover an economically plausible positive risk-return relation. Using an extremely long data set for US equity returns that spans the period 1836-2003 he finds that using long samples provides unambiguous support for a positive risk-return relation that is robust to a range of alternative GARCH-in-mean specifications. He also finds that using longer samples results in much more precise parameter estimates, which in turn gives statistical support for a positive relation between returns and conditional variance. Interestingly, we find that more precise estimates of the risk-return relation arise when information contained in the correlation between returns and volatility implied by the volatility feedback effect are exploited. In particular, we find that incorporating volatility feedback effects transforms a weak negative risk-return relation into a statistically significant positive relation. Furthermore, the standard errors of the estimated volatility-in-mean parameter in models that include volatility feedback are about one-third the size of the standard errors in models without the feedback effect. It appears that ignoring the volatility feedback effect obfuscates the positive risk-return relation.

It has long been asserted that the volatility of stock returns is too large to be explained by economic fundamentals and in particular by the low variability in dividends (see, e.g., Leroy and Porter (1981), and Shiller (1981)). Since volatility is persistent researchers hoped that a positive risk-return relation along with a volatility feedback effect could rationalize high stock return volatility. However, two important studies by Poterba and Summers (1986) and Campbell and Hentschel (1992) seemed to eliminate feedback as a potential explanation. Poterba and Summers (1986) argue that volatility is not sufficiently persistent to produce economically significant fluctuations. They argue that the size of a stock price drop following an increase in volatility depends critically on the persistence of volatility. They estimate volatility using realized monthly volatility and find that its persistence is rather low (an AR(1) coefficient of around 0.7). Their results suggest that the volatility feedback effect is insignificant.

Using an alternative methodology Campbell and Hentschel (1992) also argue that volatility feedback is economically insignificant as it is only able to explain one percent of the variance of stock returns. Although Campbell and Hentschel (1992) reach the same conclusion as Poterba and Summers (1986) they do so using a model with highly persistent volatility. Campbell and Hentschel (1992) estimate a quadratic GARCH (QGARCH) model with volatility feedback effects that forces the price of risk to be small. This is because the realized stock return is a quadratic function of the news about dividends and if the price of risk is too large then large positive returns are unattainable. As a result the QGARCH model estimates a very small price of risk. However our model, which uses an econometric approach that does not impose this restriction, produces estimates of risk aversion that are larger than their QGARCH counterparts.

In contrast to Poterba and Summers (1986) and Campbell and Hentschel (1992) we find that the volatility feedback effect is very important, explaining around 13 percent of the total variance of daily returns and about 28 percent of the variance of monthly returns. To phrase the significance of the volatility feedback slightly differently, if volatility were to double from its average level in a given month, we would expect stock prices to drop by roughly 7.77 percent. We also find that uncertainty about stock return volatility is economically significant, depressing stock prices by around 40 percent.

Perhaps the most important aspect of this model is its ability to explain the large negative correlation between returns and volatility. The most common explanation for the return-volatility correlation is changes is the leverage effect (Black (1976) and Christie (1982)) in which changes in financial leverage cause changes in equity volatility. However, it is well known that changes in leverage alone are incapable of explaining the magnitude of this asymmetric relationship (see, e.g., Schwert (1989) and Bekaert and Wu (2000)). An alternative explanation for this asymmetric effect is the so-called volatility feedback effect. In our log-AR(1) stochastic volatility model we find that feedback effects are able to generate negative correlation between returns and log-volatility comparable with the estimated correlation coefficients reported in the stochastic volatility literature.

The remainder of the paper proceeds as follows. Section 2 presents the basic stochastic volatility model and discusses our estimation strategy. In Section 3 we show how volatility feedback effects can be introduced into a stochastic volatility model. Our main empirical results are presented in Section 4. We conclude in Section 5. Details of our econometric technique are presented in the Appendix.

2 Stochastic Volatility Models

Denote by r_{t+1} the excess continuously compounded return on the stock market (over the continuously compounded risk-free return)[1] and has conditional variance $\sigma^2_{t+1} = E_t[(r_{t+1} - E_t(r_{t+1}))^2]$. We model the conditional variance as a log-AR(1) process in which log-volatility $x_{t+1} = \log(\sigma^2_{t+1})$ is an AR(1) process

$$x_{t+1} = \omega + \phi x_t + \nu_t. \tag{1}$$

[1]The continuously compounded return is given by $k_t = \log(P_t + D_t) - \log(P_t)$, then the excess return is given by $r_t = k_t - \log(1 + R_{f,t})$ were $R_{f,t}$ is the one-period risk-free return.

This type of model was originally introduced by Taylor (1986) and has been found to fit stock returns and many other financial time series very well.[2] An important feature of this modeling specification is that volatility is positive by construction.

As is common in the literature, we incorporate a risk-return relationship following Merton (1980) by modeling expected excess returns as a linear function of conditional volatility

$$E_t(r_{t+1}) = \mu + \gamma\sigma^2_{t+1}, \tag{2}$$

which is known at time t and we have included an intercept.[3] In this setting the parameter γ captures the risk-return relation and can be thought of as the representative investor's coefficient of relative risk aversion. This specification for the time-varying risk premia has become the standard model applied in finance.[4] In this stochastic volatility model stock returns evolve as:

$$\begin{aligned} r_{t+1} &= \mu + \gamma \exp(x_{t+1}) + \exp(x_{t+1}/2) z_{t+1} \\ x_{t+1} &= \omega + \phi x_t + \nu_t \end{aligned}$$

where

$$\begin{bmatrix} z_{t+1} \\ \nu_{t+1} \end{bmatrix} \sim \text{iid}MVN\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho_{z\nu}\sigma_\nu \\ \rho_{z\nu}\sigma_\nu & \sigma^2_\nu \end{bmatrix} \right).$$

When $\rho_{z\nu} \neq 0$ then the return on the stock today is conditionally correlated with future volatility.[5]

In this model log-volatility is unobserved, which complicates estimating the model's parameters. Many estimation strategies have been proposed to estimate the parameters of the discrete-time log-AR(1) volatility model, including quasi-maximum likelihood (Harvey, Ruiz, and Shephard (1994) and Harvey and Shephard (1996) for the correlated models), Bayesian-based Markov-Chain Monte Carlo methods (Jacquier, Polson, and Rossi (1994) and in correlated volatility models by Jacquier, Polson, and Rossi (2004)), Simulated Maximum Likelihood (Danielsson (1994)) and Monte-Carlo Maximum Likelihood (Sandmann and Koopman (1998)). In this paper we employ the exact non-linear filter of Kitagawa (1987), and in particular the method of Fridman and Harris (1998), which

[2] An incomplete list of papers include Harvey, Ruiz, and Shephard (1994), Jacquier and Polson and Rossi (1994, 2004), Fridman and Harris (1998), Watanabe (1999), Sandmann and Koopman (1998), Andersen, Chung, and Sorensen (1999), Ball and Torous (1999), Smith (2002). Continuous-time stochastic volatility models have been estimated by Andersen and Lund (1997), Andersen, Benzoni, and Lund (2002), Chernov, Gallant, Ghysels, and Tauchen (2003) among others.

[3] We have included an intercept but this has no effect on the volatility feedback term. It does, however, improve the overall fit of the model. The intercept can be motivated by the existence transaction costs or taxes (see also Ng (1991, p.1509)). For example, μ would be the dividend yield in Litzenberger and Ramaswamy's (1979) after-tax version of the CAPM.

[4] Note that Scruggs (1998) and Guo and Whitelaw (2006) argue that accounting for hedging demands strengthens the risk-return relation. We find that accounting for volatility feedback produces similar results.

[5] Some models allow returns to be correlated with contemporaneous volatility, e.g. Jacquier, Polson, and Rossi (2004), which produces negative skewness as the volatility of positive returns is scaled back as volatility decreases, while negative returns are compounded as volatility increases. However, Yu (2005) argues that models that allow returns to be correlated with subsequent volatility are preferable both statistically and theoretically.

integrates numerically over the latent variables. In Appendix A we present an algorithm to construct the log-likelihood numerically and which extends Fridman and Harris (1998) to allow for correlated disturbances. The idea is to recursively construct the log-likelihood for stock returns by calculating the marginal density of stock returns at each date by numerically integrating over possible values of latent volatility. A by-product of the algorithm is the density of latent volatility evaluated at N points. At each step in the algorithm we update our inference about likely values of volatility. Monte Carlo studies by Fridman and Harris (1998) and Watanabe (1999) show that such numerical integration-based algorithms have impressive finite sample properties.

Models that allow returns and volatility to be correlated invariably find a strong negative correlation, i.e., $\rho_{z\nu} < 0$, and this asymmetry is invariably interpreted as a "leverage" effect with any asymmetric stochastic volatility model being given this label (see, for example, Omori, Chib, Shephard, and Nakajima (2007), Wu (2001), and Yu (2005)). This interpretation follows the early work of Black (1976) and Christie (1982) who demonstrated that even when a firm has constant volatility it's equity volatility will increase when the firm's leverage increases. Thus negative stock returns will tend to increase volatility as the value of equity drops. This intuition is well established in the GARCH literature (see, for example, Glosten, Jagannathan, and Runkle (1993)). However, Schwert (1989) and Bekaert and Wu (2000) have shown that changes in leverage are unable to explain the magnitude of the asymmetry effect in GARCH and realized volatility models. An alternative explanation for this asymmetric relationship is the volatility feedback effect, which we discuss in the next section.

3 Modeling Volatility Feedback

In this section we extend the stochastic volatility model with a positive risk-return relationship to include feedback between shocks to volatility and returns. The volatility feedback effect is quite intuitive. When volatility increases unexpectedly, future expected returns, which are positively related to future volatility, also increase. In the absence of news about future cash flows this implies that the current stock price must drop in response to the increase in volatility, producing negative correlation between returns and volatility.

We formally develop the feedback effect using the log-linear approximation to the intertemporal budget equation of Campbell and Shiller (1988), Campbell (1991), and Campbell and Ammer (1993). In particular, the continuously compounded returns on the market portfolio k_{t+1}, which is given by $k_{t+1} = \log(P_{t+1} + D_{t+1}) - \log(P_t)$, is approximated by

$$k_{t+1} \approx \kappa + \rho p_{t+1} - p_t + (1-\rho) d_{t+1}, \tag{3}$$

where $p_{t+1} = \log(P_{t+1})$, $d_{t+1} = \log(D_t)$, $\rho = \frac{1}{1+\exp(d-p)}$, $\exp(d-p)$ is the average log-dividend-price ratio, and $\kappa = -\log(\rho) - (1-\rho)\log\left(\frac{1}{\rho}-1\right)$, and alternatively as

$$k_{t+1} \approx \kappa - \rho \cdot dp_{t+1} + dp_t + \Delta d_{t+1}, \tag{4}$$

where $dp_{t+1} = \log(D_{t+1}/P_{t+1})$ is the log dividend price ratio, and Δd_{t+1} is the log dividend growth rate. Campbell, Lo, and MacKinlay (1997, Chapter 7) note that the historical annual average log dividend price ratio is about 4% and therefore suggest using $\rho = 0.96$ for annual data, 0.997 for monthly data, and $\rho = 0.9998$ for daily data. Sequentially substituting out dividend price ratios in equation (4) and imposing a transversality condition $\lim_{j\to\infty} \rho^j dp_{t+j} = 0$ allows us to write the following accounting identity relating current excess returns (r_{t+1}) to future expected returns, dividend growth rates (Δd_{t+j+1}) and risk-free returns ($r_{f,t+j+1}$):

$$r_{t+1} - E_t(r_{t+1}) = \eta_{d,t+1} - \eta_{r,t+1}, \tag{5}$$

where:

$$\eta_{r,t+1} = (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j r_{t+j+1} \tag{6}$$

$$\eta_{d,t+1} = (E_{t+1} - E_t)\sum_{j=0}^{\infty} \rho^j \Delta d_{t+j+1} - (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j r_{f,t+j+1}. \tag{7}$$

We refer to $\eta_{r,t+1}$ as news about future expected returns, and $\eta_{d,t+1}$ as news about future cash flows. This result is nothing more than the expectation of an accounting identity that holds using the definition of returns $1 + R_{t+1} = \frac{P_{t+1}+D_{t+1}}{P_t}$ and taking a log-linearization of the budget equation.

Recall that the risk premium is linear in conditional volatility (2), and log-volatility is an AR(1) process (1). This log-AR specification for volatility implies that future volatility is lognormally distributed and admits a closed-form expression (involving an infinite sum) for future volatility forecasts. Note that future log-volatility (for $j > 0$) can be written as

$$x_{t+j+1} = \omega(1 + \phi + \cdots + \phi^{j-1}) + \phi^j x_{t+1} + \sum_{i=1}^{j} \phi^{j-i} \nu_{t+i}, \tag{8}$$

which is normally distributed because the news about log-volatility is normal. At time t we know the value of x_{t+1}, so the time t conditional mean of x_{t+j+1} for $j > 0$ is given by:

$$E_t(x_{t+j+1}) = (1 - \phi^j)\frac{\omega}{1-\phi} + \phi^j x_{t+1} \tag{9}$$

and the time t conditional variance is given by

$$V_t(x_{t+j+1}) = (1 - \phi^{2j})\frac{\sigma_\nu^2}{1-\phi^2}. \tag{10}$$

The new information at time $t+1$ is ν_{t+1}, the volatility surprise, so the updated conditional mean and variance of future volatility are given by:

$$E_{t+1}(x_{t+j+1}) = (1-\phi^j)\frac{\omega}{1-\phi} + \phi^j x_{t+1} + \phi^{j-1}\nu_{t+1} \tag{11}$$

and

$$V_{t+1}(x_{t+j+1}) = \sigma_\nu^2 \frac{1-\phi^{2(j-1)}}{1-\phi^2}. \tag{12}$$

The time $t+1$ conditional variance is smaller than the time t variance because the time $t+1$ information set is larger. Because the expected return is linear in σ^2_{t+j+1} we have

$$(E_{t+1} - E_t)\rho^j r_{t+j+1} = \gamma\rho^j (E_{t+1} - E_t)(\exp(x_{t+j+1})) \tag{13}$$

and using the log-normality of x_{t+j+1}[6] these latter two terms are given by

$$E_t(\exp(x_{t+j+1})) = \exp\left(\omega\frac{1-\phi^j}{1-\phi} + \phi^j x_{t+1} + .5\sigma_\nu^2\frac{1-\phi^{2j}}{1-\phi^2}\right) \tag{14}$$

and

$$E_{t+1}(\exp(x_{t+j+1})) = \exp\left(\omega\frac{1-\phi^j}{1-\phi} + \phi^j x_{t+1} + \phi^{j-1}\nu_{t+1} + .5\sigma_\nu^2\frac{1-\phi^{2(j-1)}}{1-\phi^2}\right). \tag{15}$$

We can simplify the expression by using the following notation

- $\sigma_0^2 = \exp(\frac{\omega}{1-\phi})$ is the exponential of the unconditional log-volatility,
- $\sigma^2_{t+1} = \exp(x_{t+1})$ is the current variance of returns
- $\xi_0^2 = \exp\left(\frac{\sigma_\nu^2}{1-\phi^2}\right) = \exp(\sigma_x^2)$ is the exponential of the unconditional variance of log-volatility.

