Center for Research in Econometric Analysis of Time Series

CREATES Research Paper 2007-16

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

Tim Bollerslev, Michael Gibson and Hao Zhou

School of Economics and Management
University of Aarhus
Building 1322, DK-8000 Aarhus C
Denmark

Aarhus School of Business
University of Aarhus
Handelshøjskolen
Aarhus Universitet

UNIVERSITY OF COPENHAGEN

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: April 2007

*The work of Bollerslev was supported by a grant from the National Science Foundation to the NBER and CREATES funded by the Danish National Research Foundation. We would also like to thank N.K. Chidambaran, Hui Guo, George Jiang, Chris Jones, Nellie Liang, Nagpurnanand R. Prabhala, Patricia White, and seminar participants at Venice Conference on Time-Varying Financial Structures 2005, Federal Reserve Conference on Financial Market Risk Premia 2005, Peking University, the Bank for International Settlement, and AFA 2006 Annual Meeting for useful comments and suggestions. The views presented here are solely those of the authors and do not necessarily represent those of the Federal Reserve Board or its staff. Matthew Chesnes and Stephen Saroki 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.

‡Risk Analysis Section, 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.

§Risk Analysis Section, 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 indicates significant temporal dependencies in the estimated stochastic volatility risk premium, which we in turn relate to a set of 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., 2004b). 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, 2005b; Carr and Wu, 2005).[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; Jones, 2003; Eraker, 2004; Aït-Sahalia and Kimmel, 2005, 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

[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.

premium. As such, the method is simple to implement and can easily be extended to allow for a time-varying volatility risk premium. Indeed, one feature of our estimation strategy is that it allows us to capture time-variation in the volatility risk premium, 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 inefficient estimates of the risk premium parameter, leading to unreliable statistical inference.

To illustrate the procedure empirically, we apply the method to estimate the volatility risk premium associated with the S&P500 market index. We extend the method to allow for time variation in the stochastic volatility risk premium. We allow the premium to vary over time and to depend on macro-finance state 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, housing start number, 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

[2]The general strategy developed here is also related to the literature on market implied risk aversion (see, e.g., Jackwerth, 2000; Aït-Sahalia and Lo, 2000; Rosenberg and Engle, 2002; Brandt and Wang, 2003; Bliss and Panigirtzoglou, 2004; Gordon and St-Amour, 2004). The closest paper to ours is arguably Garcia et al. (2001), who estimate jointly the risk-neutral and objective dynamics, using a series expansion of option-implied volatility around the Black-Scholes implied volatility rather than model-free implied volatility. A recent paper by Wu (2005) also uses model-free realized and implied volatilities to estimate a flexible affine jump-diffusion model for volatility under the risk-neutral and objective measures.

[3]For directly traded assets like equities or bonds, empirical 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, e.g., Campbell, 1987; Fama and French, 1988; Campbell and Shiller, 1988a,b), while bond risk premia may be predicted by forward rates (see, e.g., Fama and Bliss, 1987; Cochrane and Piazzesi, 2004). However, with the notable exception of the recent study by Carr and Wu (2005), academic studies on the behavior of the volatility risk premium are rare, let alone its linkage to the overall economy.

variables. As such, our findings should be of direct interest to market participants and monetary policymakers who are concerned with the links between financial markets and the overall economy.[4] Further strengthening our results, we also find that the estimated time-varying volatility risk premium predicts future stock market returns better than several 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 of the Volatility Risk Premium

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.

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

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 faces a host of market microstructure 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. (2004b)). In particular, 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,

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

[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 (2004c) have recently extended these asymptotic distributional results to allow for leverage effects.

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]

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

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 (2005b) to the case of jump diffusions. Moreover, Jiang and Tian (2005b) 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 (2006),

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

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

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

[7] When implementing the conditional moment restrictions (5) and (8), it is useful to distinguish between two information sets—the continuous sigma-algebra $\mathcal{F}_t = \sigma\{V_s; s \leq t\}$, 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 $\mathcal{G}_t$.

of equations:

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

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 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 (2005b, p. 1313). We truncate the integration at lower and upper bounds of 70 and 143 percent of the current stock price S_t. We discretize the range of integration onto a grid of 150 points.[8] 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 $\rho = -0.80$.[9]

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

[8] Jiang and Tian (2005b) show that the discretization error is negligible for 20 or more grid points.

[9] 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.

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 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.[10]

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 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

$$-\lambda V_t = \mathrm{cov}_t \left(\frac{dm_t}{m_t}, dV_t \right) \tag{10}$$

[10]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.

where m_t denotes the pricing kernel, or marginal utility of wealth for the representative investor. Moreover, if we assume that the representative agent has a power utility function

$$U_t = e^{-\delta t} \frac{W_t^{1-\gamma}}{1-\gamma}, \tag{11}$$

where δ denotes a constant subjective time discount rate, and in equilibrium the agent holds the market portfolio, marginal utility equals $m_t = e^{-\delta t} W_t^{-\gamma}$. It follows from Itô's formula that[11]

$$\text{cov}_t \left(\frac{dm_t}{m_t}, dV_t \right) = -\gamma \rho \sigma V_t. \tag{12}$$

Combining (10) and (12), the constant relative risk aversion coefficient γ 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.

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 describing asset return dynamics.[12] 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

$$U_t = e^{-\delta t} \frac{W_t^{1-\gamma_t}}{1-\gamma_t}, \tag{13}$$

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,

$$d\gamma_t = \mu(\gamma_t)dt + \sigma(\gamma_t)dB_{3t}, \tag{14}$$

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

$$-\lambda_t V_t = \text{cov}_t \left(\frac{dm_t}{m_t}, dV_t \right) = -\gamma_t \rho \sigma V_t. \tag{15}$$

[11] A similar reduced-form argument is made by Bakshi and Kapadia (2003). For a much earlier formal general equlibrium treatment, see also Bates (1988) who allows for both stochastic volatility and jumps.

[12] 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).

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

$$dV_t = \kappa_t^*(\theta_t^* - V_t)dt + \sigma_t(\cdot)dB_{2t}^*, \tag{16}$$

where now $\kappa_t^* = \kappa + \lambda_t$ and $\theta_t^* = \kappa\theta/(\kappa + \lambda_t)$.[13] 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 Approximation for the Volatility Risk Premium

The discussion in the previous section shows how our approach can accommodate a time-varying volatility risk premium. 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 are quite simple to implement.

We will explore a simple dynamic model for the risk premium parameter, λ_t. We approximate the volatility risk premium parameter as following an augmented AR(1) process,

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

where "$\text{state}_{t,k}$" are macro-finance state variables. To be consistent with an absence of 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 we estimate the time-varying risk premium specification, we add lagged squared realized volatility, lagged implied volatility, and six out of the seven macro-finance covariates (without the redundant lagged realized volatility, see Section 4.5 for details) as additional 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. The number of additional moment conditions is equal to the number of additional parameters in (17), resulting in the same $\mathcal{X}^2(1)$ asymptotic distribution for the GMM omnibus test as in the estimation with a constant λ.

[13] The option pricing model in Heston (1993) allows for time-dependent coefficients, but the closed-form solutions may be complicated by any dynamic dependencies in λ_t.

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 based on the model-free approach discussed earlier.[14] Under appropriate assumptions, the concept of CBOE's "fair value of future variance" developed by Demeterfi et al. (1999) is identical to the "model-free implied variance" by Britten-Jones and Neuberger (2000), as well as the "risk-neutral expected value of the return variance" by Carr and Wu (2005) (see Jiang and Tian, 2005a, for detailed justification). 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. Moreover, since the new VIX index is constructed for replicating the risk-neutral variance of a fixed 30 days maturity, with monthly data there are in principle no issues with telescoping option maturities.

Our realized volatilities are based on the summation of the five-minute squared returns on the S&P500 index within the month.[15] 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, including the close-to-open five-minute interval. 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 the sum of the daily squared returns.

Figure 4 plots realized volatility, implied volatility, and their difference.[16] 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

[14] 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). Historical data on both the old and new VIX are directly available from the CBOE.

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

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

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.[17] 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 Preliminary Factor Analysis on Modeling Assumptions

Our analysis relies on the stochastic volatility model (1) and its corresponding risk-neutral counterpart (2). While our approach is quite general in that the functions $\mu_t(\cdot)$ and $\sigma_t(\cdot)$ can be left unspecified, there are some assumptions embedded in equations (1)-(2) that can be tested before going further:

- Realized volatility should follow an ARMA(1,1) process.
- Model-free implied volatility should follow an AR(1) process.
- One common factor drives both integrated volatility and implied volatility.

To test whether realized volatility and model-free implied volatility follow ARMA(1,1) and AR(1) processes, respectively, we compare the ARMA(1,1) with an ARMA(2,2) model and the AR(1) with an AR(2) model.[18] We note that previous studies have typically focused on daily stock return volatility, while we work with monthly volatilities. This suggests caution in extrapolating stylized facts from other studies to our data.[19]

[17] 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.

[18] Bollerslev and Zhou (2002, Appendix B) show that a two-factor stochastic volatility model implies that the corresponding realized volatility follows an ARMA(2,2) process.

[19] In particular, we also tested for long-memory type dependencies in each of the series using the modified rescaled range statistic of Lo (1991) along with log-periodogram regressions. None of these results (available

Table 5 shows the results from estimating ARMA(1,1) and ARMA(2,2) models for realized volatility and AR(1) and AR(2) models for model-free implied volatility. A traditional time-series approach to model selection looks for the most parsimonious model that adequately captures the time-series variation in the data by producing residuals that are roughly white-noise. By this standard, the ARMA(1,1) is preferred to the ARMA(2,2) for realized volatility and the AR(1) is preferred to the AR(2) for model-free implied volatility: all four models produce white-noise residuals (the portmanteau test does not reject the hypothesis of white-noise residuals) and the ARMA(1,1) and AR(1) models minimize Schwartz's Bayesian information criterion. [20]

The third model implication listed above holds that realized and model-free implied volatility are driven by a single common factor. To informally investigate this, we performed a standard (unconditional) principal components analysis (PCA) on the two volatility series. The PCA indicates that the first principal component explains 79 percent of the variance of the two series. This is high enough to assure us that our (implicit) assumption of a single common factor is not obviously violated by the data. Of course, the remaining 21 percent of the variability could be explained, at least in part, by a time-varying volatility risk premium.

Motivated by these encouraging preliminary findings, we next turn our attention to the results from the more formal GMM-based estimation strategy.

