Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities*

Tim Bollerslev†
Michael Gibson‡
Hao Zhou§

First Draft: September 2004
This Version: January 2005

*The work of Bollerslev was supported by a grant from the National Science Foundation to the NBER. We would like to thank Nellie Liang and Patricia White for useful discussions during this project. The views presented here are solely those of the authors and do not necessarily represent those of the Federal Reserve Board or its staff. Matthew Chesnes provided excellent research assistance.

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

‡Division of Research and Statistics, Federal Reserve Board, Mail Stop 91, Washington DC 20551 USA, E-mail michael.s.gibson@frb.gov, Phone 202-452-2495, Fax 202-728-5887.

§Division of Research and Statistics, Federal Reserve Board, Mail Stop 91, Washington DC 20551 USA, E-mail hao.zhou@frb.gov, Phone 202-452-3360, Fax 202-728-5887.

Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities

Abstract

This paper proposes a method for constructing a volatility risk premium, or investor risk aversion, index. The method is intuitive and simple to implement, relying on the sample moments of the recently popularized model-free realized and option-implied volatility measures. A small-scale Monte Carlo experiment confirms that the procedure works well in practice. Implementing the procedure with actual S&P500 option-implied volatilities and high-frequency five-minute-based realized volatilities indicate significant temporal dependencies in the estimated stochastic volatility risk premium, which we in turn relate to a set of underlying macro-finance state variables. We also find that the extracted volatility risk premium helps predict future stock market returns.

JEL Classification: G12, G13, C51, C52.

Keywords: Stochastic Volatility Risk Premium, Model-Free Implied Volatility, Model-Free Realized Volatility, Black-Scholes, GMM Estimation, Return Predictability.

1 Introduction

Model-free volatility measures have figured prominently in the recent academic and financial market practitioner literatures. On one hand, several studies have argued for the use of so-called "model-free realized volatilities" computed by summing squared returns from high-frequency data over short time intervals during the trading day. As demonstrated in the literature, these types of measures afford much more accurate ex-post observations of the actual volatility than the more traditional sample variances based on daily or coarser frequency data (Andersen et al., 2001; Barndorff-Nielsen and Shephard, 2002; Meddahi, 2002; Andersen et al., 2003a,b; Barndorff-Nielsen and Shephard, 2004a; Andersen et al., 2004). On the other hand, the recently developed so-called "model-free implied volatilities" provide ex-ante *risk-neutral* expectations of the future volatilities. Importantly, and in contrast to more traditional option-implied volatilities based on the Black-Scholes pricing formula or some variant thereof, the model-free implied volatilities are computed from option prices without the use of any particular option-pricing model (Carr and Madan, 1998; Demeterfi et al., 1999; Britten-Jones and Neuberger, 2000; Lynch and Panigirtzoglou, 2003; Jiang and Tian, 2004; Carr and Wu, 2004).[1] In this paper, we combine these two new volatility measures to improve on existing estimates of the risk premium associated with stochastic volatility risk and investor risk aversion.

Because the method we present here directly uses the model-free realized and implied volatilities to extract the stochastic volatility risk premium, it is much easier to implement than other methods which rely on the joint estimation of both the underlying asset return and the price(s) of one or more of its derivatives, leading to quite complicated modeling and estimation procedures (see, e.g., Bates, 1996; Chernov and Ghysels, 2000; Jackwerth, 2000; Aït-Sahalia and Lo, 2000; Benzoni, 2002; Pan, 2002; Eraker, 2004, among many others). In contrast, the method of this paper relies on standard GMM estimation of the cross conditional moments between risk-neutral and objective expectations of integrated volatility to identify the stochastic volatility risk premium. As such, the method is simple to implement

[1]Market participants have also recently developed several new products – realized variance futures, VIX futures, and over-the-counter (OTC) variance swaps – that are based on these two model-free volatility measures. Specifically, the Chicago Board Option Exchange (CBOE) recently changed its implied volatility index (VIX) to use the model-free implied volatility approach and the more popular S&P500 index options (CBOE Documentation, 2003), while the CBOE Futures Exchange began to trade futures on the VIX on March 26, 2004 and realized variance futures on the S&P500 on May 18, 2004. Demeterfi et al. (1999) discuss OTC variance swaps.

and can easily be extended to allow for a time-varying volatility risk premium. Indeed, one feature of our estimation strategy is that it afford a characterization of the temporal variation in the volatility risk premium, or investor risk aversion, possibly driven by a set of economic state variables.[2]

To validate the performance of the new estimation strategy, we perform a small scale Monte Carlo experiment focusing directly on our ability to precisely estimate the risk premium parameter. While the estimation strategy applies generally, the Monte Carlo study focuses on the popular Heston (1993) stochastic volatility model. The results confirm that using model-free implied volatility from options with one month to maturity and realized volatility from five-minute returns, we can estimate the volatility risk premium nearly as well as if we were using the actual (unobserved and infeasible) risk-neutral implied volatility and continuous time integrated volatility. However, using Black-Scholes implied volatility and/or realized volatility from daily returns generally results in biased and (highly) inefficient estimates of the risk premium parameter and corresponding unreliable statistical inference.

To illustrate the procedure empirically, we apply the method in estimating the volatility risk premium associated with the S&P500 market index. We extend the method to allow for two types of time variation in the stochastic volatility risk premium. In the first, the premium follows a specific autoregressive process. In the second, it varies over time with other macro-finance variables. We find statistically significant effects on the volatility risk premium from several macro-finance variables, including the market volatility itself, the price-earnings (P/E) ratio of the market, a measure of credit spread, industrial production, the producer price index, and nonfarm employment.[3] Our results give structure to the intuitive notion that the difference between implied and realized volatilities reflects a volatility risk premium that responds to economic state variables. As such, our findings should be of direct interest to market participants and monetary policymakers alike concerned with the links between

[2]The general strategy developed here is also related to the literature on market implied risk aversion (see, Jackwerth, 2000; Aït-Sahalia and Lo, 2000; Rosenberg and Engle, 2002; Brandt and Wang, 2003; Bliss and Panigirtzoglou, 2004; Gordon and St-Amour, 2004, e.g.). The closest to ours is arguably that of Garcia et al. (2001), who estimate jointly the risk-neutral and objective dynamics, using a series expansion of option implied volatilities around the Black-Scholes formula.

[3]For directly traded assets like equities or bonds, the links between the risk premium—expected excess return—and macro-finance state variables are already well established. For example, the equity risk premium is predicted by the dividend–price ratio and short-term interest rates (see, Campbell, 1987; Fama and French, 1988; Campbell and Shiller, 1988a,b, e.g.), while bond risk premia may be predicted by forward rates (see, Fama and Bliss, 1987; Cochrane and Piazzesi, 2004, e.g.). However, with the notable exception of the recent study by (Carr and Wu, 2004), academic studies on the behavior of the volatility risk premium are rare, let alone its linkage to the overall economy.

the financial markets and the overall economy.[4] Further strengthening our results, we also find that the estimated time-varying volatility risk premium better predicts future stock market returns than several other well-established predictor variables.

The rest of the paper is organized as follows. Section 2 outlines the basic theory behind our simple GMM estimation procedure, while Section 3 provides finite sample simulation evidence on the performance of the estimator. Section 4 applies the estimator to the S&P500 market index, explicitly linking the temporal variation in the volatility risk premium to a set of underlying macro-finance variables. This section also documents our findings related to return predictability. Section 5 concludes.

2 Identification and Estimation

Consider the general continuous-time stochastic volatility model for the logarithmic stock price process ($p_t = \log S_t$),

$$\begin{aligned} dp_t &= \mu_t(\cdot)dt + \sqrt{V_t}dB_{1t}, \\ dV_t &= \kappa(\theta - V_t)dt + \sigma_t(\cdot)dB_{2t}, \end{aligned} \tag{1}$$

where the instantaneous $\text{corr}(dB_{1t}, dB_{2t}) = \rho$ denotes the familiar leverage effect, and the functions $\mu_t(\cdot)$ and $\sigma_t(\cdot)$ must satisfy the usual regularity conditions. Assuming no arbitrage and a linear volatility risk premium, the corresponding risk-neutral distribution then takes the form

$$\begin{aligned} dp_t &= r_t^* dt + \sqrt{V_t}dB_{1t}^*, \\ dV_t &= \kappa^*(\theta^* - V_t)dt + \sigma_t(\cdot)dB_{2t}^*, \end{aligned} \tag{2}$$

where $\text{corr}(dB_{1t}^*, dB_{2t}^*) = \rho$, and r_t^* denotes the risk-free interest rate. Importantly, the risk-neutral parameters in (2) are directly related to the parameters of the actual price process in equation (1) by the relationships, $\kappa^* = \kappa + \lambda$ and $\theta^* = \kappa\theta/(\kappa + \lambda)$, where λ refers to the volatility risk premium parameter of interest. Note that the functional forms of $\mu_t(\cdot)$ and $\sigma_t(\cdot)$ are completely flexible as long as they avoid arbitrage.

2.1 Model-Free Volatility Measures and Moment Restrictions

The point-in-time volatility V_t entering the stochastic volatility model above is latent and its consistent estimation through filtering is complicated by a host of market microstructure

[4] See e.g., Tarashev et al. (2003) and Liang and Zhou (2003) for a discussion from the perspective of central bank policy makers.

complications. Alternatively, the model-free realized volatility measures afford a simple approach for quantifying the integrated volatility over non-trivial time intervals. In our notation, let $\mathcal{V}^n_{t,t+\Delta}$ denote the realized volatility computed by summing the squared high-frequency returns over the $[t, t+\Delta]$ time-interval:

$$\mathcal{V}^n_{t,t+\Delta} \equiv \sum_{i=1}^{n} \left[p_{t+\frac{i}{n}(\Delta)} - p_{t+\frac{i-1}{n}(\Delta)} \right]^2 \tag{3}$$

It follows then by the theory of quadratic variation (see, e.g., Andersen et al. (2003a), for a recent survey of the realized volatility literature),

$$\lim_{n\to\infty} \mathcal{V}^n_{t,t+\Delta} \xrightarrow{a.s.} \mathcal{V}_{t,t+\Delta} \equiv \int_t^{t+\Delta} V_s \, ds \tag{4}$$

In other words, when n is large relative to Δ, the realized volatility should be a good approximation for the unobserved integrated volatility $\mathcal{V}_{t,t+\Delta}$.[5]

Moments for the integrated volatility for the model in (1) have previously been derived by Bollerslev and Zhou (2002) (see also Meddahi (2002) and Andersen et al. (2004)). In particular, it follows that the first conditional moment under the physical measure satisfies

$$\mathrm{E}(\mathcal{V}_{t+\Delta,t+2\Delta}|\mathcal{F}_t) = \alpha_\Delta \, \mathrm{E}(\mathcal{V}_{t,t+\Delta}|\mathcal{F}_t) + \beta_\Delta \tag{5}$$

where the coefficients $\alpha_\Delta = e^{-\kappa\Delta}$ and $\beta_\Delta = \theta\left(1 - e^{-\kappa\Delta}\right)$ are functions of the underlying parameters κ and θ of (1).

Using option prices, it is also possible to construct a model-free measure of the risk-neutral expectation of the integrated volatility. In particular, let $\mathrm{IV}^*_{t,t+\Delta}$ denote the time t implied volatility measure computed as a weighted average, or integral, of a continuum of Δ-maturity options,

$$\mathrm{IV}^*_{t,t+\Delta} = 2 \int_0^\infty \frac{C(t+\Delta, K) - C(t, K)}{K^2} \, dK \tag{6}$$

where $C(t, K)$ denotes the price of a European call option maturing at time t with strike price K. As formally shown by Britten-Jones and Neuberger (2000), this model-free implied volatility then equals the true risk-neutral expectation of the integrated volatility,[6]

$$\mathrm{IV}^*_{t,t+\Delta} = \mathrm{E}^* \left(\mathcal{V}_{t,t+\Delta} | \, \mathcal{F}_t \right), \tag{7}$$

[5] The asymptotic distribution (for $n \to \infty$ and Δ fixed) of the realized volatility error has been formally characterized by Barndorff-Nielsen and Shephard (2002) and Meddahi (2002). Also, Barndorff-Nielsen and Shephard (2004b) have recently extended these asymptotic distributional results to allow for leverage effects.