With this notation we can write the change in expected returns as

$$(E_{t+1} - E_t)\rho^j r_{t+j+1} = \gamma\rho^j \left(\sigma_0^2\right)^{1-\phi^j} \left(\sigma_t^2\right)^{\phi^j} \left[\xi_0^{1-\phi^{2(j-1)}} \exp\{\phi^{j-1}\nu_{t+1}\} - \xi_0^{1-\phi^{2j}}\right] \tag{16}$$

and the feedback term $\eta_{r,t+1}$ is then obtained by summing (16), discounted by ρ^j, over all $j > 0$, which gives the feedback effect induced by a shock to log-volatility ν_{t+1}:

$$\eta_{r,t+1} := \eta_r(\nu_{t+1}; x_{t+1}) = \gamma\sum_{j=1}^{\infty} \rho^j \left(\sigma_0^2\right)^{1-\phi^j} \left(\sigma_{t+1}^2\right)^{\phi^j} \left[\xi_0^{1-\phi^{2(j-1)}} \exp\{\phi^{j-1}\nu_{t+1}\} - \xi_0^{1-\phi^{2j}}\right]. \tag{17}$$

The first two terms $(\sigma_t^2)^{\phi^j}(\sigma_0^2)^{1-\phi^j}$ represent a geometric average of two forecasts of future volatility: the current level of volatility σ^2_{t+1} and its long run average $\sigma_0^2 = \exp(\frac{\omega}{1-\phi})$ and the weights depend on ϕ which determines the speed of mean reversion. Forecasts further into the future place less

[6] If x is normal with mean μ and variance σ^2, then $E(\exp(x)) = \exp(\mu + \sigma^2/2)$.

weight on the current value of log-volatility and more weight on its unconditional expected value. The quantity in the square brackets accounts for the revision of future volatility forecasts that arises from learning the true value of ν_{t+1}. It is easy to verify that the expected value of the feedback term $E(\eta_{r,t+1}) = 0$ since the expected value of each element in the square bracket is zero.

Volatility feedback effects are capable of generating asymmetry in returns when both $\phi > 0$ and $\gamma > 0$. The unexpected return generated by news about volatility is given in equation (17). When $\gamma > 0$ an unexpected increase in future volatility $\nu_{t+1} > 0$ implies a positive feedback realization $\eta_{r,t+1} > 0$, which, in turn generates a negative unexpected return since returns are given by $r_{t+1} - E_t(r_{t+1}) = \eta_{d,t+1} - \eta_{r,t+1}$. This captures the intuition of the volatility feedback effect.

Combining these results, the stochastic volatility model with feedback effects we estimate is given by

$$\begin{aligned} r_{t+1} &= \mu + \gamma\sigma_{t+1}^2 - \gamma \sum_{j=1}^{\infty} \rho^j \left(\sigma_0^2\right)^{1-\phi^j} \left(\sigma_{t+1}^2\right)^{\phi^j} \left[\xi_0^{1-\phi^{2(j-1)}} \exp\{\phi^{j-1}\nu_{t+1}\} - \xi_0^{1-\phi^{2j}}\right] \\ &\quad + \delta_{t+1} z_{t+1} \qquad (18) \\ x_{t+1} &= \omega + \phi x_t + \nu_t, \qquad (19) \end{aligned}$$

where the iid innovations $z_{t+1} \sim N(0,1)$ and $\nu_{t+1} \sim N(0,\sigma_\nu^2)$, and $\delta_{t+1} = \delta^2(x_{t+1}) = \sigma_{t+1}^2 - E_t(\eta_{r,t+1}^2|x_{t+1})$ is the conditional variance (conditional on x_{t+1}) of the news about future cash flows and is defined in equation (50) in Appendix B. The conditional expected return is $\mu + \gamma\sigma_{t+1}^2$ as in the basic SV model. Also, using the definition of δ_{t+1} we have the conditional variance is σ_{t+1}^2. Finally, the correlation between returns and future volatility arises because the innovation in future volatility x_{t+1} is ν_{t+1} and appears in both expressions. The correlation between returns and future volatility is determined endogenously by x_{t+1}, ϕ and γ, increasing in magnitude with ϕ and γ, and will have the opposite sign as γ. We assume that although the econometrician does not observe conditional volatility, prices are set as if the representative investor does know the true value of log-volatility x_{t+1} at time t. In Appendix B we present the details of an algorithm to construct the quasi-log-likelihood for the stochastic volatility model with feedback.

This model is the first to incorporate volatility feedback when the risk premium depends on the variance of stock returns. There are several models in the literature that are similar to our model but do not simultaneously include both these features. In particular, there are two basic approaches to estimating the risk-return relation with volatility feedback. The first approach is to model the risk premium as a function of the volatility of dividends. Prominent examples of this type of model are Campbell and Hentschel (1992) and Wu (2001). A second approach is to model the risk premium using an estimate of stock return volatility (e.g., using realized volatility estimated from daily returns or implied volatility extracted from option prices) along with a model of volatility dynamics to generate a feedback effect. An important example of this approach is Guo and Whitelaw (2006) who model realized volatility dynamics using a VAR model including standard financial instrumental variables and implied volatility. Returns are decomposed into news about future returns, whose variance is given by the dynamics of volatility, and the news about future cash flows, whose variance is estimated as a free parameter. However, the variance of stock returns depends on the variance of these two news components and generally does not equal the "variance" used in the risk premium. Our model has a risk-premium that depends on stock return volatility and connects the variance of the news about cash flows and future returns to the conditional volatility used in the risk-premium. Interestingly we find that closing the loop between these news components and the volatility of returns results in a strong and statistically significant risk-return relation.

Our model also has some other important advantages. By modeling log-volatility directly we ensure that volatility is strictly positive, while some previous models do not require conditional variance to be positive. For example, Wu (2001) models volatility as an autoregressive process and Guo and Whitelaw (2006) models volatility as a linear function of conditioning variables that follow a VAR process. Also, we estimate the model parameters by quasi-maximum likelihood, which is an important advantage over Wu's (2001) model, which must be estimated by EMM and has some econometric concerns.[7] Perhaps the most important advantage of this setup is that it allows us to relate the conditional volatility used in specifying expected returns to the variance of unexpected returns. The current literature either specifies returns as a function of the variance of dividends (e.g., Campbell and Hentschel (1992) and Wu (2001)) or treats stock return "volatility" as a state variable that drives the conditional mean but is not related to the conditional variance of unexpected returns (e.g., Guo and Whitelaw (2006)). We find that using conditional variance to specify both the first and second moments of stock returns produces a model that generates a positive risk-return relationship and is able to explain several other important characteristics of stock returns.

4 Empirical Results

We estimate the models using both daily and monthly continuously compounded returns on the value-weighted portfolio of all stocks on the NYSE, AMEX and NASDAQ taken from CRSP. We

[7] In particular, Wu's (2001) model has 7 parameters to be estimated and these are chosen so the structural stochastic volatility model will generate data that is close to the AR(1)-GARCH(1,1) SNP auxiliary model fitted to actual data. This auxiliary model has eight parameters, of which only two govern the mean dynamics and three describe the volatility dynamics. The other three parameters explain the tail behavior of stock returns. To statistically identify the model Wu (2001) had to arbitrarily fix two of the parameters (the unconditional mean of dividend volatility and the unconditional mean of the log dividend growth rate) as these are only weakly identified. Because some of the parameters have been fixed we must maintain a healthy degree of scepticism about the standard errors, which are likely understated because the uncertainty about these two key parameters is ignored. Furthermore, given that the structural model is only capable of producing symmetric returns we have serious reservations about the informativeness of the three tail parameters in the auxiliary model. It seems that we have only five reliable moments to estimate seven structural parameters. In contrast we are able to reliably estimate all the model's parameters by using quasi-maximum likelihood.

report results using monthly data for the CRSP value-weighted index from 1926 to 2004 and, following Campbell and Hentschel (1992), we also consider a shorter sample period starting in January 1952. The post-1951 data starts after the Fed-Treasury Accord since we are dealing with excess returns, and does not include the Great Depression. We also report results using the CRSP value-weighted daily daily for the period July 1962 to December 2004. To allay concerns that our daily results may be contaminated by microstructural issues, we repeated the analysis using both weekly data and daily data purged of non-synchronous trading-induced autocorrelation with a moving-average filter. The results are materially unchanged and are available on request.[8]

We report the parameter estimates and standard errors for a range of stochastic volatility models in Tables 1 through 3 for each of the three samples. These models include both correlated and uncorrelated innovations, and both with and without a risk-return relation. Models 1 through 4 do not include the volatility feedback effect, which is incorporated in Model 5. We note that because of the feedback effects only Model 1 is nested as a special case of Model 5 (i.e. when $\gamma = 0$). Conditional volatility in all four non-feedback models exhibit very high degrees of volatility persistence. The autocorrelation coefficients for the full sample results are all around 0.95 or higher, implying a half-life of volatility shocks (i.e., $\log(0.5)/\log(\phi)$) of between 12 and 20 months. Volatility shocks in the post-1951 sample are slightly less persistent but still have half-lives of between 6 to 10 months, and the half-life of shocks to daily volatility implied by Table 3 are around 50 trading days. This high persistence is common for stochastic volatility and GARCH models.

Models 1 and 2 assume zero correlation between returns and volatility, while Models 3 and 4 relax this assumption. There is strong evidence of an asymmetric relationship between returns and volatility as we find estimates a correlation between returns and volatility of between -0.33 and -0.59 in monthly data, and around -0.36 in the daily data. The log-likelihoods for all asymmetric volatility models improve dramatically over the symmetric volatility models. Formal likelihood ratio tests of the null hypothesis $\rho = 0$ are easily rejected in all cases.

We next consider the statistical evidence of a risk-return relation. Although theory predicts a positive relation (i.e. $\gamma > 0$) we find a negative risk-return relation (i.e. $\gamma < 0$) in every specification without feedback effects in all three samples. The point estimates of γ are between 1.5 and 3 standard errors less than zero in Models 2 and 4 in all three samples. Interestingly, the evidence for a variance-in-mean effect is weakest in the correlated model. For example, in the post-1951 data the point estimate is only 1.47 standard errors below zero. Furthermore, the quasi log-likelihood is relatively flat around $\gamma = 0$. In the monthly data the maximum log-likelihood restricting $\gamma = -10$ and $\gamma = 1$ are virtually identical at 1150.8 and 1150.6 and the maximum unrestricted value is only at 1151.5–hardly compelling evidence.

[8]All our sample periods include the stock market crash of October 1987. To avoid having our results driven by this highly influential "outlier", we Windsorize the data by truncating the single largest absolute return to have the same magnitude as the second largest absolute return. This does not have a material impact on the parameter estimates, but the standard errors are slightly higher when using the non-truncated data. These results are available upon request.

However, when we include volatility feedback this weak evidence of a negative risk-return relation is dramatically overturned.[9] We find a statistically significant positive risk-return relation in all three data sets and find $\gamma \approx 3.5$ in both the monthly post-1951 and daily data. Models that omit volatility feedback find a weak negative relationship but after accounting for volatility feedback we find a strong positive relation. These estimates of γ are statistically significant. A robust Wald test in the full sample monthly data is 3.5252 and the post-1951 data is 4.4723, and the daily data is larger still at 49.1092 (these tests are significant respectively at the 6%, 3% and 0% levels). To formally compare Model 5 with Model 4 we use the non-nested likelihood ratio test of Vuong (1989),[10] which has a standard normal asymptotic distribution, though we note that Model 5 is based on a quasi-log likelihood. The Vuong (1989) test for the monthly post-1951 data is 1.6028 (with a p-value of 0.109), and the daily data is 1.6836 (with a p-value of 0.093). Thus although the more heavily parameterized Model 4, which estimates the correlation as an extra free parameter, has a higher log-likelihood, we cannot distinguish it statistically from Model 5, in which the correlation is endogenously determined from the other model parameters through the feedback effect. The volatility feedback effect does a remarkable job of explaining the negative correlation between returns and future volatility. Also, failing to account for the volatility feedback obfuscates the risk-return relation and the previously identified negative relationship is spurious.