4.5 GMM Estimation Results

Table 6 reports the GMM estimation results for two volatility risk premium specifications: (i) a constant λ; (ii) a time-varying λ_t driven by shocks to macro-finance variables as in equation (17).[21]

As seen in the first column of the table, when we restrict the risk premium to be constant, the estimated λ is negative and statistically significant. This finding is consistent with other papers that have found a negative risk premium on stochastic volatility. However, the chi-

upon request) produced any strong evidence for long memory.

[20]While a standard likelihood ratio (LR) test based on the reported maximized values for the Gaussian log likelihoods results in the ARMA(1,1) for the realized volatility being rejected in favor of the ARMA(2,2) model, the errors from both models (the volatility-of-volatility) are heteroskedastic, so the standard critical values overstates the difference in fit between the two models.

[21]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 4, point toward a high degree of volatility persistence in the (latent) V_t process.

square omnibus test of overidentifying restrictions rejects the overall specification at the 10% (although not at the 5%) level.

The second column presents the results obtained by explicitly including macro-finance covariates. To select the macro-finance variables in the time-varying risk premium specification, we did an extensive search over 29 monthly data series. The series we used are listed in Table 10. 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.[22] 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.[23]

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 0.68. 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).[24] The impact of AAA bond spread over Treasuries (0.19) likely reflects a business cycle effect

[22] 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 10. 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.

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

[24] This result contradicts the finding 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 risk aversion coefficient; or caused by their particular methodology of assuming a constant risk aversion coefficient, but splitting the sample into periods of high and low volatilities.

(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.

4.6 Robustness Checks

The consistency of the realize volatility hinges on the idea of ever finer sampled observations over a fixed-length time interval. Yet, it is well-known that a host of market microstructure frictions, including price discreteness and staleness (see, e.g., Stoll and Whaley, 1990), invalidates the basic underlying martingale assumption at the ultra high frequencies.[25] In order to investigate the robustness of our findings based on the 5-minute returns with respect to this issue, we re-estimate our model with realized volatilities constructed from coarser sampled 30-minute returns. As seen from the first column in Table 7, the sign and significance of the parameter estimates are qualitatively and quantitatively very similar to the previous findings, and the p-values for the overall goodness-of-fit tests for the model are also remarkably close (0.697 versus 0.681).

The Monte Carlo experiment in Section 3 indicates that the use of Black-Scholes as opposed to model-free implied volatilities do not result in materially different estimates for the (constant) volatility risk premium, when the sample size is relatively small (150 months). Of course, in large samples (600 months), the BSIV based estimate is ultimately rejected. To investigate the sensitivity of our results to the specific volatility measure, the second column in Table 7 reports the GMM estimation results obtained by using BSIV in place of the MFIV measure. Compared to the original results in the last column in Table 6, the parameter estimates are generally close, as are their standard errors. Again, this is completely in line

[25] A large, and rapidly growing, literature have sought different ways in which to best deal with these complications in the construction of improved realized volatility measures; see, e.g., Aït-Sahalia et al. (2005); Bandi and Russell (2005); Hansen and Lunde (2005).

with the Monte Carlo evidence presented earlier, which suggest that a much larger sample size is needed to effectively distinguish between the two volatility measures. Of course, only the approach using the model-free implied volatilities is formally justified.

It is been widely argued in the literature that most major market indices contain jumps, or price discontinuities (see, e.g., Bates, 1996; Bakshi et al., 1997; Pan, 2002; Chernov et al., 2003). This suggests that it may be important to separately consider jump risk when estimating the stochastic volatility risk premium. However, the moment condition in equation (8) only identifies a single risk premium parameter. Thus, to check the robustness of our estimates with respect to jumps, we simply fix the jump risk premium and re-estimate the resulting volatility risk premium. To identify the jumps, we follow Barndorff-Nielsen and Shephard (2004b,d) in using the difference between the realized volatility,

$$\mathcal{V}^n_{t,t+\Delta} \equiv \sum_{i=1}^{n} \left[p_{t+\frac{i}{n}(\Delta)} - p_{t+\frac{i-1}{n}(\Delta)} \right]^2 \rightarrow \int_t^{t+\Delta} V_s ds + \int_t^{t+\Delta} J_s^2 ds, \tag{18}$$

and the so-called *bi-power variation*,

$$\mathcal{BV}^n_{t,t+\Delta} \equiv \sum_{i=2}^{n} \left[p_{t+\frac{i}{n}(\Delta)} - p_{t+\frac{i-1}{n}(\Delta)} \right] \left[p_{t+\frac{i-1}{n}(\Delta)} - p_{t+\frac{i-2}{n}(\Delta)} \right] \rightarrow \int_t^{t+\Delta} V_s ds, \tag{19}$$

for separating the monthly diffusive and jump volatility, $\int_t^{t+\Delta} V_s ds$ and $\int_t^{t+\Delta} J_s^2 ds$, respectively, under the objective measure (see also Andersen et al., 2004a; Huang and Tauchen, 2005). Following Jiang and Tian (2005b) and Carr and Wu (2005) the model-free implied volatility may be similarly decomposed under the risk-neutral expectation,

$$\mathrm{IV}^*_{t,t+\Delta} \approx \mathrm{E}^* \left(\int_t^{t+\Delta} V_s ds \middle| \mathcal{F}_t \right) + \mathrm{E}^* \left(\int_t^{t+\Delta} J_s^2 ds \middle| \mathcal{F}_t \right). \tag{20}$$

Since it isn't possible to separately identify a volatility risk premium and a jump risk premium, we instead perform a counter-factual experiment and assume that the risk-neutral and objective expectations of the jump contribution differ by a constant multiple. In particular, under Jump Scenario (h),

$$\mathrm{E}^* \left(\int_t^{t+\Delta} J_s^2 ds | \mathcal{F}_t \right) \;=\; h \cdot \mathrm{E} \left(\int_t^{t+\Delta} J_s^2 ds | \mathcal{F}_t \right) \tag{21}$$

The corresponding estimation results, reported in the last three columns in Table 7, show that the level, persistence, and macro-finance sensitivities of the volatility risk premium are all largely unaffected. Interestingly, on comparing the three jumps scenarios, the overall goodness-fit appears to improve as the price of jump risk increases. Nonetheless, overall the results clearly confirm the robustness of our previous findings with respect to the specific jump dynamics and risk prices entertained here.

4.7 Comparing Alternative Estimates of Time-Varying Risk Premia

Several alternative procedures for estimating the time-varying volatility risk premium have previously been implemented in the literature. One 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. In the context of our modeling framework, such an approach would produce the time-varying risk premium shown in the middle panel of Figure 6. The general shape of Figure 6 matches the simple difference between implied and realized volatilities shown in the bottom panel of Figure 4. Papers that have taken this approach include Rosenberg and Engle (2002, p. 363) and Tarashev et al. (2003, p. 62). Their risk premium estimates are shown in Figure 7.

As previously noted, because this approach attributes every wiggle in the data to changes in the risk premium, it produces a very 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 that monthly volatility risk premiums estimated in this way are implausibly volatile.

A second approach for estimating investors' "risk appetite," popular among market participants, is to construct a simple average of macro-finance variables.[26] The bottom panel of Figure 6 shows such a average index constructed from the 29 (standardized) macro-finance variables listed in Table 10, 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.

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

A third approach to estimating risk premium parameters comes from the consumption-based asset pricing literature. This approach typically assumes that risk premia are constant, or if risk preferences 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.[27] Such a modeling strategy explicitly prevents the risk premia from being excessively variable in the short-run.[28]

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 series in that same figure. The estimated risk premium 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. Nearly all of the peaks in the series are readily identifiable with major macroeconomic or financial market developments, including the 1994 monetary policy tightening 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". The chart suggests that the risk premium often rises sharply but declines only gradually.

4.8 Stock Return Predictability

Because the volatility risk premium can be related to investor risk aversion, it may be 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. The top panel of Table 8 reports the results of simple univariate regressions of monthly S&P500 excess returns on the volatility risk premium and on the most significant individual variables from the pool of covariates listed in Table 10. The extracted volatility

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

[28] Along these lines, the estimated risk aversions of Brandt and Wang (2003, p. 1481) and Gordon and St-Amour (2004, p. 249) do not pick out most recessions (as reproduced in Figure 8).

risk premium has the highest predictive power with an adjusted R^2 of 4.4%.[29] The second best predictor is the S&P500 P/E ratio with an adjusted R^2 of 2.2%. Next in order are industrial production and nonfarm payrolls with adjusted R^2's of 1.0% and 0.5%, respectively. Dividend yield — a significant predictor according to many other studies — only explains 0.3% of the monthly return variation. These results are consistent with previous findings that macroeconomic state variables do predict returns, though the predictability measured by adjusted 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 highest adjusted 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 8. Interestingly, none of the macro-finance variables 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 6), so the finding that these variables are "driven out" when included together with the premium is not necessarily surprising. However, the macro variables entering the model for λ_t only impact the returns indirectly through the temporal variation in the premium, and the volatility risk premium itself is also estimated from a different set of moment conditions involving only the model-free realized and options implied volatilities.

Table 9 examines stock return predictability over a quarterly horizon. In addition to the volatility risk premium and the P/E ratio from the last month of the previous quarter, which were the two most important predictor variables in the monthly regressions shown in table 8, we now add the quarterly consumption-wealth ratio. The consumption-wealth ratio, termed CAY, has previously been found by Lettau and Ludvigson (2001) to be significant in explaining longer-horizon returns.

The first three regressions in Table 9 show that each of the three predictor variables are statistically significant in univariate regressions. The volatility risk premium results in the highest individual adjusted R^2 of 15.6%, higher than its monthly R^2 of 4.4%. The next three lines of the table show results for two right-hand side predictor variables. Adding the P/E ratio or CAY to the volatility risk premium does not help predict excess returns: the

[29]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. Also, the persistence of the right-hand-side variables in predictive regressions can cause biased coefficient estimates and inferences (Stambaugh, 1999; Amihud and Hurvich, 2004).

adjusted R^2 actually falls and only the risk premium is statistically significant. Combining the P/E ratio and CAY in the same regression renders both insignificant.

These results for quarterly excess returns reinforce the earlier findings for the monthly returns in Table 8. The estimated volatility risk premium appears to be a new and powerful stock market predictor over longer quarterly horizons. The fact that the volatility risk premium and CAY or the P/E ratio partially crowd each other out suggests that these three variables capture some common time-variation of risk premia in the economy.