[6] Carr and Madan (1998) and Demeterfi et al. (1999) have previously derived a closely related expression.

where $\mathrm{E}^*(\cdot)$ refers to the expectation under the risk-neutral measure. Although the original derivation of this important result in Britten-Jones and Neuberger (2000) assumes that the underlying price path is continuous, this same result has recently been extended by Jiang and Tian (2004) to the case of jump diffusions. Moreover, Jiang and Tian (2004) also demonstrates that the integral in the formula for $\mathrm{IV}^*_{t,t+\Delta}$ may be accurately approximated from a finite number of options in empirically realistic situations.

Combining these results, it now becomes possible to directly and analytically link the expectation of the integrated volatility under the risk-neutral dynamics in (2) with the objective expectation of the integrated volatility under (1). As formally shown by Bollerslev and Zhou (2004),

$$\mathrm{E}\left(\mathcal{V}_{t,t+\Delta} | \mathcal{F}_t\right) = \mathcal{A}_\Delta \mathrm{IV}^*_{t,t+\Delta} + \mathcal{B}_\Delta, \tag{8}$$

where $\mathcal{A}_\Delta = \frac{(1-e^{-\kappa\Delta})/\kappa}{(1-e^{-\kappa^*\Delta})/\kappa^*}$ and $\mathcal{B}_\Delta = \theta[\Delta - (1-e^{-\kappa\Delta})/\kappa] - \mathcal{A}_\Delta \theta^*[\Delta - (1-e^{-\kappa^*\Delta})/\kappa^*]$ are functions of the underlying parameters κ, θ, and λ. This equation, in conjunction with the moment restriction in (5), provides the necessary identification of the risk premium parameter, λ.[7]

2.2 GMM Estimation and Statistical Inference

Using the moment conditions (5) and (8), we can now construct a standard GMM type estimator. To allow for overidentifying restrictions, we augment the moment conditions with a lagged instrument of realized volatility, resulting in the following four dimensional system of equations:

$$f_t(\xi) = \begin{bmatrix} \mathcal{V}_{t+\Delta,t+2\Delta} - \alpha_\Delta \mathcal{V}_{t,t+\Delta} - \beta_\Delta \\ (\mathcal{V}_{t+\Delta,t+2\Delta} - \alpha_\Delta \mathcal{V}_{t,t+\Delta} - \beta_\Delta)\mathcal{V}_{t-\Delta,t} \\ \mathcal{V}_{t,t+\Delta} - \mathcal{A}_\Delta \mathrm{IV}^*_{t,t+\Delta} - \mathcal{B}_\Delta \\ (\mathcal{V}_{t,t+\Delta} - \mathcal{A}_\Delta \mathrm{IV}^*_{t,t+\Delta} - \mathcal{B}_\Delta)\mathcal{V}_{t-\Delta,t} \end{bmatrix} \tag{9}$$

where $\xi = (\kappa, \theta, \lambda)'$. By construction $\mathrm{E}[f_t(\xi_0)|\mathcal{G}_t] = 0$, and the corresponding GMM estimator is defined by $\hat{\xi}_T = \arg\min g_T(\xi)' W g_T(\xi)$, where $g_T(\xi)$ refers to the sample mean of the

[7] When implementing the conditional moment restrictions (5) and (8), it is useful to distinguish between two information sets—the continuous sigma-algebra $\mathcal{F}_t = \sigma\{V_s; s \le t\}$, generated by the point-in-time volatility process, and the discrete sigma-algebra $\mathcal{G}_t = \sigma\{\mathcal{V}_{t-s-1,t-s}; s = 0,1,2,\cdots,\infty\}$, generated by the integrated volatility series. Obviously, the coarser filtration is nested in the finer filtration (i.e., $\mathcal{G}_t \subset \mathcal{F}_t$), and by the Law of Iterated Expectations, $\mathrm{E}[\mathrm{E}(\cdot|\mathcal{F}_t)|\mathcal{G}_t] = \mathrm{E}(\cdot|\mathcal{G}_t)$. The GMM estimation method implemented later is based on the coarser information set $\mathcal{G}_t$.

moment conditions, $g_T(\xi) \equiv 1/T \sum_{t=2}^{T-2} f_t(\xi)$, and W denotes the asymptotic covariance matrix of $g_T(\xi_0)$ (Hansen, 1982). Under standard regularity conditions, the minimized value of the objective function $J = \min_\xi g_T(\xi)' W g_T(\xi)$ multiplied by the sample size should be asymptotically chi-square distributed, allowing for an omnibus test of the overidentifying restrictions. Moreover, inference concerning the individual parameters is readily available from the standard formula for the asymptotic covariance matrix, $(\partial f_t(\xi)/\partial \xi' W \partial f_t(\xi)/\partial \xi)/T$. Further, since the lag structure in the moment conditions in equations (5) and (8) entails a complex dependence, we use a heteroscedasticity and autocorrelation consistent robust covariance matrix estimator with a Bartlett-kernel and a lag length of five in implementing the estimator (Newey and West, 1987).

3 Finite Sample Distributions

3.1 Monte Carlo Design

To determine the finite sample performance of the GMM estimator based on the moment conditions described above, we conducted a small scale Monte Carlo study for the specialized Heston (1993) version of the model in (1) and (2) with $\sigma_t(\cdot) = \sigma\sqrt{V_t}$. To illustrate the advantage of the new model-free volatility measures, we estimated the model using three different implied volatilities:

1. **RNIV:** risk-neutral expectation of integrated volatility (this is, of course, not observable in practice but can be calculated inside the simulations where we know both the latent volatility state V_t and the risk neutral parameters κ^* and θ^*);

2. **MFIV:** model-free implied volatility computed from one-month maturity option prices using a truncated and discretized version of equation (6);

3. **BSIV:** Black-Scholes implied volatility from a one-month maturity, at-the-money option as a (misspecified) proxy for RNIV.

We also use three different realized volatility measures to assess how the mis-measurement of realized volatility affects the estimation:

1. **Integrated Volatility:** The monthly true integrated volatility $\int_t^{t+\Delta} V_s \, ds$ (again, this is not observable in practice but can be calculated inside the simulations);

2. **Realized Volatility, 5-minute:** monthly realized volatilities computed from five-minute returns;

3. **Realized Volatility, daily:** monthly realized volatilities computed from daily returns.

The dynamics of (1) are simulated with the Euler method. We calculate model-free implied volatility for a given level of V_t with the discrete version of (6) presented by Jiang and Tian (2004). The call option prices needed to compute model-free implied volatility are computed with the Heston (1993) formula. The Black-Scholes implied volatility is generated by calculating the price of an at-the-money call and then inverting the Black-Scholes formula to extract the implied volatility.

The accuracy of the asymptotic approximations are illustrated by contrasting the results for sample sizes of 150 and 600. The total number of Monte Carlo replications is 500. To focus on the volatility risk premium, the drift of the stock return in (1) and the risk-free rate in (2) are both set equal to zero. The benchmark scenario is labeled (a) and sets $\kappa = 0.10$, $\theta = 0.25$, $\sigma = 0.10$, $\lambda = -0.20$, $\rho = -0.50$. Three additional variations we consider are (b) high volatility persistence, or $\kappa = 0.03$; (c) high volatility-of-volatility, or $\sigma = 0.20$; and (d) pronounced leverage, or $\rho = -0.80$.[8]

3.2 Simulation Results

Tables 1-3 summarize the parameter estimation for the volatility risk premium. The use of model-free implied volatility (MFIV) achieves a similar root-mean-squared error (RMSE) and convergence rate as the true infeasible risk-neutral implied volatility (RNIV). On the other hand, the misspecified Black-Scholes implied volatility (BSIV) shows slow convergence in estimating the volatility risk premium. Also, using realized volatility from five-minute returns (over a monthly horizon) has virtually the same small bias and high efficiency as the estimates based on the (infeasible) integrated volatility. In contrast, using the realized volatility from daily returns generally results in a larger bias and significantly lower efficiency.

Figures 1-3 report the Wald test for the risk premium parameter, which should be asymptotically $\mathcal{X}^2(1)$ distributed. In the cases of (infeasible) integrated volatility and five-minute realized volatility, the test statistics for the MFIV and RNIV measures are generally indistinguishable and closely approximated by the asymptotic distribution, the only exception

[8] The first three designs are the same as in Bollerslev and Zhou (2002), and the estimation results for the κ and θ parameters (available upon request) mirror the results reported therein based on the moment conditions for the model in (1) only.

being the high volatility persistence scenario (b) for which the MFIV measure results in slight over-rejection. In contrast, the (misspecified) BSIV measure is clearly biased for all of the different scenarios. When the realized volatility is constructed from daily squared returns, the Wald test systematically loses power to detect any misspecification, and the RNIV and MFIV measures now both show some under-rejection bias.[9]

In a sum, the Monte Carlo results clearly demonstrate that it is possible to accurately estimate the volatility risk premium from the model-free implied volatilities and the five-minute based realized volatilities. On the other hand, the use of Black-Scholes implied volatilities and/or realized volatilities from daily squared returns both produce biased and inefficient estimates, and generally do not allow for reliable inference concerning the true value of the risk premium parameter.

4 Estimates for the Market Volatility Risk Premium

4.1 Time-Varying Volatility Risk Premium and Relative Risk Aversion

There is an intimate link between the stochastic volatility risk premium and the coefficient of risk aversion for the representative investor within the standard intertemporal asset pricing framework. In particular, assuming a linear volatility risk premium along with an affine version of the stochastic volatility model corresponding to $\sigma_t(\cdot) = \sigma\sqrt{V_t}$ in (1), as in Heston (1993), it follows that $-\lambda V_t = \text{cov}_t\left(\frac{dm_t}{m_t}, dV_t\right)$, where m_t denotes the pricing kernel, or marginal utility of wealth for the representative investor. Moreover, assuming that the representative agent has a power utility function

$$U_t = e^{-\delta t}\frac{W_t^{1-\gamma}}{1-\gamma}, \tag{10}$$

where δ denotes a constant subjective time discount rate, and that in equilibrium the agent holds the market portfolio. The marginal utility then equals $m_t = e^{-\delta t}W_t^{-\gamma}$, and it follows

[9] The GMM omnibus test also has the correct size for the RNIV and MFIV measures, but often cannot reject for the misspecified BSIV. This is because even for BSIV the objective moment (5) is still correctly specified, only the cross moment (8) is misspecified. These additional graphs are omitted to conserve space but available upon request.

by Itô's formula that[10]

$$-\lambda V_t = \text{cov}_t\left(\frac{dm_t}{m_t}, dV_t\right) = -\gamma\rho\sigma V_t. \tag{11}$$

Thus, in this situation the constant relative risk aversion coefficient is directly proportional to the volatility risk premium $\gamma = \lambda/(\rho\sigma)$. Moreover, given the estimated values of $\rho = -0.8$ and $\sigma = 1.2$ for the S&P500 data analyzed below, $-\lambda$ is approximately equal to the representative investor's risk aversion, γ.

Meanwhile, a number of studies have argued that the assumption of constant risk aversion, or by the equivalence discussed above a constant volatility risk premium parameter, is too restrictive for satisfactorily describing asset return dynamics.[11] The development of a formal preference-based model for explaining temporal variation in the risk aversion coefficient is beyond the scope of the present paper. Instead, suppose simply that the utility function for the representative investor may be expressed as

$$U_t = e^{-\delta t}\frac{W_t^{1-\gamma_t}}{1-\gamma_t}, \tag{12}$$

where γ_t now represents a possibly *time-varying* relative risk aversion coefficient. Moreover, assume that the evolution in γ_t may be described by the separate diffusion process,

$$d\gamma_t = \mu(\gamma_t)dt + \sigma(\gamma_t)dB_{3t}, \tag{13}$$

where importantly the preference shocks are exogenous, in the sense that the dB_{3t} innovation process is uncorrelated with the two Brownian motions driving the log-price and volatility processes, dB_{1t} and dB_{2t}, respectively. On applying Itô's formula, it follows then by similar arguments to the ones in Gordon and St-Amour (2004) that in equilibrium

$$-\lambda_t V_t = \text{cov}_t\left(\frac{dm_t}{m_t}, dV_t\right) = -\gamma_t\rho\sigma V_t. \tag{14}$$

In particular, the no-arbitrage requirement implies the following modification to the risk-neutral distribution for the volatility in equation (2),

$$dV_t = \kappa_t^*(\theta_t^* - V_t)dt + \sigma_t(\cdot)dB_{2t}^*, \tag{15}$$

[10] A similar argument is made by Bakshi and Kapadia (2003).

[11] Constant relative risk aversion is also not consistent with more general utility functions, like the habit persistence model of Campbell and Cochrane (1999) or the relative social status model of Bakshi and Chen (1996).

where now $\kappa_t^* = \kappa + \lambda_t$ and $\theta_t^* = \kappa\theta/(\kappa + \lambda_t)$.[12] This expression directly motivates our estimation of a time-varying volatility risk premium λ_t, or equivalently a time-varying risk aversion coefficient, $\gamma_t = \lambda_t/(\rho\sigma)$. With the caveat that more generally this equivalence is at best an approximation, we will continue to use the phrases volatility risk premium and investor risk aversion interchangeably in the following discussion.

4.2 Empirical Approximations for the Volatility Risk Premium

The discussion in the previous section suggests that the volatility risk premium is likely time-varying. We will explore two different empirical approximations which explicitly link the period-by-period variation in the estimated risk premium parameter through a simple dynamic model for λ_{t+1}.

Our first approximation is based on the AR(1) model

$$\lambda_{t+1} = a + b\lambda_t + cu_t, \tag{16}$$

where we allow the time-variation in the risk premium to be driven by the fitted error in the cross moment between the realized and implied volatility, $u_t = \mathcal{V}_{t-1,t} - \mathcal{A}_1 \mathrm{IV}^*_{t-1,t} - \mathcal{B}_1$. This formulation has a precedent in ARCH-GARCH type modeling, where the shock to the volatility equation comes from the previous period's mean equation error. Importantly, it is also consistent with the no-arbitrage requirement that the preference parameter λ_{t+1} be a pre-determined function of time-t state variables. To identify the additional two parameters a and b, we add the lag squared realized volatility as an instrument to the moment conditions in (5) and (8), leaving us with the same single degree of freedom for the chi-square omnibus test.

Our second approach explores whether the observed difference between the implied and realized volatility can further be explained by a set of macro-finance state variables in a manner consistent with the underlying theoretical option pricing framework. Specifically, we approximate the volatility risk premium parameter as the following augmented AR(1) process,

$$\lambda_{t+1} = a + b\lambda_t + \sum_{k}^{K} c_k \times \mathrm{state}_{t,k} \tag{17}$$

[12] The option pricing formula in Heston (1993) also allow for time-dependent coefficients, although the closed-form solutions are complicated by any dynamic dependencies in λ_t.

where $\text{state}_{t,k}$ will be chosen from around thirty popular candidate variables. Previous efforts to explain time-varying volatility risk premia with economic variables have been rare and challenging at best. In contrast, the model and GMM estimation procedure that we use here is quite simple to implement. Again, to be consistent with no-arbitrage, the macro-finance shocks "$\text{state}_{t,k}$" must be interpreted either as fixed covariates or predetermined functions of the time-t state variables, S_t and V_t. When actually implementing the estimation below we used the lagged realized volatility, the lagged squared realized volatility, and the lagged implied volatility as instruments for the cross moment in (8), while leaving the moment for the realized volatility in (5) the same as in the constant risk premium case. This in turn results in the identical $\mathcal{X}^2(1)$ asymptotic distribution for the GMM omnibus test.

4.3 Data Sources and Summary Statistics

Our empirical analysis is based on monthly implied and realized volatilities for the S&P500 index from January 1990 through May 2004. For the risk-neutral implied volatility measure, we rely on the VIX index provided by the Chicago Board of Options Exchange (CBOE). The VIX index, available back to January 1990, is based on the liquid S&P500 index options, and more importantly, it is calculated with the model-free approach advocated by Carr and Madan (1998), Demeterfi et al. (1999), and Britten-Jones and Neuberger (2000).[13] As shown in the Monte Carlo study, the model-free implied volatility should be a good approximation to the true (unobserved) risk-neutral expectation of the integrated volatility, and, in particular, a much better approximation than the one afforded by the Black-Scholes implied volatility.

Our realized volatilities are based on the summation of the five-minute squared returns on the S&P500 index within the month.[14] Thus, for a typical month with 22 trading days, we have $22 \times 78 = 1,716$ five-minute returns, where the 78 five-minute subintervals cover the normal trading hours from 9:30am to 4:00pm. Again, as indicated by the Monte Carlo simulations in the previous section, the monthly realized volatilities based on these five-minute returns should provide a very good approximation to the true (unobserved) continuous-time integrated volatility, and, in particular, a much better approximation than the one based on

[13]In September 2003, CBOE replaced the old VIX index, based on S&P100 options and Black-Scholes implied volatility, with the new VIX index based on S&P500 options and model-free implied volatilities involving a discrete approximation to the theoretical result in Carr and Madan (1998), Demeterfi et al. (1999), and Britten-Jones and Neuberger (2000). Historical data on both the old and new VIX are directly available from the CBOE.

[14]The high-frequency data for the S&P500 index is provided by the Institute of Financial Markets.

the sum of the daily squared returns.

Figure 4 plots realized volatility, implied volatility, and their difference.[15] Both of the volatility measures were generally higher during the latter half of the sample, although they have also both decreased more recently. Summary statistics are reported in Table 4. Realized volatility is systematically lower than implied volatility, and its unconditional distribution deviates more from the normal. Both measures exhibit pronounced serial correlation with extremely slow decay in their autocorrelations.

There is a long history of market participants (and some academic researchers) using the level of the VIX implied volatility as a gauge of market fear or, in the economists' jargon, investor risk aversion. Along similar lines, the difference between the implied and realized volatilities are also sometimes associated with the market-implied risk aversion.[16] Unfortunately, the raw difference, as depicted in the bottom panel in Figure 4, is typically very noisy and uninformative, and essentially just follows the level of the volatility. A more structured approach for extracting the volatility risk premium (or implied risk aversion), as discussed in the previous sections, thus holds the promise of revealing a deeper understanding of the way in which the volatility risk premium evolves over time, and its relationship to the macroeconomy. We next turn to a discussion of our pertinent estimation results.

4.4 GMM Estimation Results

Table 5 reports the GMM estimation results for the three volatility risk premium specifications: (i) a constant λ; (ii) a time-varying $\lambda_{t+\Delta}$ driven by the time-t error from the cross moment as in equation (16); and (iii) a time-varying $\lambda_{t+\Delta}$ determined by the time-t macro-finance variables as in equation (17).[17]

First, restricting the risk premium to be constant results in a highly significant estimate of -1.79. However, the chi-square omnibus test of overidentifying restrictions rejects the overall specification at the 10% (although not at the 5%) level.

The second column of the table presents the result allowing for temporal variation in the risk premium driven by the error from the cross moment. The corresponding estimated

[15] Here and throughout the paper, monthly standard deviations are "annualized" by multiplying by $\sqrt{12}$.

[16] In support of this, Rosenberg and Engle (2002) also find that their empirical risk aversion measure is positively related to the difference between implied and objective volatility.

[17] In order to conserve space, we only report the results pertaining to the parameters for the volatility risk premium. The results for the other parameters in the model are directly in line with previous results reported in the literature, and consistent with the summary statistics in Table 3, point toward a high degree of volatility persistence in the (latent) V_t process.

coefficient ($c = 0.02$) is highly statistically significant. Interestingly, the estimates for this specification also point toward a high degree of persistence ($b = 0.80$) in the volatility risk premium, with an implied average value for the full sample of $a/(1-b) = -1.99$. Yet, the overall specification test continues to (barely) reject the model at the 10% (but not at the 5%) level.

To circumvent these shortcomings, the third column presents the results obtained by explicitly including the macro-finance covariates. To select the macro-finance variables in the time-varying risk premium specification, we did an extensive search with 29 monthly data series (listed in Table 8). If part of the temporal variation in investor risk aversion reflects investors focusing on different aspects of the economy at different points in time, as seems likely, some flexibility in specifying the set of covariates seems both appropriate and unavoidable. Hence, we select the group of variables that jointly achieves the highest p-value of the GMM omnibus specification test and that are significant (at the 5% level) based on their individual t-test statistics.[18] To facilitate the subsequent discussion, the resulting seven variables have all been standardized to have mean zero and variance one so that their marginal contribution to the time-varying risk premium are directly comparable.[19]

The results for the autoregressive part of the specification implies an average risk premium of $a/(1-b) = -1.82$, and, without figuring in the dynamic impact of the macro state variables, an even higher degree of own persistence, $b = 0.93$. As necessitated by the specification search, all of the individual parameters for the macro-finance covariates are statistically significant at the 5% level, and the overall GMM specification test is greatly improved, with a p-value of 0.92. The resulting estimate for the volatility risk premium, along with the seven macro-finance input variables, are plotted in Figure 5.

Both the signs and magnitudes of the macro-finance shock coefficients are important in understanding the time-variation of the volatility risk premium. Sticking to the convention that $(-\lambda)$ represents the risk premium, or risk aversion, the realized volatility has the biggest contribution (-0.32) and a positive impact (i.e., when volatility is high so is risk aversion).[20]

[18] We are, of course, aware of the danger of data mining that such a specification search presents. However, we have attempted to limit the degree of data mining by choosing a limited set of candidate macro-finance covariates, as listed in Table 8. Also, it is not the case that adding more covariates in the GMM estimation automatically improves the fit of the model, as judged by the p-value for the over-identifying restrictions.

[19] For stationary variables the unit is the level, while for non-stationary variables the unit is the logarithmic change for the past twelve months.

[20] This result contradicts the counter-intuitive findings in Bliss and Panigirtzoglou (2004) that risk aversion appears to be lower when volatility is higher. However, this finding may possibly be explained by their omission of other important macro-finance variables for jointly describing the time-variation in the estimated

The impact of AAA bond spread over Treasuries (0.19) likely reflects a business cycle effect (i.e., credit spreads tend to be high before a downturn which usually coincides with low risk aversion). Conversely, housing starts have a positive impact on the risk premium (-0.19) (i.e., a real estate boom usually precedes higher risk aversion). The S&P 500 P/E ratio is the fourth most important factor (0.14), and impacts the premium negatively (i.e., everything else equal, higher P/E ratios lowers the degree of risk aversion). The fifth variable in the table is industrial production growth (0.10), which also has a negative impact (i.e., higher growth leads to a lower volatility risk premium). On the contrary, the sixth PPI inflation variable leads to higher risk aversion (-0.05). Finally, the last significant macro state variable, payroll employment, marginally raises the volatility risk premium (-0.04), possibly as a result of wage pressure.

4.5 Alternative Estimation Strategies

Several alternative procedures for estimating the time-varying volatility risk premium have previously been implemented in the literature. One such approach is to vary the risk premium parameter each time period to best match that period's market data. In the context of volatility modeling, that approach would vary the risk premium parameter to match each month's difference between realized and implied volatility. Such an approach would produce the time-varying risk premium shown in the middle panel of Figure 6, the general shape of which matches the earlier plot in the bottom panel of Figure 4 (the simple difference between implied and realized volatilities).

As previously noted, the problem with this approach is that by attributing each wiggle in the data to changes in the risk premium, it produces an excessively volatile time series of monthly risk premia. Economic theory argues that an asset's risk premium should depend on deep structural parameters. For example, in the consumption CAPM (C-CAPM), an asset's risk premium varies with investors' risk aversion and the asset's covariance with investors' consumption. By definition, deep structural parameters should be relatively stable over time. Yet the approach of period-by-period estimation of a time-varying risk premia forces the parameters to vary (almost independently) from one period to the next. As such, we find the monthly volatility risk premium series shown in the middle panel of Figure 6 to be implausibly volatile.[21]

risk aversion coefficient.

[21]Several recent papers have charts that look similar to the middle panel of Figure 6. For example, see

A second approach for estimating investors' "risk appetite," more popular among market participants, is to construct a weighted-average of macro-finance variables.[22] The bottom panel of Figure 6 shows such a weighted-average index constructed from the 29 macro-finance variables listed in Table 8, all standardized to have mean zero and variance one. In addition to concerns that such indexes are too ad hoc to be reliable, indexes constructed in this way also tend to be excessively and implausibly volatile.

A third approach to estimating risk premium parameters comes from the macroeconomic, or consumption-based asset pricing literature. This approach typically assumes that risk premia are constant, or if the risk preference are allowed to vary over time, they end up being implausibly smooth and possibly nonstationary. For example, Campbell and Cochrane (1999) generate time variation in risk aversion through habit formation in which the level of habit reacts only gradually to changes in consumption.[23] Such a modeling strategy explicitly prevents the risk premia from being excessively variable in the short-run.[24]

In contrast, consider the top panel of Figure 6 which plots our estimated volatility risk premium parameter based on the model involving the seven macro-finance covariates. Peaks and troughs in the series are generally multiple years apart, and reassuringly the series is void of the excessive month-by-month fluctuations that plague both of the other two series in that same figure. The estimated risk premium also rises sharply during the two NBER-dated macroeconomic recessions (the shaded areas in the plots), as well as the periods of slow recovery and job growth after the 1991 and 2001 recessions. Moreover, the peaks in the series are readily identifiable with major macroeconomic or financial market developments, including the 1994 rate hike and soft landing, the 1998 Russian debt crisis, and the bursting of the stock market "bubble" in 2000. There is also a peak in the risk premium in 1996 that does not appear to directly line up with any major economic event, except perhaps the worry about over-valuation in the stock market sometimes labeled as the period of "Irrational Exuberance". Lastly, the estimates also suggest that the risk premium often rises fairly sharply but in general declines only gradually.

Rosenberg and Engle (2002) page 363, Tarashev et al. (2003) page 62, and Gordon and St-Amour (2004) page 249.

[22] Chaboud (2003) discusses several such indexes constructed by J.P. Morgan, State Street Bank, and Credit Suisse First Boston.

[23] In a similar vein, Cochrane and Piazzesi (2004) model a slowly varying risk premia on Treasury bonds as a function of current forward rates.

[24] Along these lines, note that except for the 1978-82 monetary experiment period, the estimated risk aversions of Brandt and Wang (2003) do not pick out most recessions.

4.6 Stock Return Predictability

Our characterization of the volatility risk premium has the potential of being informative about other risk premia in the economy. To illustrate, we compare its predictive power for aggregate stock market returns with that of other traditionally-used macro-finance variables. To that end, the top panel of Table 6 reports the results of simple univariate regressions of the monthly S&P500 excess returns on the most significant individual variables from the pool of covariates listed in Table 8. As evidenced by the results, the extracted volatility risk premium has the highest predictive power with an R^2 of 3.67%.[25] The next best predictor is the S&P500 P/E ratio with an R^2 of 2.80%. Next in order are industrial production and nonfarm payrolls with R^2's of 1.53% and 1.06%, respectively. Dividend yields - a significant predictor according to many other studies - only explains 0.85% of the monthly return variation. All-in-all, these results are consistent with previous findings that macroeconomic state variables do predict returns, though the predictability measured by R^2 is usually in the low single digits. Nonetheless, it is noteworthy that of all the predictor variables, the volatility risk premium results in the single highest R^2.

Combining all of the marginally significant variables into a single multiple regression, results in the estimates shown in the bottom panel of Table 6. Interestingly, none of the macro-finance variable remains significant when the volatility risk premium is included, while only the P/E ratio is significant in the regression excluding the premium. Of course, the estimate for the volatility risk premium already incorporates some of the same macroeconomic variables (see Table 5), so the finding that these variables are "driven out" when included together with the premium is not necessarily that surprising. However, the macro variables entering the model for $\lambda_{t+\Delta}$ only impact the returns indirectly through the temporal variation in the premium, and the volatility risk premium itself is also estimated from a very different set of moment conditions involving only the model-free realized and options implied volatilities.

Our examination of the monthly stock excess return in Table 6 singles out the volatility risk premium and the stock market P/E ratio as the two most important predictor variables. Table 7 augments these results with regressions involving longer-run quarterly excess returns spanning 1990Q1 through 2003Q2. In addition to the volatility risk premium and the P/E ratio from the last month of the previous quarter, we now also include the quarterly

[25]The use of the volatility risk premium as a second-stage regressor suffers from a standard errors-in-variables type problem, resulting in too large a standard error for the estimated slope coefficient.

consumption-wealth ratio popularized by Lettau and Ludvigson (2001) in the regressions. The consumption-wealth ratio, termed CAY, has previously been found to be significant in explaining longer horizon returns. The first three regressions in the table show that each of the three variables are indeed individually significant. At the same time, it is noteworthy that the risk premium results in the highest individual R^2 of 11.6%, much higher than the monthly R^2 of 3.7%. Adding the P/E ratio and/or the CAY variable further increase the quarterly R^2's in excess of 14%. The risk premium remains significant in all of the multiple regressions, while combining the P/E ratio and the CAY variable in the same regression renders both insignificant, and does not increase the R^2 by much. As such, these results further reinforce the earlier findings for the monthly returns in Table 6 and the role of the estimated volatility risk premium as a new and powerful stock market predictor over longer quarterly horizons.

5 Conclusion

This paper develops a simple consistent approach for estimating the volatility risk premium. The approach exploits the linkage between the objective and risk-neutral expectations of the integrated volatility. The estimation is facilitated by the use of newly available model-free realized volatilities based on high-frequency intraday data along with model-free option-implied volatilities. The approach allows us to explicitly link any temporal variation in the risk premium to underlying state variables within an internally consistent and simple-to-implement GMM estimation framework. A small scale Monte Carlo experiment indicates that the procedure performs well in estimating the volatility risk premium in empirically realistic situations. In contrast, the estimates based on the Black-Scholes implied volatilities and/or monthly sample variances based on daily squared returns result in highly inefficient and statistically unreliable estimates of the risk premium. Applying the methodology to the S&P500 market index, we find significant evidence for temporal variation in the volatility risk premium, which we directly link to a set of underlying macro-finance state variables. Interestingly, the extracted volatility risk premium also appears to be helpful in predicting the return on the market itself.

The volatility risk premium (or risk aversion index) extracted in our paper differs sharply from other approaches in the literature. In particular, earlier estimates relying directly on period-by-period differences in the estimated risk-neutral and objective distributions tend to produce implausibly volatile estimates. On the other hand, earlier procedures based on structural macroeconomic/consumption-type pricing models typically result in implausibly smooth estimates. In contrast, the model-free realized and implied volatility-based procedure developed here results in an estimated premium that avoids the excessive period-by-period random fluctuations, yet responds to recessions, financial crises, and other economic events in an empirically realistic fashion.

It would be interesting to more closely compare and contrast the risk aversion index estimated here to other popular gauges of investor fear or market sentiment. Also, how do the estimated volatility risk premium for the S&P500 compare to that of other markets? The results in the paper show that the extracted volatility risk premium for the current month is useful in predicting next month's aggregate S&P500 return. It would interesting to further explore the cross sectional pricing implications of this finding. Does the volatility risk premium represent a systematic priced risk factor?[26] Also, what is the link between stock and bond market volatility risk premia? Lastly, better estimates for the volatility risk premium is, of course, of direct importance for derivatives pricing. We leave further work along these lines for future research.

[26]The recent results in Ang et al. (2005) suggest that volatility risk may indeed be a priced factor.

References

Aït-Sahalia, Yacine and Andrew W. Lo (2000), "Nonparametric Risk Management and Implied Risk Aversion," *Journal of Econometrics*, vol. 94, 9–51.

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

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

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys (2003b), "Modeling and Forecasting Realized Volatility," *Econometrica*, vol. 71, 579–625.

Andersen, Torben G., Tim Bollerslev, and Nour Meddahi (2004), "Analytical Evaluation of Volatility Forecasts," *International Economic Review*, forthcoming.

Ang, Andrew, Robert Hodrick, Yuhang Xing, and Xiaoyan Zhang (2005), "The Cross-Section of Volatility and Expected Returns," *Journal of Finance*, forthcoming.

Bakshi, Gurdip and Nikunj Kapadia (2003), "Delta-Hedged Gains and the Negative Market Volatility Risk Premium," *Review of Financial Studies*, vol. 16, 527–566.

Bakshi, Gurdip S. and Zhiwu Chen (1996), "The Spirit of Capitalism and Stock-Market Prices," *American Economic Review*, vol. 86, 133–157.

Barndorff-Nielsen, Ole and Neil Shephard (2002), "Econometric Analysis of Realised Volatility and Its Use in Estimating Stochastic Volatility Models," *Journal of Royal Statistical Society*, vol. Series B 64.

Barndorff-Nielsen, Ole and Neil Shephard (2004a), "Econometric Analysis of Realised Covariation: High Frequency Based Covariance, Regression and Correlation," *Econometrica*, vol. 72, 885–925.

Barndorff-Nielsen, Ole and Neil Shephard (2004b), "A Feasible Central Limit Theory for Realised Volatility under Leverage," *manuscript*, Nuffield College, Oxford University.

Bates, David S. (1996), "Jumps and Stochastic Volatility: Exchange Rate Process Implicit in Deutsche Mark Options," *The Review of Financial Studies*, vol. 9, 69–107.

Benzoni, Luca (2002), "Pricing Options under Stochastic Volatility: An Empirical Investigation," *Working Paper*, University of Minnesota.

Bliss, Robert R. and Nikolaos Panigirtzoglou (2004), "Option-Implied Risk Aversion Estimates," *Journal of Finance*, vol. 59, 407 – 446.

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

Bollerslev, Tim and Hao Zhou (2004), "Volatility Puzzles: A Simple Framework for Gauging Return-Volatility Regressions," *Journal of Econometrics*, forthcoming.

Brandt, Michael W. and Kevin Q. Wang (2003), "Time-varying Risk Aversion and Expected Inflation," *Journal of Monetary Economics*, vol. 50, 1457–1498.

Britten-Jones, Mark and Anthony Neuberger (2000), "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance*, vol. 55, 839–866.

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

Campbell, John Y. and John H. Cochrane (1999), "By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, vol. 107, 205–251.

Campbell, John Y. and Robert J. Shiller (1988a), "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, vol. 1, 195–228.

Campbell, John Y. and Robert J. Shiller (1988b), "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance*, vol. 43, 661–676.

Carr, Peter and Dilip Madan (1998), *Towards a Theory of Volatility Trading*, chap. 29, pages 417–427, Risk Books.

Carr, Peter and Liuren Wu (2004), “Variance Risk Premia,” *Working Paper*, NYU and CUNY.

CBOE Documentation (2003), “VIX: CBOE Volatility Index,” *White Paper*.

Chaboud, Alain (2003), “Indexes of Risk Appetite in International Financial Markets,” *memo*, Federal Reserve Board.

Chernov, Mikhail and Eric Ghysels (2000), “A Study towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation,” *Journal of Financial Economics*, vol. 56, 407–458.

Cochrane, John H. and Monika Piazzesi (2004), “Bond Risk Premia,” *NBER Working Paper 9178*.

Demeterfi, Kresimir, Emanuel Derman, Michael Kamal, and Joseph Zou (1999), “A Guide to Volatility and Variance Swaps,” *Journal of Derivatives*, vol. 6, 9–32.

Eraker, Bjørn (2004), “Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices,” *Journal of Finance*, vol. 59, 1367–1403.

Fama, Eugene F. and Robert R. Bliss (1987), “The information in Long-Maturity Forward Rates,” *The American Economic Review*, vol. 77, 680–692.

Fama, Eugene F. and Kenneth R. French (1988), “Dividend Yields and Expected Stock Returns,” *Journal of Financial Economics*, vol. 22, 3–25.

Garcia, René, Marc-André Lewis, and Éric Renault (2001), “Estimation of Objective and Risk-Neutral Distributions Based on Moments of Integrated Volatility,” *Working Paper*, CRDE, Université de Montréal.

Gordon, Stephen and Pascal St-Amour (2004), “Asset Returns and State-Dependent Risk Preferences,” *Journal of Business and Economic Statistics*, vol. 22, 241–252.

Hansen, Lars Peter (1982), “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, vol. 50, 1029–1054.

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

Jackwerth, Jens Carsten (2000), "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, vol. 13, 433–451.

Jiang, George and Yisong Tian (2004), "Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, forthcoming.

Lettau, Martin and Sydney Ludvigson (2001), "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance*, vol. 56, 815–849.

Liang, Nellie and Hao Zhou (2003), "Implied Volatility, Realized Volatility, and Risk Aversion," *memo*, Federal Reserve Board.