The standard errors on γ decrease dramatically after including volatility feedback. In the non-feedback models the only information about γ comes from the average returns conditional on various levels of volatility through time. However, when we include the volatility feedback effect, information about γ also comes from the covariance between returns and volatility, since the correlation is determined endogenously by the other model parameters including γ. This extra information results in a large drop in the standard errors of γ. This result compliments Lundblad (2005) who searches for more information about the risk-return by using a longer sample period. We find that information from the covariance between returns and volatility viewed through the lens of the feedback effect also provides information about the risk-return relation. The only way returns can be negatively correlated with future volatility is if returns and volatility are positively related.

We plot the conditional volatility and the conditional risk premium implied by the volatility feedback model estimated with monthly data in Figure 1 for the full sample and in Figure 2 for

[9] The volatility dynamics implied by Model 5 are virtually identical to the non-feedback models.

[10] Vuong (1989) shows that under certain regularity conditions the appropriately normalized likelihood ratio statistic converges to a standard normal random variable:

$$T^{-1/2} LR_T / \hat{\omega}_T \xrightarrow{D} N(0,1)$$

where $LR_T = L_T^A - L_T^B$, and L_T^A and L_T^B are the log-likelihoods of two non-nested models, T is the number of observations, and

$$\hat{\omega}_T^2 = \frac{1}{T} \sum_{t=1}^{T} \left[\log \frac{f_A(r_t)}{f_B(r_t)} \right]^2 - \left[\frac{1}{T} \sum_{t=1}^{T} \log \frac{f_A(r_t)}{f_B(r_t)} \right]^2$$

is a consistent estimator for the variance of the likelihood ratio statistic.

the post-1951 data, and the daily estimates in Figure 3. The well known episodes of high volatility, the great depression and the crash of October 1987 are clearly evident. The other high volatility episodes are the 1937-1938 recession, the recessions in 1970 and 1974, and the internet bubble at the end of the sample. The high variability in expected returns through time is also clear. To avoid having our results being driven by the Great Depression, which is clearly dominant, we report only the plots for the post-1951 results for the remainder of the paper.

4.1 Economic Significance

Poterba and Summers (1986) use the elasticity of volatility to assess the economic significance of changes in volatility on prices. They find an elasticity of prices with respect to volatility with an AR(1) model of realized volatility (using data from 1930-1984) suggesting a 2.2 percent drop in stock prices following a doubling in volatility. Our model suggests that when volatility doubles from its unconditional level ($E(\sigma^2) = \exp(\omega/(1-\phi) + 0.5\sigma_\nu^2/(1-\phi^2)) = 0.0017$) stock returns will drop by 7.77 percent (using monthly post-1951 estimates). This is a three standard deviation event. A more modest two standard deviation shock to log-volatility increases volatility from its unconditional value will still result in an impressive 4.79 percent drop in stock prices.

Campbell and Hentschel (1992) argue that volatility feedback is only capable of explaining one percent of the variance of stock returns. Given a level of log-volatility $x_{t+1} = x_i$ the fraction of total returns volatility that is due to volatility feedback is given by

$$\chi(x_i) = \frac{E(\eta_{r,t+1}^2 | x_i)}{\delta^2(x_i) + E(\eta_{r,t+1}^2 | x_i)}, \tag{20}$$

where $\delta^2(x_i) = E(\eta_{d,t+1}^2 | x_{t+1} = x_i)$ is the conditional variance of the news about future cash flows and is defined in equation (50) in Appendix B. Given the conditional density of log-volatility given information available at time t, the conditional expected value of χ is computed as follows:

$$E_t(\chi_{t+1}) \approx \sum_{i=1}^{N} w_i \chi(x_i) f(x_{t+1} = x_i | I_t). \tag{21}$$

We plot the time series of the conditional fraction of total variance due to the volatility feedback effect in Figure 4. Volatility feedback is clearly economically significant, explaining about 13 percent of the total variance of daily stock returns and around 28 percent of the variance of monthly stock returns.

The economically significant feedback effect we find stands in stark contrast with the insignificant feedback effect reported by Poterba and Summers (1986) and Campbell and Hentschel (1992). To generate a significant feedback effect two things are needed: 1) a large price of risk (γ), and 2) highly persistent volatility (ϕ). We find that both of these necessary conditions are met in our model, which consequently generates an economically meaningful role for volatility feedback.

In contrast, Poterba and Summers (1986) suggest that volatility is not terribly persistent, and Campbell and Hentschel (1992) find a small risk aversion coefficient, leading them to conclude that volatility feedback is insignificant. We demonstrate in Sections 4.4 and 4.5 that the conclusions of both papers are driven by their choice of either econometric specification or technique, which obfuscates the economically significant feedback effect.

4.2 Volatility Asymmetry

There are two common explanations for the asymmetric relationship between returns and volatility: changes in financial leverage and volatility feedback. Previous research has demonstrated that changes in leverage alone is incapable of explaining the asymmetry in volatility (see, e.g., Schwert (1989) and Bekaert and Wu (2000)). However, we demonstrate that this negative asymmetry can be explained by volatility feedback. The correlation between returns and log-volatility conditional on log-volatility being equal to x_i is computed as

$$\rho(x_i) = -\int \eta_r(\nu; x_i)\nu\phi(\nu|0, \sigma_\nu^2)d\nu \tag{22}$$

$$\approx -\sum_{i=1}^{N} w_i'\eta_r(\nu_i; x_i)\nu_i\phi(\nu_i|0, \sigma_\nu^2) \tag{23}$$

where $\phi(.|a, b)$ denotes the probability density function of a normal random variable with mean a and variance b, and w_i' and ν_i are the weights and nodes for Gauss-Legendre integrating over ν.[11] Using this the conditional correlation is calculated as

$$\rho_{t+1|t} \approx \sum_{i=1}^{N} w_i\rho(x_i)f(x_{t+1} = x_i|I_t), \tag{24}$$

and if we use the unconditional density of x_t we obtain the unconditional correlation.

We plot the time series of conditional correlation coefficients in Figure 5 for both daily and monthly stock returns. We also report the unconditional correlation coefficient in the last column (Model 5) of Tables 1 through 3. The main result to take from these results is how close the conditional and unconditional correlation coefficients are to the point estimates of the empirical correlation for Models 3 and 4 of about −0.53 for the monthly data (in Table 2) and about −0.37 for the daily data (in Table 3). The volatility feedback effect can explain the negative return-volatility asymmetry we observe in the data that cannot be explained by changes in financial leverage alone.

4.3 Volatility Uncertainty Discount

Another way to view the economic significance of conditional volatility on returns is the discount in prices induced by the uncertainty in conditional volatility. Our approach adapts Campbell and

[11] In particular we integrate over -6 to +6 standard deviations.

Hentschel (1992) who evaluate the economic importance of time variation in the risk premium from time-varying volatility. The idea is to compare the level of prices we would observe under two very different specifications of expected excess returns. The approximation to the budget equation given in equation (3) can also be used to express the level of stock prices (Campbell, Lo, and MacKinlay 1997, Eq. 7.1.22, p.263) including the risk-free return as

$$p_t = \frac{\kappa}{1-\rho} + E_t\left(\sum_{j=1}^{\infty} \rho^j [(1-\rho)d_{t+j+1} - r_{f,t+j+1} - r_{t+j+1})]\right), \tag{25}$$

which can alternatively be expressed as

$$p_t = \frac{\kappa}{1-\rho} + E_t\left(\sum_{j=1}^{\infty} \rho^j [(1-\rho)d_{t+j+1} - r_{f,t+j+1})]\right) - \zeta_t, \tag{26}$$

where

$$\zeta_t = \sum_{j=1}^{\infty} \rho^j E_t(r_{t+j+1}). \tag{27}$$

Different specifications of the risk premium $E_t(r_{t+j+1})$ will produce different levels of stock prices for the same sequence of dividends. One way to assess the economic significance of, for example, uncertainty about volatility σ_ν^2, is to compare the level of stock prices that would prevail in a world with no volatility uncertainty (i.e., $\sigma_\nu^2 = 0$) holding other parameters constant, with the level of stock prices in an economy with the current level of volatility uncertainty. This quantification of economic significance is made without reference to the dynamics of dividends as it only requires that we specify the discounted sum of future expected returns under different specifications for the price of risk. In particular, the difference between the current level of log-stock prices p_t, and the level of log-stock prices we would observe under an alternative specification for expected returns $p_t^\star$, depends only on the discounted conditional expected returns in the two economic models:

$$p_t - p_t^\star = \zeta_t^\star - \zeta_t, \tag{28}$$

since the first two terms in (26) do not depend on the dynamics of expected returns. Using the log-normality of conditional expected returns we can write

$$\zeta_t = \gamma \sum_{j=1}^{\infty} \rho^j E_t e^{x_{t+j+1}} = \gamma \sum_{j=1}^{\infty} \rho^j (\sigma_0^2)^{1-\phi^j} (\sigma_{t+1}^2)^{\phi^j} \xi_0^{1-\phi^{2j}}. \tag{29}$$

We can express the discount in stock prices caused by volatility uncertainty by setting $\sigma_\eta^2 = 0$ in $\zeta_t^\star$, which gives

$$p_t - p_t^\star = \zeta_t^\star - \zeta_t = \gamma \sum_{j=1}^{\infty} \rho^j (\sigma_0^2)^{1-\phi^j} (\sigma_{t+1}^2)^{\phi^j} (1 - \xi_0^{1-\phi^{2j}}). \tag{30}$$

We plot the time series of volatility uncertainty discounts in Figure 6. Uncertainty about stock return volatility depresses stock prices by about 40 percent in monthly data and by about 44 percent in daily data. This stock price depression is higher during periods of high volatility.

4.4 Comparison with QGARCH Feedback Models

Campbell and Hentschel (1992) present a model of volatility feedback in which the conditional expected return on the stock market (rather than excess returns) is modeled as a linear function of the conditional volatility of news about dividends[12]

$$E_t(k_{t+1}) = \mu + \gamma\sigma_{t+1}^2 \tag{31}$$

where the variance of dividends $\sigma_{t+1}^2 = E_t(\eta_{d,t+1}^2)$ is modeled as a QGARCH process:

$$\sigma_{t+1}^2 = \beta_0 + \beta_1(\eta_{d,t} - b)^2 + \beta_2\sigma_t^2.$$

The advantage of using the QGARCH specification is that the feedback effect $\eta_{h,t+1}$ (the revision of the discounted expected value of future total returns analogous to $\eta_{r,t+1}$) is available in closed form and depends on the news about dividends

$$\eta_{h,t+1} = \lambda(\eta_{d,t+1}^2 - \sigma_t^2 - 2b\eta_{d,t+1}) \tag{32}$$

where $\lambda = \frac{\gamma\rho\beta_1}{1-\rho(\beta_1+\beta_2)}$ so the unexpected stock return is a known quadratic function of only $\eta_{d,t+1}$ and the actual realized stock return is given by

$$k_{t+1} = \mu + \gamma\sigma_t^2 + \kappa\eta_{d,t+1} - \lambda(\eta_{d,t+1}^2 - \sigma_t^2) \tag{33}$$

where $\kappa = 1 + 2\lambda b$.

Because the unexpected return is a quadratic function of the news about dividends $\eta_{d,t+1}$ there is a maximum unexpected return that can be generated by this model. Given a realized unexpected return k_{t+1}, the news about dividends at time $t+1$, $\eta_{d,t+1}$, is found by solving the quadratic equation

[12] The reason for using total rather than excess returns is that the functional form of the QGARCH model forces the choice between modeling expected returns as a linear function of the variance of dividends, as do Campbell and Hentschel (1992), or modeling expected excess returns as a linear function of the variance of the news about both future dividends and future risk-free returns. The log-linear expansion models expected returns using a volatility-in-mean specification that depends on the variance of future cash flows (the first term in equation (7)), which is $E_t(\eta_{d,t+1}^{\star 2})$ for excess returns, and $E_t(\eta_{d,t+1}^2)$ for total returns. The key is that the news about "cash flows" is different when we model total rather than excess returns. When we decompose total returns (k_{t+1}) we have $\eta_{d,t+1}^\star = k_{t+1} - E_t(k_{t+1}) + \eta_{k,t+1}$, and $\eta_{d,t+1}^\star$ is news about future dividends. However, when we decompose excess returns (r_{t+1}) in the same fashion we have $\eta_{d,t+1} = r_{t+1} - E_t(r_{t+1}) + \eta_{r,t+1}$, and $\eta_{d,t+1}$ represents news about both 1) future dividends and 2) future risk-free rates of return. (Note that $\eta_{k,t+1}$ and $\eta_{r,t+1}$ both refer to news about future expected total and excess returns and are computed in the same way.) Thus the QGARCH model implies that expected stock returns are a quadratic function of the news about dividends (total returns) or the news about dividends and risk-free returns (excess returns).

in (33) and is denoted by $\hat{\eta}_{t+1}$. To be well defined $\eta_{d,t+1}$ must be real, which occurs if the inequality $\kappa \leq \lambda\hat{\eta}_{t+1}$ holds. This places an upper bound on the price of risk λ. Because the quadratic nature of the QGARCH model restricts the price of risk to be small it also limits the magnitude of the volatility feedback effect that can be generated.

We report the parameter estimates and robust standard errors using total returns for the QGARCH model of Campbell and Hentschel (1992) along with a range of alternative GARCH-in-mean models as a reference in Tables 4 (monthly post-1951 data) and 5 (daily data). All the alternative GARCH-in-mean models can be written as special cases of the following recursive conditional volatility specification

$$\sigma_{t+1}^2 = \beta_0 + (\beta_1 + \delta 1_{e_t<0})(e_t - b)^2 + \beta_2 \sigma_t^2, \tag{34}$$

which includes both symmetric and asymmetric (following Glosten, Jagannathan, and Runkle (1993)) GARCH and GARCH-in-mean models. It is clear that the QGARCH model proposed by Campbell and Hentschel (1992) (Model 4) dominates all the other specifications. There is evidence of a positive risk-return relation as suggested by the estimates of γ: a strongly significant risk-return relation in the daily data (γ is 4 standard errors from zero), but only marginally significant in the post-1951 monthly data (γ is only 1.88 standard errors from zero).