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 does 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 be interesting to further explore the cross sectional pricing implications of this finding. Does the volatility risk premium represent a systematic priced risk factor?[30] Also, what is the link between stock and bond market volatility risk premia? Lastly, better estimates for the volatility risk premium are, of course, of direct importance for derivatives pricing. We leave further work along these lines for future research.

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

References

Adrian, Tobias and Joshua Rosenberg (2005), "Stock Returns and Volatility: Pricing the Long-Run and Short-Run Components of Market Risk," *Working Paper*, Federal Reserve Bank of New York.

Aït-Sahalia, Yacine and Bob Kimmel (2005), "Maximum Likelihood Estimation of Stochastic Volatility Models," *Journal of Financial Economics*, forthcoming.

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

Aït-Sahalia, Yacine, Per Mykland, and Lan Zhang (2005), "How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise," *Review of Financial Studies*, vol. 18, 351–416.

Amihud, Yakov and Clifford Hurvich (2004), "Predictive Regressions: A Reduced-Bias Estimation Method," *Journal of Financial and Quantitative Analysis*.

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, and Francis X. Diebold (2004a), "Some Like it Smooth, and Some Like it Rough: Untangling Continuous and Jump Components in Measuring, Modeling and Forecasting Asset Return Volatility," *Working Paper*, Duke University.

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 (2004b), "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, Charles Cao, and Zhiwu Chen (1997), "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, vol. 52, 2003–2049.

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.

Bandi, Federico and Jeffrey Russell (2005), "Separating Microstructure Noise from Volatility," *Journal of Financial Economics*, forthcoming.

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), "Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation," *Working Paper*, Nuffield College, Oxford University.

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

Barndorff-Nielsen, Ole and Neil Shephard (2004d), "Power and Bipower Variation with Stochastic Volatility and Jumps," *Journal of Financial Econometrics*, vol. 2, 1–48.

Bates, David S. (1988), "Pricing Options on Jump-Diffusion Process," *Working Paper*, Rodney L. White Center, Wharton School.

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 (2006), "Volatility Puzzles: A Simple Framework for Gauging Return-Volatility Regressions," *Journal of Econometrics*, vol. 131, 123–150.

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 (2005), "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, A. Ronald Gallant, Eric Ghysels, and George Tauchen (2003), "Alternative Models for Stock Price Dynamics," *Journal of Econometrics*, forthcoming.

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.

Hansen, Peter Reinhard and Asger Lunde (2005), "Realized Variance and Market Microstructure Noise," *Journal of Business and Economic Statistics*, forthcoming.

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.

Huang, Xin and George Tauchen (2005), "The Relative Contribution of Jumps to Total Price Variance," *Journal of Financial Econometrics*, vol. 3, 456–499.

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 (2005a), "Gauging the "Investor Fear Gauge": Problems in the CBOE's New Volatility Index and a Simple Solution," *Working Paper*, Finance Department, University of Arizona.

Jiang, George and Yisong Tian (2005b), "Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, pages 1305–1342.

Jones, Christopher S. (2003), "The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets," *Journal of Econometrics*, vol. 116, 181–224.

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.

Lo, Andrew (1991), "Long-Term Memory in Stock Market Prices," *Econometrica*, vol. 59, 1279–1313.

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.

Stambaugh, Robert F. (1999), "Predictive Regressions," *Journal of Financial Economics*, vol. 54, 375–421.

Stoll, Hans and Robert Whaley (1990), "The Dynamics of Stock Index and Stock Index Futures Returns," *Journal of Financial and Quantitative Analysis*, vol. 25, 441–468.

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.

Wu, Liuren (2005), "Variance Dynamics: Joint Evidence from Options and High-Frequency Returns," *Working Paper*, Baruch College of CUNY.

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: Diagnosing Alternatives for Realized and Implied Volatilities

For realized volatility, we estimate the ARMA(2,2) model

$$\mathcal{V}_{t+\Delta,t+2\Delta} = \beta + \alpha_1 \mathcal{V}_{t,t+\Delta} + \alpha_2 \mathcal{V}_{t-\Delta,t} + e_{t+\Delta,t+2\Delta} + \theta_1 e_{t,t+\Delta} + \theta_2 e_{t-\Delta,t}$$

where the ARMA(1,1) model restricts $\alpha_2 = \theta_2 = 0$. For implied volatility, we estimate the AR(2) model

$$\mathrm{IV}^*_{t+\Delta,t+2\Delta} = \beta + \alpha_1 \mathrm{IV}^*_{t,t+\Delta} + \alpha_2 \mathrm{IV}^*_{t-\Delta,t} + e_{t+\Delta,t+2\Delta}$$

where the AR(1) model restricts $\alpha_2 = 0$. The columns "Q(10)", "BIC" and "log L" show the Ljung-Box portmanteau test statistic for white-noise residuals using 10 lags, Schwartz's Bayesian information criterion, and the maximized value of the Gaussian log-likelihood, respectively. Q(10) is asymptotically distributed as $\chi^2(10)$ under the null hypothesis of white-noise residuals; the 95 percent critical value of a $\chi^2(10)$ random variable is 18.3. Standard errors are shown in parentheses. All regressions are estimated over January 1990 to May 2004.

Realized Volatility	β	α_1	α_2	θ_1	θ_2	Q(10)	BIC	$\log L$
ARMA(1,1)	.052	.70	—	-.32	—	14.5	2951.7	-1465.5
	(.050)	(.07)		(.10)				
ARMA(2,2)	.052	.21	.45	.27	-.42	7.9	2952.4	-1460.7
	(.053)	(.26)	(.12)	(.26)	(.11)			
Implied Volatility	β	α_1	α_2	θ_1	θ_2	Q(10)	BIC	$\log L$
AR(1)	.044	.76	—	—	—	13.4	2323.0	-1153.8
	(.01)	(.04)						
AR(2)	.044	.80	-.05	—	—	12.5	2327.8	-1153.6
	(.01)	(.09)	(.08)					

Table 6: Estimation of Volatility Risk Premium

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.

	Constant	Macro-Finance
λ	-1.793 (0.216)	
a		-0.122 (0.051)
b		0.933 (0.030)
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)	0.169 (0.681)

Table 7: Robustness Checks

All of the macro-finance variables are the same as in the previous table. 30-Minute stands for sampling the returns every 30-minute instead of every 5-minute. BSIV is the Black-Scholes implied volatility, replacing the model-free implied volatility. Jump (1), (2), and (3) represent the cases where the risk-neutral expectation of jump-squared is assumed to be the same as, double, and tipple the objective expectation.

Macro-Finance Specification	30-Minute	BSIV	Jump (1)	Jump (2)	Jump (3)
a	-0.156 (0.040)	-0.129 (0.065)	-0.158 (0.055)	-0.145 (0.048)	-0.129 (0.050)
b	0.891 (0.028)	0.933 (0.033)	0.916 (0.032)	0.919 (0.029)	0.925 (0.030)
c_1 Realized Volatility	-0.262 (0.102)	-0.281 (0.033)	-0.352 (0.049)	-0.365 (0.063)	-0.386 (0.070)
c_2 Moody AAA Bond Spread	0.107 (0.065)	0.151 (0.029)	0.225 (0.039)	0.229 (0.054)	0.236 (0.056)
c_3 Housing Start	-0.175 (0.054)	-0.142 (0.056)	-0.199 (0.056)	-0.203 (0.056)	-0.209 (0.057)
c_4 S&P500 P/E Ratio	0.136 (0.021)	0.129 (0.014)	0.144 (0.013)	0.147 (0.012)	0.152 (0.011)
c_5 Industrial Production	0.064 (0.054)	0.056 (0.026)	0.110 (0.028)	0.117 (0.032)	0.124 (0.033)
c_6 Producer Price Index	-0.037 (0.019)	-0.032 (0.021)	-0.048 (0.024)	-0.046 (0.022)	-0.043 (0.021)
c_7 Payroll Employment	-0.019 (0.043)	-0.013 (0.020)	-0.045 (0.020)	-0.051 (0.023)	-0.059 (0.024)
$\mathcal{X}^2$(d.o.f. = 1) (p-Value)	0.151 (0.697)	0.373 (0.541)	0.150 (0.698)	0.059 (0.807)	0.004 (0.951)

Table 8: Monthly Stock Market Return Predictability

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.

Variables	Intercept	(s.e.)	Slope	(s.e.)	Adj. R-Square
Volatility Risk Premium	-22.283	(10.120)	14.567	(5.045)	0.044
S&P500 PE Ratio	35.939	(13.750)	-1.272	(0.566)	0.022
Industrial Production	-0.926	(5.495)	1.992	(1.226)	0.010
Nonfarm Payroll Employment	-0.311	(5.478)	3.643	(2.624)	0.005
26 Other Macro-Finance Variables					<0.005

Joint Estimation	Including λ_t		Excluding λ_t	
Variables	Parameter	(s.e.)	Parameter	(s.e.)
Intercept	-9.279	(21.245)	32.619	(15.306)
Volatility Risk Premium	11.575	(5.438)		
S&P500 PE Ratio	-0.430	(0.594)	-1.259	(0.585)
Industrial Production	1.841	(2.245)	2.719	(2.204)
Nonfarm Payroll Employment	-1.388	(4.721)	-3.306	(4.657)
Adj. R-Square	0.037		0.022	

Table 9: Quarterly Stock Market Return Predictability

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.