Lynch, Damien and Nikolaos Panigirtzoglou (2003), "Option Implied and Realized Measures of Variance," *Working Paper*, Monetary Instruments and Markets Division, Bank of England.

Meddahi, Nour (2002), "Theoretical Comparison Between Integrated and Realized Volatility," *Journal of Applied Econometrics*, vol. 17, 479–508.

Newey, Whitney K. and Kenneth D. West (1987), "A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, vol. 55, 703–708.

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

Rosenberg, Joshua V. and Robert F. Engle (2002), "Empirical Pricing Kernels," *Journal of Financial Economics*, vol. 64, 341–372.

Tarashev, Nikola, Kostas Tsatsaronis, and Dimitrios Karampatos (2003), "Investors' Attitude toward Risk: What Can We Learn from Options?" *BIS Quarterly Review*, Bank of International Settlement.

Table 1: Monte Carlo Result for λ with Risk-Neutral Implied Volatility

	Mean Bias		Median Bias		Root-MSE	
	T = 150	T = 600	T = 150	T = 600	T = 150	T = 600
	Scenario (a), Benchmark Case					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.10$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	-0.0046	-0.0015	-0.0041	-0.0013	0.0202	0.0091
Realized, 5-min.	-0.0043	-0.0014	-0.0027	-0.0014	0.0201	0.0090
Realized, 1-day	-0.0129	-0.0036	-0.0169	-0.0040	0.0576	0.0260
	Scenario (b), High Volatility Persistence					
	$\kappa = 0.03$, $\theta = 0.20$, $\sigma = 0.10$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	-0.0097	-0.0029	-0.0079	-0.0017	0.0244	0.0099
Realized, 5-min.	-0.0088	-0.0026	-0.0059	-0.0014	0.0237	0.0098
Realized, 1-day	-0.0172	-0.0051	-0.0187	-0.0039	0.0615	0.0275
	Scenario (c), High Volatility-of-Volatility					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.20$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	-0.0166	-0.0054	-0.0127	-0.0049	0.0463	0.0193
Realized, 5-min.	-0.0162	-0.0054	-0.0119	-0.0048	0.0457	0.0190
Realized, 1-day	-0.0278	-0.0089	-0.0288	-0.0085	0.0804	0.0342
	Scenario (d), High Leverage					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.20$, $\lambda = -0.20$, $\rho = -0.80$					
Integrated Vol.	-0.0046	-0.0016	-0.0040	-0.0015	0.0200	0.0093
Realized, 5-min.	-0.0042	-0.0014	-0.0043	-0.0012	0.0200	0.0092
Realized, 1-day	-0.0130	-0.0032	-0.0165	-0.0025	0.0569	0.0253

Table 2: Monte Carlo Result for λ with Model-Free Implied Volatility

	Mean Bias		Median Bias		Root-MSE	
	T = 150	T = 600	T = 150	T = 600	T = 150	T = 600
	Scenario (a), Benchmark Case					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.10$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	0.0013	0.0044	0.0015	0.0048	0.0199	0.0101
Realized, 5-min.	0.0017	0.0045	0.0030	0.0045	0.0199	0.0101
Realized, 1-day	-0.0068	0.0021	-0.0103	0.0017	0.0569	0.0258
	Scenario (b), High Volatility Persistence					
	$\kappa = 0.03$, $\theta = 0.20$, $\sigma = 0.10$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	-0.0005	0.0064	0.0000	0.0071	0.0248	0.0130
Realized, 5-min.	0.0003	0.0066	0.0020	0.0068	0.0244	0.0130
Realized, 1-day	-0.0081	0.0036	-0.0093	0.0053	0.0598	0.0276
	Scenario (c), High Volatility-of-Volatility					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.20$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	-0.0034	0.0075	-0.0008	0.0078	0.0475	0.0221
Realized, 5-min.	-0.0030	0.0077	-0.0018	0.0086	0.0471	0.0219
Realized, 1-day	-0.0166	0.0029	-0.0170	0.0041	0.0796	0.0341
	Scenario (d), High Leverage					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.20$, $\lambda = -0.20$, $\rho = -0.80$					
Integrated Vol.	0.0016	0.0045	0.0021	0.0046	0.0198	0.0103
Realized, 5-min.	0.0020	0.0047	0.0016	0.0048	0.0198	0.0104
Realized, 1-day	-0.0068	0.0029	-0.0101	0.0035	0.0561	0.0253

Table 3: Monte Carlo Result for λ with Black-Scholes Implied Volatility

	Mean Bias		Median Bias		Root-MSE	
	T = 150	T = 600	T = 150	T = 600	T = 150	T = 600
	Scenario (a), Benchmark Case					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.10$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	0.0089	0.0119	0.0094	0.0122	0.0209	0.0147
Realized, 5-min.	0.0092	0.0120	0.0106	0.0121	0.0211	0.0148
Realized, 1-day	0.0010	0.0100	-0.0019	0.0094	0.0562	0.0276
	Scenario (b), High Volatility Persistence					
	$\kappa = 0.03$, $\theta = 0.20$, $\sigma = 0.10$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	0.0045	0.0107	0.0065	0.0120	0.0214	0.0139
Realized, 5-min.	0.0055	0.0111	0.0079	0.0118	0.0214	0.0142
Realized, 1-day	-0.0015	0.0094	-0.0007	0.0105	0.0601	0.0285
	Scenario (c), High Volatility-of-Volatility					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.20$, $\lambda = -0.20$, $\rho = -0.50$					
Integrated Vol.	0.0215	0.0321	0.0247	0.0324	0.0444	0.0361
Realized, 5-min.	0.0220	0.0321	0.0258	0.0322	0.0443	0.0361
Realized, 1-day	0.0136	0.0312	0.0144	0.0311	0.0742	0.0450
	Scenario (d), High Leverage					
	$\kappa = 0.10$, $\theta = 0.20$, $\sigma = 0.20$, $\lambda = -0.20$, $\rho = -0.80$					
Integrated Vol.	0.0127	0.0156	0.0134	0.0156	0.0227	0.0179
Realized, 5-min.	0.0131	0.0158	0.0128	0.0160	0.0230	0.0181
Realized, 1-day	0.0041	0.0141	0.0002	0.0153	0.0555	0.0288

Table 4: Summary Statistics for Monthly Implied and Realized Volatilities

Statistics	Realized Volatility	Implied Volatility
Mean	12.68	20.08
Std. Dev.	5.84	6.39
Skewness	1.21	0.84
Kurtosis	4.63	3.87
Minimum	4.73	10.63
5% Qntl.	5.92	11.73
25% Qntl.	7.93	14.79
50% Qntl.	11.56	19.52
75% Qntl.	15.42	24.19
95% Qntl.	24.62	31.17
Maximum	36.61	44.28
ρ_1	0.81	0.83
ρ_2	0.68	0.69
ρ_3	0.61	0.60
ρ_4	0.54	0.56
ρ_5	0.55	0.55
ρ_6	0.55	0.53
ρ_7	0.52	0.50
ρ_8	0.53	0.49
ρ_9	0.53	0.52
ρ_{10}	0.53	0.54

Table 5: Estimation of Volatility Risk Premium

	Constant	Time-Varying	Macro-Finance
λ	-1.793 (0.216)		
a		-0.394 (0.107)	-0.122 (0.051)
b		0.803 (0.070)	0.933 (0.030)
c		0.022 (0.003)	
c_1 Realized Volatility			-0.319 (0.042)
c_2 Moody AAA Bond Spread			0.194 (0.034)
c_3 Housing Start			-0.191 (0.055)
c_4 S&P500 P/E Ratio			0.140 (0.015)
c_5 Industrial Production			0.097 (0.026)
c_6 Producer Price Index			-0.047 (0.023)
c_7 Payroll Employment			-0.040 (0.019)
$\mathcal{X}^2$(d.o.f. = 1) (p-Value)	2.889 (0.089)	2.722 (0.099)	0.169 (0.919)

All of the macro-finance variables are standardized to have mean zero and variance one. The growth variables (Industrial Production, Producer Price Index, and Payroll Employment) are expressed in terms of the logarithmic difference over the past twelve months. The lag length in the Newey-West weighting matrix employed in the estimation is set at 25.

Table 6: Monthly Stock Market Return Predictability

Variables	Intercept	(s.e.)	Slope	(s.e.)	R-Square
Volatility Risk Premium	-18.51	(10.58)	12.49	(5.18)	0.04
S&P500 PE Ratio	35.94	(13.75)	-1.27	(0.57)	0.03
Industrial Production	-0.93	(5.50)	1.99	(1.23)	0.02
Nonfarm Payroll Employment	-0.31	(5.48)	3.64	(2.62)	0.01
26 Other Macro-Finance Variables					<0.01

Joint Estimation	Including λ_t		Excluding λ_t	
Variables	Parameter	(s.e.)	Parameter	(s.e.)
Intercept	3.58	(20.30)	32.62	(15.31)
Volatility Risk Premium	8.34	(5.49)		
S&P500 PE Ratio	-0.71	(0.57)	-1.26	(0.59)
Industrial Production	2.07	(2.26)	2.72	(2.20)
Nonfarm Payroll Employment	-1.94	(4.70)	-3.31	(4.66)
R-Square	0.05		0.04	

The table reports predictive regressions for the monthly excess return on S&P500 index measured in annualized percentage term. Industrial Production and Payroll Employment numbers represent the past year logarithmic changes in annualized percentages.

Table 7: Quarterly Stock Market Return Predictability

Intercept (s.e.)	Risk Premium (s.e.)	PE Ratio (s.e.)	CAY (s.e.)	R-Square
-22.039 (12.840)	13.63 (6.28)			0.12
41.044 (16.289)		-1.54 (0.69)		0.10
2.412 (4.307)			5.38 (2.03)	0.09
8.740 (17.576)	9.51 (5.45)	-0.95 (0.59)		0.14
-16.755 (13.545)	10.49 (6.83)		3.19 (1.79)	0.14
29.736 (27.887)		-1.10 (1.14)	2.48 (3.35)	0.11
2.403 (23.700)	9.03 (5.64)	-0.66 (0.96)	1.75 (3.04)	0.15

The quarterly data range from 1990Q1 to 2003Q2. The consumption-wealth-ratio, or CAY, variable is defined in Lettau and Ludvigson (2001), and the data is downloaded from their website.

Table 8: List of Macro-Finance Variables

Macro-Finance Variables	Data Source
S&P500 Realized Volatility	Constructed from IFM (CME)
S&P500 Implied Volatility	CBOE
S&P500 Market Return	Standard & Poors
S&P500 PE Ratio	Standard & Poors
S&P500 Dividend Yield	Standard & Poors
NYSE Trading Volume	NYSE
Unemployment Rate	Bureau of Labor Statistics
Nonfarm Payroll Employment	Bureau of Labor Statistics
Industrial Capacity Utilization	Federal Reserve
Industrial Production	Federal Reserve
CPI Inflation	Bureau of Labor Statistics
Producer Price Index	Bureau of Labor Statistics
Expected CPI Inflation	Michigan Survey
Treasury Spread 5yr-6mn	Federal Reserve
Treasury Spread 10yr-6mn	Federal Reserve
Mortgage Spread (over 10yr Treasury)	Federal Reserve
Swap Spread (over 10yr Treasury)	Bloomberg
AAA Corporate Spread (over 10yr Treasury)	Moody
BAA Corporate Spread (over 10yr Treasury)	Moody
AA Corporate Spread (over 10yr Treasury)	Merrill Lynch
BBB Corporate Spread (over 10yr Treasury)	Merrill Lynch
Consumer Sentiment	Michigan Survey
Consumer Sentiment (Expected)	Michigan Survey
Consumer Confidence	Conference Board
Consumer Confidence (Expected)	Conference Board
Housing Permit Number	HUD
Housing Start Number	HUD
Money Supply (M2)	Federal Reserve
Business Cycle Indicator	NBER

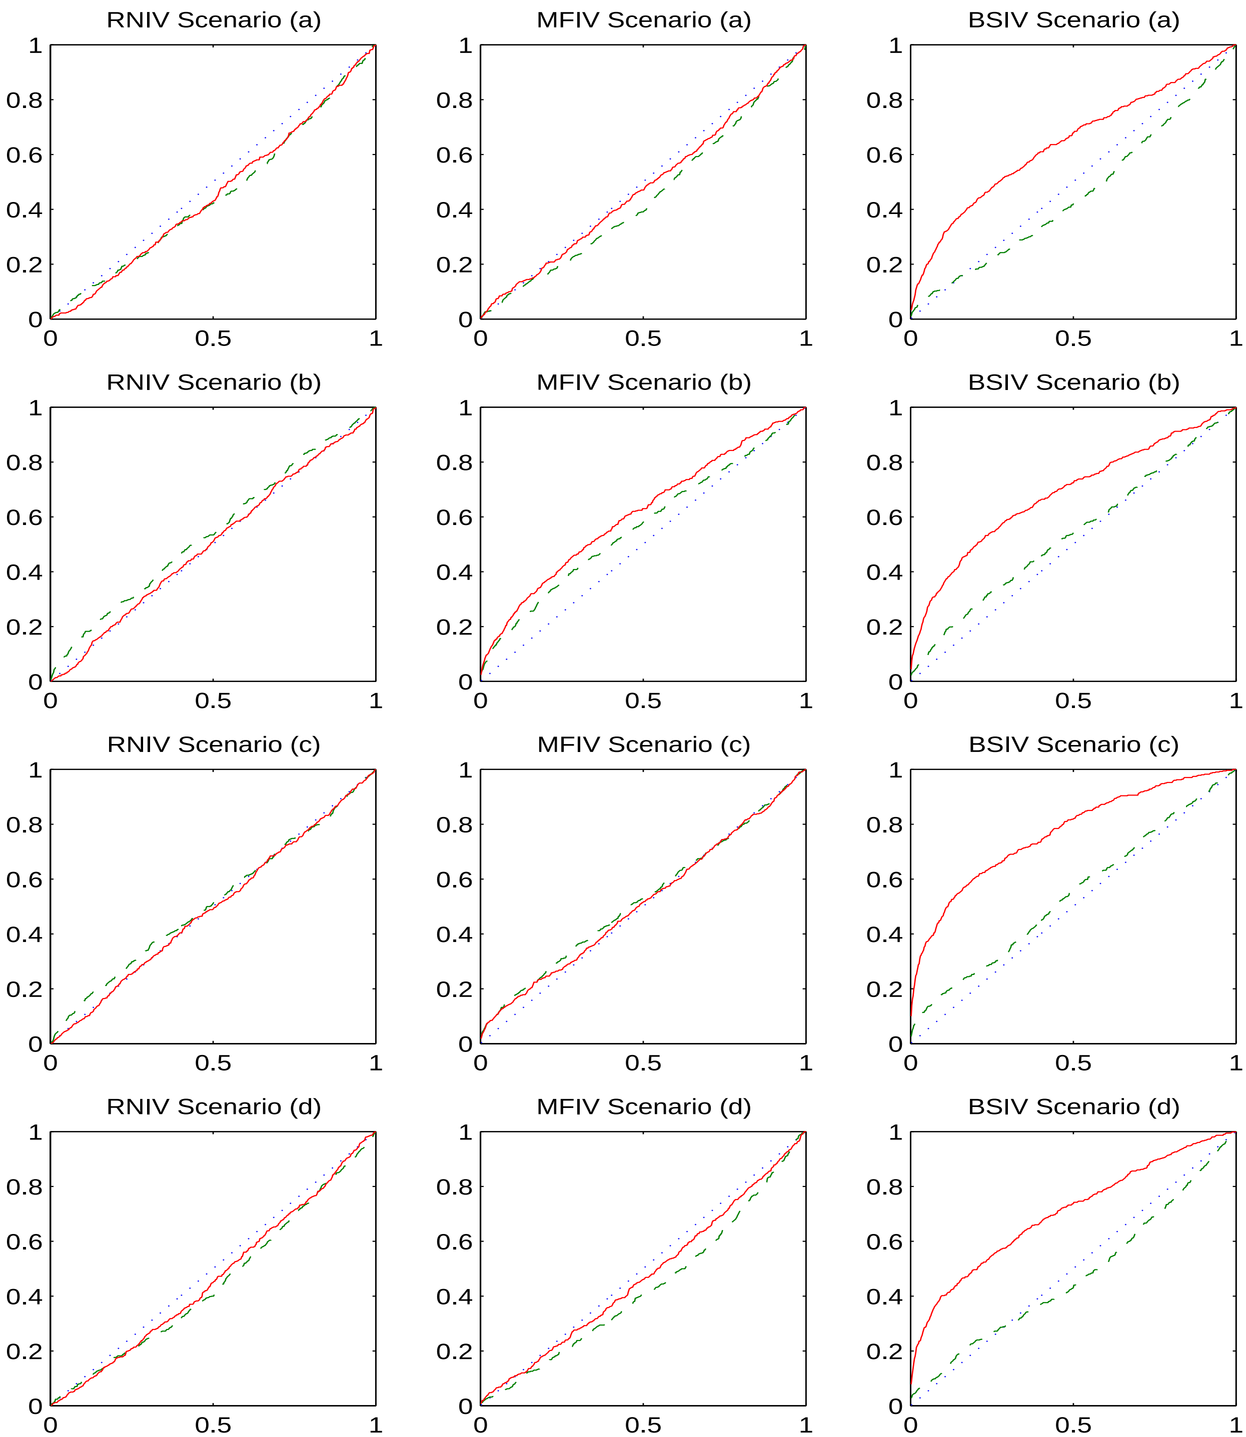

Figure 1: Wald Test for Risk Premium with True Integrated Volatility.
The X-axis gives the nominal level of the test and Y-axis the probability of rejection. The dotted line represents the uniform reference distribution, the dash line is T = 150, and the solid line is T = 600.

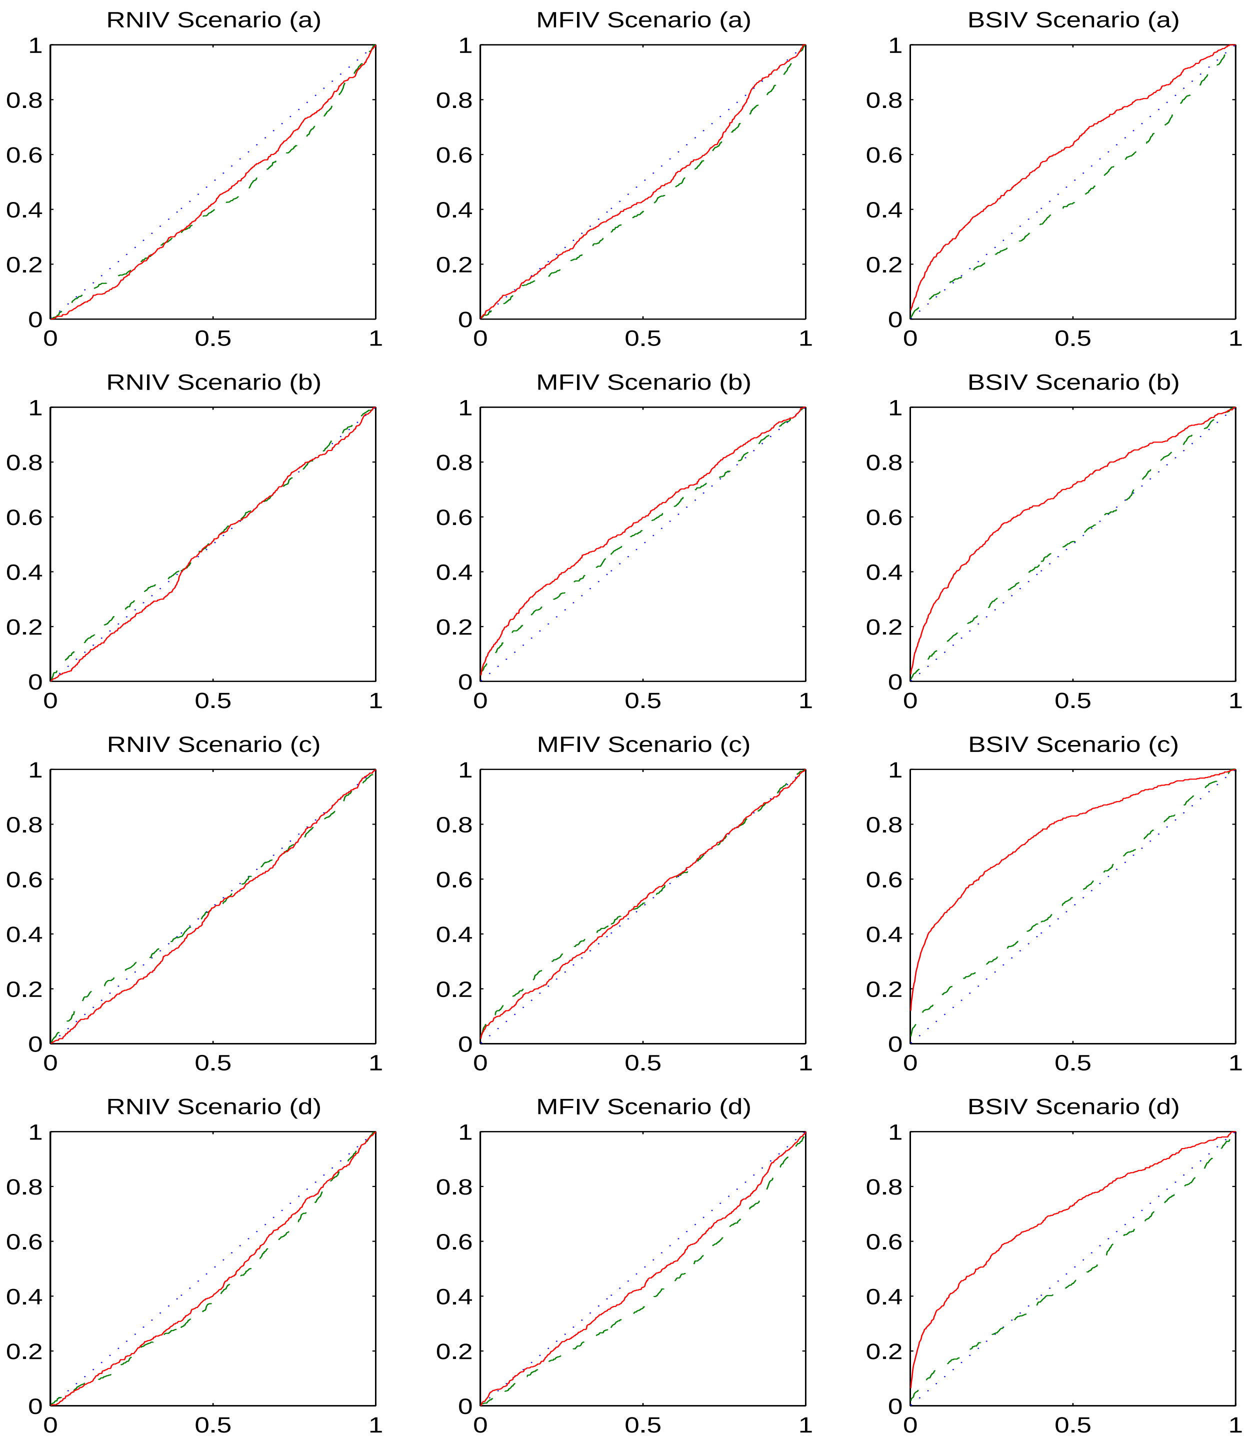

Figure 2: Wald Test for Risk Premium with Five-Minute Return Realized Volatility. The X-axis gives the nominal level of test and the Y-axis the probability of rejection. The dotted line represents the uniform reference distribution, the dash line is $T = 150$, and the solid line is $T = 600$.

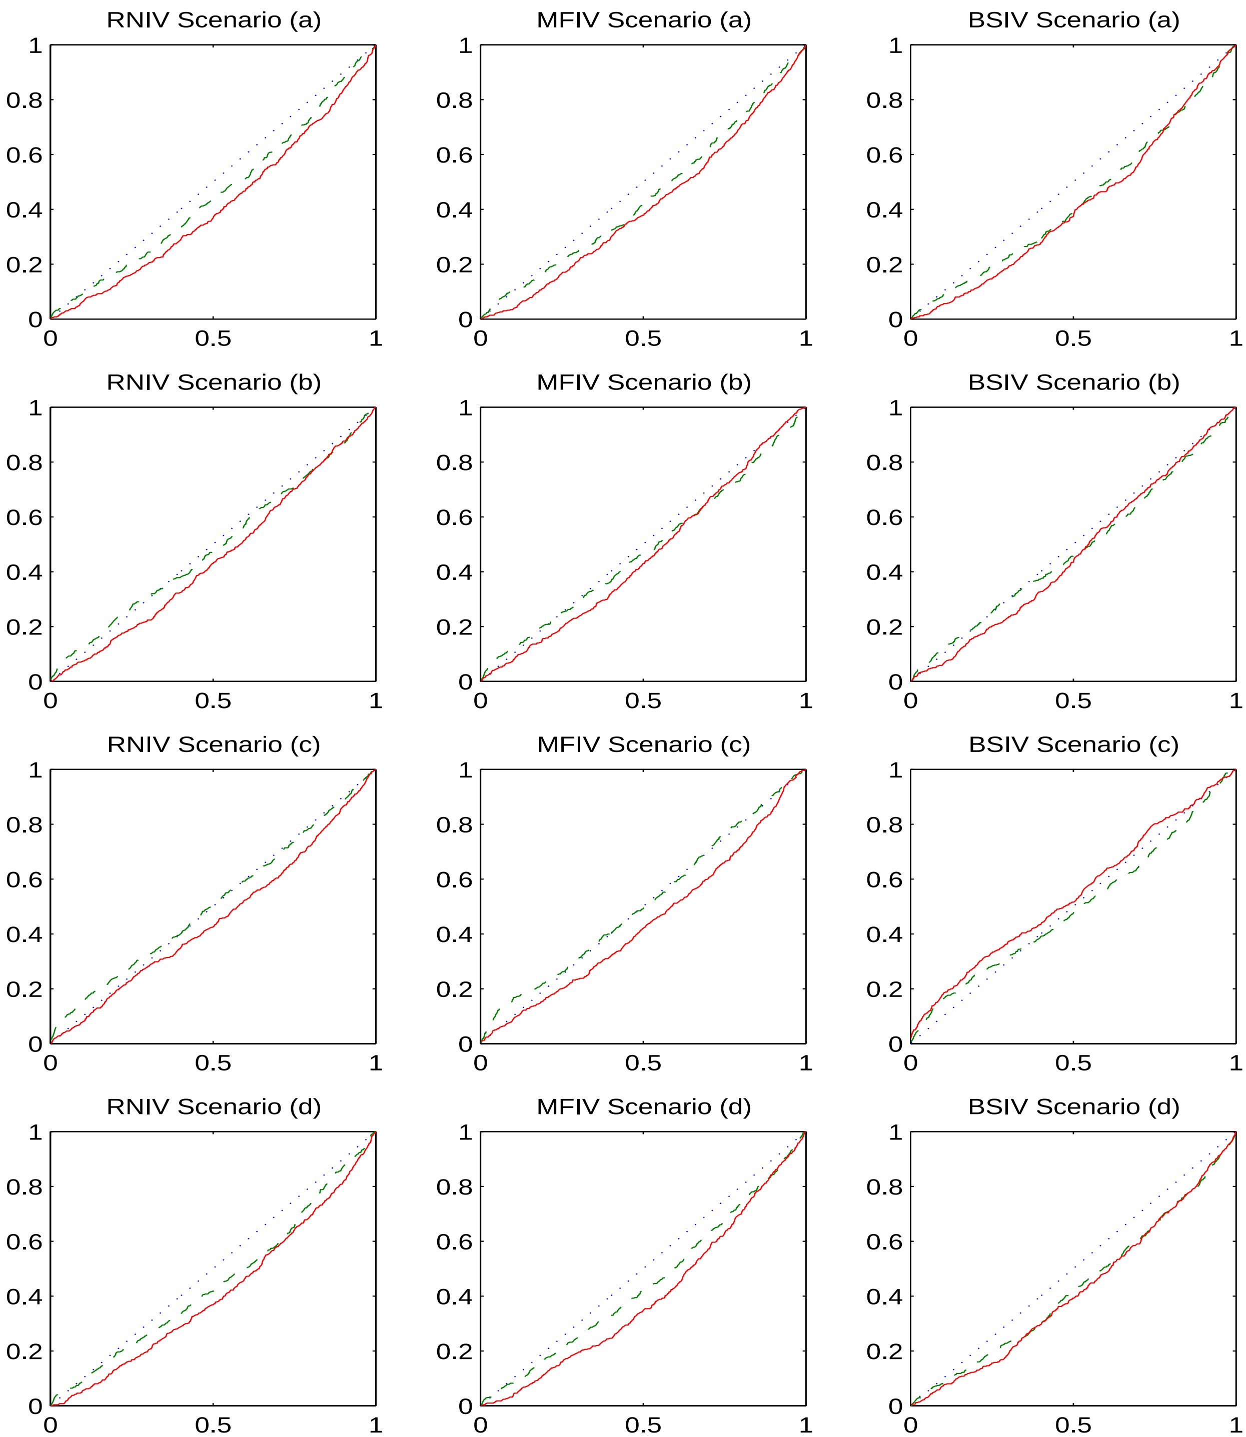

Figure 3: Wald Test for Risk Premium with Daily Return Realized Volatilities.
The X-axis gives the nominal level of the test and the Y-axis the probability of rejection. The dotted line represents the uniform reference distribution, the dash line is T = 150, and the solid line is T = 600.

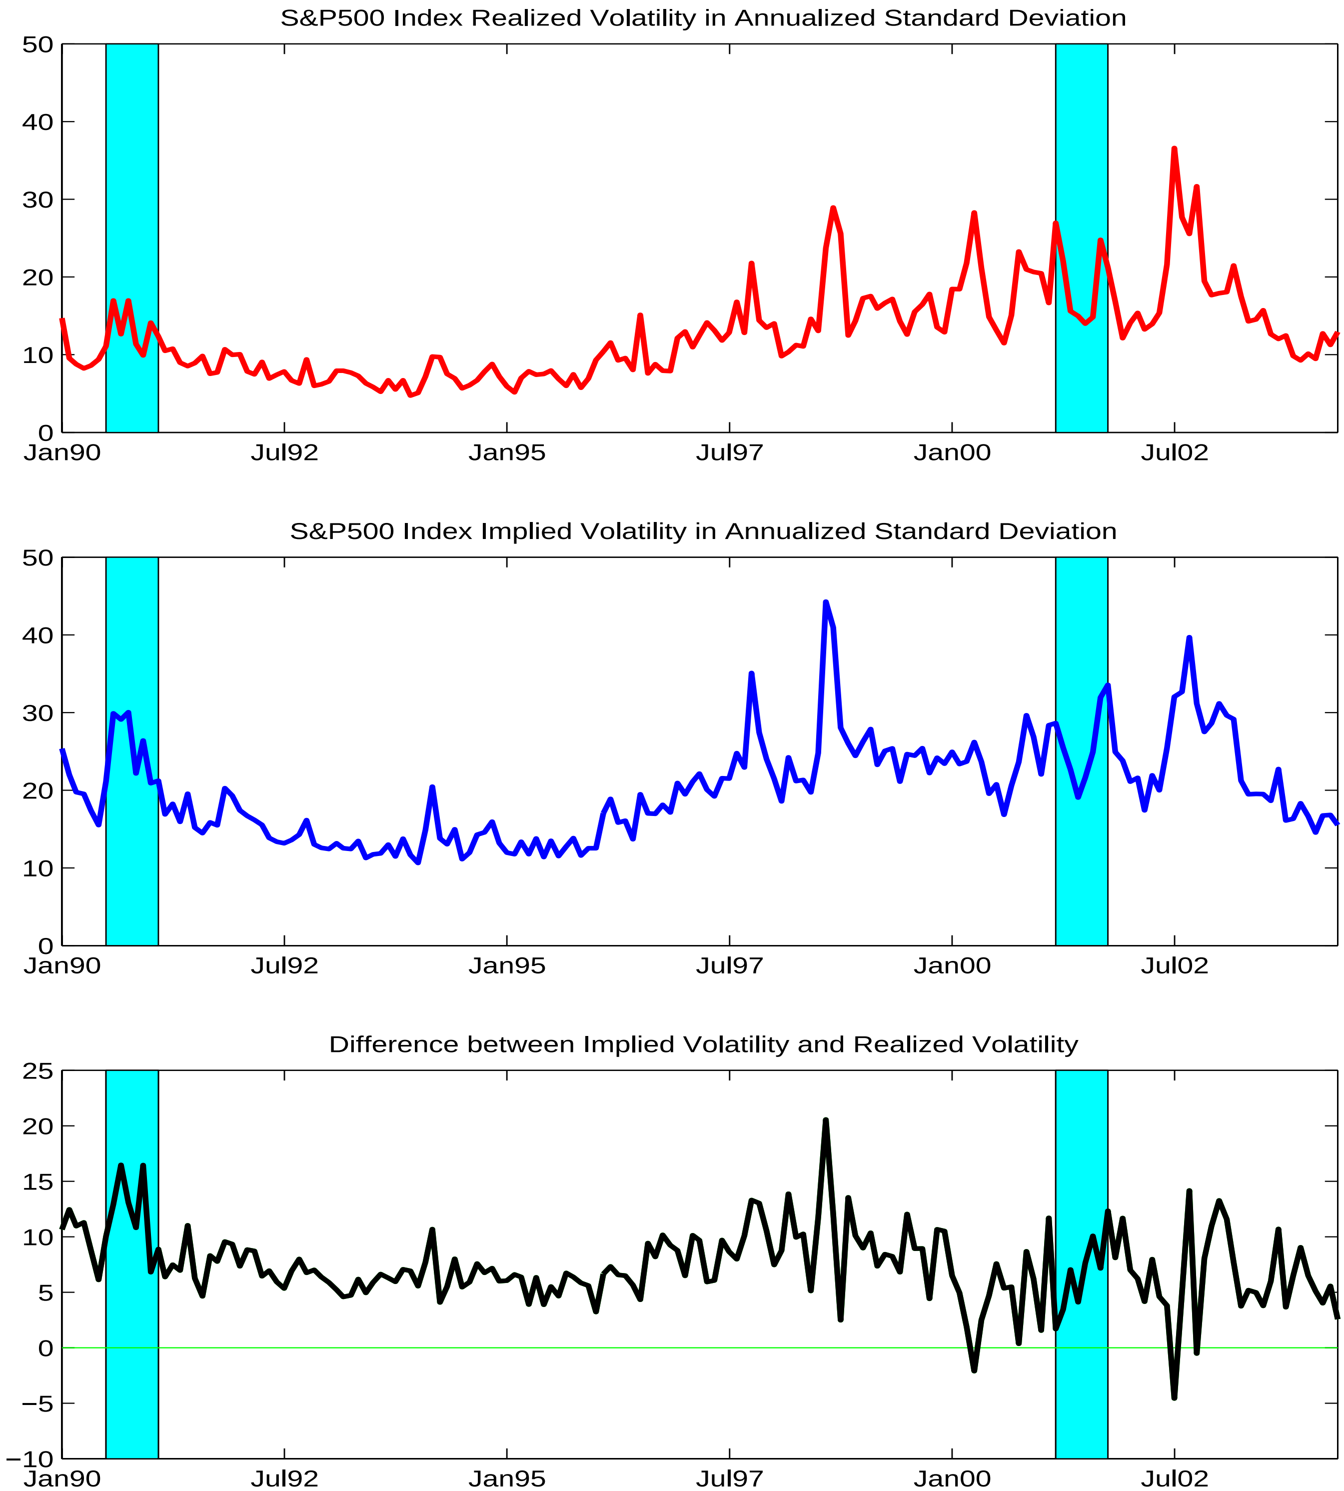

Figure 4: Model-Free Realized and Implied Volatilities

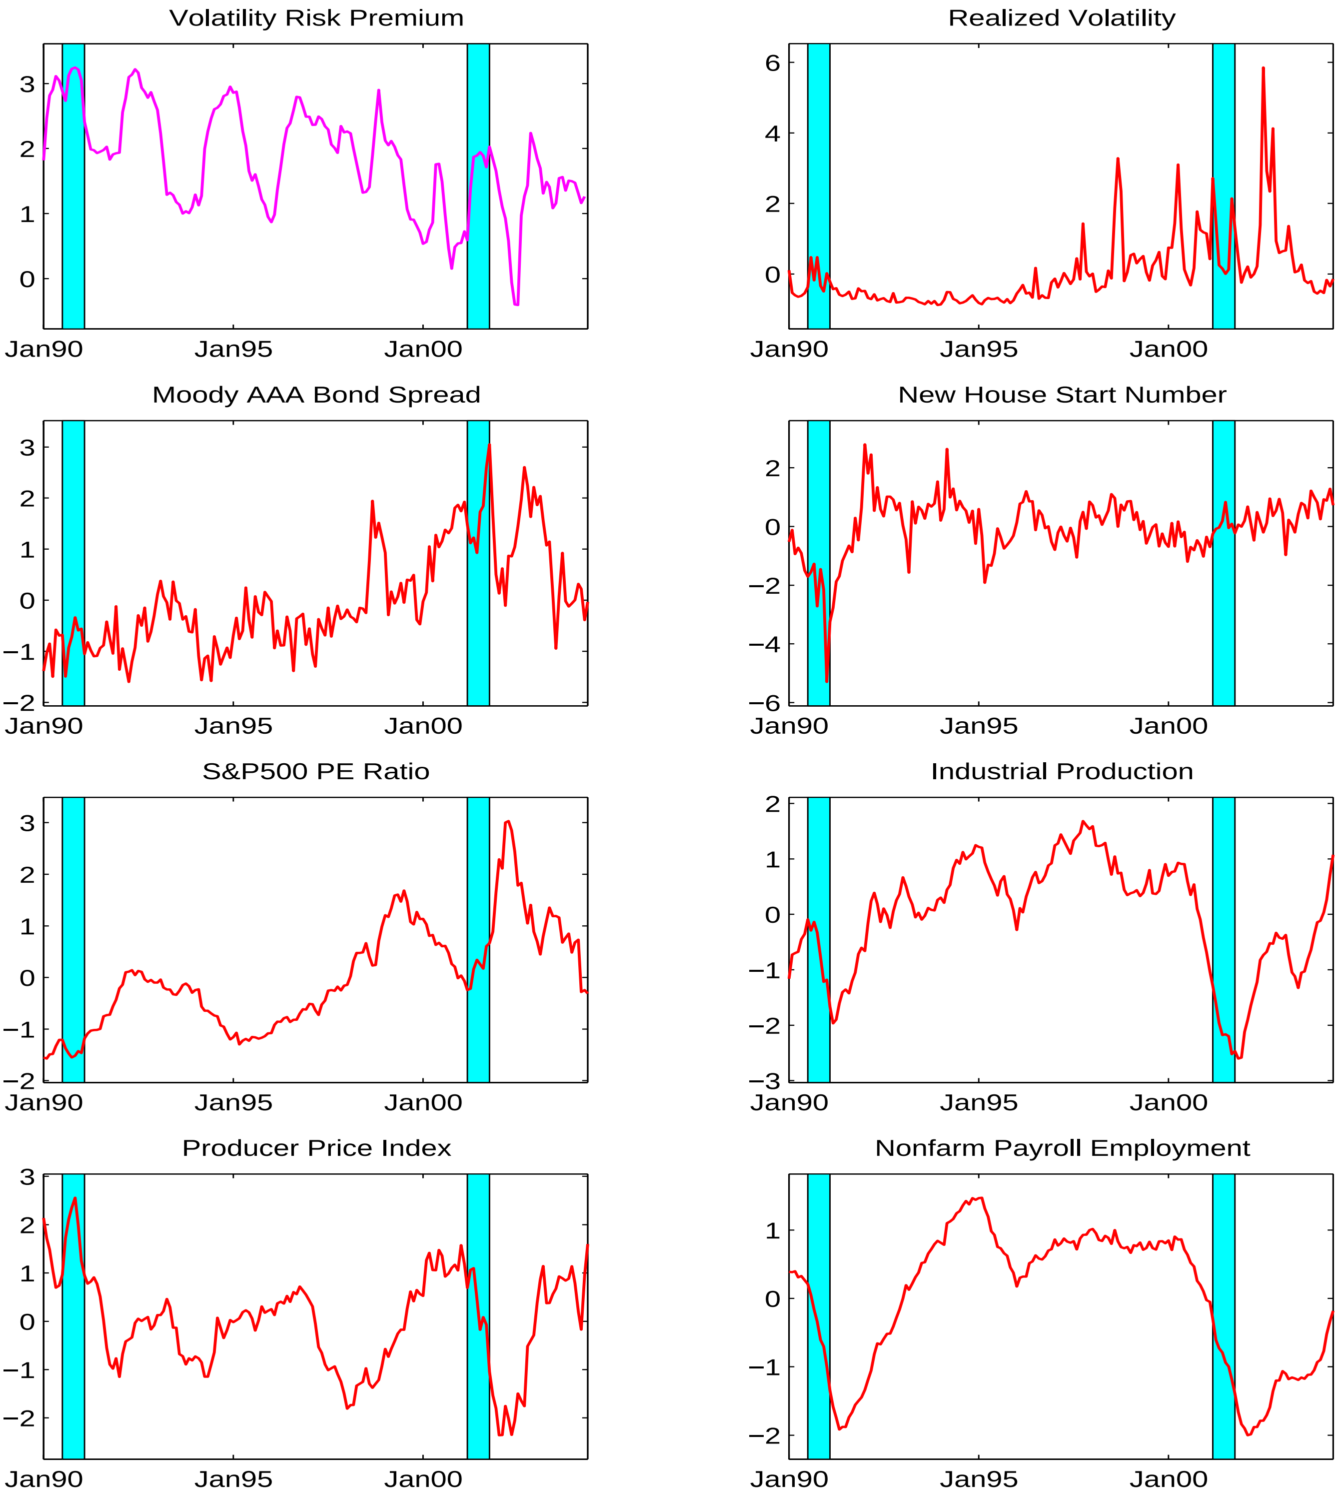

Figure 5: Standardized Macro-Finance Covariates

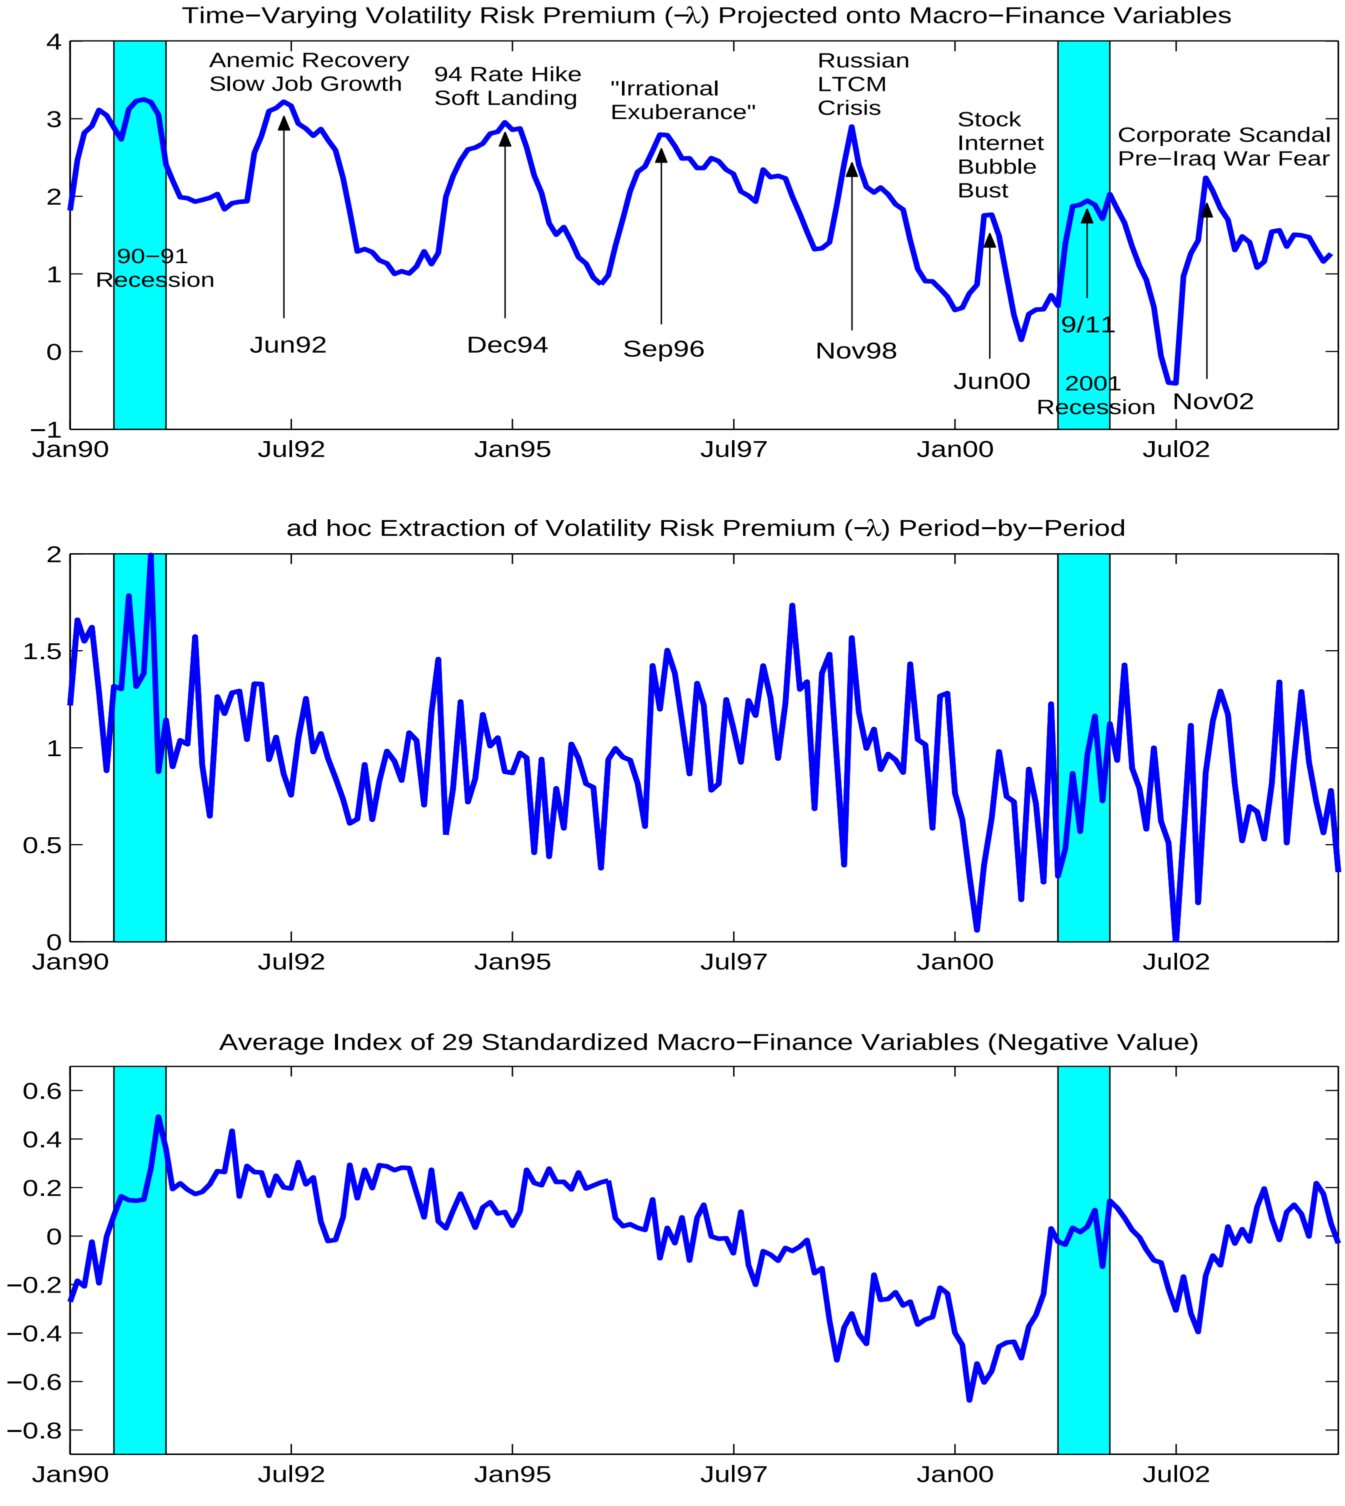

Figure 6: Time-Varying Volatility Risk Premium and Other Indices