We cannot directly compare the log-likelihoods for the QGARCH models with the stochastic volatility models because the GARCH-based models are fitting total returns while the stochastic volatility-based models use excess returns. To compare the models we fit the stochastic volatility feedback model directly to total returns using $E_t(k_{t+1}) = \mu + \gamma\sigma_{t+1}^2$. We report the estimates of γ and the log-likelihood for this model in the notes to Tables 4 and 5. Note that the estimated price of risk is larger in the stochastic volatility model than in the QGARCH model: in the monthly post-1951 data it is twice as large, and is 10 times larger in the daily data. This supports our earlier claim that the QGARCH model finds a small feedback effect because it mechanically places a limit on the size of the price of risk. Interestingly the implied risk-return relation in the stochastic volatility model is slightly stronger in the total returns than in excess returns. We also compare the model fit as measured by quasi log-likelihood between the QGARCH and stochastic volatility models. The QGARCH feedback model with $b = 0$ (i.e., Model 3) has the same number of parameters as the stochastic volatility feedback model but has a lower log-likelihood in both the monthly and daily data. The most general QGARCH feedback model has one extra parameter, yet has a lower log-likelihood in the daily data. In the monthly post-1951 data the QGARCH model 4 has a slightly higher log-likelihood than the stochastic volatility model, but note that it has an extra parameter (after accounting for this extra parameter the Bayesian Information Criteria are virtually identical -2266.72 (stochastic volatility model) and -2267.00 (QGARCH model).

4.5 Realized Volatility-Based Persistence

Poterba and Summers (1986) find that volatility is not sufficiently persistent to generate an economically meaningful feedback effect. Poterba and Summers report that the first-order autocorrelation coefficient for realized S&P500 volatility[13] from 1950-1984 is 0.57 and over 1930-1984 it is slightly higher at 0.73, both estimates are rather low. This low persistence in realized volatility is a common finding, e.g. French, Schwert, and Stambaugh (1987) report first order autocorrelation coefficients for realized volatility (again for S&P500 returns) over various subperiods of between 0.62 to 0.71. This is clearly at odds with the GARCH and stochastic volatility literatures in which estimated persistence (the AR coefficient in stochastic volatility models and the sum of the GARCH parameters in GARCH models) is typically around 0.95. Why are these estimates of volatility persistence so disparate? The difference is driven by two biases in the estimates of persistence endemic to realized volatility. The first bias is the well known bias from OLS estimates of the autocorrelation coefficient ϕ, which Kendall (1954) shows is (when estimated along with an intercept) approximately $-(1+3\phi)/T$ (see also Stambaugh, 1999). The second and most important bias is an errors-in-variables induced downward bias, which arises because monthly realized volatility is calculated using on average only 22 daily returns. This sample estimate is therefore quite noisy and its use as a regressor biases the estimated persistence toward zero (the estimate is also inconsistent, Greene (1990, p.375)).

To assess the extent of the downward bias induced by these two problems we undertake a Monte Carlo experiment in which we simulate $T = \{60, 120, 480, 960, 1920\}$ observations of monthly volatility, which evolve as a log-AR(1) process with the same parameters as reported in Table 1 (in particular we use the parameters from Model 1):[14]

$$x_{t+1} = -0.4402 + 0.9327 \cdot x_t + 0.2254 \cdot z_t$$

where $z_t \sim N(0,1)$. We then simulate $M = \{22, 50, 100, 200\}$ intra-monthly returns for each month, which we denote by $r_s(M)$ and are indexed as $s = 1, \ldots, TM$. The realized volatility for month t is then computed using the M intra-monthly returns as $\hat{\sigma}_t^2 = \sum_{s=(t-1)\times M+1}^{t\times M} r_s^2$. The next step is to estimate an AR(1) regression (including a constant) and we report the average AR coefficient and R^2 across $N = 1000$ simulations in Table 6 for $\hat{\sigma}_t^2$ (columns 2 and 3) and $\log \hat{\sigma}_t^2$ (columns 4 and 5). We also estimate a regression using true monthly log-volatility $\log \sigma_t^2$, which allows us to determine the finite-same AR bias alone. By comparing this average coefficient with the realized volatility-based estimates we can gauge the importance of the errors-in-variables bias.

The main point to take from this experiment is that both biases are important. The magnitude of the finite sample bias in the autocorrelation coefficient can be gauged from the $M = \infty$ row for

[13] Realized volatility for any given month is computed as the sum of squared daily returns within that month.

[14] Note that Anderson, Bollerslev, Diebold, and Ebens (2001) find that the lognormal distribution adequately describes the dynamics of daily realized stock return volatility, and Anderson, Bollerslev, Diebold, and Labys (2001) provide similar evidence for daily exchange rate volatility.

various sample sizes. For samples of 10 years or less the bias is important, but for very large samples the bias declines. We also find that the approximation $-(1+3\phi)/T$ is quite accurate. The errors-in-variables bias can be seen by comparing the size of the average log-AR coefficient for various M holding the sample size constant. When $T = 60$ using 22 daily observations to estimate realized volatility, we find a 35 percent downward bias in the estimated persistence in log-volatility (relative to the estimate obtained using true log-volatility) and the estimated persistence is almost 40 percent less than the true persistence. Increasing the number of observations used to calculate realized volatility to 100 mitigates the errors-in-variables problem somewhat but the total bias is still around 13 percent. Increasingly the total sample size marginally improves the bias but even with 160 years of monthly data using 22 daily observations still results in a downward bias of more than 17 percent. The problem is exacerbated when using realized volatility rather than log-realized volatility. Using daily realized volatility (rather than log-realized volatility) to estimate volatility produces even lower persistence with average persistence measures anywhere between 0.54 (with 5 years of returns) and 0.70 (with 80 and 160 years of returns), suggesting that model specification error is yet another source of downward bias on volatility persistence. Recall that Poterba and Summers (1986) report estimates of the autocorrelation in realized volatility estimated using daily returns of between 0.57 and 0.73. These low estimates of volatility persistence are actually reasonably high given the biases inherent in the estimation method.

This experiment suggests that estimates of volatility persistence that use daily returns to estimate monthly realized volatility significantly understate the persistence of volatility. Sampling more frequently improves matters, but non-synchronous trading and other microstructural issues complicate the issue when using stock indices.[15] So we can place more confidence in the estimated persistence of volatility that come from studies that use many observations, though we must be concerned with microstructural issues.

5 Conclusions

In this paper we have demonstrated that a very simple model of stock return volatility is able to reproduce a rich set of stock return behavior. When the expected excess return on the market portfolio depends positively on the conditional level of stock return volatility, which is modeled as a log-AR(1) process, there are economically significant volatility feedback effects. An increase in volatility raises future expected returns through the positive risk-return relation since volatility is persistent. Stock prices drop so investors can earn these higher future required returns. This feedback effect is large, explaining around 13 percent of the variance of daily returns and 28 percent of

[15] This point is made by Aït-Sahalia, Mykland, and Zhang (2005) and Zhang, Mykland, and Aït-Sahalia (2005) who demonstrate that the optimal choice of sampling frequency in ultra-high frequency settings (i.e., when using intra-daily data) trades off the benefits of sampling very frequently and using more information against the noise induced by microstructural issues. They show how to calculate the optimal sampling frequency in the presence of micro-structural complications. Aït-Sahalia, Mykland, and Zhang (2006) present a strategy to estimate realized volatility using all tick-by-tick trade data and explicitly takes into account market microstructure effects.

the variance of monthly stock returns. This effect is time-varying and is orders of magnitude larger than reported in the previous literature. The economic importance of this feedback-induced return volatility somewhat offsets concerns about the seemingly excessive volatility of stock returns given relatively smooth consumption and dividend process. Uncertainty about stock return volatility is also important, causing stock prices to be depressed by around 40 percent.

Our model is also able to generate the negative asymmetric relationship between returns and volatility. Previous research has shown that changes in financial leverage cannot explain this negative correlation. However, we find that volatility feedback is able to explain the large negative correlation between returns and volatility. In fact we find a remarkable correspondence between the estimated correlation coefficient and the average conditional correlation implied by our feedback model.

When we estimate a stochastic volatility model that ignores the feedback effect we find a negative risk aversion coefficient. However, after we account for volatility feedback effects we find a positive and statistically significant risk-return relation. Accounting for volatility feedback dramatically improves the precision with which we estimate the risk aversion coefficient. It is apparent that the correlation between returns and volatility provides useful information we can exploit to estimate the risk-return relation.

We reconcile our empirical findings with previous conclusions in the literature that volatility feedback is insignificant. We demonstrate that by using realized volatility to estimate volatility persistence, Poterba and Summers (1986) introduce a severe downward bias in estimated persistence, which spuriously muted the magnitude of the volatility feedback effect. On the other hand the quadratic nature of the QGARCH model used by Campbell and Hentschel (1992) produces artificially low estimates of the risk aversion coefficient because of the quadratic functional form of volatility. This serves to depresses the size of the volatility feedback effect. When we employ a specification that avoids these two restrictions we find that volatility is persistent and a strong risk-return relation which combine to generate our significant volatility feedback effect.

Appendix

A Estimating Stochastic Volatility Models

We introduce our estimation strategy by considering the basic stochastic volatility-in-mean model

$$\begin{aligned} r_{t+1} &= \mu + \gamma \exp(x_{t+1}) + \exp(x_{t+1}/2) z_{t+1} \\ x_{t+1} &= \omega + \phi x_t + \nu_t \end{aligned}$$

where $z_t \sim N(0,1)$ and $\nu_t \sim N(0, \sigma_\nu^2)$. Estimating stochastic volatility models is complicated because conditional volatility is latent and the joint distribution of the observable variable r_{t+1} and log-volatility is not jointly normal, which confounds the use of standard filtering schemes. We will estimate the parameters of this model with an algorithm that implements brute force numerical integration over the unobserved latent volatility using the conditional probability formulae following Kitagawa (1987) and Fridman and Harris (1998). To introduce the algorithm suppose that at time t the conditional density of log-volatility is known to be $f(x_t|I_t)$, where I_t denotes the information available to the econometrician at time t. Using the conditional normality of x_{t+1} implied by the AR(1) model we can construct the forecast density of x_{t+1}:

$$f(x_{t+1}|I_t) = \int f(x_{t+1}|x_t) f(x_t|I_t) dx_t, \tag{35}$$

where $f(x_{t+1}|x_t) = \phi(x_{t+1}; \omega + \phi x_t, \sigma_\nu^2)$ and $\phi(.; a, b)$ is the pdf of a normal random variable with mean a and variance b. The joint density of x_{t+1} and the data r_{t+1} is given by

$$f(r_{t+1}, x_{t+1}|I_t) = f(r_{t+1}|x_{t+1}) f(x_{t+1}|I_t), \tag{36}$$

where $f(r_{t+1}|x_{t+1}) = \phi(r_{t+1}; \mu + \gamma \exp(x_{t+1}), \exp(x_{t+1}))$. The log-likelihood function requires the conditional density be known

$$f(r_{t+1}|I_t) = \int f(r_{t+1}, x_{t+1}|I_t) dx_{t+1} \tag{37}$$

and is then calculated as

$$\mathcal{L}(r; \theta) = \sum_{t=0}^{T-1} \log f(r_{t+1}|I_t). \tag{38}$$

We can update our inference about the distribution of x_{t+1} by conditioning on the new information r_{t+1} using

$$f(x_{t+1}|I_{t+1}) = \frac{f(r_{t+1}, x_{t+1}|I_t)}{f(r_{t+1}|I_t)}. \tag{39}$$

The difficulty we now face is that we cannot summarize the whole conditional density from one iteration to the next. In the standard Kalman filter the entire distribution is described by the filtered estimate and its mean squared error. We overcome this limitation following Kitagawa (1987) and Fridman and Harris (1998) and use a numerical integration scheme that only requires that we keep track of the conditional densities at a finite number of (N) points. In particular we employ a numerical scheme to calculate the general integral

$$\int f(x)dx \approx \sum_{i=1}^{N} w_i f(x_i) \tag{40}$$

where w_i are a set of weights and $f(x_i)$ denotes the value of the function f evaluated at only N different points x_i for $i = 1, \ldots, N$. There are many schemes that approximate integrals with finite summation including the mid-point, trapizoidal, Simpson's rule as well as the various quadrature rules. In particular, we follow Fridman and Harris (1998) and use the Gauss-Legendre integration that approximates f over a finite interval using N orthonormal polynomials. In this scheme the weights and nodes at which f is approximated solve a set of linear equations and are thus prespecified and common to all integrals (see Judd, 1998). This scheme is quite efficient and to ensure the accuracy of our results we use 75 nodes and verified that including further nodes does not cause the results to change.

This numerical approach requires that we evaluate the joint density N times at x_i for $i = 1, \ldots, N$

$$f(r_{t+1}, x_{t+1} = x_i | I_t) = \phi(r_{t+1}; \mu + \gamma \exp(x_i), \exp(x_i)) f(x_{t+1} = x_i | I_t). \tag{41}$$

Given this joint density we can then calculate the forecast of the marginal density of conditional volatility

$$f(x_{t+1} = x_i | I_t) \approx \sum_{j=1}^{N} w_j \phi(x_{t+1} = x_i; \omega + \phi x_j, \sigma_\nu^2) f(x_t = x_j | I_t), \tag{42}$$

and the density of returns

$$f(r_{t+1} | I_t) \approx \sum_{i=1}^{N} w_i f(r_{t+1}, x_{t+1} = x_i | I_t), \tag{43}$$

which can then be used to construct the log-likelihood as in (38). We then update the conditional density of log-volatility as

$$f(x_{t+1} = x_i | I_{t+1}) = \frac{f(r_{t+1}, x_{t+1} = x_i | I_t)}{f(r_{t+1} | I_t)}, \tag{44}$$

which is then used as the input for the next iteration in the algorithm.