Intercept (s.e.)	Risk Premium (s.e.)	PE Ratio (s.e.)	CAY (s.e.)	Adj. R-Square
-27.087 (12.215)	16.477 (6.145)			0.156
41.044 (16.289)		-1.543 (0.692)		0.086
2.413 (4.308)			5.378 (2.031)	0.068
-6.927 (21.494)	13.619 (5.966)	-0.609 (0.649)		0.151
-23.338 (13.413)	14.253 (6.970)		2.042 (1.946)	0.149
29.737 (27.888)		-1.098 (1.139)	2.476 (3.352)	0.078
-10.698 (25.008)	13.204 (6.180)	-0.432 (0.926)	1.146 (2.860)	0.136

Table 10: 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

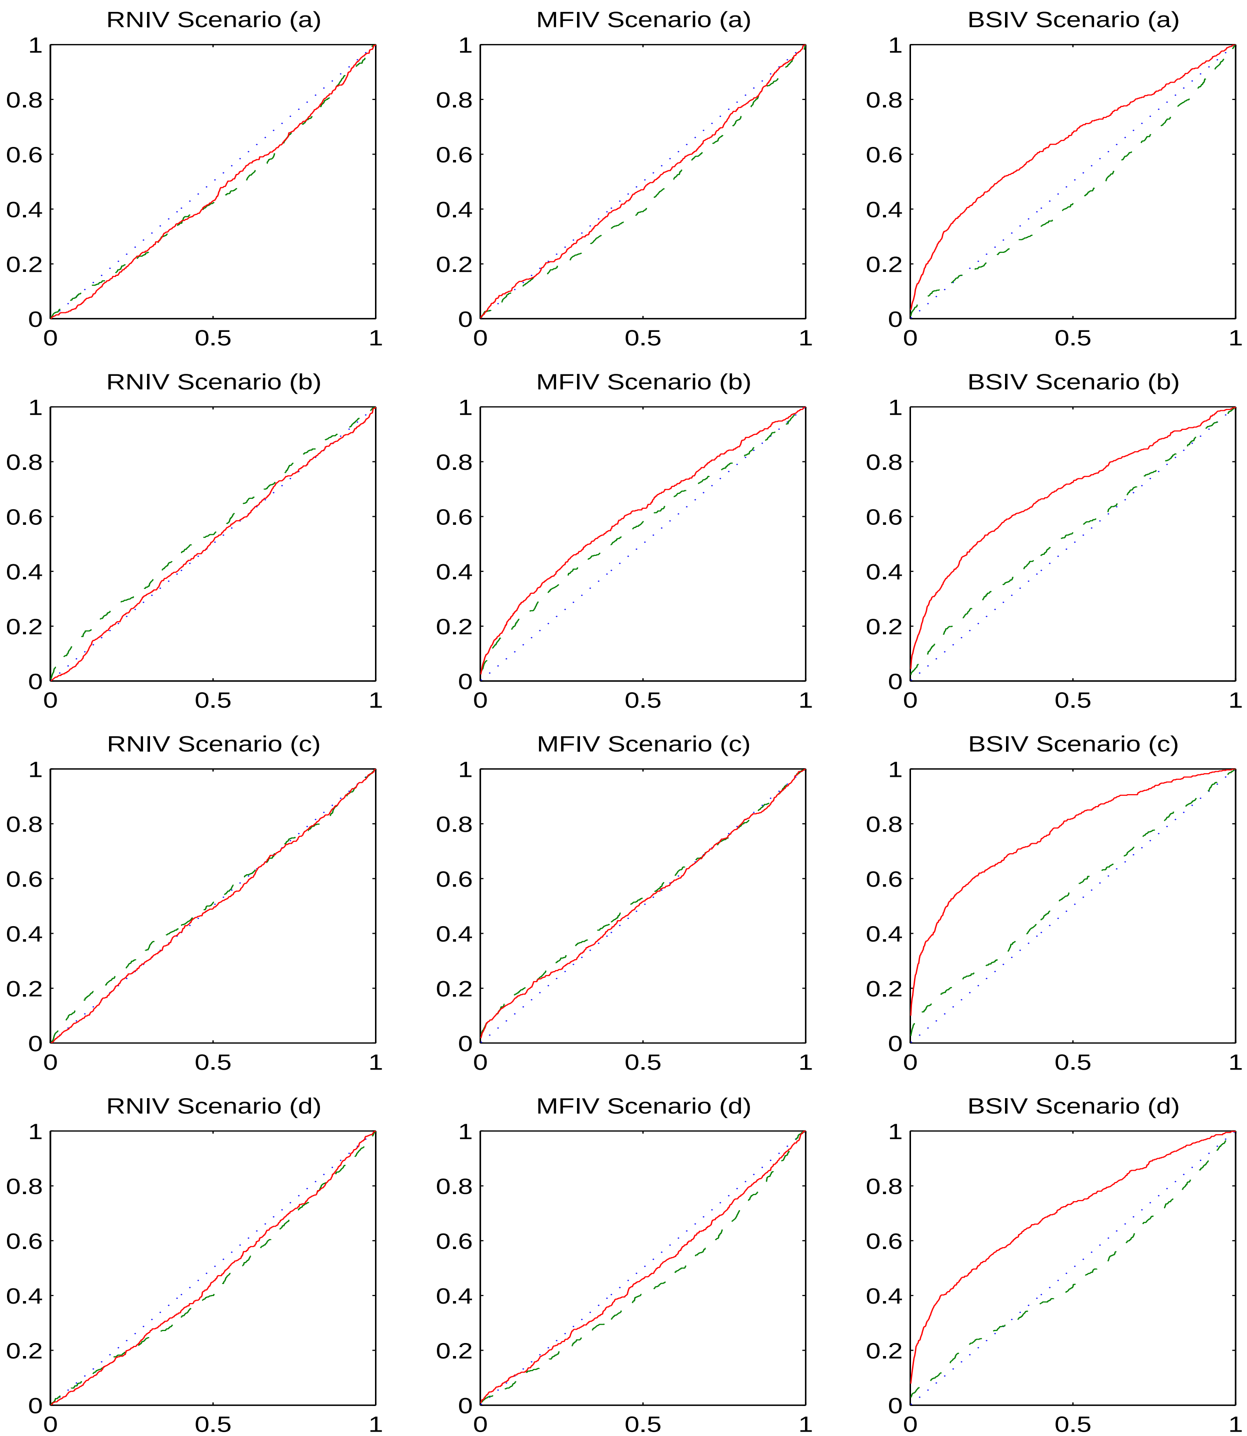

Figure 1: Wald Test for Risk Premium with True Integrated Volatility.
The X-axis gives the nominal level of the test and Y-axis the probability of rejection. The dotted line represents the uniform reference distribution, the dash line is T = 150, and the solid line is T = 600.

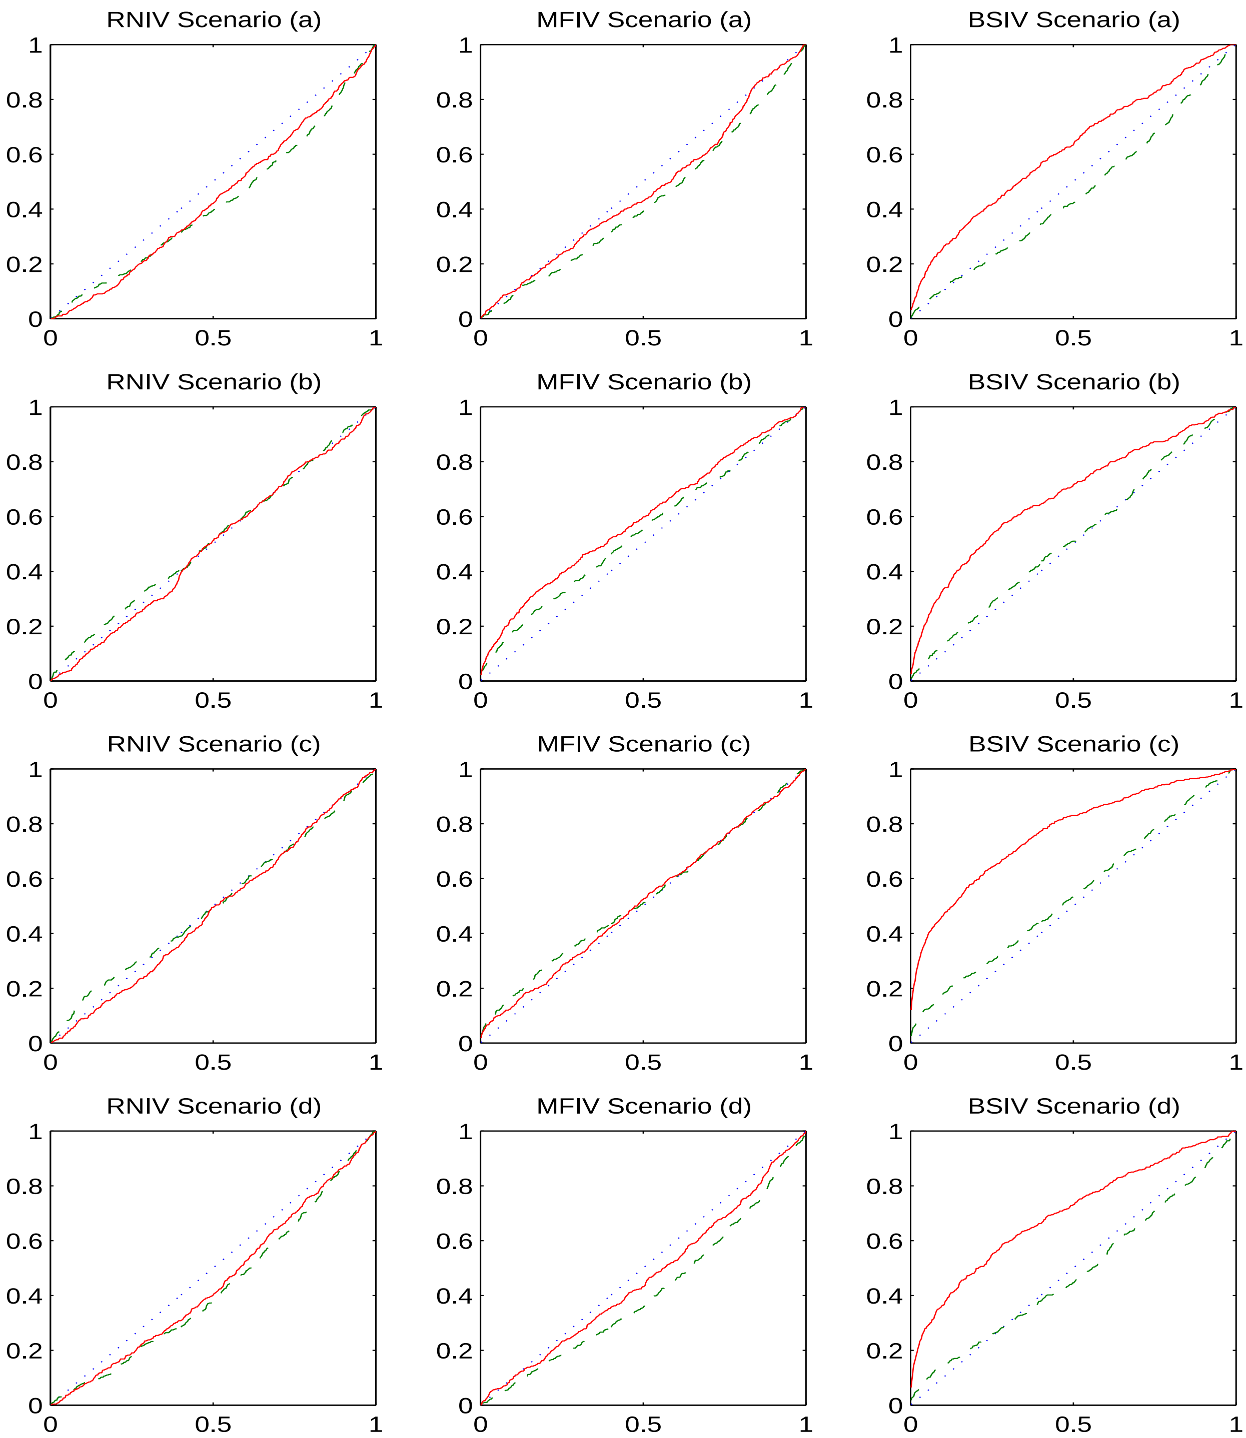

Figure 2: Wald Test for Risk Premium with Five-Minute Return Realized Volatility. The X-axis gives the nominal level of test and the Y-axis the probability of rejection. The dotted line represents the uniform reference distribution, the dash line is T = 150, and the solid line is T = 600.