There is significant empirical evidence that the relationship between returns and volatility is asymmetric and in particular that returns are correlated with volatility. We model correlation between returns and subsequent volatility (see Yu, 2005).[16] We can incorporate asymmetry by

[16] An alternative modeling strategy is to allow for contemporaneous correlation between returns and volatility. This

modifying the algorithm outlined above as follows. We need to model the joint density of returns and subsequent log-volatility

$$f(r_{t+1}, x_{t+2} = x_i | I_t) \approx \sum_{j=1}^{N} w_j f(r_{t+1}, x_{t+2} = x_i | x_{t+1} = x_j) f(x_{t+1} = x_j | I_t), \tag{45}$$

which requires $f(r_{t+1}, x_{t+2} = x_i | x_{t+1} = x_j)$ the pdf of a multivariate normal random variable with mean $[\mu + \gamma e^{x_j} : \omega + \phi x_j]'$ and covariance matrix $\begin{bmatrix} e^{x_j} & \rho_{r\nu}\sigma_\nu e^{x_j/2} \\ \rho_{r\nu}\sigma_\nu e^{x_j/2} & \sigma_\nu^2 \end{bmatrix}$ with $\rho_{r\nu} = E(z_t, \nu_t)/\sigma_\nu$ being the correlation between returns and log-volatility. We calculate updated inference about volatility using

$$f(x_{t+2} = x_i | I_{t+1}) = \frac{f(r_{t+1}, x_{t+2} = x_i | I_t)}{f(r_{t+1} | I_t)} \tag{46}$$

where

$$f(r_{t+1} | I_t) \approx \sum_{i=1}^{N} w_i f(r_{t+1}, x_{t+2} = x_i | I_t) \tag{47}$$

is used to construct the log-likelihood function that we maximize when estimating the parameters.

B Estimating the Volatility Feedback Model

We will use Quasi-Maximum Likelihood in our parameter estimation that does not require that we know the exact distribution of $\eta_{d,t+1}$. We are directly modeling the conditional volatility of returns $\sigma^2_{t+1|t} = E_t(r_{t+1} - E_t(r_{t+1}))^2$. Excess returns are a function of two news innovations, which we assume are independent,[17] so we can write

$$\sigma^2_{t+1} = E_t(r_{t+1} - E_t(r_{t+1}))^2 = E_t(\eta^2_{r,t+1}) + E_t(\eta^2_{d,t+1}). \tag{48}$$

The news about expected returns $\eta_{r,t+1} = \eta_r(\nu_{t+1}, x_{t+1})$ is a nonlinear function of the normally distributed innovation in volatility ν_{t+1} whose variance we can calculate numerically

$$\begin{aligned} E_t(\eta^2_{r,t+1} | x_{t+1}) &= \int \eta_r(\nu, x_{t+1})^2 \phi(\nu; 0, \sigma^2_\nu) d\nu \\ &\approx \sum_{i=1}^{N} w_i \eta_r(\nu_i, x_{t+1}) \phi(\nu_i; 0, \sigma^2_\nu) \end{aligned} \tag{49}$$

class of model has the advantage of producing negative skewness in returns (see Jacquier, Polson, and Rossi (2004)) but has been criticized by Yu (2005) who argues that models should allow returns to be correlated with subsequent volatility as we have considered above.

[17] Interestingly, in a standard predictive regression framework (using the term spread, price earnings ratio and small-stock value spread as predictors) Campbell and Vuolteenaho (2004) find that the news about future returns and cash-flows are nearly uncorrelated.

using $\nu_{t+1} \sim N(0, \sigma_\nu^2)$ independent of $\eta_{d,t+1}$, and $\phi(\cdot; m, s^2)$ is the *pdf* of a normal random variable with mean m and variance m^2. We can now relate the variance of the news about dividends to the conditional variance of returns σ_{t+1}^2:

$$\delta^2(x_{t+1}) := E_t(\eta_{d,t+1}^2 | x_{t+1}) = \sigma_{t+1}^2 - E_t(\eta_{r,t+1}^2 | x_{t+1}). \tag{50}$$

In general we will have that $E_t(\eta_{r,t+1}^2 | x_{t+1}) < \sigma_{t+1}^2$ and the volatility is well defined, however to ensure that volatility is well defined we let

$$\delta^2(x_i) = \max(\exp(x_i) - E_t(\eta_{r,t+1}^2 | x_i), 10^{-8}) \tag{51}$$

(or a minimum standard deviation of 10^{-4}). It is worth noting that this never occurs for the parameter values we observe in practice.

To implement our QML procedure we assume that news about dividends is normally distributed, which allows us to express the density of stock returns as the product of the conditional densities of two densities conditional on knowing the value of the unobserved news about returns, which in turn depends on the unobserved news about volatility:

$$f(r_{t+1} | \sigma_{t+1}^2; \theta) = f(\eta_{r,t+1} | \sigma_{t+1}^2; \theta) \cdot \phi(r_{t+1} - E_t(r_{t+1}) + \eta_{r,t+1} | 0, \delta^2(x_i)). \tag{52}$$

In our algorithm we integrate over values of the unobserved latent state variable σ_{t+1}^2 and σ_t^2. These two terms with the parameter vector θ identify ν_{t+1} and hence $\eta_{r,t+1}$. The density of returns thus depends on both these two values of conditional volatility.

We need the joint density of returns, current and future log-volatility. The dependence of the density of returns on current log-volatility is immediately apparent, and the subsequent log-volatility is required to capture the volatility feedback effect $\eta_r(\nu_{t+1}, x_{t+1})$ since $\nu_{t+1} = x_{t+2} - \omega - \phi x_{t+1}$. The joint density is given by

$$f(r_{t+1}, x_{t+1}, x_{t+2} | I_t) = f(r_{t+1} | x_{t+1}, x_{t+2}) f(x_{t+1}, x_{t+2} | I_t). \tag{53}$$

Once the current and future values of log-volatility are realized the distribution of r_{t+1} depends only on the normally distributed news shock and the density is therefore given by

$$f(r_{t+1} | x_{t+1}, x_{t+2}) = f(r_{t+1} - \mu - \gamma \exp(x_{t+1}) + \eta_r(x_{t+2} - \omega - \phi x_{t+1}, x_{t+1}), 0, \delta^2(x_{t+1})) \tag{54}$$

where the variance of the dividend new component $\delta^2(x_i) = V_t(\eta_{d,t+1})$ is given by

$$\delta^2(x_i) = \exp(x_i) - \int \eta_r(u, x_i)^2 f(u) du + \left(\int \eta_r(u, x_i) f(u) du \right)^2 \tag{55}$$

and the integrals are evaluated numerically. We evaluate this conditional density at a finite number of points

$$f(r_{t+1}|x_{t+1} = x_i, x_{t+2} = x_j) = f(r_{t+1} - \mu - \gamma \exp(x_i) + \eta_r(x_j - \omega - \phi x_i, x_i), 0, \delta^2(x_i)). \quad (56)$$

The joint density of current and future log-volatility is given

$$f(x_{t+1}, x_{t+2}|I_t) = f(x_{t+2} - \omega - \phi x_{t+1}; 0, \sigma_v^2) f(x_{t+1}|I_t) \quad (57)$$

and in particular we focus on the N^2 values

$$f(x_{t+1} = x_i, x_{t+2} = x_j|I_t) = f(x_j - \omega - \phi x_i; 0, \sigma_v^2) f(x_i|I_t) \quad (58)$$

for $i, j = 1, \ldots, N$. The marginal density of returns is obtained by integrating out the underlying volatility which is again done numerically:

$$\begin{aligned} f(r_{t+1}|I_t) &= \int\int f(r_{t+1}, x_{t+1}, x_{t+2}|I_t) dx_{t+2} dx_{t+1} \\ &\approx \sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j f(r_{t+1}, x_{t+1} = x_i, x_{t+2} = x_j|I_t). \end{aligned} \quad (59)$$

We can now update our inference about the latent log-volatility:

$$f(x_{t+1} = x_i, x_{t+2} = x_j|I_{t+1}) = \frac{f(r_{t+1}, x_{t+1} = x_i, x_{t+2} = x_j|I_t)}{f(r_{t+1}|I_t)}. \quad (60)$$

The final step in each iteration is to integrate out x_{t+1}

$$f(x_{t+2} = x_j|I_{t+1}) = \sum_{i=1}^{N} w_i f(x_{t+1} = x_i, x_{t+2} = x_j|I_{t+1}), \quad (61)$$

which is used as the input to the next iteration. We use the conditional density $f(r_{t+1}|I_t)$ to construct the quasi-log likelihood which is in turn maximized to estimate the parameters. We are using QML because although the density of $\eta_{r,t+1}$ is an extremely nonlinear function of a normally distributed random variable we can easily calculate its conditional mean and variance and the resulting QML estimates will be consistent and asymptotically normal under suitable regularity conditions. The final step is to initialize the algorithm with $f(x_1|I_0)$ for which we use the ergodic distribution of the AR(1) log-volatility process.

References

Able, A. B., 1988. Stock Prices Under Tim-Varying Dividend Risk: An Exact Solution in an Infinite-Horzon General Equilibrium Model. Journal of Financial Economics 22, 375–393.

Aït-Sahalia, Y., P. A. Mykland, L. Zhang, 2005. How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise. Review of Financial Studies 18, 351–416.

——— , 2006. Ultra High Frequency Volatility Estimation with Dependent Microstructure Noise. Unpublished working paper. Princeton University.

Andersen, T. G., L. Benzoni, J. Lund, 2002. An Empirical Investigation of Continuous-Time Equity Returns. Journal of Finance 57, 1239–1284.

Andersen, T. G., H.-J. Chung, B. E. Sorensen, 1999. Efficient Method of Moments Estimation of a Stochastic Volatility Model: A Monte Carlo Study. Journal of Econometrics 91, 61–87.

Andersen, T. G., J. Lund, 1997. Estimating Continuous-Time Stochastic Volatility Models of the Short Term Interest Rate. Journal of Econometrics 77, 343–377.

Anderson, T., T. Bollerslev, F. Diebold, H. Ebens, 2001. The Distribution of Realized Stock Return Volatility. Journal of Financial Economics 61, 43–76.

Anderson, T., T. Bollerslev, F. Diebold, P. Labys, 2001. The Distribution of Realized Exchange Rate Volatility. Journal of the American Statistical Association 96, 42–55.

Backus, D., A. W. Gregory, 1993. Theoretical Relations Between Risk Premiums and Conditional Variances. Journal of Business and Economic Statistics 11, 173–185.

Baillie, R. T., R. P. DeGennaro, 1990. Stock Returns and Volatility. Journal of Financial and Quantitative Analysis 25, 203–214.

Ball, C. A., W. N. Torous, 1999. The Stochastic Volatility of Short-term Interest Rates: Some International Evidence. Journal of Finance 56, 2339–2359.

Bekaert, G., G. Wu, 2000. Asymmetric Volatility and Risk in Equity Markets. Review of Financial Studies 13, 1–42.

Black, F., 1976. Studies of Stock Price Volatility Changes. Proceedings of the Business and Economic Statistical Section, American Statistical Association pp. 177–181.

Cambpell, J., 1987. Stock Returns and the Term Structure. Journal of Financial Economics 18, 157–179.

Campbell, J., L. Hentschel, 1992. No News is Good News. Journal of Financial Economics 31, 281–318.

Campbell, J. Y., 1991. A Variance Decomposition for Stock Returns. Economic Journal 101, 157–179.

Campbell, J. Y., J. Ammer, 1993. What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns. Journal of Finance 48, 3–37.

Campbell, J. Y., A. W. Lo, A. C. MacKinlay, 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton, USA.

Campbell, J. Y., R. Shiller, 1988. The Dividend-Price Ratio and Expectations of Fuure Dividends and Discount Factors. Review of Financial Studies 1, 195–228.

Campbell, J. Y., T. Vuolteenaho, 2004. Bad Beta, Good Beta. American Economic Review 94, 1249–1275.

Chernov, M., A. R. Gallant, E. Ghysels, G. Tauchen, 2003. Alternative Models for Stock Price Dynamics. Journal of Econometrics 116, 225–257.

Christie, A. A., 1982. The Stochastic Behavior of Common Stock Variances: Value, Leverage, and Interest Rate Effects. Journal of Financial Economics 19, 407–432.

Danielsson, J., 1994. Stochastic Volatility in Asset Prices: Estimation with Simulated Maximum Likelihood. Journal of Econometrics 61, 375–400.

French, K., G. W. Schwert, R. Stambaugh, 1987. Expected Stock Returns and Volatility. Journal of Financial Economics 19, 3–30.

Fridman, M., L. Harris, 1998. A Maximum Likelihood Approach for Non-Gaussian Stochastic Volatility Models. Journal of Business and Economic Statistics 16, 284–291.

Ghysels, E., P. Santa-Clara, R. Valkanov, 2005. There is a Risk-Return Tradeoff After All. Journal of Financial Economics 76, 509–548.

Glosten, L., R. Jagannathan, D. Runkle, 1993. On The Relation Between Expected Value and the Volatility of the Nominal Excess Return on Stocks. Journal of Finance 48, 1779–1801.

Greene, W. H., 1990. Econometric Analysis. MacMillan, New York, NY.

Guo, H., R. Whitelaw, 2006. Uncovering the Risk-Return Relation in the Stock-Market. Journal of Finance 63, 1433–1463.

Harvey, A., E. Ruiz, N. Shephard, 1994. Multivariate stochastic variance models'. Review of Economic Studies 61, 247–264.