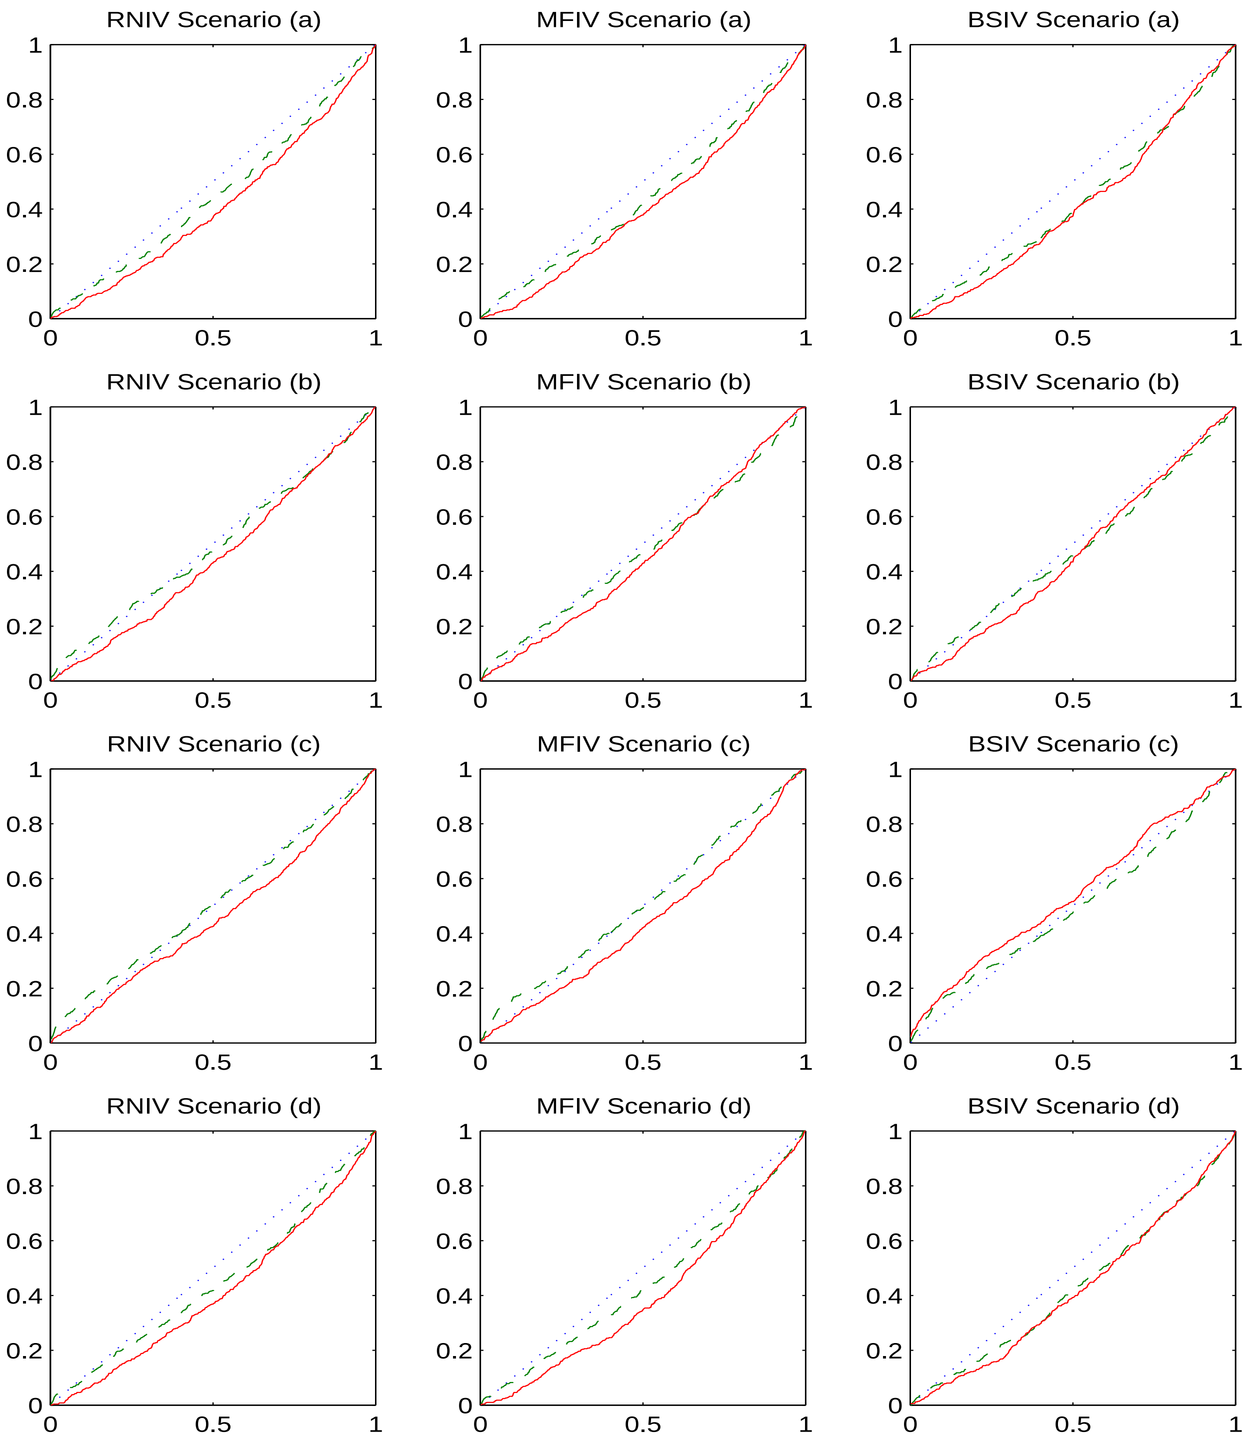

Figure 3: Wald Test for Risk Premium with Daily Return Realized Volatilities.
The X-axis gives the nominal level of the test and the Y-axis the probability of rejection. The dotted line represents the uniform reference distribution, the dash line is T = 150, and the solid line is T = 600.