Harvey, A., N. Shephard, 1996. Estimation of an Asymmetric Stochastic Volatility Model for Asset Returns. Journal of Business and Economic Statistics 14, 429–434.

Harvey, C., 2001. The Specification of Conditional Expectations. Journal of Empirical Finance 8, 573–638.

Jacquier, E., N. G. Polson, P. E. Rossi, 1994. Bayesian Analysis of Stochastic Volatility Models. Journal of Business and Economic Statistics 12, 371–389.

——— , 2004. Bayesian Analysis of Stochastic Volatility Models with Fat-Tails and Correlated Errors. Journal of Econometrics 122, 185–212.

Kendall, M. G., 1954. Note on Bias in the Estimation of Autocorrelation. Biometrika 41, 403–404.

Kitagawa, G., 1987. Non-Gaussian State-Space Modeling of Nonstationary Time Series. Journal of the American Statistical Association 82, 1032–1041.

Leroy, S., R. Porter, 1981. The Present Value Relation: Tests Based on Variance Bounds. Econometrica 49, 555–577.

Litzenberger, R., K. Ramaswamy, 1979. The Effect of Personal Taxes and Dividends on Capital Asset Prices. Journal of Financial Economics 7, 163–196.

Lundblad, C., 2005. The Risk Return Tradeoff in the Long-Run: 1836-2004. Unpublished working paper. Kelley School of Business, Indianna University.

Merton, R. C., 1980. On Estimating the Expected Return on the Market: An Exploratory Investigation. Journal of Financial Economics 8, 323–361.

Nelson, D., 1991. Conditional Heteroscedasticity in Stock Returns: A New Approach. Econometrica 59, 347–370.

Ng, L., 1991. Tests of the CAPM with Time-Varying Covariances: A Multivariate GARCH Approach. Journal of Finance 46, 2507–1521.

Omori, Y., S. Chib, N. Shephard, J. Nakajima, 2007. Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference. Journal of Econometrics 140, 425–449.

Pastor, L., M. Sinha, B. Swaminathan, 2006. Estimating the Intertemporal Risk-Return Tradeoff Using the Implied Cost of Capital. Unpublished working paper. w11941 NBER.

Poterba, J. M., L. H. Summers, 1986. The Persistence of Volatility and Stock Market Fluctuations. American Economic Review 76, 1142–1151.

Sandmann, G., S. J. Koopman, 1998. Estimation of Stochastic Volatility Models via Monte Carlo Maximum Likelihood. Journal of Econometrics 87, 271–301.

Schwert, W. G., 1989. Why Does Stock Market Volatility Change Over Time?. Journal of Finance 44, 1115–1153.

Scruggs, J. T., 1998. Resoving the Puzzling Intertemporal Relationship Between the Market Risk Premium and Conditional Market Variance: A Two-Factor Approach. Journal of Finance 53, 575–603.

Shiller, R., 1981. Do Stock Prices Change Too Much to Be Justified by Subsequent Changes in Dividends. American Economic Review 71, 421–436.

Smith, D. R., 2002. Markov-Switching and Stochastic Volatility Diffusion Models of Short-Term Interest Rates. Journal of Business and Economic Statistics 20, 183–197.

Stambaugh, R. F., 1999. Return Predictability. Journal of Financial Economics 54, 375–421.

Taylor, S., 1986. Modelling Financial Time Series. John Wiley and Sons:, Chichester, UK.

Turner, C. M., R. Starz, C. R. Nelson, 1989. A Markov Model of Heteroskedasticity, Risk and Learning in the Stock Market. Journal of Financial Economics 25, 3–22.

Vuong, Q. H., 1989. Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses. Econometrica 57, 307–333.

Watanabe, T., 1999. A Nonlinear Filtering Approach to Stochastic Volatility Models with an Application to Daily Stock Returns. Journal of Applied Econometrics 16, 521–536.

Whitelaw, R., 2000. Stock Market Risk and Return: An Equilibrium Approach. Review of Financial Studies 13, 521–547.

Whitelaw, R. F., 1994. Time Variations and Covariations in the Expectation and Volatility of Stock Market Returns. Journal of Finance 49, 515–541.

Wu, G., 2001. The Determinants of Asymmetric Volatility. Review of Financial Studies 14, 837–859.

Yu, J., 2005. On Leverage in a Stochastic Volatility Model. Journal of Econometrics 127, 165–178.

Zhang, L., P. A. Mykland, Y. Aït-Sahalia, 2005. A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High Frequency Data. Journal of the American Statistical Association, Forthcoming.

Table 1: Parameter Estimates for the Stochastic Volatility Model: Monthly Data 1926-2004
We report the parameter estimates and robust standard errors (in parenthesis) for five models. For models 1 through 4, returns are modeled as $r_{t+1} \sim N(\mu_{t+1|t}, \sigma^2_{t+1})$ the conditional mean return on all models is given by $\mu_{t+1|t} = \mu + \gamma\sigma^2_{t+1}$, with log-volatility $x_{t+1} = \log(\sigma^2_{t+1})$ evolving as a latent AR(1) process $x_{t+1} = \omega + \phi x_t + \nu_t$ with $E(\nu_t^2) = \sigma^2_\nu$, $E((r_{t+1} - \mu_{t+1|t})(x_{t+2}) = \rho_{z\nu}\sigma_\nu\sigma_{t+1}$. Model 5 includes the feedback between returns and volatility. All parameters are estimated by maximum likelihood implemented using Gauss-Legendre numerical integration.

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5
μ	0.0085	0.0159	0.0085	0.0133	0.0059
	(0.0014)	(0.0027)	(0.0014)	(0.0029)	(0.0021)
γ	–	-4.6534	–	-3.4099	1.1900
		(1.4899)		(1.4690)	(0.6338)
ω	-0.2227	-0.2735	-0.3584	-0.3146	-0.2421
	(0.0726)	(0.0989)	(0.1097)	(0.1126)	(0.0862)
ϕ	0.9650	0.9572	0.9440	0.9507	0.9620
	(0.0115)	(0.0156)	(0.0173)	(0.0177)	(0.0137)
σ_η	0.1926	0.0453	0.0501	0.0525	0.0411
	(0.0275)	(0.0146)	(0.0153)	(0.0176)	(0.0122)
$\rho_{z\nu}$	–	–	-0.4325	-0.3329	-0.2995*
			(0.0885)	(0.1110)	
LL	1594.68	1601.65	1601.56	1605.04	1597.07
Feedback:	No	No	No	No	Yes

NOTES: $*$ the correlation coefficient for the feedback model is the unconditional expected correlation between returns and log volatility and is computed following equation (24).

Table 2: Parameter Estimates for the Stochastic Volatility Model: Monthly Data 1952-2004
We report the parameter estimates and robust standard errors (in parenthesis) for five models. For models 1 through 4, returns are modeled as $r_{t+1} \sim N(\mu_{t+1|t}, \sigma^2_{t+1})$ the conditional mean return on all models is given by $\mu_{t+1|t} = \mu + \gamma\sigma^2_{t+1}$, with log-volatility $x_{t+1} = \log(\sigma^2_{t+1})$ evolving as a latent AR(1) process $x_{t+1} = \omega + \phi x_t + \nu_t$ with $E(\nu_t^2) = \sigma_\nu^2$, $E((r_{t+1} - \mu_{t+1|t})(x_{t+2}) = \rho_{z\nu}\sigma_\nu\sigma_{t+1}$. Model 5 includes the feedback between returns and volatility. All parameters are estimated by maximum likelihood implemented using Gauss-Legendre numerical integration.

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5
μ	0.0076	0.0177	0.0076	0.0122	0.0008
	(0.0016)	(0.0042)	(0.0016)	(0.0044)	(0.0036)
γ	–	-7.4450	–	-4.0722	3.8747
		(2.9025)		(2.7739)	(1.8322)
ω	-0.4402	-0.5965	-0.7294	-0.6980	-0.4753
	(0.1401)	(0.2672)	(0.2371)	(0.3019)	(0.1978)
ϕ	0.9327	0.9094	0.8889	0.8936	0.9275
	(0.0215)	(0.0406)	(0.0364)	(0.0462)	(0.0304)
σ_η	0.0508	0.0699	0.0736	0.0818	0.0559
	(0.0158)	(0.0286)	(0.0240)	(0.0346)	(0.0173)
$\rho_{z\nu}$	–	–	-0.5891	-0.5178	-0.5336*
			(0.0940)	(0.1258)	
LL	1141.49	1147.15	1151.55	1153.19	1146.17
Feedback:	No	No	No	No	Yes

NOTES: $*$ the correlation coefficient for the feedback model is the unconditional expected correlation between returns and log volatility and is computed following equation (24).

Table 3: Parameter Estimates for the Stochastic Volatility Model: Daily Data 1962-2004
We report the parameter estimates and robust standard errors (in parenthesis) for five models. For models 1 through 4, returns are modeled as $r_{t+1} \sim N(\mu_{t+1|t}, \sigma^2_{t+1})$ the conditional mean return on all models is given by $\mu_{t+1|t} = \mu + \gamma\sigma^2_{t+1}$, with log-volatility $x_{t+1} = \log(\sigma^2_{t+1})$ evolving as a latent AR(1) process $x_{t+1} = \omega + \phi x_t + \nu_t$ with $E(\nu_t^2) = \sigma_\nu^2$, $E((r_{t+1} - \mu_{t+1|t})(x_{t+2}) = \rho_{z\nu}\sigma_\nu\sigma_{t+1}$. Model 5 includes the feedback between returns and volatility. All parameters are estimated by maximum likelihood implemented using Gauss-Legendre numerical integration.

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5
μ	0.0005	0.0008	0.0005	0.0005	0.0002
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
γ	–	-6.4955	–	-3.7600	3.5065
		(1.8325)		(1.7509)	(0.5004)
ω	-0.1364	-0.1426	-0.1509	-0.1416	-0.1316
	(0.0245)	(0.0258)	(0.0256)	(0.0252)	(0.0007)
ϕ	0.9862	0.9855	0.9847	0.9856	0.9867
	(0.0025)	(0.0026)	(0.0026)	(0.0026)	(0.0002)
σ_η	0.0223	0.0233	0.0229	0.0233	0.0213
	(0.0037)	(0.0039)	(0.0038)	(0.0039)	(0.0024)
$\rho_{z\nu}$	–	–	-0.3653	-0.3547	-0.3690*
			(0.0290)	(0.0289)	
LL	37048.67	37056.97	37116.17	37118.95	37109.80
Feedback:	No	No	No	No	Yes

NOTES: $*$ the correlation coefficient for the feedback model is the unconditional expected correlation between returns and log volatility and is computed following equation (24).

Table 4: Parameter Estimates for GARCH Model on Gross Returns: Monthly Data 1952-2004
In this table we report the parameter estimates and standard errors (in parenthesis) for a range of GARCH models. Model 1 and 2 are standard GARCH (Model 1) and GARCH-in-mean (Model 2), Models 5 and 6 allow the impact of positive and negative innovations to have different impacts on subsequent volatility (following Glosten, Jagannathan, and Runkle (1993), Model 3 and 4 are Campbell and Hentschel (1992) QGARCH models. In all models conditional volatility follows a GARCH process:

$$\sigma_{t+1}^2 = \beta_0 + (\beta_1 + \delta 1_{e_t<0})(e_t - b)^2 + \beta_2 \sigma_t^2$$

with $e_t = k_t - E_{t-1}k_t$ for the GARCH and GJR models and $e_t = \eta_{d,t}$ in the QGARCH model. For reference, in the stochastic volatility model with feedback effects fitted to total returns we find $\gamma = 4.2386$ (s.e. 2.0996) with log-likelihood 1149.5.

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
μ	0.0061	0.0061	0.0051	0.0029	0.0058	0.0058
	(0.0016)	(0.0044)	(0.0020)	(0.0022)	(0.0016)	(0.0044)
γ	–	-0.0025	0.8483	2.1435	–	-0.0236
		(2.7785)	(0.4872)	(1.1343)		(2.6076)
β_0	0.0018	0.0018	0.1059	0.1283	0.0016	0.0016
	(0.0003)	(0.0003)	(0.0418)	(0.0489)	(0.0002)	(0.0002)
β_1	0.0959	0.0959	0.1026	0.1226	-0.0520	-0.0519
	(0.0240)	(0.0293)	(0.0263)	(0.0212)	(0.0346)	(0.0373)
β_2	0.8389	0.8389	0.8364	0.7322	0.7851	0.7855
	(0.0315)	(0.0350)	(0.0325)	(0.0625)	(0.0497)	(0.0525)
δ	–	–	–	–	0.2295	0.2293
					(0.0659)	(0.0671)
b	–	–	–	0.0251	–	–
				(0.0074)		
LL	1134.09	1134.09	1143.92	1152.87	1143.01	1143.00
Feedback:	No	No	Yes	Yes	No	No

Table 5: Parameter Estimates for GARCH Model with Gross Returns: Daily Data 1962-2004
In this table we report the parameter estimates and standard errors (in parenthesis) for a range of GARCH models. Model 1 and 2 are standard GARCH (Model 1) and GARCH-in-mean (Model 2), Models 5 and 6 allow the impact of positive and negative innovations to have different impacts on subsequent volatility (following Glosten, Jagannathan, and Runkle (1993), Model 3 and 4 are Campbell and Hentschel (1992) QGARCH models. In all models conditional volatility follows a GARCH process:

$$\sigma_{t+1}^2 = \beta_0 + (\beta_1 + \delta 1_{e_t<0})(e_t - b)^2 + \beta_2 \sigma_t^2$$

with $e_t = k_t - E_{t-1}k_t$ for the GARCH and GJR models and $e_t = \eta_{d,t}$ in the QGARCH model. For reference, in the stochastic volatility model with feedback effects fitted to total returns we find $\gamma = 5.6246$ (s.e. 0.4796) with log-likelihood $37,421$.

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
μ	0.0006	0.0006	0.0006	0.0005	0.0005	0.0005
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
γ	–	0.0009	0.0194	0.4870	–	-0.0011
		(1.3427)	(0.0648)	(0.1239)		(1.5680)
β_0	0.0001	0.0001	0.0005	0.0002	0.0001	0.0001
	(0.0000)	(0.0000)	(0.0002)	(0.0001)	(0.0000)	(0.0000)
β_1	0.0806	0.0806	0.0830	0.0804	0.0421	0.0421
	(0.0102)	(0.0102)	(0.0110)	(0.0075)	(0.0046)	(0.0046)
β_2	0.9143	0.9142	0.9139	0.9061	0.9105	0.9105
	(0.0111)	(0.0112)	(0.0095)	(0.0085)	(0.0056)	(0.0057)
δ	–	–	–	–	0.0703	0.0703
					(0.0075)	(0.0075)
b	–	–	–	0.0031	–	–
				(0.0003)		
LL	37128.33	37128.33	37133.96	37246.21	37180.73	37180.73
Feedback:	No	No	Yes	Yes	No	No

Table 6: Bias in Estimated Persistence Using Realized Volatility.

For each month $t = 1, \ldots, T$ for $T = \{60, 120, 480, 960, 1920\}$ we simulate $M = \{22, 50, 100, 200\}$ normally distributed iid returns each with variance $\exp(x_{t+1})/M$ for each month, and model x_{t+1} as a log-AR(1) process

$$x_{t+1} = -0.4402 + 0.9327 \cdot x_t + 0.2254 \cdot z_t$$

for $z_t \sim N(0, 1)$. For each month we estimate the realized volatility as the sum of all M squared daily returns $\hat{\sigma}^2_{t+1}$. We then estimate an AR(1) (columns 2 and 3) and log-AR(1) (columns 4 and 5) process with the realized variances and report the average AR(1) parameter and R^2 across $N = 1000$ simulation. The first row (identified with $M = \infty$ uses the true variance $\exp(x_{t+1})$.

M	$\hat{\beta}$	R^2	$\hat{\beta}_{\log}$	$R^2_{\log}$
		T= 60		
∞			0.8973	0.8071
22	0.5442	0.3234	0.5805	0.3652
50	0.6726	0.4711	0.7050	0.5173
100	0.7466	0.5695	0.7772	0.6176
200	0.7994	0.6429	0.8246	0.6861
		T= 120		
∞			0.8973	0.8071
22	0.6169	0.3942	0.6643	0.4549
50	0.7406	0.5551	0.7780	0.6132
100	0.8066	0.6559	0.8368	0.7051
200	0.8355	0.7022	0.8622	0.7470
		T= 480		
∞			0.8973	0.8071
22	0.6840	0.4715	0.7271	0.5316
50	0.7946	0.6335	0.8285	0.6878
100	0.8464	0.7175	0.8728	0.7627
200	0.8763	0.7686	0.8980	0.8071
		T= 960		
∞			0.8973	0.8071
22	0.7009	0.4932	0.7374	0.5452
50	0.8049	0.6489	0.8358	0.6993
100	0.8555	0.7323	0.8798	0.7744
200	0.8838	0.7815	0.9040	0.8175
		T= 1920		
∞			0.8973	0.8071
22	0.7090	0.5036	0.7446	0.5551
50	0.8130	0.6617	0.8413	0.7081
100	0.8615	0.7425	0.8839	0.7814
200	0.8877	0.7883	0.9064	0.8217

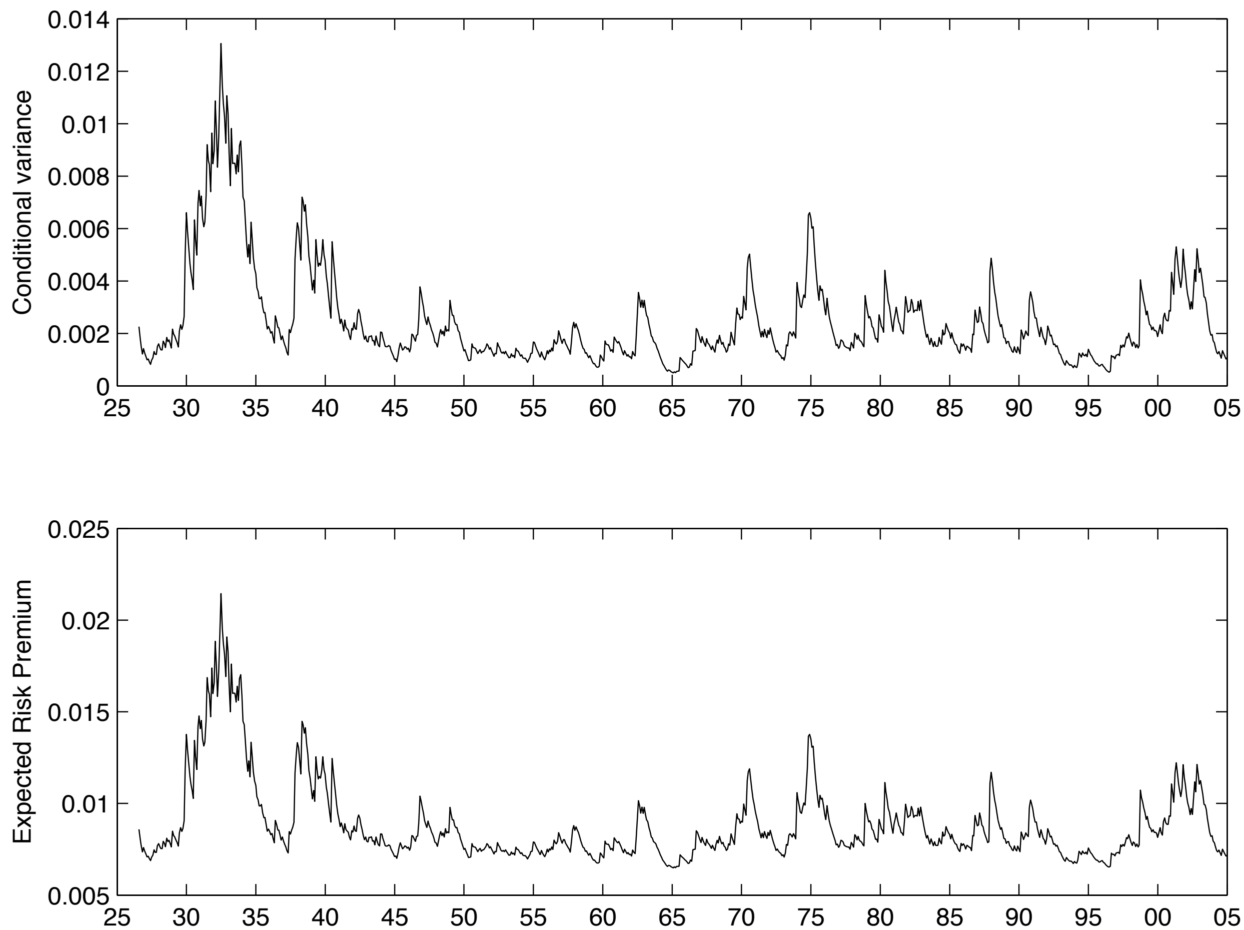

Figure 1: Conditional volatility and risk premium: Monthly Returns 1926-2004. The conditional variance of stock returns and the conditional expected return in excess of the risk-free rate is plotted.

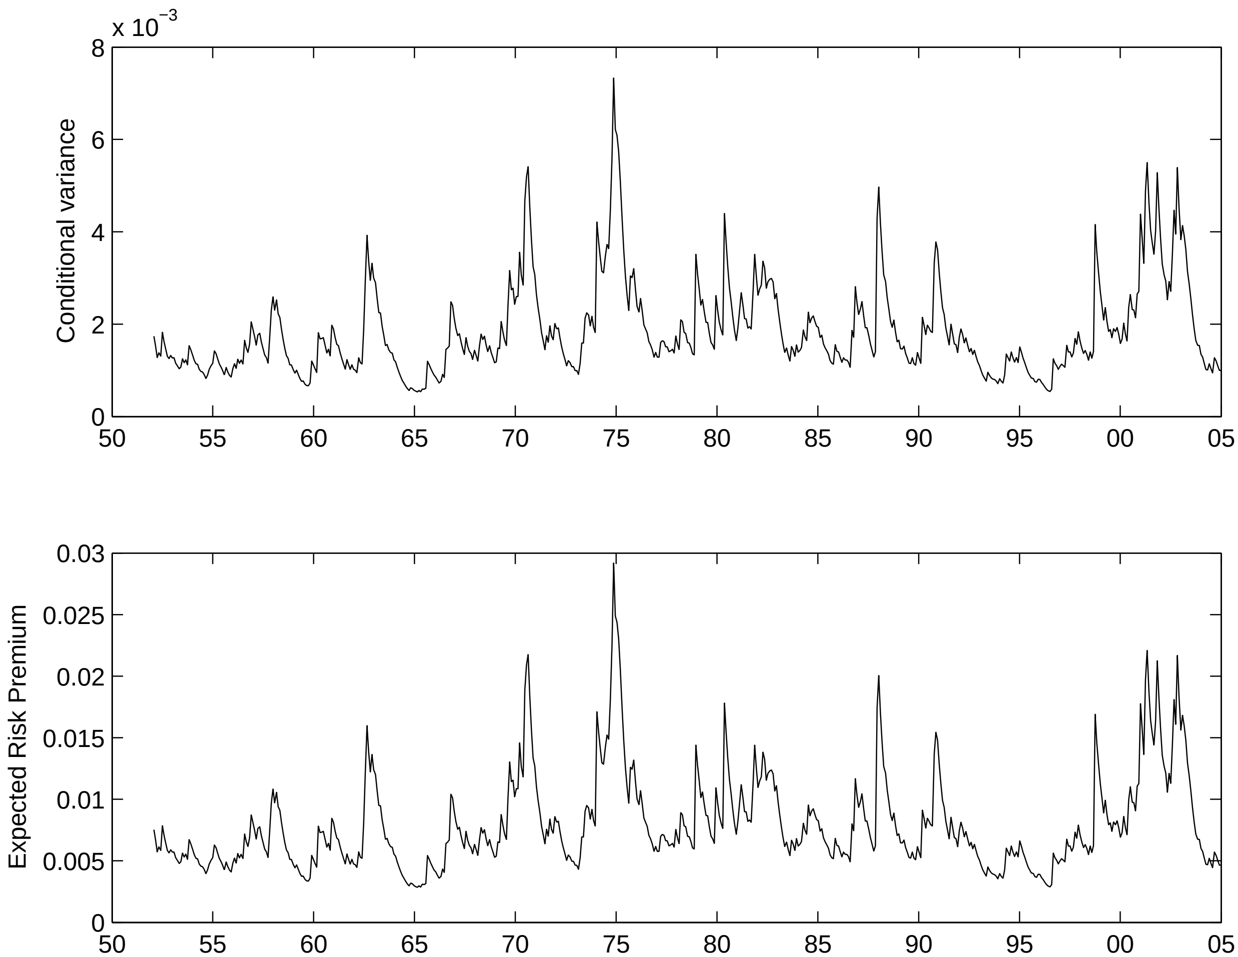

Figure 2: Conditional volatility and risk premium: Monthly Returns 1952-2004. The conditional variance of stock returns and the conditional expected return in excess of the risk-free rate is plotted.

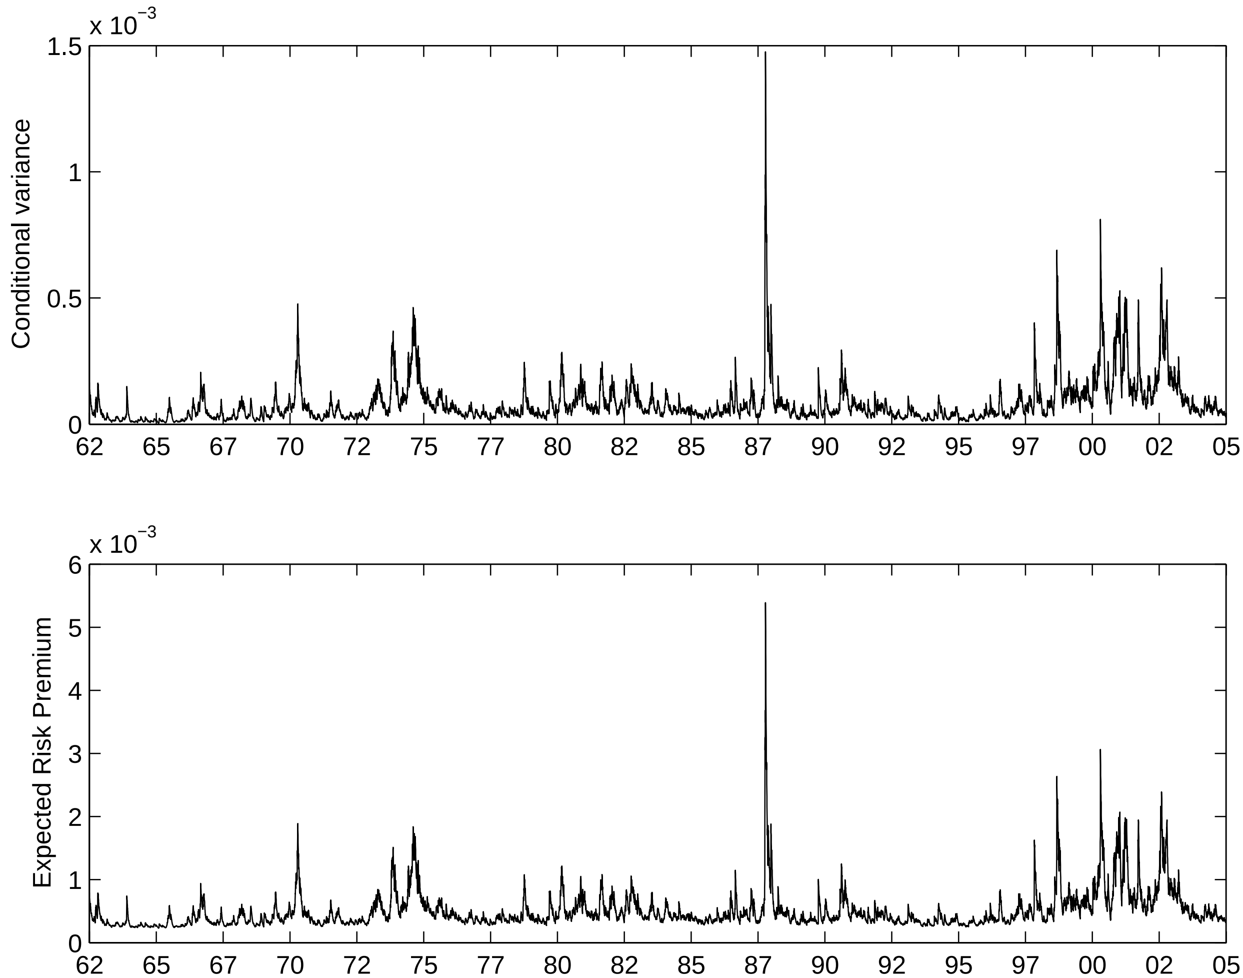

Figure 3: Conditional volatility and risk premium: Daily Returns 1962-2004. The conditional variance of stock returns and the conditional expected return in excess of the risk-free rate is plotted.

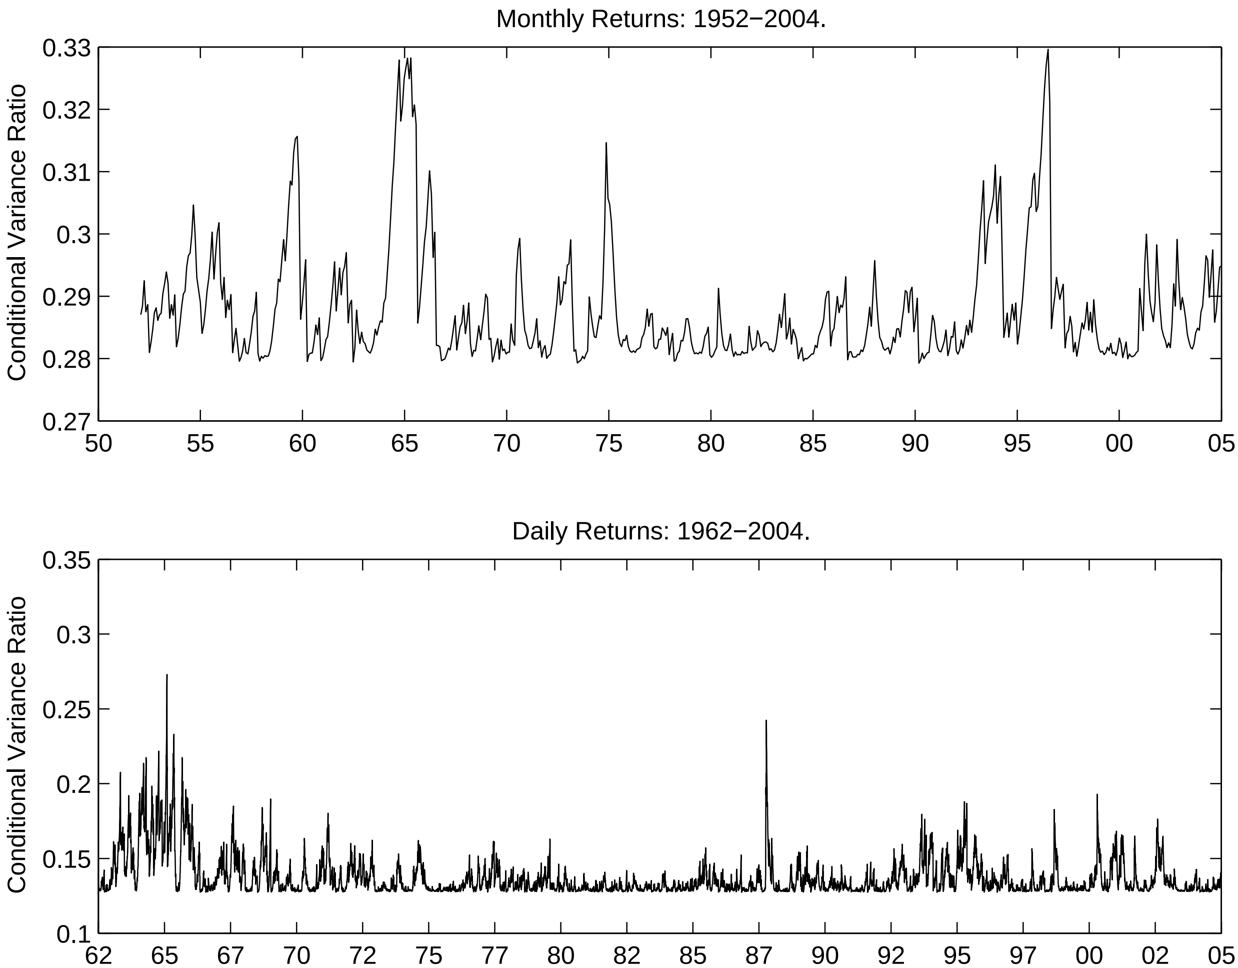

Figure 4: Fraction of return volatility due to volatility feedback. We plot the fraction of total conditional volatility that is due to volatility feedback is calculated as in equation (20).

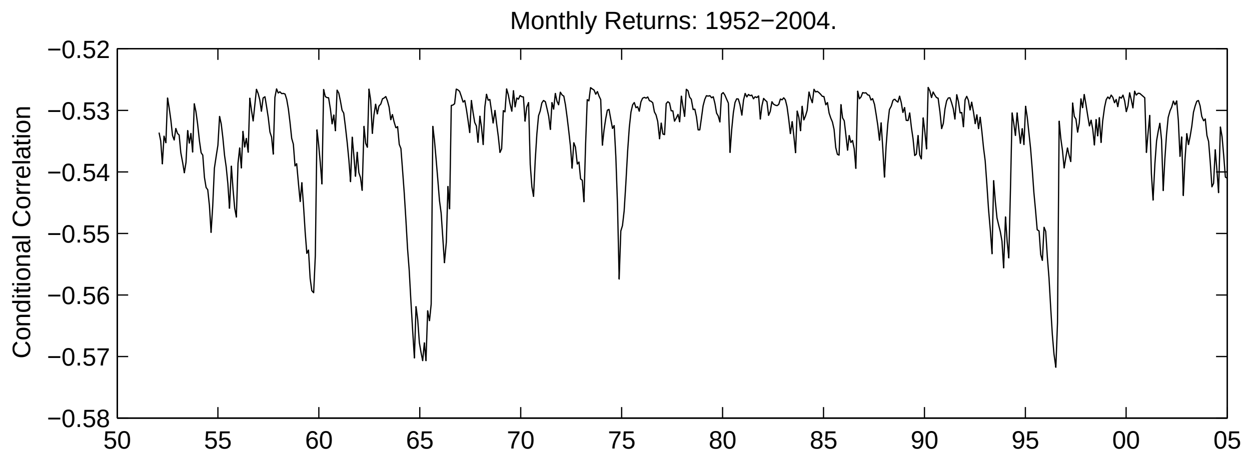

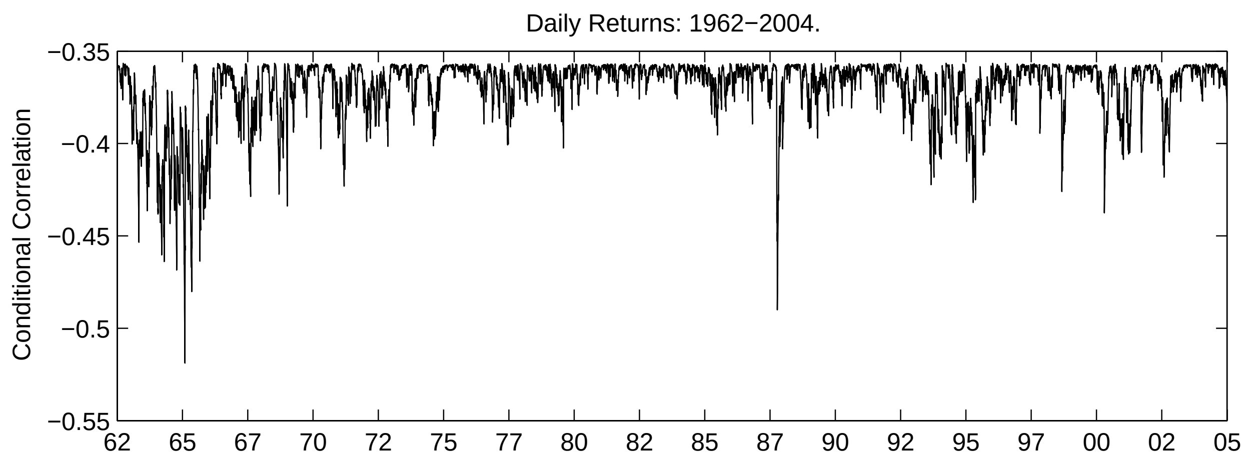

Figure 5: Plot of the feedback-induced conditional correlation between returns and volatility. We plot the conditional correlation between returns and future log-volatility induced by the feedback effect following (22).

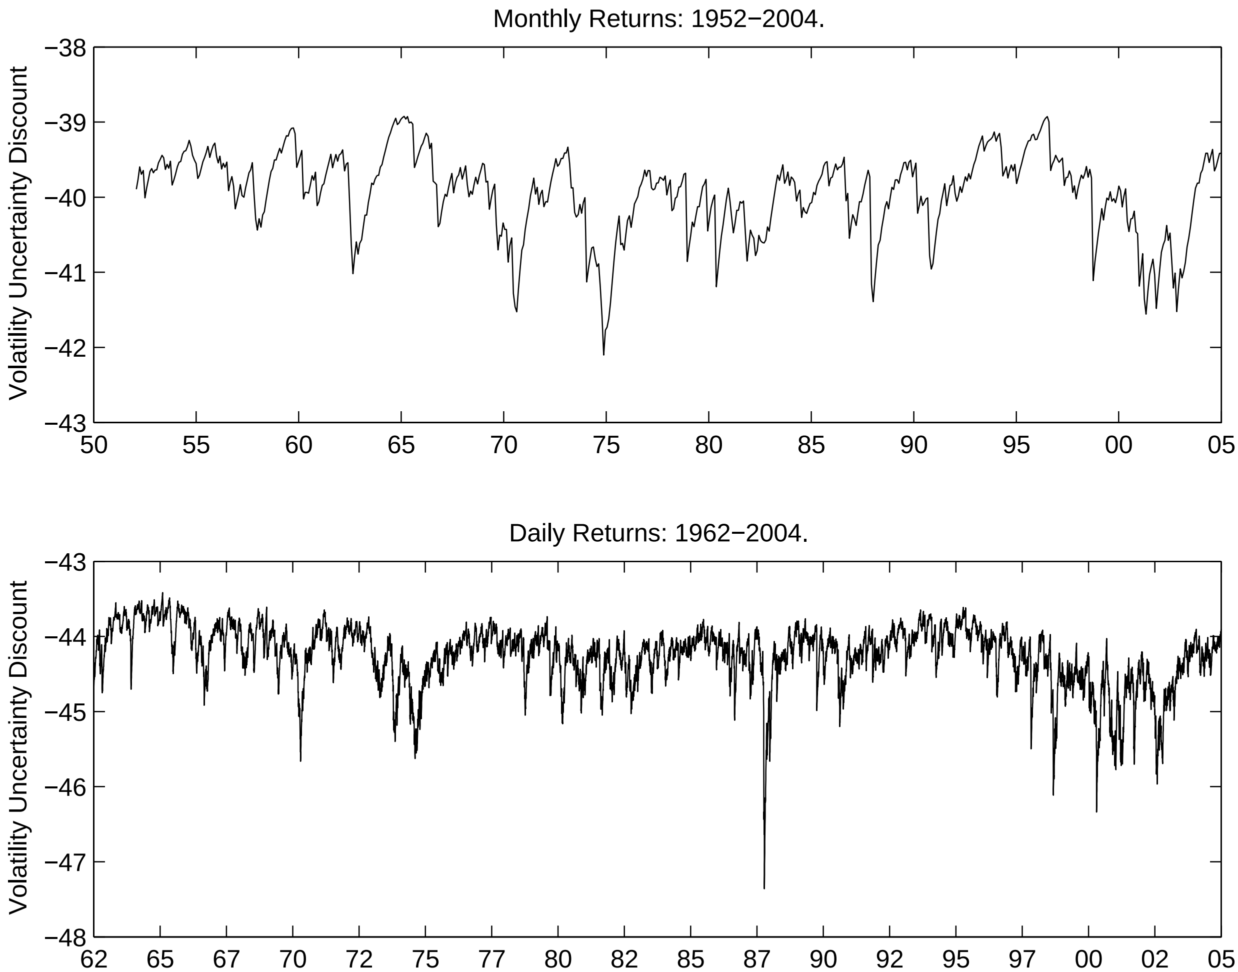

Figure 6: Historical Volatility Uncertainty Discount (in percent). The volatility discount ζ_t (measured in percent) is defined in equation (30) and represents the increase in the level of log-stock prices that would occur if there was no uncertainty about stock return volatility (i.e., $\sigma_\eta^2 = 0$).