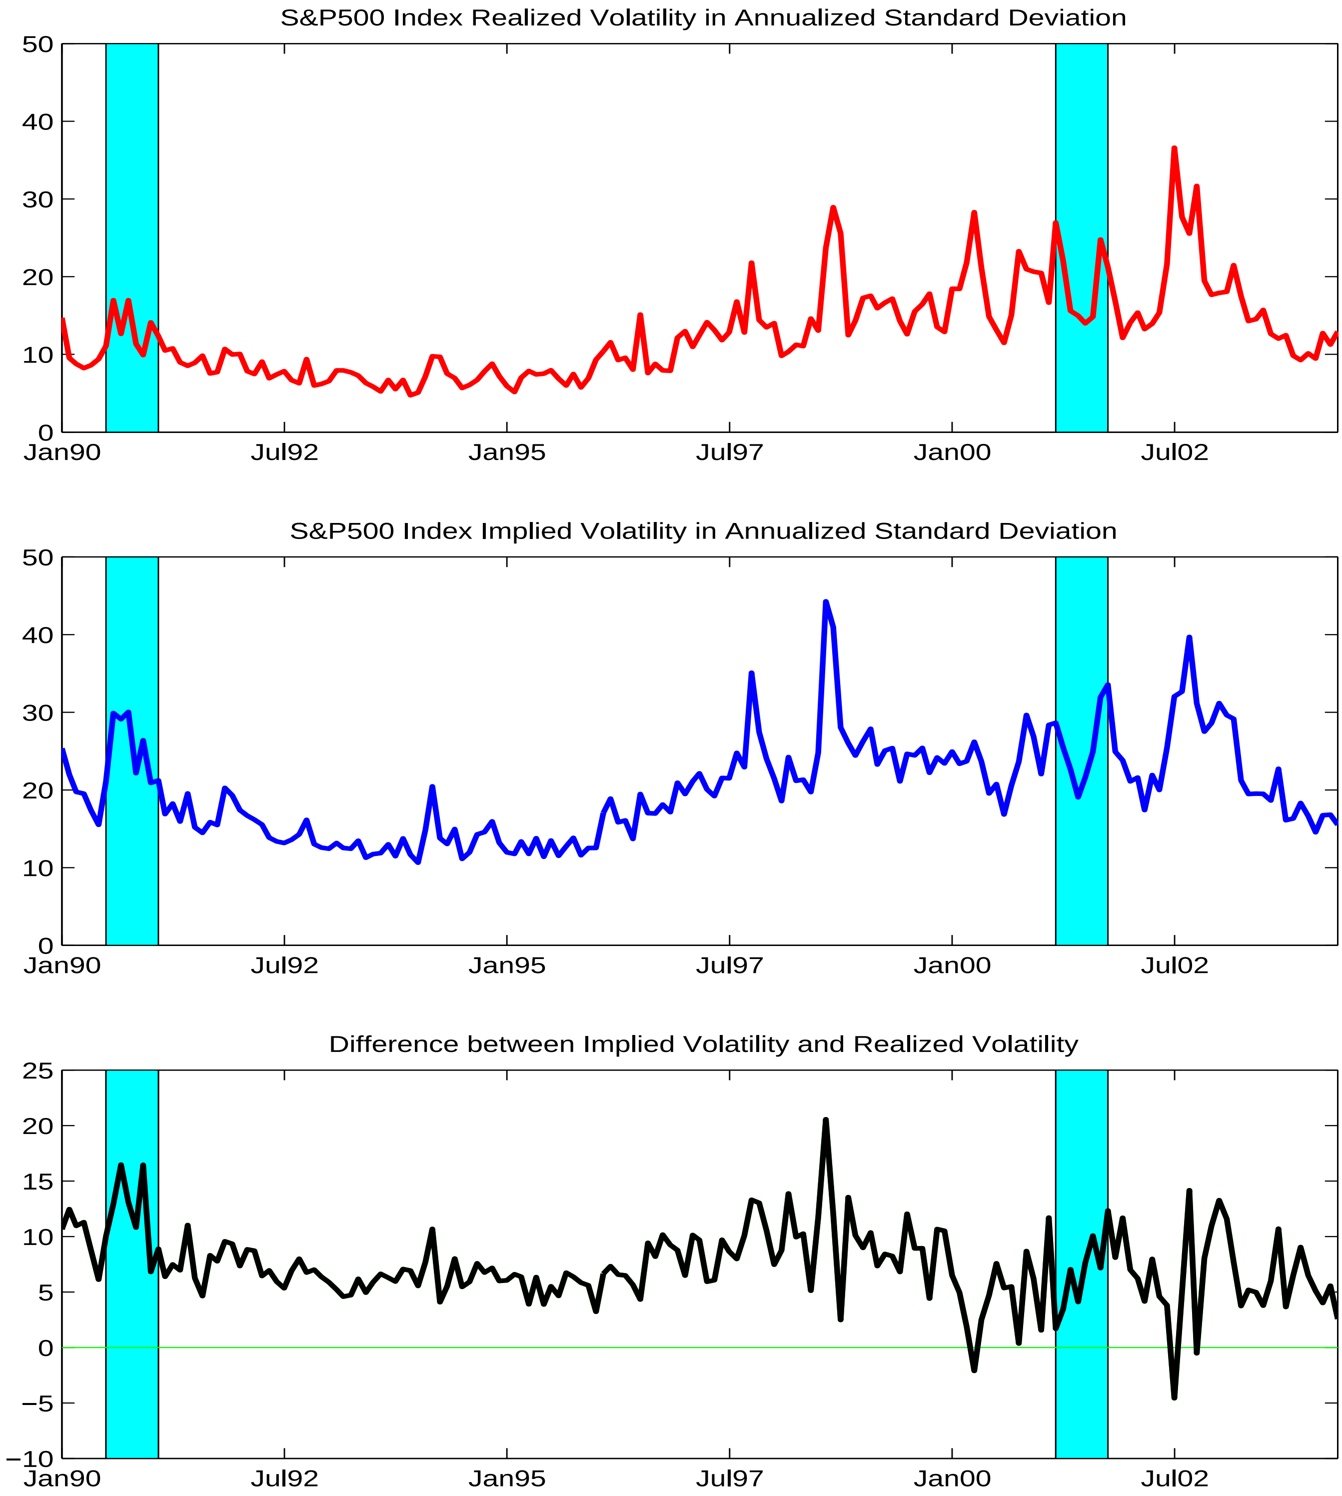

Figure 4: Model-Free Realized and Implied Volatilities

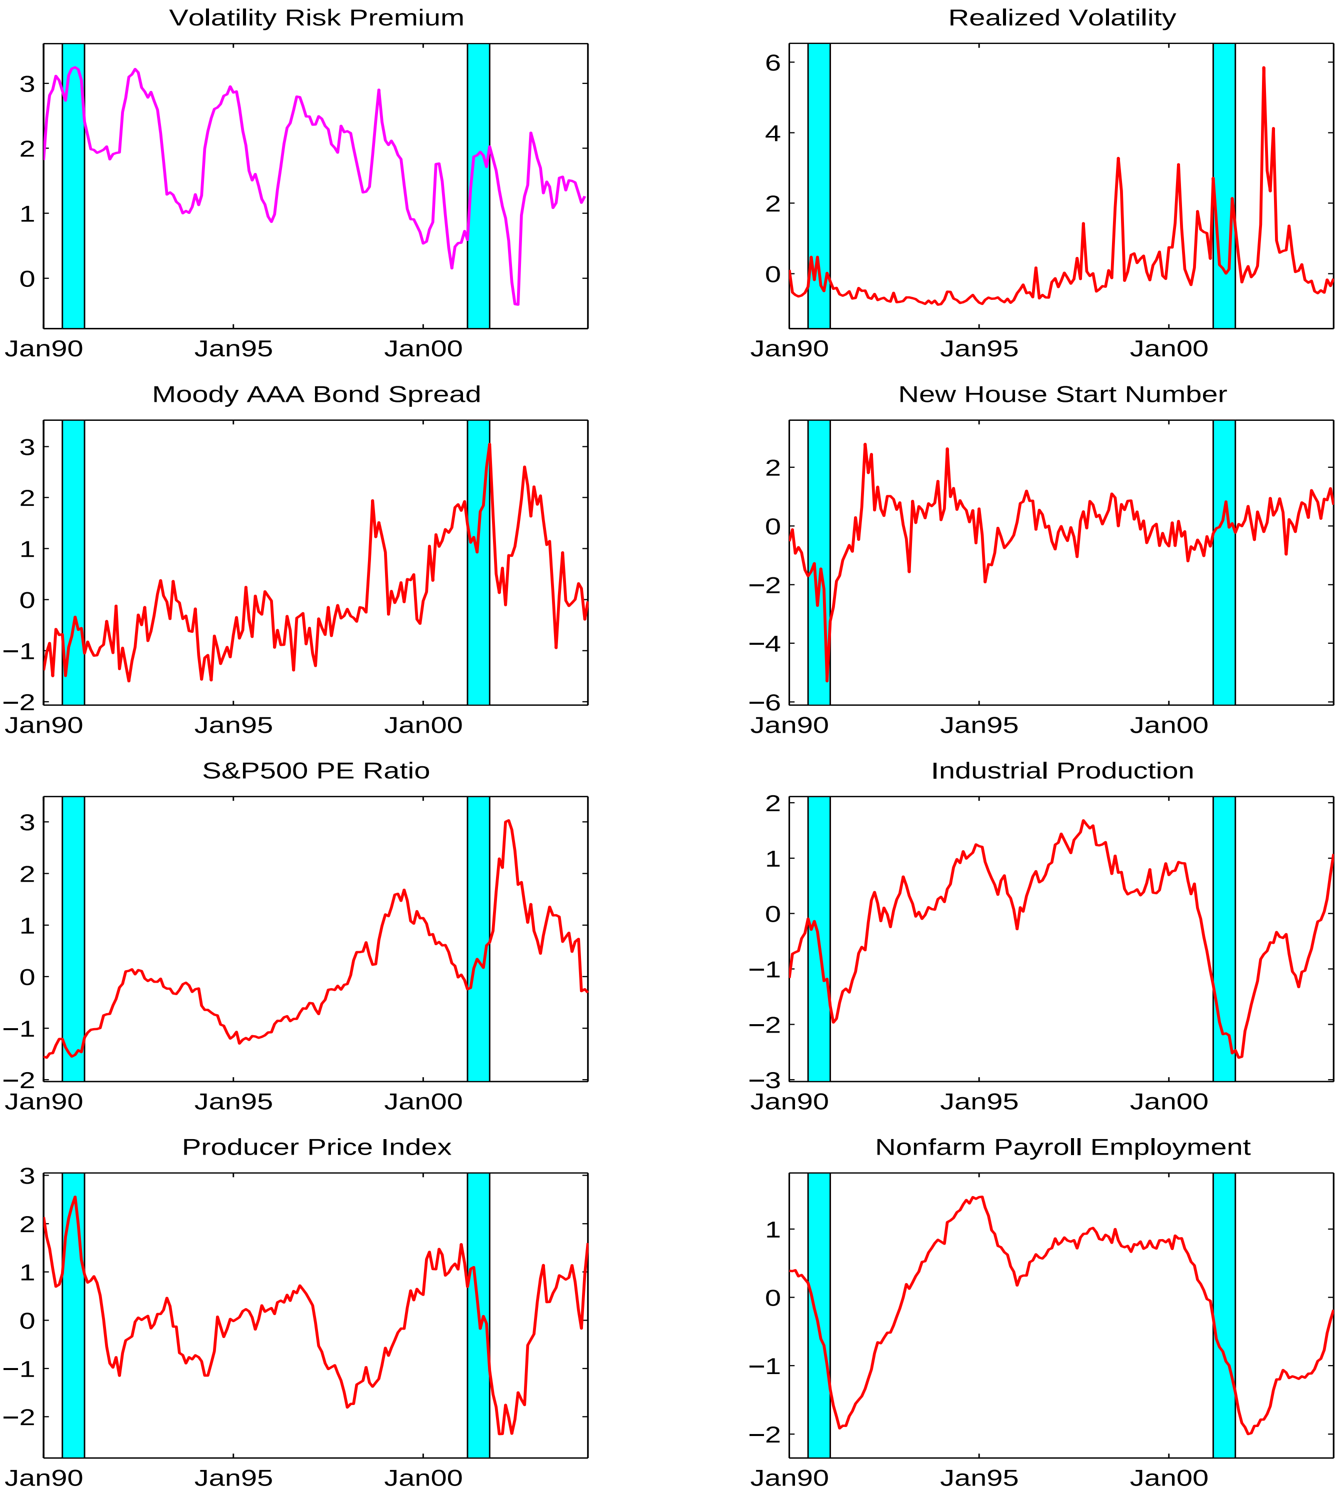

Figure 5: Standardized Macro-Finance Covariates

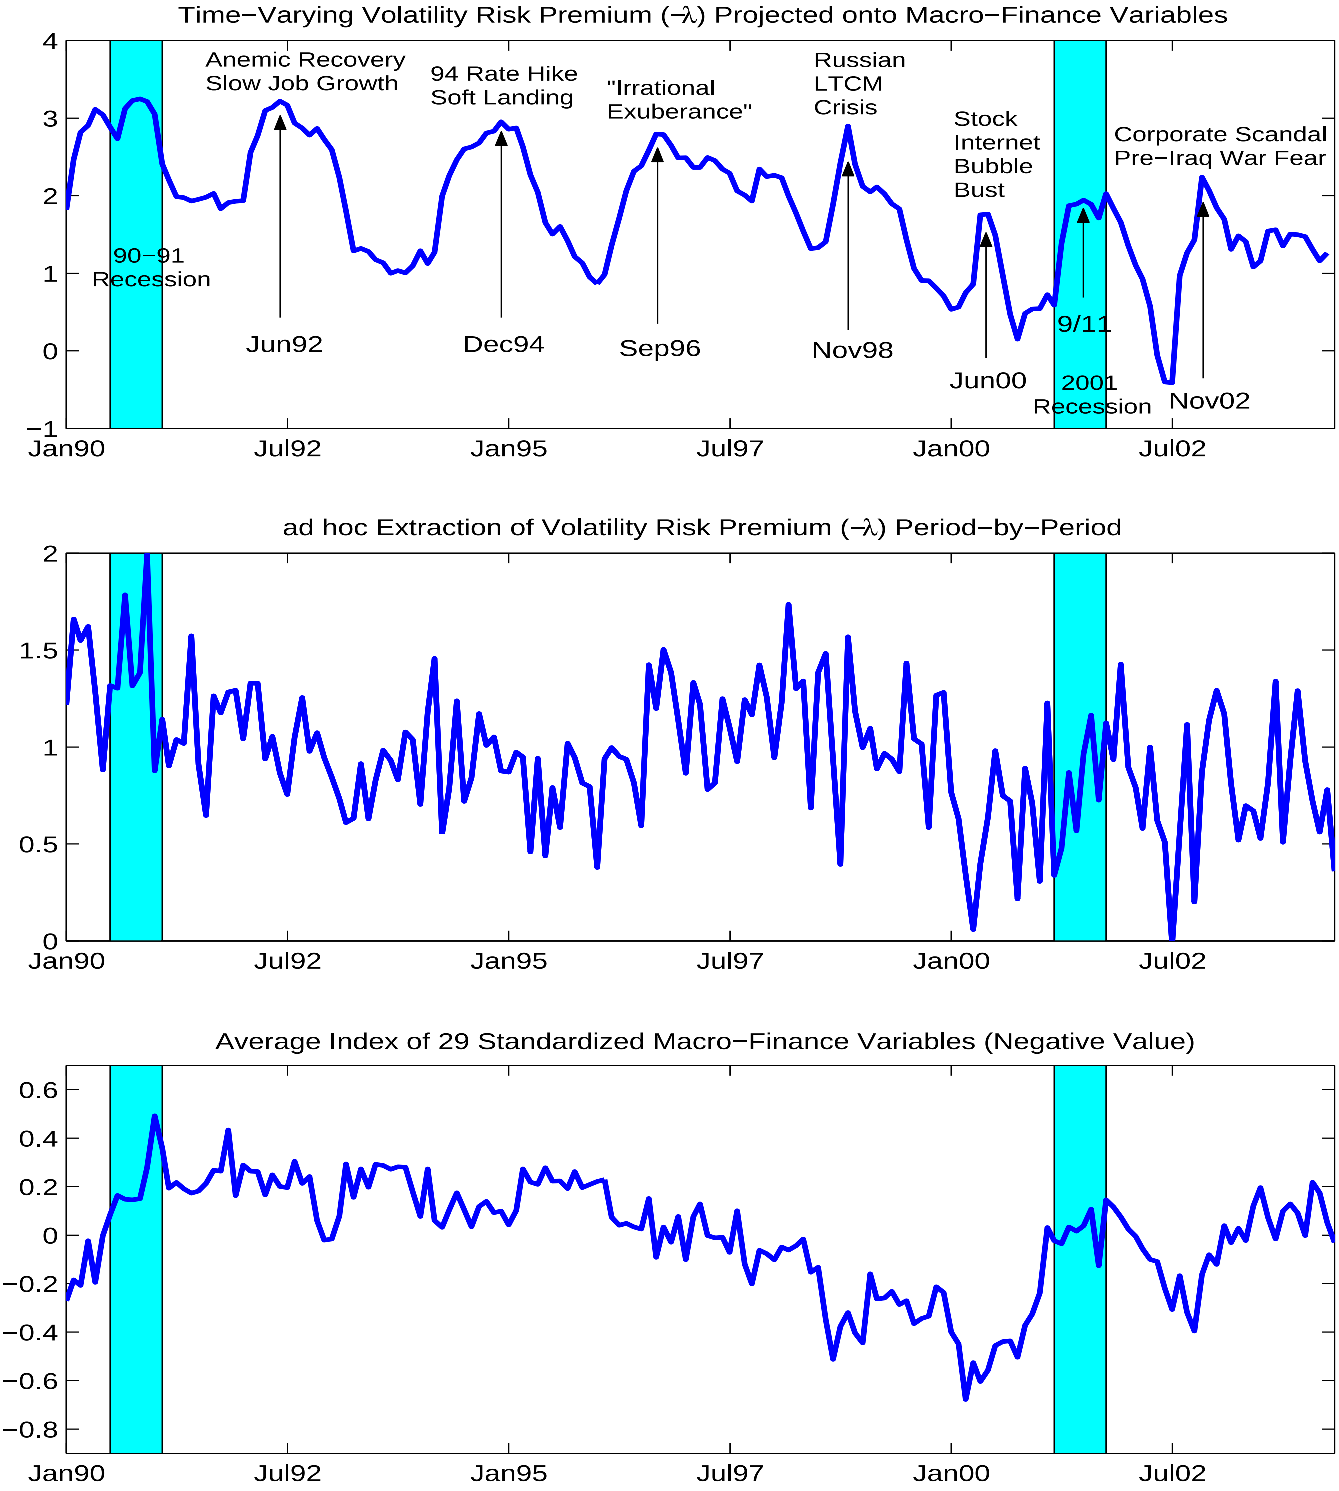

Figure 6: Time-Varying Volatility Risk Premium and Other Indices

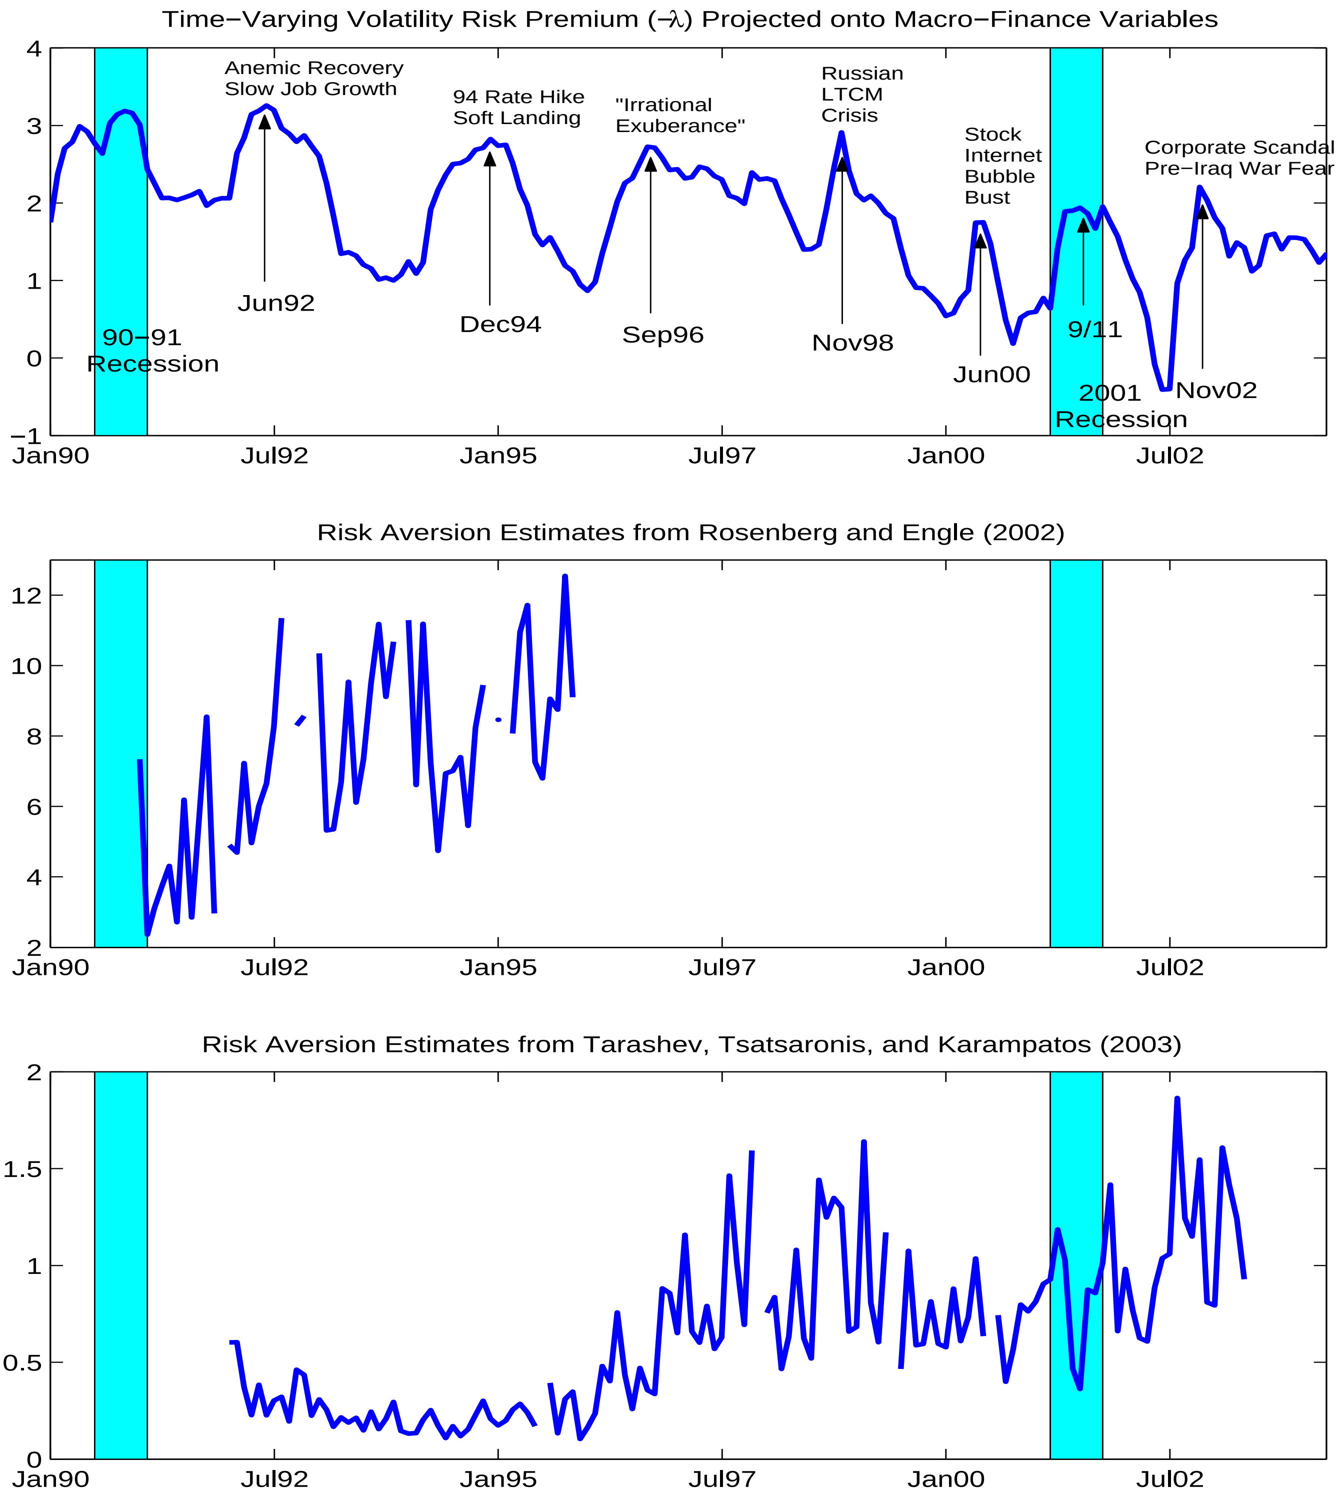

Figure 7: Comparison with Estimates Using Risk-Neutral and Objective Distributions

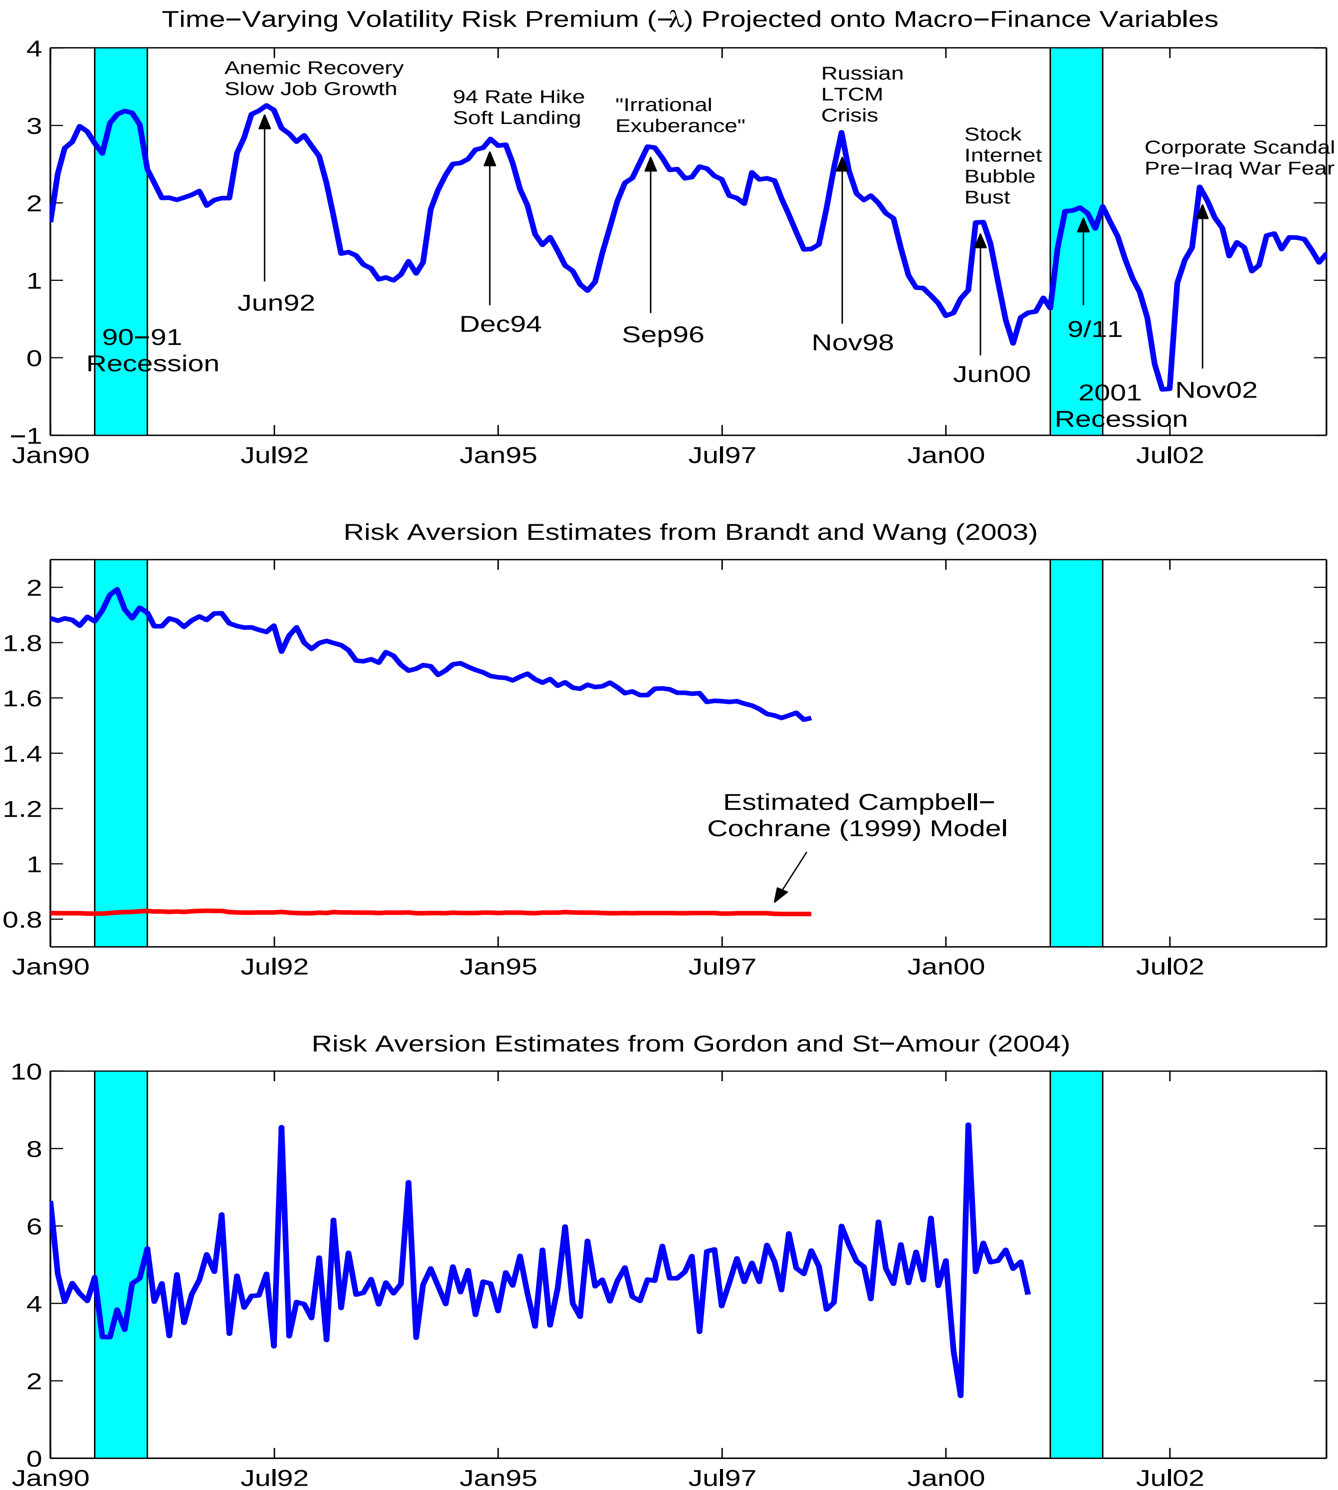

Figure 8: Comparison with Estimates from Consumption-Based Asset Pricing Models

Research Papers 2007

2007-3	Bent Jesper Christensen and Morten Ørregaard Nielsen: The Effect of Long Memory in Volatility on Stock Market Fluctuations
2007-4	Amber Anand, Carsten Tanggaard and Daniel G. Weaver: Paying for Market Quality
2007-5	Charlotte Christiansen: Level-ARCH Short Rate Models with Regime Switching: Bivariate Modeling of US and European Short Rates
2007-6	Charlotte Christiansen: Decomposing European Bond and Equity Volatility
2007-7	Stig Møller: Habit persistence: Explaining cross sectional variation in returns and time-varying expected returns
2007-8	Charlotte Christiansen, Juanna Schröter Joensen and Jesper Rangvid: Are Economists More Likely to Hold Stocks?
2007-9	Thomas Busch, Bent Jesper Christensen and Morten Ørregaard Nielsen: The Role of Implied Volatility in Forecasting Future Realized Volatility and Jumps in Foreign Exchange, Stock, and Bond Markets
2007-10	Bent Jesper Christensen, Morten Ørregaard Nielsen and Jie Zhu: Long Memory in Stock Market Volatility and the Volatility-in-Mean Effect: The FIEGARCH-M Model
2007-11	Matias D. Cattaneo, Richard K. Crump and Michael Jansson: Optimal Inference for Instrumental Variables Regression with non-Gaussian Errors
2007-12	Michael Jansson: Semiparametric Power Envelopes for Tests of the Unit Root Hypothesis
2007-13	Jens Perch Nielsen, Carsten Tanggaard and M.C. Jones: Local Linear Density Estimation for Filtered Survival Data,with Bias Correction
2007-14	Torben G. Andersen, Tim Bollerslev and Xin Huang: A Reduced Form Framework for Modeling Volatility of Speculative Prices based on Realized Variation Measures
2007-15	Viktor Todorov and Tim Bollerslev: Jumps and Betas: A New Framework for Disentangling and Estimating Systematic Risks
2007-16	Tim Bollerslev, Michael Gibson and Hao Zhou: Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities