Variance Risk Dynamics, Variance Risk Premia, and Optimal Variance Swap Investments*

DANIEL EGLOFF,[†] MARKUS LEIPPOLD,[‡] AND LIUREN WU[§]

First draft: December 16, 2005; This version: September 17, 2007

ABSTRACT

With more than a decade worth of variance swap quotes at five maturities, we exploit the information in both the time series and the term structure of the variance swap rates to analyze the return variance rate dynamics and market pricing of variance risk. We then study how the investor's optimal portfolio choice is altered by the presence of variance swap contracts across different maturities to span the variance risk. Our empirical study indicates that the gains from including variance swaps into the portfolio mix are large.

JEL CLASSIFICATION CODES: G12, G13, C52.

KEY WORDS: Return variance swap; equity index options; term structure.

*We welcome comments, including references we have inadvertently missed. We are grateful to Dave Backus, Cliff Ball, Craig Brown, Peter Carr, Jaksa Cvitanic, Stephen Figlewski, Gikas A. Hardouvelis, Robert Kosowski, Craig Lewis, Dimitrios Malliaropoulos Loriano Mancini, Carsten Murawski, Bob Schwartz, Clemens Sialm, George Skiadopoulos, Hans Stoll, Anders Trolle, Jun Wang, Bob Whaley, Rudi Zagst, and the seminar participants at Baruch College, Imperial College London, University of Piraeus, Vanderbilt University, Technische Universität München (TMU) and University of Munich (LMU), the 2006 NCCR Research Conference in Gerzensee, and the 2006 EFA Annual Meetings in Zurich for comments. The authors acknowledge the financial support of INQUIRE Europe and of the Swiss National Science Foundation (NCCR FINRISK). Part of this work was completed when Markus Leippold was a visiting researcher at the Capital Markets Research Group, Federal Reserve Bank of New York.

†Zurich Cantonalbank, Josefstrasse 222, 8010 Zurich, Switzerland; tel: (+41) 44-292-4533; daniel.egloff@zkb.ch.

‡Imperial College London - Tanaka Business School, South Kensington Campus, London SW7 2AZ, United Kingdom; m.leippold@imperial.ac.uk.

§Zicklin School of Business, Baruch College, One Bernard Baruch Way, Box B10-225, New York, NY 10010; tel: (646) 312-3509; fax: (646) 312-3451; Liuren_Wu@baruch.cuny.edu; http://faculty.baruch.cuny.edu/lwu.

Contents

The financial market is becoming increasingly aware of the fact that return variance on stock indexes is stochastic and the variance risk is heavily priced.[1] Associated with this recognition is the development of a large number of variance-related derivative products. The most actively traded among them is the variance swap contract. The contract has zero value at inception. At maturity, the long side of the variance swap contract receives the difference between a standard measure of the realized variance and a fixed rate, called the variance swap rate determined at the inception of the contract. Although traditional derivative contracts such as calls, puts, and straddles have variance risk exposure, entering a variance swap contract represents the most direct way of achieving exposure to or hedging against variance risk.

Variance swap contracts on major equity indexes are actively traded over the counter. Accordingly, variance swap rate quotes on such indexes are now readily available from several broker dealers. In this paper, we obtain more than a decade worth of variance swap rate quotes from a major investment bank on the S&P 500 index at five fixed maturities from two months to two years. With the data, we first propose a class of models on variance risk dynamics and then estimate the variance dynamics and the market pricing of different sources of variance risk by exploiting the rich information embedded in both the time series and the term structure of the variance swap rate quotes. Furthermore, based on the estimated dynamics and market pricing, we also study both theoretically and empirically how investors can use variance swap contracts across different maturities to span the variance risk and to revise their dynamic asset allocation decisions.

Despite the well-appreciated importance of understanding variance risk dynamics and variance risk premia, it remains an unsettled issue on how to model and estimate variance risk dynamics and variance risk premia mainly because return variance is not directly observable. Previous literature mostly relies on the information in time series returns and option prices on the underlying security to make the inference. Yet, such inference is almost always joint inference with the underlying return dynamics. Misspecification on one leads to erroneous conclusions on the other. With variance swap rate quotes, we show that we can directly study and estimate the variance dynamics and variance risk premia without specifying the underlying return dynamics, and hence without interference from the potential misspecifications on the return dynamics.

Our model design and estimation show that we need two stochastic variance risk factors to explain the variation in the variance swaps at different maturities, with one factor controlling the instantaneous variance

[1] See Engle (2004) for a recent review of variance risk and modeling, and Bakshi and Kapadia (2003), Carr and Wu (2007), and Bondarenko (2004) for evidence on the market price of variance risk.

rate while the other controlling the central tendency of the variance rate movements. The variance rate factor is much more transient than the central tendency factor under both the risk-neutral and the statistical measures, thus generating different loading patterns across the term structure from the two risk factors and different autocorrelation patterns for variance swap rates of different maturities. We also find that the market prices for both variance risk factors must be negative to match the upward-sloping mean term structure observed from the variance swap data.

With the estimated dynamics and market pricing, we consider a dynamic asset allocation problem, where an investor equipped with a CRRA utility function trades in the S&P 500 stock index, a riskless bond, and a series of variance swap contracts to maximize her utility on the terminal wealth. Compared to option contracts, variance swap contracts provide a more direct approach in spanning the variance risk. Since the variance swap is a linear contract in variance risk, by trading these contracts, the investor does not build up additional delta exposures to the underlying stock, as would be the case for strategies involving vanilla options.

We first derive the allocation decision in analytical form, and then calibrate the decision to the estimated variance dynamics. It has been argued in the literature (see, e.g., Chacko and Viceira (2005)) that the intertemporal hedging demand caused by the presence of stochastic volatility does not lead to significant changes in the dynamic portfolio allocation. We find that by making the variance swap contract available to span the variance risk, an investor drastically alters the investment allocation to stocks, not because of the hedging demand but merely due to the availability of the variance swap market.

When we perform an out-of-sample study to investigate the impact of variance swap investment on the overall performance of portfolio strategies, we find that incorporating variance swap contracts significantly increases the portfolio performance. We interpret our results as evidence in favor of profitable trading strategies that may occur to the investor who possesses the ability to go both long and short in variance swap contracts.

Our study is related to three strands of literature. The first strand includes all traditional studies that estimate the variance dynamics joint with the return dynamics (see Engle (2004) for a review). The second strand studies the market pricing of variance risk by comparing time-series returns and/or realized variances to options or option portfolios. Prominent examples include Bakshi and Kapadia (2003), Bondarenko (2004), Carr and Wu (2007), and Coval and Shumway (2001). By using the time series of variance swap rates across different maturities in this paper, we infer both the variance risk dynamics and the market prices of different sources of variance risk without resorting to the specification of the return dynamics. The third strand of the literature studies the asset allocation problem in the presence of derivative securities. Carr and Madan (2001) and Carr, Jin, and Madan (2001) focus on how to use vanilla options across different strikes to span the jump risk with random jump size, in the absence of stochastic variance. Complementing to their study, we focus on how to use variance swap contracts of different maturities to span the multi-dimensional stochastic variance risk, while secluding ourselves from jump risk. Liu and Pan (2003) analyze investments in vanilla options in the presence of both jumps and stochastic variance. In this case, the allocation to a vanilla option at a given strike and maturity is a result of mixed effects from spanning the jump risk and the stochastic variance risk. To disentangle the effects, they assume a constant jump size and hence effectively seclude themselves from the strike dimension analyzed in Carr and Madan (2001) and Carr, Jin, and Madan (2001).

The remainder of the paper is organized as follows. Section I introduces a class of affine stochastic variance models and shows how to price variance swaps under this setting and how to identify the variance dynamics and market prices of variance risks using both the time series and the term structure of variance swap rates. Section II presents the estimation results for the variance risk dynamics and variance risk premia. Section III derives the optimal portfolio allocation policies. Section IV calibrates the allocation decisions to the estimated dynamics and analyzes how incorporating variance swap contracts alters an investor's strategic decision and improves her welfare. Section V concludes.

I. Affine Market Models of Stochastic Variance

A. Basic setup

Formally, let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ be a complete stochastic basis and let $\mathbb{P}$ be the statistical probability measure. Let $\mathbb{Q}$ denote a risk-neutral measure that is absolutely continuous with respect to $\mathbb{P}$. No arbitrage guarantees that there exists at least one such measure that prices all traded securities (Duffie (1992)). Let $V_{t,T} \equiv \int_t^T v_s ds$ denote the aggregate return variance during the period $[t,T]$ with $\tau = T - t$ denoting the length of the horizon.

We assume that the dynamics of the instantaneous variance rate v_t is controlled by a k-dimensional Markov process X, which starts at X_0 and satisfies the following stochastic differential equation under the risk-neutral measure $\mathbb{Q}$:

$$dX_t = \mu(X_t)dt + \Sigma^X(X)dB_t^X + \left(qdN^X(\lambda(X_t)) - \bar{q}\lambda(X_t)dt\right), \tag{1}$$

where $\mu(X_t) \in \mathbb{R}^k$ denotes the instantaneous drift function, B^X denotes a k-dimensional independent Brownian motion with $\Sigma^X(X)\Sigma^X(X)^\top \in \mathbb{R}^{k\times k}$ being the symmetric and positive definite instantaneous covariance matrix, and N^X denotes k independent Poisson jump components with intensities $\lambda(X_t) \in \mathbb{R}^k$ and with the random jump magnitudes q being a diagonal $(k \times k)$ matrix, characterized by its two-sided Laplace transform $\mathcal{L}_q(\cdot)$ and with $\bar{q} = \mathbb{E}^{\mathbb{Q}}[q]$. The last two terms in equation (1) form a k-dimensional jump martingale.

To analyze the variance risk dynamics using variance swap rates, we adopt the affine framework of Duffie, Pan, and Singleton (2000) and model the term structure of variance swaps within the affine class.

Definition 1 *In **affine stochastic variance models**, the Laplace transform of the quadratic variation,* $V_{t,T} = \int_t^T v_s ds$*, under the risk-neutral measure* $\mathbb{Q}$ *is an exponential-affine function of the state vector* X_t*:*

$$\mathcal{L}_V(u) \equiv \mathbb{E}^{\mathbb{Q}}\left[e^{-uV_{t,T}} \,\middle|\, \mathcal{F}_t\right] = \exp\left(-\mathbf{b}(\tau)^\top X_t - c(\tau)\right), \tag{2}$$

where $\mathbf{b}(\tau) \in \mathbb{R}^k$ *and* $c(\tau)$ *is a scalar.*

The definition implicitly limits us to time-homogeneous models since the coefficients depend only on the horizon $\tau = T - t$, but not on the calendar time t. The following proposition presents a set of sufficient conditions for the affine definition in equation (2) to hold.

Proposition 1 *If under the risk-neutral measure* $\mathbb{Q}$*, the instantaneous variance rate* v_t*, the drift vector* $\mu(X)$*, the diffusion covariance matrix* $\Sigma^X(X)\Sigma^X(X)^\top$*, and the jump arrival rate* $\lambda(X)$ *of the Markov process* X *are all affine in* X*, then the Laplace transform* $\mathcal{L}_V(u)$ *is exponential-affine in* X_t*.*

The above process specifications are directly adopted from Duffie, Pan, and Singleton (2000) on general asset pricing modeling. In particular, let the $\mathbb{Q}$-dynamics be defined as[2]

$$\begin{aligned} v_t &= \mathbf{b}_v^\top X_t + c_v, \quad \mathbf{b}_v \in \mathbb{R}^k, c_v \in \mathbb{R}, \\ \mu(X_t) &= \kappa(\theta - X_t), \quad \kappa \in \mathbb{R}^{k\times k}, \theta \in \mathbb{R}^k, \\ \Sigma^X(X)\Sigma^X(X)^\top &= \mathrm{diag}\left[\alpha + \beta X_t\right], \quad \alpha \in \mathbb{R}^k, \beta \in \mathbb{R}^{k\times k}, \\ \lambda(X_t) &= \alpha_\lambda + \beta_\lambda X_t, \quad \alpha_\lambda \in \mathbb{R}^k, \beta_\lambda \in \mathbb{R}^{k\times k}. \end{aligned} \tag{3}$$

We further constrain β and β_λ to be diagonal matrices. Given the above specification, the coefficients $\{\mathbf{b}(\tau), c(\tau)\}$ for the Laplace transform in (2) are determined by the following ordinary differential equations:

$$\begin{aligned} \mathbf{b}'(\tau) &= u\mathbf{b}_v - (\kappa + \bar{q}\beta_\lambda)^\top \mathbf{b}(\tau) - \frac{1}{2}\beta\,\mathrm{diag}\left[\mathbf{b}(\tau)\mathbf{b}(\tau)^\top\right] - \beta_\lambda\left(\mathcal{L}_q(\mathbf{b}(\tau)) - \mathbf{1}\right), \\ c'(\tau) &= uc_v + (\kappa\theta - \bar{q}\alpha_\lambda)^\top \mathbf{b}(\tau) - \frac{1}{2}\alpha^\top\,\mathrm{diag}\left[\mathbf{b}(\tau)\mathbf{b}(\tau)^\top\right] - \alpha_\lambda^\top\left(\mathcal{L}_q(\mathbf{b}(\tau)) - \mathbf{1}\right), \end{aligned} \tag{4}$$

with the boundary conditions $\mathbf{b}(0) = 0$, and $c(0) = 0$.

B. Pricing variance swaps

The terminal payoff of a variance swap contract is the difference between the realized variance over a certain time period and a fixed variance swap rate, determined at the inception of the contract. The variance difference is multiplied by a notional amount that converts the difference into dollar terms. Since the contract is worth zero at inception, no-arbitrage dictates that the value of the variance swap rate is equal to the risk-neutral expected value of the realized variance over the relevant time period. Formally, the time-t value of an annualized variance swap rate with expiry date T is given by,

$$VS_t(T) = \frac{1}{\tau}\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^T v_s ds\right], \quad \tau = T - t. \tag{5}$$

[2]We use the convention that $\mathrm{diag}[v]$ maps the vector v onto a diagonal matrix with elements from the vector v. For a matrix M, $\mathrm{diag}[M]$ is the vector consisting of the diagonal elements of M.

In the affine stochastic variance framework, the variance swap rate can be solved from the Laplace transform in equation (2):

$$\begin{aligned}\mathbb{E}_t^{\mathbb{Q}}[V_{t,T}] &= -\left.\frac{\partial \mathbb{E}_t^{\mathbb{Q}}\left[e^{-uV_{t,T}}\right]}{\partial u}\right|_{u=0} = \mathcal{L}_V(u)\left.\left(\left[\frac{\partial \mathbf{b}(\tau)}{\partial u}\right]^\top X_t + \frac{\partial c(\tau)}{\partial u}\right)\right|_{u=0} \\ &= \mathbf{B}(\tau)^\top X_0 + C(\tau),\end{aligned} \tag{6}$$

which is affine in the current level of the state vector X_t. Note that $\mathcal{L}_{V_t}(u)|_{u=0} = 1$ and the coefficients $\mathbf{B}(t)$ and $C(t)$ are defined as the partial derivatives of $\mathbf{b}(\tau)$ and $c(\tau)$ with respect to u. Plugging the derivatives into the ordinary differential equations in (4) and setting $u = 0$, we obtain a new set of ordinary different equations that determine the coefficients of the variance swap rates:

$$\begin{aligned}\mathbf{B}'(\tau) &= \mathbf{b}_v - (\kappa + \bar{q}\beta_\lambda)^\top \mathbf{B}(\tau) - \beta\,\mathrm{diag}\left[\mathbf{b}(\tau)\mathbf{B}(\tau)^\top\right] - \beta_\lambda\,\mathrm{diag}\left[\nabla\mathcal{L}_q(\mathbf{b}(\tau))\mathbf{B}(\tau)\right], \\ C'(\tau) &= c_v + (\kappa\theta - \bar{q}\alpha_\lambda)^\top \mathbf{B}(\tau) - \alpha^\top\,\mathrm{diag}\left[\mathbf{b}(\tau)\mathbf{B}(\tau)^\top\right] - \alpha_\lambda^\top\,\mathrm{diag}\left[\nabla\mathcal{L}_q(\mathbf{b}(\tau))\mathbf{B}(\tau)\right],\end{aligned} \tag{7}$$

with the boundary conditions $\mathbf{B}(0) = \mathbf{b}(0) = 0$ and $C(0) = 0$. With $u = 0$, the ordinary differential equation for $\mathbf{b}(\tau)$ becomes

$$\mathbf{b}'(\tau) = -(\kappa + \bar{q}\beta_\lambda)^\top \mathbf{b}(\tau) - \frac{1}{2}\beta\,\mathrm{diag}\left[\mathbf{b}(\tau)\mathbf{b}(\tau)^\top\right] - \beta_\lambda\left(\mathcal{L}_q(\mathbf{b}(\tau)) - \mathbf{1}\right). \tag{8}$$

Starting at $\mathbf{b}(0) = 0$, we have $\mathbf{b}'(0) = 0$. Thus, $\mathbf{b}(t) = 0$, for all $t = 0$, is a solution. The equations in (7) simplify to

$$\begin{aligned}\mathbf{B}'(\tau) &= \mathbf{b}_v - (\kappa + \bar{q}\beta_\lambda)^\top \mathbf{B}(\tau) - \beta_\lambda\,\mathrm{diag}\left[\nabla\mathcal{L}_q(\mathbf{b}(0))\mathbf{B}(\tau)\right], \\ C'(\tau) &= c_v + (\kappa\theta - \bar{q}\alpha_\lambda)^\top \mathbf{B}(\tau) - \alpha_\lambda^\top\,\mathrm{diag}\left[\nabla\mathcal{L}_q(\mathbf{b}(0))\mathbf{B}(\tau)\right],\end{aligned} \tag{9}$$

where $\nabla\mathcal{L}_q(\mathbf{b}(0))$ denotes the gradient of $\mathcal{L}_q(\mathbf{b}(\tau))$ with respect to $\mathbf{b}(\tau)$, evaluated at $\mathbf{b}(0) = 0$. Therefore,

$$\nabla\mathcal{L}_q(\mathbf{b}(0)) = -\bar{q} \tag{10}$$

and the equations in (9) reduce to

$$\mathbf{B}'(\tau) = \mathbf{b}_v - \kappa^\top \mathbf{B}(\tau), \quad C'(\tau) = c_v + \mathbf{B}(\tau)^\top \kappa\theta. \tag{11}$$

The two ordinary differential equations in (11) can be solved analytically.

Proposition 2 *In the affine stochastic variance framework as specified in (3), the time-t variance swap rate with expiry date T is given by*

$$VS_t(T) = \frac{1}{\tau}\left[\mathbf{B}(\tau)^\top X_t + C(\tau)\right], \tag{12}$$

with

$$\mathbf{B}(\tau) = \left(I - e^{-\kappa^\top \tau}\right)\left(\kappa^\top\right)^{-1}\mathbf{b}_v; \quad C(\tau) = \left(c_v + \mathbf{b}_v^\top \theta\right)\tau - \mathbf{B}(\tau)^\top \theta, \tag{13}$$

where $e^{(\cdot)}$ denotes the matrix exponential operator, defined as:

$$e^M \equiv \sum_{n=0}^{\infty} \frac{M^n}{n!}. \tag{14}$$

From Proposition 2, we obtain the following corollary:

Corollary 1 *Under the affine stochastic variance framework defined in (2), the term structure of the return variance swap rate only depends upon the specification of the drift of the state vector, but does not depend upon the type and specification of the martingale component of these factors. Holding constant the long run mean and the reverting speed, the term structure remains the same whether the martingale component is a pure diffusion, a pure jump martingale, or a mixture of both.*

The corollary follows readily by inspecting the solutions of the coefficients $\{\mathbf{B}(\tau), C(\tau)\}$ in equation (13) that determine the variance swap rate in equation (6). Note in particular that the parameters controlling the covariance matrix of the diffusion component (α, β) and the parameters for the jump component $(\alpha_\lambda, \beta_\lambda, \mathcal{L}_q(\cdot))$ do not enter the solutions of the coefficients.

Corollary 1 implies that from the term structure of the return variance swap, one can identify the risk-neutral drift of the state vector that controls the dynamics of the return quadratic variation. Nevertheless, the innovation (martingale) specifications of the instantaneous variance rate play little role in determining the term structure of the return variance swap rate, although they affect the time-series behavior of the swap rates.

It is worth noting that although we use the affine specification to illustrate the corollary, the variance swap rate only depends on the risk-neutral drift of the instantaneous variance rate $\mu(v)$ under any instantaneous variance rate dynamics:

$$VS_t(T) = \frac{1}{\tau}\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^T v_s ds\right] = \frac{1}{\tau}\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^T \mu(v_s) ds\right], \quad (15)$$

as the expectation of the martingale component equals zero. This conclusion is a result of the linear relation between the variance swap rate and future instantaneous variance rates.

C. Model design

Proposition 1 identifies a set of conditions that generates the affine stochastic variance class. Based on these conditions, we design both a one-factor and a two-factor model for the variance risk dynamics and compare their empirical performance in matching the time-series and term structure behaviors of the variance swap rates.

C.1. A one-factor variance rate model

In the one-factor setting, we let the variance rate follow the square-root dynamics as in Heston (1993). Under the risk-neutral measure $\mathbb{Q}$, the instantaneous variance rate dynamics are:

$$dv_t = \kappa_v(\theta_v - v_t)\,dt + \sigma_v\sqrt{v_t}dB_t^v. \quad (16)$$

Comparing equation (16) to the general conditions in (3), we have $b_v = 1, c_v = 0, b_v = 1, \alpha = 0, \beta = \sigma_v^2, \lambda = 0$. Plugging these parameterizations into (13) and rearranging, we have the variance swap rate as,

$$VS_t(T) = \phi_v(\tau)v_t + (1 - \phi_v(\tau))\,\theta_v, \quad (17)$$

with

$$\phi_v(\tau) = \frac{1 - e^{-\kappa_v\tau}}{\kappa_v\tau}. \quad (18)$$

With a stationary risk-neutral variance rate dynamics ($\kappa_v > 0$), the coefficient $\phi_v(\tau)$ is between zero and one. Thus, the variance swap rate is a weighted average of the instantaneous variance rate v_t and its risk-neutral long-run mean θ_v. The weight depends on the maturity (τ) of the variance swap contract and the risk-neutral mean-reversion speed of the variance rate (κ_v). The linear structure further dictates that under this one-factor setting, variance swap rates of all maturities show the same statistical persistence as that for the instantaneous variance rate v_t.

Holding the maturity fixed as κ_v declines and the variance rate becomes more persistent, $\phi_v(\tau)$ increases and the current variance rate v_t has a larger impact on the variance swap rate. In the limit $\kappa \to 0$ and $\phi_v(\tau) \to 1$, the variance swap rates of all maturities are equal to the instantaneous variance rate level. On the other hand, as the mean-reversion speed increases, the long-run mean imposes a heavier weight on the variance swap rate. As $\kappa \to \infty$ and the instantaneous variance rate shows zero persistence, the variance swap rate is always equal to the long-run mean.

Holding a fixed mean-reversion speed $\kappa_v > 0$, the coefficient $\phi_v(\tau)$ starts at one at $\tau = 0$ and declines to zero with increasing maturity. Hence, the variance swap rate converges to the instantaneous variance rate as the maturity goes to zero and converges to the risk-neutral long-run mean as the maturity goes to infinity.

Taking expectations on both sides of equation (17) under the statistical measure $\mathbb{P}$, we obtain the mean term structure of variance swap rates as

$$\mathbb{E}^{\mathbb{P}}[VS_t(T)] = \phi_v(\tau)\theta_v^{\mathbb{P}} + (1 - \phi_v(\tau))\,\theta_v, \tag{19}$$

which is a weighted average of the statistical mean $\theta_v^{\mathbb{P}} \equiv \mathbb{E}^{\mathbb{P}}[v_t]$ and the risk-neutral mean θ_v of the instantaneous variance rate. Since $\phi_v(\tau)$ declines monotonically with increasing maturity, the risk-neutral mean has an increasing weight at longer maturities. Therefore, to generate an upward or downward sloping mean term structure for the variance swap rates, we need the statistical mean and the risk-neutral mean of the instantaneous variance rate to be different. The difference between the two mean values dictates the sign of the variance risk premium.

C.2. A two-factor variance rate model

We also consider a two-factor variance rate dynamics, which satisfy the following risk-neutral dynamics,

$$\begin{array}{rcl} dv_t & = & \kappa_v\,(m_t - v_t)\,dt + \sigma_v\sqrt{v_t}dB_t^v, \\ dm_t & = & \kappa_m\,(\theta_m - m_t)\,dt + \sigma_m\sqrt{m_t}dB_t^m, \end{array} \qquad dB_t^v dB_t^m = 0, \tag{20}$$

where the instantaneous variance rate (v_t) reverts to a stochastic mean level (m_t). The mean level follows another square-root process. Analogous to Balduzzi, Das, and Foresi (1998) for interest rate modeling, we label m_t as the stochastic central tendency of the instantaneous variance rate, with θ_m being the unconditional long-run mean for both v_t and m_t.

Under this specification, the variance swap rates are given by,

$$VS_t(T) = \phi_v(\tau)v_t + \phi_m(\tau)m_t + (1 - \phi_v(\tau) - \phi_m(\tau))\,\theta_m, \tag{21}$$

with

$$\phi_v(\tau) = \frac{1 - e^{-\kappa_v \tau}}{\kappa_v \tau}, \quad \phi_m(\tau) = \frac{1 + \frac{\kappa_m}{\kappa_v - \kappa_m} e^{-\kappa_v \tau} - \frac{\kappa_v}{\kappa_v - \kappa_m} e^{-\kappa_m \tau}}{\kappa_m \tau}. \tag{22}$$

The variance swap rate in equation (21) is a weighted average of the instantaneous variance rate v_t, its stochastic central tendency m_t, and the risk-neutral long-run mean θ. The weight on the instantaneous variance rate is the same as in the one-factor case. The weight converges to one as the maturity goes to zero and converges to zero as the maturity goes to infinity. The weight on m_t also converges to zero as the maturity goes to infinity. Hence, the variance swap rate starts at the instantaneous variance rate level at zero maturity and converges to the long-run mean θ_m as maturity goes to infinity. The stochastic central tendency factor plays a role in the intermediate maturities, with the weighting coefficient $\phi_m(\tau)$ exhibiting a hump-shaped term structure.

Under this two-factor structure, swap rates at different maturities can show different degrees of persistence. In particular, if the central tendency factor is more persistent than the instantaneous variance rate, the short-term swap rate will become less persistent than the long-term swap rate. For the variance process v_t to be stationary under $\mathbb{Q}$, we furthermore require that $\kappa_v, \kappa_m > 0$.

D. Market prices of variance risks

For both models, we assume that the market price on each source of risk is proportional to the square root of the risk level:

$$\gamma(B_t^v) = \gamma_v \sqrt{v_t}, \quad \gamma(B_t^m) = \gamma_m \sqrt{m_t}. \tag{23}$$

Under the one-factor model, the statistical dynamics of the variance rate become,

$$dv_t = \kappa_v^{\mathbb{P}} \left(\theta_v^{\mathbb{P}} - v_t\right) dt + \sigma_v \sqrt{v_t} dB_t^v, \tag{24}$$

with $\kappa_v^{\mathbb{P}} = (\kappa_v - \gamma_v \sigma_v)$ and $\theta_v^{\mathbb{P}} = \kappa_v \theta / \kappa_v^{\mathbb{P}}$. Thus, a negative market price for the variance risk makes the statistical variance rate process more mean reverting than its risk-neutral counterpart ($\kappa_v^{\mathbb{P}} > \kappa_v$) and makes the statistical mean variance rate lower than its risk-neutral counterpart ($\theta_v^{\mathbb{P}} < \theta_v$). A positive market price for the variance risk generates the opposite effects.

Under the two-factor model, the statistical dynamics of the variance rate become,

$$\begin{aligned} dv_t &= \kappa_v^{\mathbb{P}} \left(\tfrac{\kappa_v}{\kappa_v^{\mathbb{P}}} m_t - v_t\right) dt + \sigma_v \sqrt{v_t} dB_t^v, \\ dm_t &= \kappa_m^{\mathbb{P}} \left(\theta_m^{\mathbb{P}} - m_t\right) dt + \sigma_m \sqrt{m_t} dB_t^m, \end{aligned} \tag{25}$$

with $\kappa_m^{\mathbb{P}} = \kappa_m - \gamma_m \sigma_m$ and $\theta_m^{\mathbb{P}} = \kappa_m \theta_m / \kappa_m^{\mathbb{P}}$. In this case, the long-run statistical mean of the variance rate becomes $\theta_v^{\mathbb{P}} = \kappa_v \theta_m^{\mathbb{P}} / \kappa_v^{\mathbb{P}} = \frac{\kappa_v \kappa_m}{\kappa_v^{\mathbb{P}} \kappa_m^{\mathbb{P}}} \theta_m$. Similar to the one-factor case, negative market price on the variance risk (B^v) increases the mean-reversion speed ($\kappa_v^{\mathbb{P}}$) and reduces the long-run mean ($\theta_v^{\mathbb{P}}$) under the statistical measure. Negative market price on the central tendency has similar effects on the central tendency dynamics. It also further reduces the long-run statistical mean of the instantaneous variance rate. Positive market prices for both risk factors generate the opposite effects.

II. Estimating Variance Risk Dynamics and Variance Risk Premia

We estimate the variance risk dynamics and variance risk premia using over-the-counter quotes on variance swap rates on the S&P 500 index. From a major broker dealer, we obtain daily closing quotes on variance swap rates with fixed time to maturities at two, three, six, 12, and 24 months starting January 10, 1996, and ending March 30, 2007, spanning over 12 years. To avoid the effect of weekday patterns on the dynamics estimation, we sample the data weekly on every Wednesday. When Wednesday is a holiday, we use the quotes from the previous business day. The data sample contains 586 weekly observations for each series.

To enhance the identification of variance risk premia, we also obtain the time series data on the S&P 500 index and compute the ex-post annualized realized variance based on daily log returns,

$$RV_{t,T} = \frac{365}{D} \sum_{d=1}^{D} \left(\ln S_{t+d}/S_{t+d-1}\right)^2$$

where S_t denotes the index level at time t and D denotes the number of days between time t and T. Following industry standard in variance swap payoff calculations, we compute the realized variance using un-demeaned log daily returns. At each date t, we compute realized variance over fixed horizons of 7, 30, 60, 90, 120, and 150 days, and use these realized variance series together with the variance swap rates to identify the variance risk dynamics and variance risk premia. Ideally, we want to include realized variance over horizons that match the whole spectrum of the variance swap maturity. Yet incorporating realized variance of longer horizons would severely shorten our sample length. The choice of horizon from seven to 150 days is a trade off between the two consideration. Since we have obtained the S&P 500 index series up to September 4, 2007, we can compute the realized variance for the full variance swap sample period up to 150 day horizons.

A. Summary statistics of variance swap rates and realized variance

The left panel of Figure 1 plots the time series of the variance swap rate quotes at three selected maturities of two (solid line), six (dashed line), and 24 (dotted line) months. The right panel plots the mean term structure (solid line) and various representative term structure shapes at different dates. From the time series plots, we observe that the variance swap rates started at relatively low levels, but experienced a spike during the 1997 Asian crisis, and another even larger spike during the hedge fund crisis in late 1998. The series witnessed another two spikes between 2001 and 2003, but otherwise have been declining to very low levels. Over the course of the past ten years, the variance swap rate level has varied greatly from as low as 0.0108 to as high as 0.2548. In volatility terms, the variation is from 10.39% to 50.48%.

[Figure 1 about here.]

The right panel shows that the term structure of variance swap rates can take a wide variety of shapes, including upward sloping, downward sloping, and hump-shaped term structures. A successful model of variance risk dynamics and variance risk premium specification must capture not only the large variation in the volatility levels, but also the different shapes of the term structure.

Table I reports the summary statistics of the variance swap rates in panel A. The mean variance swap rates increase with maturity from 0.048 at two-month maturity to 0.056 at two-year maturity. The standard deviation estimates decline as the maturity increases. The variance swap rates show positive skewness and kurtosis. The weekly autocorrelation estimates for the variance swap rates are high and increasingly so as the swap maturity increases.

The summary statistics of the realized variance are reported in panel B of Table I. In theory, the population mean of the annualized realized variance should be the same regardless of the measurement horizons. The estimates show some small sample variation around 0.031. The mean estimates on the realized variance are markedly lower than the mean variance swap rates. The standard deviation estimates on the realized variance are comparable to the estimates on the variance swap rates. The non-overlapping seven-day realized variance shows much larger skewness and excess kurtosis, and much lower weekly autocorrelation. As the horizon increases and the degree of overlapping increases, the skewness and kurtosis estimates become smaller and the autocorrelation estimates become larger.

B. Profit and loss from long variance swaps and holding the contracts to maturity

By comparing the variance swap rates to the corresponding realized variance, we can learn the historical profit and loss on investing in variance swaps. The average profit and loss reflects the average variance risk premium. To gain a preliminary understanding of the variance risk premium behavior, we estimate the ex post realized variance corresponding to variance swap rate quote and we measure the profit and loss both in dollar terms $100 \times (RV_{t,T} - VS_{t,T})$ and in log excess returns $\ln(RV_{t,T}/VS_{t,T})$. As discussed in Carr and Wu (2007), the first measure represents the dollar capital gain from going long a \$100 par variance swap contract and holding it to maturity. The second measure represents the log excess return by regarding the forward price (i.e., the swap rate) as the forward cost. Table II reports the summary statistics of the dollar capital gains in panel A and the statistics of the log excess returns in panel B. We can match the two- and three-month variance swap rates with their corresponding ex post realized variance over the whole sample period. For swaps at longer maturities, the sample length for the profit and loss is shorten accordingly given the constraint on the ex post realized variance calculation.

During our sample period, a \$100 par investment in the two-month variance swap contracts generate an average loss of the \$1.72. The average log excess return is -54.9%. As the investment horizon increases, both the dollar gain and the log excess return increase moderately. The raw standard deviation estimates decline with the investment horizon mainly due to increasing overlapping in the sample. When we adjust for serial dependence in the time series according to Newey and West (1987) with the number of lags optimally chosen following Andrews (1991), the adjusted standard deviation increases with the investment horizon. The dollar capital gains show negative skewness and large excess kurtosis, but skewness and excess kurtosis estimates are smaller for the log excess returns.

The negative average return on the long variance swap investments suggests that the variance risk premium is negative on average. Shorting variance swaps generate positive average returns. In the last column, we compute the annualized Sharpe ratios for long variance investments, defined as the mean divided by the serial-dependence-adjusted standard deviation and then annualized by the investment horizon. The Sharpe ratio estimates are relatively high for shorting short-term variance swap contracts, but the Sharpe ratio estimates decline for shorting long-term variance swaps and holding them to maturity.

Figure 2 contrasts the time series of the variance swap rates (solid lines) to the corresponding ex post realized variance (dashed lines) at maturities of two, six, 12, and 24 months, with one maturity for each panel. The variance swap rates almost always stay above the ex post realized variance, with only occasional cross overs. The realized variance looks much smoother at longer horizons. Nevertheless, it is important to caution that with the longer investment horizon, the effective sample size is also shorter.

[Figure 2 about here.]

Figure 3 plots the time series of the profit and loss from long variance swaps of various maturities and holding the contracts to maturities. Solid lines in the left panels measure the dollar capital gains. Solid lines in the right panels measure the log excess returns. In all panels, the dashed lines represent the sample mean and the dotted line represent the zero-profit benchmark. Plots in the left panels show that the variation of the dollar gain increases with the variance swap rate level. During the most recent four years, as the variance swap rate level declines, the magnitude of the capital gain has also become smaller. However, when we scale the dollar gain by the variance swap rate level and compute the log excess return in the right panels, the variations become much more uniform over time, suggesting that the variance risk premium is relatively proportional to the variance risk level.

[Figure 3 about here.]

C. Estimation methodology

To estimate the variance dynamics, we first cast the model into a state-space form and extract the variance risk factors (v_t and m_t) from the observed variance swap rates using Kalman filter.

We build the state propagation equation based on the statistical dynamics of the variance rates. Under the two factor model, we set $X_t = [v_t, m_t]^\top$, and construct the state propagation equation based on the Euler approximation of the statistical dynamics in (25) as,

$$X_{t+1} = A + \Phi X_t + \Sigma(X_t)\sqrt{\Delta t}\varepsilon_t, \tag{26}$$

with ε denoting a two-dimensional iid standard normal innovation vector,

$$A = (I - \Phi)\theta^{\mathbb{P}}, \quad \Phi = e^{-\kappa^{\mathbb{P}}\Delta t}, \theta^{\mathbb{P}} = \begin{bmatrix} \theta_v^{\mathbb{P}} \\ \theta_m^{\mathbb{P}} \end{bmatrix}, \kappa^{\mathbb{P}} = \begin{bmatrix} \kappa_v^{\mathbb{P}} & -\kappa_v \\ 0 & \kappa_m^{\mathbb{P}} \end{bmatrix}, \Sigma(X_t) = \begin{bmatrix} \sigma_v\sqrt{v_t} & 0 \\ 0 & \sigma_m\sqrt{m_t} \end{bmatrix},$$

and $\Delta t = 1/52$ being the weekly time interval of the discretization. The one-dimensional state propagation equation for the one-factor model is defined analogously.

We construct the measurement equation based on the observed variance swap rates:

$$y_t \quad = \quad VS_t(T, X_t) + e_t, \quad T - t = 2, 3, 6, 12, 24 \quad \text{months}. \tag{27}$$

where y_t denotes the observed variance swap series, $VS_t(T, X_t)$ denotes their corresponding model values as a function of the variance risks X_t, and e_t denotes the measurement error. We assume that the measurement error is independent of the state vector and that the measurement error on each of the five series is also mutually independent but with distinct variance.

Since the state propagation equation is Gaussian linear and the measurement equation is linear in the state vector, the Kalman filter provides the efficient forecasts and updates on the state vector and the observed variance swap rates. We build the likelihood on the variance swap rates based on the forecasting errors from the Kalman filter. Specifically, let $(\overline{y}_t, \overline{Q}_t)$ denote the time-$(t-1)$ Kalman filter forecasts on the conditional mean and the conditional variance of the variance swap rates at time t, the time-t likelihood on the variance swap rates is,

$$l_t^{VS}(\Theta) = -\frac{1}{2}\left[\log\left|\overline{Q}_t\right| + \left((y_t - \overline{y}_t)^\top \left(\overline{Q}_t\right)^{-1}(y_t - \overline{y}_t)\right)\right], \tag{28}$$

where Θ denotes the set of model parameters.

Given the variance risk factors (X_t) extracted from the Kalman filter, we can predict the annualized realized variance based on the statistical dynamics of the risk factors,

$$\mathbb{E}^{\mathbb{P}}\left[RV_{t,T}|X_t\right] = \frac{1}{\tau}\left[\left(\mathbf{B}(\tau)^{\mathbb{P}}\right)^{\top} X_t + C(\tau)^{\mathbb{P}}\right], \tag{29}$$

where the coefficients are analogous to those defined as in Proposition 2,

$$\mathbf{B}(\tau)^{\mathbb{P}} \quad = \quad \left(I - e^{-(\kappa^{\mathbb{P}})^{\top}\tau}\right)\left((\kappa^{\mathbb{P}})^{\top}\right)^{-1}\mathbf{b}_v, \quad C(\tau)^{\mathbb{P}} = \left(c_v + \mathbf{b}_v^{\top}\theta^{\mathbb{P}}\right)\tau - (\mathbf{B}(\tau)^{\mathbb{P}})^{\top}\theta^{\mathbb{P}}. \tag{30}$$

Under the two-factor model, we have $\mathbf{b}_v = [1,0]^{\top}$ and $c_v = 0$. Under the one-factor model, we have $\mathbf{b}_v = 1$, $c_v = 0$, $\theta^{\mathbb{P}} = \theta_v^{\mathbb{P}}$, and $\kappa^{\mathbb{P}} = \kappa_v^{\mathbb{P}}$.

Given the forecasts, we build the likelihood function on the realized variances assuming that the forecasting errors on the realized variance, $e_t^{RV} = RV_{t,t+D} - \mathbb{E}^{\mathbb{P}}\left[RV_{t,t+D}|X_t\right]$ with $D = 7, 30, 60, 90, 120, 150$ days, are normally distributed with constant covariance matrix Q_{RV}. Thus, the time-t likelihood on the realized variance becomes,

$$l_t^{RV}(\Theta) = -\frac{1}{2}\left[\log|Q_{RV}| + \left(e_t^{RV}\right)^{\top}(Q_{RV})^{-1}\left(e_t^{RV}\right)\right]. \tag{31}$$

For estimation, we assume that the forecasting errors on variance swaps and realized variances are independent but with distinct variance. Thus, the aggregate likelihood becomes the summation of the likelihoods on the two sets of data series. We numerically maximize the aggregate likelihood to estimate the model parameters.

D. Model performance

Table III reports the summary statistics of the pricing errors, defined as the difference between the variance swap rate quotes and the model-implied values. To make more economic sense of the pricing errors, we follow industry convention and represent both the market quotes and model values in volatility percentage points. The one-factor model fits the six-month variance swap to near perfection, but the pricing errors increase at other maturities. In contrast, the performance of the two-factor model is more uniform across different maturities. The explained variations range from 98.77% to 99.99%. The root mean squared pricing errors range from practically zero to 0.8 volatility percentage points, no larger than the average bid-ask spreads for the over-the-counter variance swap rate quotes.[3]

The last row of Table III reports the maximized likelihood values for the two models. The two-factor model also performs significantly better than the one-factor model in terms of the log likelihood values. A formal likelihood ratio test rejects the one-factor model over any reasonable confidence level.

E. Variance risk dynamics and variance risk premia

Exploiting the information in the time series and term structure of variance swap rates and the realized variance, we can accurately identify the variance risk dynamics and the market prices of different sources of variance risks. Table IV reports the parameter estimates and the absolute magnitudes of the t-statistics in parentheses. Focusing on the two-factor model specification, we observe that the risk-neutral mean-reversion speed for the instantaneous variance rate ($\kappa_v = 4.373$) is much higher than that for the central tendency risk factor ($\kappa_m = 0.1022$). To gain more intuition, we define the half life of each series as the number of weeks for the autocorrelation of the series to decay to half of its weekly autocorrelation level, $\text{Half-life (in weeks)} = \ln(\phi/2)/\ln(\phi)$, with $\phi = \exp(-\kappa\Delta t)$ denoting the weekly autocorrelation of a series. Under the risk-neutral measure, the mean-reversion speed estimates imply a half life of less than ten weeks for the instantaneous variance rate v_t but almost seven years for the central tendency factor m_t.

The different risk-neutral persistence dictates that the two risk factors have different impacts across the term structure of the variance swap rates. In particular, we can convert the risk-neutral persistence estimates into factor loadings coefficients $\phi_v(\tau)$ and $\phi_m(\tau)$ as in (22):

$$\phi_v(\tau) = \frac{1 - e^{-\kappa_v \tau}}{\kappa_v \tau}, \quad \phi_m(\tau) = \frac{1 + \frac{\kappa_m}{\kappa_v - \kappa_m} e^{-\kappa_v \tau} - \frac{\kappa_v}{\kappa_v - \kappa_m} e^{-\kappa_m \tau}}{\kappa_m \tau}, \tag{32}$$

with which the variance swap rates of different maturities are linked to the two risk factors:

$$VS_t(T) = \phi_v(\tau) v_t + \phi_m(\tau) m_t + (1 - \phi_v(\tau) - \phi_m(\tau))\, \theta_m. \tag{33}$$

The loading coefficients measure the contemporaneous response of the variance swap term structure to unit shocks in the two variance risk factors v_t and m_t. Figure 4 plots the term structure of the two responses in the

[3]Currently, the bid-ask spreads on variance swap rate quotes from major broker dealers average around half to one volatility percentage point for the S&P 500 index.

left panel, with the solid line denoting the response of v_t, the dashed line denoting the response of m_t, and the dotted line capturing the remaining weight on the common risk-neutral mean of the variance rate and central tendency θ_m. The impact of the transient variance rate factor (v_t) is mainly at short maturities. Its impact declines as maturity increases. On the other hand, the contribution of the persistent central tendency factor (m_t) starts at zero, but increases progressively as the variance swap maturity increases. The remaining weight on the risk-neutral mean also starts at zero and increases monotonically with increasing maturity. This increasing weight on a constant is responsible in generating a downward sloping term structure on the standard deviation in the variance swap rate innovations.

[Figure 4 about here.]

In line with the strongly negative P&L from holding variance swaps reported in Table II, the market prices of the instantaneous variance rate is strongly negative. The market price of the central tendency factor is also negative, but to a much lesser degree. The negative market prices on the two risk factors make the statistical mean-reversion speeds much larger and the statistical long-run mean much lower than their risk-neutral counterparts. In particular, the three long-run means exhibit the following ranking: $\theta_v^{\mathbb{P}} < \theta_m^{\mathbb{P}} < \theta_m$. The statistical mean of the variance rate is lower than the statistical mean of the central tendency factor. Both are lower than their common risk-neutral mean θ_m.

Taking unconditional expectations on both sides of equation (17) under the statistical measure $\mathbb{P}$, we obtain the mean term structure of variance swap rates as

$$\mathbb{E}^{\mathbb{P}}[VS_t(T)] = \phi_v(\tau)\theta_v^{\mathbb{P}} + \phi_m(\tau)\theta_m^{\mathbb{P}} + (1 - \phi_v(\tau) - \phi_v(\tau))\,\theta_m, \tag{34}$$

which is a weighted average of the statistical mean of the instantaneous variance rate $\theta_v^{\mathbb{P}}$, the statistical mean of the central tendency factor $\theta_m^{\mathbb{P}}$, and the common risk-neutral mean of the variance rate and the central tendency θ_m. The factor loading in Figure 4 suggests that $\theta_v^{\mathbb{P}}$ has the highest weighting at short maturities whereas both $\theta_m^{\mathbb{P}}$ and θ_m have increasing weights as the swap maturity increases. The factor loadings and the ranking of the three long-run mean values generate a mean upward sloping term structure, as shown in the right panel of Figure 4.

The two-factor structure in (33) also allows us to generate different statistical persistence for swap rates at different maturities. Table IV shows that the statistical mean-reversion speed for the instantaneous variance rate is also much larger than that for the central tendency factor. The statistical half life of the instantaneous variance rate is just about four weeks whereas the statistical half life of the central tendency factor is over three years. The increasing weight on the central tendency factor suggests that the variance swap rate becomes increasingly persistent as the swap maturity increases. This observation is consistent with the weekly autocorrelation estimates reported in Table I.

Given the model parameter estimates, we use the Kalman filter to extract the variance risk factors from the variance swap rate series. Figure 5 plots the time series of the extracted variance risk factors in the left panel, with the solid line denoting the instantaneous variance rate v_t and the dashed line denoting the central tendency m_t. The instantaneous variance rate moves around the central tendency factor, which shows more persistence and less instantaneous volatility than the variance rate.

[Figure 5 about here.]

Combining the market price of risk coefficient estimates with the extracted risk factors, we can compute the instantaneous risk premium on the two risk factors, $\gamma^v \sigma_v v_t$ for the instantaneous variance rate and $\gamma^m \sigma_m m_t$ for the central tendency factor. The right panel of Figure 5 plots the time series of the risk premia for the two risk factors, with the solid line for the variance rate and the dashed line for the central tendency. Since the risk premium for the central tendency is about 50 times smaller than the risk premium for the instantaneous variance rate, the plot uses different scales for the two lines, with the scale on the left hand side for the variance rate risk premium and the scale on the right hand side for the central tendency risk premium.

F. Model stability and out-of-sample performance

To gauge the stability of the model and its out-of-sample performance, we divide the sample into two subsample periods. The first subsample is from January 10, 1996 to June 27, 2001, 286 weekly observations for each series. The second subsample is the remaining sample period from July 4, 2001 to March 28, 2007, 300 weekly observations for each series. Each sample contains about five and a half years of data.

Table V report the subsample model parameter estimates. The estimates are largely in line with the whole-sample estimates. Comparing the parameter estimates of the two-factor variance risk model during the two subsamples, we observe that the risk-neutral mean-reversion speed for the instantaneous variance rate is smaller during the first subsample than during the second sample. The opposite is true for the central tendency factor as the estimate for its mean-reversion speed is no longer significantly different from zero during the second subsample. The different estimates suggest that the roles of the two risk factors become more separated during the second sample: The impacts of the instantaneous variance rate are mainly at short maturities and die out quickly as the maturity increases. The impacts of the central tendency factor become even more persistent across the variance swap term structure. On the other hand, the market price estimates are relatively stable over the two sample periods.

To gauge the out-of-sample pricing performance of the two models, we use the parameters estimated from the first subsample to price the variance swap rates during the second subsample. Table VI reports the summary statistics on the out-of-sample pricing errors. Compared to the in-sample pricing error statistics in Table III, the out-of-sample errors are not much different. The average explained variation for the one-factor model is 97.36% compared to the in-sample performance of 95.81%. The average explained variation for the two-factor model is 99.49% compared to the in-sample estimate of 99.54%. These statistics show that the model generates very stable performance over time.

III. Optimal Portfolio Choice

The availability of variance swap contracts makes it convenient for institutional investors to either hedge away variance risk or to achieve additional exposures in it and receive variance risk premium. How does this availability alter an investor's strategic allocation in the stock market? How does it influence the investor's welfare? To answer these questions, we solve the optimal portfolio choice problem for an institutional investor who has access to bonds, a stock index, and return variance swaps on the index across different maturities.

We assume that under the statistical measure $\mathbb{P}$, the stock index dynamics are governed by the following stochastic differential equation,

$$dS_t/S_t \quad = \quad \left(r+\gamma^S v_t\right)dt+\sqrt{v_t}dB_t^S, \tag{35}$$

where r denotes the instantaneous riskfree interest rate, which we assume constant, B_t^S denotes a Brownian motion that describes the stock index return risk, and $\gamma^S v_t$ denotes the instantaneous risk premium on the index return, which we assume is proportional to the instantaneous variance rate level. We analyze the allocation problem under both the one-factor and two-factor variance risk specifications that we have estimated in the previous section.

A. The one-factor stochastic variance model

The statistical dynamics of the one-factor variance rate is specified in (24). To capture the well-known correlation between the index return and index variance, we can decompose the Brownian motion in the variance dynamics into two components,

$$dB_t^v = \rho dB_t^S + \sqrt{1-\rho^2} dB_t^z, \tag{36}$$

where ρ measures the instantaneous correlation between the return risk and the variance risk and B_t^z denotes the independent component of the variance risk. If we let the market price of the independent variance risk also proportional to the square root of the instantaneous variance rate level,

$$\gamma(B_t^z) = \gamma^z \sqrt{v_t}, \tag{37}$$

we can decompose the instantaneous total variance risk premium into two component, one from the return risk and the other from the independent variance risk,

$$\gamma^v v_t = \rho\gamma^S v_t + \sqrt{1-\rho^2}\gamma^z v_t. \tag{38}$$

Under the one-factor variance risk dynamics, variance swap rates of all maturities are affine functions of the instantaneous variance rate. Thus, we can use one variance swap contract of any maturity to completely span the variance risk.

Let W^S, W^B denote the amount of money invested in stock and bond, and let N denote the notional of the variance swap contract. Consider now an investor who allocates her initial wealth $W_0 > 0$ among the stock index S_t, the variance swap contract $VS_t(T)$, and a money market account with a riskfree rate of return r. Then, the budget constraint becomes

$$W_t = W_t^S + W_t^B + N(VS_t - K),$$

where K is the delivery price of the variance swap contract. At inception, we have $K = VS_t = 0$. If we use share prices and portfolio fractions instead of money amounts, we can write the portfolio dynamics as

$$\frac{dW}{W} = w_t dS_t/S_t + (1-w_t)\, rdt + n_t dVS_t$$

where w denotes the fraction of wealth in stock and n the fraction of wealth used as notional for the variance swap investment. Therefore, in the one-factor model, the investor's wealth evolves according to

$$\frac{dW}{W} = rdt + \left(w_t\gamma^S + n_t\phi\sigma_v\gamma^v\right) v_t dt + w_t\sqrt{v}dB_t^S + n_t\phi_v\sigma_v\sqrt{v}dB_t^v \tag{39}$$

The investor chooses the portfolio weights to maximize her utility of terminal wealth $W_{\mathcal{T}}$ at time $\mathcal{T}$. Assuming constant relative risk aversion (CRRA) utility with relative risk aversion coefficient $\eta > 0$, we can write the indirect utility function as

$$J(t,W,v) = \sup_{(w_t,n_t)} \mathbb{E}\left(\frac{W_{\mathcal{T}}^{1-\eta}}{1-\eta}\middle| W_t = W, v_t = v\right), \quad \eta \neq 1, \tag{40}$$

subject to the budget constraint in (39). The corresponding Hamilton-Jacobi-Bellman (HJB) equation is

$$\begin{aligned} 0 \;=\; & \sup_{(w_t,n_t)}\left\{J_t + J_v\kappa_v^{\mathbb{P}}\left(\theta_v^{\mathbb{P}} - v_t\right) + J_W W r + J_W W\left(w_t\gamma^S + n_t\phi_v\sigma_v\gamma^v\right)v_t \right.\\ & +\tfrac{1}{2}J_{WW}W^2\left[w_t^2 + n_t^2\phi_v^2\sigma_v^2 + 2w_t n_t\phi_v\sigma_v\rho\right]v_t \\ & \left. +J_{Wv}W\left[\sigma_v\rho w_t + \sigma_v^2 n_t\phi_v\right]v_t\right\}, \end{aligned} \tag{41}$$

where J_t, J_v, J_W, J_{Wv}, and J_{WW} denote the first, second, and cross derivatives with respect to t, v, and W. From the first order conditions with respect to (w_t, n_t), we obtain the optimal portfolio allocation in following generic form.

Proposition 3 *Given the optimization problem in* (40) *under the budget constraint* (39)*, the stock and variance dynamics in (35) and* (24)*, the optimal allocation in the stock (w_t) and the variance swap contract (n_t) are*

$$w_t = -\frac{J_W}{W_t J_{WW}}\left(\gamma^S - \frac{\rho}{\sqrt{1-\rho^2}}\gamma^z\right), \tag{42}$$

$$n_t = -\frac{1}{\phi_v(\tau)}\left(\frac{J_W}{W_t J_{WW}}\frac{\gamma^z}{\sqrt{1-\rho^2}\sigma_v} + \frac{J_{Wv}}{W_t J_{WW}}\right). \tag{43}$$

For comparison, we also derive analogously the generic optimal allocation decisions when the investor can only invest in one of the two risky assets:

Proposition 4 *If the investor can only invest in bonds and the stock index, the optimal fraction of wealth invested in the stock index,*

$$w_t = -\frac{J_W}{W_t J_{WW}}\gamma^S - \rho\sigma_v \frac{J_{Wv}}{W_t J_{WW}}. \tag{44}$$

If the investor can only invest in bond and the spot variance contract, the optimal fraction of wealth invested in the spot variance contract is

$$n_t = -\frac{1}{\phi_v(\tau)}\left(\frac{J_W}{W_t J_{WW}}\frac{\gamma^v}{\sigma_v} + \frac{J_{Wv}}{W_t J_{WW}}\right). \tag{45}$$

As shown in the seminal paper of Merton (1971), the optimal allocation to risky assets includes two components: a myopic component that is proportional to the mean excess return and an intertemporal hedging demand that is proportional to the covariance between the risky asset returns and the state variables that govern the stochastic investment opportunity, both scaled by the covariance matrix of the asset return. In our application, the stochastic variance risk represents the stochastic investment opportunity, which induces an intertemporal hedging demand when we invest in either the stock index alone (equation (44)) or the variance swap contract alone (equation (45)).

Interestingly, when we invest in both the stock index and the variance swap contract to span both the return risk and the variance risk, the optimal investment in the stock index no longer includes an intertemporal hedging demand. Stochastic investment opportunity asks for intertemporal hedging demand; yet an appropriately designed derivative contract can be used to span the risk inherent in the stochastic investment opportunity and to eliminate the need for intertemporal hedging using the primary securities.

To better understand what type of derivative contracts can help us eliminate the need for intertemporal hedging, we consider a generic example of two risky assets, with the first security as the primary security (a stock index) and the second security as a specially designed derivative contract to hedge against the stochastic investment opportunity, which is governed by a one-factor state variable X_t.

If we use $\Sigma_{SS}, \mu_S, \Sigma_{SX}$ to denote the return covariance matrix of the two risky assets, the mean excess return vector, and the covariance between the two risky assets and the state variable, with

$$\Sigma_{SS} = \begin{bmatrix} \sigma_1^2 & \sigma_1\sigma_2\rho_{12} \\ \sigma_1\sigma_2\rho_{12} & \sigma_2^2 \end{bmatrix}, \quad \mu_S = \begin{bmatrix} \gamma_1\sigma_1 \\ \gamma_2\sigma_2 \end{bmatrix}, \quad \Sigma_{SX} = \begin{bmatrix} \sigma_1\sigma_x\rho_{1x} \\ \sigma_2\sigma_x\rho_{2x} \end{bmatrix},$$

we can write the optimal allocation weights as,

$$w = -\frac{J_W}{W_t J_{WW}} \Sigma_{SS}^{-1} \mu_S - \frac{J_{WX}}{W_t J_{WW}} \Sigma_{SS}^{-1} \Sigma_{SX}. \tag{46}$$

Plugging in the elements of the matrices, we have

$$\begin{aligned} w_1 &= -\tfrac{J_W}{W_t J_{WW}} \tfrac{1}{\left(1-\rho_{12}^2\right)\sigma_1} \left[\gamma_1 - \rho_{12}\gamma_2\right] - \tfrac{J_{WX}}{W_t J_{WW}} \tfrac{1}{\left(1-\rho_{12}^2\right)\sigma_1} \left[\rho_{1x} - \rho_{12}\rho_{2x}\right] \sigma_x, \\ w_2 &= -\tfrac{J_W}{W_t J_{WW}} \tfrac{1}{\left(1-\rho_{12}^2\right)\sigma_2} \left[\gamma_2 - \rho\gamma_1\right] - \tfrac{J_{WX}}{W_t J_{WW}} \tfrac{1}{\left(1-\rho_{12}^2\right)\sigma_2} \left[\rho_{2x} - \rho_{12}\rho_{1x}\right]. \end{aligned} \tag{47}$$

To exclude the investment in the first risky asset from having an intertemporal hedging demand, we need

$$\rho_{1x} - \rho_{12}\rho_{2x} = 0. \tag{48}$$

In our application, the stock index return has an instantaneous correlation of ρ with the stochastic variance risk. The variance swap contract is an affine function of only the variance rate and depends on no other variables. Thus, the variance swap contract has perfect instantaneous correlation with the stochastic variance risk factor. The condition in (48) is satisfied. By having $\rho_{2x} = 1$, we have $\rho_{12} = \rho_{1x} = \rho$ and hence $\rho_{1x} - \rho_{12}\rho_{2x} = 0$.

Therefore, when we design a derivative instrument that shows perfect instantaneous correlation with the state variable, it eliminates the intertemporal hedging need from the primary security investment. Under the one-factor stochastic variance risk model, variance swap contracts play this role perfectly as they are only a function of the variance rate itself. By contrast, although an option contract can also be used to span the variance risk, it cannot be used to exclude the intertemporal hedging demand from the underlying primary security, because the variation of the option value not only depends on the variance risk, but also depends on the stock index level and hence the return risk.

Chacko and Viceira (2005) argue that the presence of stochastic volatility has only a small effect on the portfolio allocation. Our results show that the picture can be dramatically different when we have access to variance swap contracts, which can be used to hedge against the variance risk and alleviate the underlying stock index investment from the need of intertemporal hedging. In the presence of strongly negative variance risk premia, the variance swap contracts can also be used to gain exposures to the variance risk for enhanced returns.

Under the one-factor variance risk dynamics, we can further solve the indirect utility function and the optimal allocation weights in analytical forms, which are summarized in the following proposition.

Proposition 5 *In the one-factor stochastic variance model, the indirect utility function has the following analytical representation*

$$J(t,W,v) = \frac{W_t^{1-\eta}}{1-\eta}\exp\left(h_v(u)v_t + k(u)\right), \tag{49}$$

with u denoting the investment horizon $\mathcal{T}-t$. The optimal portfolio weights in the stock index and the variance swap contract are

$$w_t^S = \frac{1}{\eta}\left(\gamma^S - \frac{\rho}{\sqrt{1-\rho^2}}\gamma^z\right), \tag{50}$$

$$w_t^{SV} = \frac{1}{\eta\phi_v(\tau)}\left(\frac{\gamma^z}{\sqrt{1-\rho^2}\sigma_v} + h_v(u)\right), \tag{51}$$

where the intertemporal hedging demand is determined by the coefficient $h_v(u)$, which solves the following partial differential equation,

$$h_v'(u) = \alpha h_v(u)^2 + \beta h_v(u) + \zeta, \tag{52}$$

with

$$\alpha = \frac{1}{2}\frac{(1-\eta)}{\eta}\sigma_v^2, \quad \beta = \frac{(1-\eta)}{\eta}\sigma_v\gamma^v - \kappa^{\mathbb{P}}, \quad \zeta = \frac{1}{2}\frac{(1-\eta)}{\eta}\left((\gamma^S)^2 + (\gamma^z)^2\right), \tag{53}$$

starting at $h_v(0) = 0$.

The derivation and the expression for the coefficient $k(u)$ in the indirect utility function are given in Appendix A.

For a myopic investor, the intertemporal hedging demand is zero. The investments in the stock index and the variance contract depend crucially on the risk premiums on the two sources of risks. In the absence of a risk premium on the independent variance risk $\gamma^z = 0$, the investor invests positively in the stock index (given that the index generates positive risk premium, $\gamma^S > 0$) but does not invest in the variance contract. A positive risk premium on the independent variance risk component induces a positive investment in the variance contract and also increases the investment in the stock index. On the other hand, when the risk premium on the independent risk component is negative, the investor short sells the variance contract and also reduces her investment in the stock index accordingly. In the case when the market price of the independent variance risk is highly negative and its magnitude dominates the market price of the return risk, the negative variance risk premium effect can dominate the positive return risk premium and the optimal investment in the stock index can also become negative.

Under the one-factor variance risk model, our estimate for the market price of the total variance risk is highly negative at $\gamma^v = -17.0141$. Given the extracted variance rate series, we can estimate the return risk dynamics and its instantaneous correlation with variance risk through a simple likelihood specification on the index return data and the extracted variance rate series. We have done so over the whole sample period and obtain an estimate for the market price of return risk at $\gamma^S = 1.5326$ and the instantaneous correlation at $\rho = -0.8069$.[4] Combining these estimates with γ^v, we can compute the market price of the independent variance risk as $\gamma^z = (\gamma^v - \rho\gamma^S)/\sqrt{1-\rho^2} = -26.71$. Under such highly negative variance risk premium estimates, the investor not only shorts the variance contract, but also shorts the stock index, as $\left(\gamma^S - \frac{\rho}{\sqrt{1-\rho^2}}\gamma^z\right) = -34.95$.

When the investor has non-myopic investment horizon, her investment in the variance contract varies with the investment horizon, but her investment in the stock index does not. The hedging coefficient $h_v(u)$ starts at zero when the investment horizon u is zero. With $\eta > 1$ and hence $\zeta < 0$ in equation (53), the hedging demand becomes negative as the investment horizon increases.

B. The two-factor stochastic variance model

Under the two-factor stochastic variance model, both the instantaneous variance rate and its central tendency are stochastic. To fully span these two sources of variance risk, we need to invest in two variance contracts with different maturities T_1 and T_2, with $T_1 \neq T_2$.

[4]The estimate for γ^S is not statistically significant, with a t-statistic close to zero. The t-static for ρ is significant at 2.26.

Solving an analogous utility optimization problem,

$$J(t, W, v, m) = \sup_{(w_t, n_t^{\tau_1}, n_t^{\tau_2})} \mathbb{E}\left(\frac{W_T^{1-\eta}}{1-\eta} \middle| W_t = W, v_t = v, m_t = m \right), \quad \eta \neq 1, \tag{54}$$

with the budget constraint

$$\frac{dW_t}{W_t} = w_t dS_t / S_t + n_t^{\tau_1} dVS_t(T_1) + n_t^{\tau_2} dVS_t(T_2) + (1 - w_t) r dt, \tag{55}$$

we obtain the following optimal allocations.

Proposition 6 *Under the two-factor variance risk model, the optimal fractions of wealth invested in the stock index and the two spot variance contracts are*

$$\begin{aligned} w_t^S &= \tfrac{1}{\eta}\left(\gamma^S - \tfrac{\rho}{\sqrt{1-\rho^2}}\gamma^z\right), \\ n_t^{\tau_1} &= \tfrac{1}{\eta D}\left[\left(\tfrac{\gamma^z}{\sigma_v\sqrt{1-\rho^2}} + h_v(u)\right)\phi_m(\tau_2) - \left(\tfrac{\gamma^m}{\sigma_m} + h_m(u)\right)\phi_v(\tau_2)\right], \\ n_t^{\tau_2} &= \tfrac{1}{\eta D}\left[-\left(\tfrac{\gamma^z}{\sigma_v\sqrt{1-\rho^2}} + h_v(u)\right)\phi_m(\tau_1) + \left(\tfrac{\gamma^m}{\sigma_m} + h_m(u)\right)\phi_v(\tau_1)\right], \end{aligned} \tag{56}$$

where

$$D = \phi_v(\tau_1)\phi_m(\tau_2) - \phi_v(\tau_2)\phi_m(\tau_1), \tag{57}$$

and the intertemporal hedging demand coefficients $h_v(u)$ *and* $h_m(u)$ *solve the ordinary differential equations* (A.7) *and* (A.8) *in Appendix B.*

Even in the presence of two variance risk factors, investing in the variance swap contracts can still eliminate the intertemporal hedging demand on the stock index, further illustrating that variance swap contacts are the simplest and most direct contracts for spanning variance risks.

At short investment horizons, the optimal allocations in the two variance contracts depend on the market prices of the independent variance rate risk (γ^z) and the central tendency risk (γ^m), as well as the exposures of the two contracts on the two sources of risk, $\phi_v(\tau)$ and $\phi_m(\tau)$. If we assume $T_1 < T_2$ such that the first is a short-term variance swap contract whereas the second is a long-term contract, we have $\phi_v(\tau_1) > \phi_v(\tau_2)$ whereas $\phi_m(\tau_1) < \phi_m(\tau_2)$ and the common scaling coefficient D on the two variance contracts is positive.

Thus, the long-term contract always puts more emphasis on the market price of the central tendency factor, and the short-term contract on the market price of the independent short-term variance risk.

Given a large τ_2 for the variance swap with the longer maturity, the first variance swap contract is mainly exposed to the instantaneous variance rate risk, since $\phi_m(\tau_2)$ dominates $\phi_v(\tau_2)$. As such, investment in the short-term contract is mainly determined by the market price of the independent variance rate risk γ^z. Analogously, with a very small τ_1, the investment in the long-term contract is mainly determined by the market price of the central tendency factor γ^m. The investor may short both contracts when the market prices of the two sources of risks are sufficiently negative. Indeed, as long as the variance risk premia on the independent variance risk and the central tendency are negative, the net position in variance swaps $n_t^{\tau_1} + n_t^{\tau_2}$ is always negative.

On the other hand, when the maturity separation is moderate, optimal allocations to the two variance contracts also depend on the relative magnitude of the two market prices γ^z and γ^m. In particular, since our estimates for γ^z are much larger in absolute magnitude than those for γ^m, the γ^z term can dominate the investment decision. In this case, the optimal allocation decision for the investor is to short the short-term variance swap contract and long the long-term variance swap contract, $n^{\tau_1} < 0$ and $n^{\tau_2} > 0$.

The intertemporal hedging demands for the two swap contracts start at zero for a myopic investor. At longer investment horizons, the hedging demands are determined by the two coefficients $h_v(u)$ and $h_m(u)$ for the two variance risk exposures, which start at zero at $u = 0$ and propagate according to the two ordinary differential equations in (A.7) and (A.8) as the investment horizon increases. Since the constant terms in the two ordinary differential equations are both negative, the hedging demand coefficients grow increasingly negative as the investment horizon increases.

IV. Empirical Analysis on Variance Swap Investments

In this section, we combine the theoretical results from Section III with the parameter estimates from Section II to compare the historical out-of-sample performance. For the optimal strategy, we choose a variance swap with the shortest maturity ($\tau_1 = 2/12$) and with the longest maturity ($\tau_2 = 2$).

For the out-of-sample study, we proceed as follows. For the estimation of our model, we use the estimates from our first subsample in Table V, i.e., we start January 10, 1996, and use the data until June 27, 2001, for estimation. For our investment strategy, we start on January 10, 1996. By doing so, we can split up the whole period into an in-sample and an out-of-sample subperiod.

We set the investment horizon equal to two months, i.e., the investment horizon corresponds to the shortest maturity of the variance swap contracts available. After two months, we sell the whole portfolio to the current available market prices. For the calculation of the return on the long-term variance swap, we use the mark-to-market value of the variance swap at time s, $t \leq s \leq T$:

$$VS_s(T) = \lambda (RV(t,s) - K_t T) + (1+\lambda)(K_s(T) - K_t(T)), \tag{58}$$

where $RV(t,s)$ is the realized variance between t and s. $K_t(T)$ and $K_s(T)$ are the strikes of variance swap contracts maturing at T and initiated at time t and s, respectively. The weight λ is given as

$$\lambda = \frac{s-t}{T-t}.$$

To obtain weekly returns, we repeat this exercise every Wednesday throughout the whole sample period.

In Figure 6, we plot the optimal portfolio weights for the stock and variance swaps (upper panel), and the stock only (lower panel). We observe that the presence of the variance swap market together with the highly negative and significant variance risk premia gives rise to substantial changes in the optimal asset allocation in the stock. While the classical stock-bond portfolio optimally gives a substantial long-position in the index ($w_p = 0.7063$), the optimal position in stocks becomes slightly negative ($w = -0.1176$). At the same time, the investor goes long the variance swap contract with the longer maturity $(n^1 = 0.1912)$ and short the variance swap with the shorter maturity $(n^1 = -0.8076)$. Hence, the portfolio allocation strongly react to the high (and negative) variance risk premia. Nevertheless, the investor uses both the long-term variance swap and the stock to hedge her position in the short-term variance swap. In particular, the long-term variance swap helps to dampen the effect of persistent changes in the volatility level. Due to the large negative correlation, the stock serves as a hedging instrument when volatilities increase, which might be particularly useful when volatility spikes.

[Figure 6 about here.]

In the right panels of Figure 6, we plot the cumulative wealth path for the optimal variance swap strategy (upper panel) and the stock-bond strategy (lower panel), together with the S&P 500 index level. The dash-dotted vertical line indicates the beginning of the out-of-sample period. Not surprisingly, we observe that the stock-bond strategy is highly correlated with the market. In contrast, the strategy based on variance swaps is able to provide a smooth evolution of wealth. The downturn after 2001 has practically no influence on its success. Since we have reported very stable estimates for the return dynamics, the out-of-sample performance is almost the same as the in-sample performance.

Table VII reports the summary statistics for the cumulative wealth of three different strategies: the optimal strategy with stock, bond, and two variance swaps, $\mathcal{S}_1$, the stock-bond strategy $\mathcal{S}_2$, and $\mathcal{S}_3$, where we only invest in two variance swaps, but not in the stock. We see that the stock-bond strategy yields only a small Sharpe ratio, which moreover varies a lot for the two subsample periods. In contrast, the strategies based on variance swaps gives a high Sharpe ratio not only for the in-sample period, but also for the whole sample and out-of-sample period.

V. Concluding Remarks

Using more than a decade worth of weekly quotes on variance swap rates, we design and estimate models for the S&P 500 index return variance dynamics and market prices of variance risks. We find that we need two stochastic variance risk factors to explain the variation in the variance swaps at different maturities, with one factor controlling the instantaneous variance rate while the other controlling the central tendency of the variance rate movements. The variance rate factor is much more transient than the central tendency factor under both the risk-neutral and the statistical measures, thus generating different loading patterns across the term structure from the two risk factors and different autocorrelation patterns for variance swap rates of different maturities. We also find that the specification of market prices for both the equity and variance risk factors is crucial to provide enough flexibility for portfolio allocation.

Embedding variance swaps into an optimal investment strategy, we find that with the variance swap contract to span the stochastic variance risk, an investor drastically changes her portfolio allocation. The intertemporal hedging demand is completely removed from the stock investment. Moreover, the large and negative variance risk premia forces stock investments to become negative. Analyzing different strategies, we find that our variance swap strategy can provide smooth returns with high Sharpe ratios. Since the parameter estimates of the model are consistent for different subperiods, we are able to generate high Sharpe ratios also out-of-sample.

Compared to traditional mean variance analysis, the modern financial industry has recognized the important impacts that stock price jumps and stochastic variance can have on an investor's welfare. Accordingly, derivative securities such as options and variance swaps have been developed to span risks along these two dimensions. To simplify the problem and to gain a clearer picture of their separate effects, the academic literature either assumes constant volatility and focuses on how to choose options at different strikes to span the random jump risk in the stock price (Carr and Madan (2001), Carr and Wu (2002)), or assumes purely continuous dynamics (or jumps of fixed size) and focuses on how to choose options (Liu and Pan (2003)) or, as what we do in this paper, variance swap contracts at different maturities to span the stochastic variance risk. Integrating these two dimensions can be a challenging but interesting direction for future research.

Appendix

A. Proof of Proposition 5

We can exploit the homogeneity structure of the optimization problem and make the Ansatz

$$J(t,W,v) = \frac{e^{g(u,v)}W_t^{1-\eta}}{1-\eta}, \tag{A.1}$$

and

$$g(t,v) = h_v(u)v_t + k(u),$$

with the boundary condition $g(0,X) = 0$, the corresponding HJB reduces to

$$\begin{aligned} 0 &= -h_v'(u)v_t - k'(u) + r(1-\eta) + \left(\kappa_v^{\mathbb{P}}\theta_v^{\mathbb{P}} - \frac{\eta\kappa_v^{\mathbb{P}} + (\eta-1)\sigma_v\gamma^v}{\eta}v_t\right)h_v(u) \\ &\quad + \frac{1-\eta}{2\eta}v_t\left((\gamma^S)^2 + (\gamma^z)^2 + \sigma_v^2 h_v(u)^2\right). \end{aligned}$$

Comparing coefficients, we obtain the ordinary differential equations for $h_v(t)$ and $c_{gv}(t)$ as

$$-h_v'(u) = \frac{\eta-1}{2\eta}\sigma_v^2 h_v(u)^2 + \frac{\eta\kappa_v^{\mathbb{P}} + (\eta-1)\sigma_v\gamma^v}{\eta}h_v(u) + \frac{\eta-1}{2\eta}\left((\gamma^S)^2 + (\gamma^z)^2\right) \tag{A.2}$$

$$k'(u) = r(\eta-1) - \tilde{\theta}_v\tilde{\kappa}_v h_v(u), \tag{A.3}$$

with terminal condition $h_v(0) = k(u) = 0$. We can solve for $h_v(u)$ explicitly. Writing the ODE as

$$h_v'(u) = \alpha h_v(u)^2 + c_1 h_v(u) + c_0. \tag{A.4}$$

Then, the solution of $h_v(t)$ can be written as

$$h_v(t) = \frac{\alpha\left(e^{-D(T-t)} - 1\right)(\beta - D)}{\beta^2 - 2\alpha\zeta(1 + e^{-Du}) - \alpha D}, \tag{A.5}$$

where $D = \sqrt{\beta^2 - 4\alpha\zeta}$. ■

B. Proof of Proposition 6

The proof follows the same line as for the one-factor model. For the indirect utility function, we can make the Ansatz

$$J(t,W,v) = \frac{e^{g(u,v,m)} W_t^{1-\eta}}{1-\eta}, \tag{A.6}$$

and

$$g(t,v) = h_v(u)v_t + h_m(u) + k(u),$$

with the boundary condition $g(0,X) = 0$. Solving for the optimal weights triplet $(w_t, n_t^{\tau_1}, n_t^{\tau_2})$, plugging them back into the HJB, and collecting terms, we get the ordinary differential equations for $h_v(u), h_m(u), k(u)$:

$$-h_v'(u) = \frac{\eta-1}{2\eta}\sigma_v^2 h_v(u)^2 + \frac{\eta\kappa_v^{\mathbb{P}} + (\eta-1)\sigma_v\gamma^v}{\eta} h_v(u) + \frac{\eta-1}{2\eta}\left((\gamma^S)^2 + (\gamma^z)^2\right), \tag{A.7}$$

$$-h_m'(u) = \frac{\eta-1}{2\eta}\sigma_m^2 h_m(u)^2 - \kappa_v h_v(u) + \frac{\eta\kappa_m^{\mathbb{P}} + (\eta-1)\sigma_m\gamma^m}{\eta} h_m(u) + \frac{\eta-1}{2\eta}(\gamma^m)^2, \tag{A.8}$$

$$-k'(u) = r(\eta-1) - \theta_m\tilde{\kappa}_m h_m(u). \tag{A.9}$$

Note that due to the correlation structure, the solution for $h_v(u)$ is exactly the same as in the one-factor model. ■

References

Andrews, D., 1991, “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica*, 59, 817–858.

Bakshi, G., and N. Kapadia, 2003, “Delta-Hedged Gains and the Negative Market Volatility Risk Premium,” *Review of Financial Studies*, 16(2), 527–566.

Balduzzi, P., S. Das, and S. Foresi, 1998, “The Central Tendency: A Second Factor in Bond Yields,” *Review of Economics and Statistics*, 80(1), 62–72.

Bondarenko, O., 2004, “Market Price of Variance Risk and Performance of Hedge Funds,” working paper, University of Illinois at Chicago.

Carr, P., X. Jin, and D. Madan, 2001, “Optimal Investment in Derivative Securities,” *Finance and Stochastics*, 5(1), 33–60.

Carr, P., and D. Madan, 2001, “Optimal Positioning in Derivative Securities,” *Quantitative Finance*, 1(1), 19–37.

Carr, P., and L. Wu, 2002, “Static Hedging of Standard Options,” working paper, New York University and Baruch College.

——— , 2007, “Variance Risk Premia,” *Review of Financial Studies*, forthcoming.

Chacko, G., and L. Viceira, 2005, “Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets,” *Review of Financial Studies*, 18(4).

Coval, J. D., and T. Shumway, 2001, “Expected Option Returns,” *Journal of Finance*, 56(3), 983–1009.

Duffie, D., 1992, *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, New Jersey, 2nd edn.

Duffie, D., J. Pan, and K. Singleton, 2000, “Transform Analysis and Asset Pricing for Affine Jump Diffusions,” *Econometrica*, 68(6), 1343–1376.

Engle, R., 2004, “Risk and Volatility: Econometric Models and Financial Practice,” *American Economic Review*, 94(3), 405–420.

Heston, S., 1993, “Closed-Form Solution for Options with Stochastic Volatility, with Application to Bond and Currency Options,” *Review of Financial Studies*, 6(2), 327–343.

Liu, J., and J. Pan, 2003, “Dynamic Derivative Strategies,” *Journal of Financial Economics*, 69, 401–430.

Merton, R. C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3(4), 373–413.

Newey, W. K., and K. D. West, 1987, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55(3), 703–708.

Table I
Summary statistics of the variance swap rates

Entries report the mean, standard deviation (Std), skewness (Skew), excess kurtosis (Kurt) and weekly autocorrelation (Auto) of the variance swap rates (panel A) and annualized realized variance (panel B) on S&P 500 index at different maturities. Data are weekly (every Wednesday) from January 10, 1996, to March 30, 2007 (586 observations for each series).

Maturity	Mean	Std	Skew	Kurt	Auto
A. Variance swap rates, Maturity in months					
2	0.048	0.033	1.861	5.734	0.920
3	0.048	0.031	1.750	5.253	0.943
6	0.050	0.031	1.825	6.338	0.959
12	0.053	0.030	1.385	3.421	0.971
24	0.056	0.029	1.146	1.780	0.977
B. Annualized realized variance, Maturity in days					
7	0.031	0.044	4.900	33.500	0.382
30	0.032	0.031	2.524	7.937	0.925
60	0.030	0.026	1.958	4.188	0.970
90	0.031	0.025	1.651	2.924	0.982
120	0.031	0.024	1.498	2.480	0.989
150	0.031	0.022	1.315	1.805	0.992

Table II
Summary statistics of profit and loss from long variance swaps and holding the contracts to maturity

Entries report the mean, standard deviation (Std), skewness (Skew), excess kurtosis (Kurt), Newey-West serial dependence adjusted standard deviation (Newey), and annualized Sharpe ratios (IR) on the profit and loss from long variance swaps and holding the contracts to maturity. The profit and loss are measured in terms of the dollar capital gains per \$100 par investment, $\$100 \times (RV - VS)$, in panel A and in terms of log excess returns, $\ln(RV/SV)$, in panel B. The investment horizon is measured in months. Data are weekly (every Wednesday) from January 10, 1996, to March 30, 2007, 586 observations for each full sample data series. The samples on six-, 12-, and 24-month investments are shortened by one, seven, and 19 months, respectively, due to the constraint on ex post realized variance calculation.

Horizon	Mean	Std	Skew	Kurt	Newey	IR
A. Dollar capital gain per \$100 par investment, $100(RV - VS)$						
2	-1.721	2.771	-1.141	8.754	7.694	-0.548
3	-1.684	2.778	-1.006	7.259	8.988	-0.375
6	-1.919	2.778	-1.485	8.770	11.399	-0.238
12	-2.296	2.525	-1.754	7.316	13.174	-0.174
24	-2.712	2.448	-1.431	3.566	13.836	-0.139
B. Log excess return, $\ln(RV/VS)$						
2	-0.549	0.497	0.658	0.322	1.322	-1.017
3	-0.529	0.505	0.593	0.221	1.630	-0.649
6	-0.551	0.485	0.521	0.450	2.078	-0.375
12	-0.603	0.445	0.222	-0.232	2.690	-0.224
24	-0.636	0.487	-0.216	-1.278	3.278	-0.137

Table III
Summary statistics of the pricing errors on the variance swap rates

Entries report the summary statistics of the model pricing errors on the variance swap rates, including the sample average (Mean), root mean squared error (Rmse), weekly autocorrelation (Auto), maximum absolute error (Max), and explained percentage variation (R^2), defined as one minus the variance of the pricing error to the variance of the original swap rate quotes. The pricing errors are defined as the difference between the variance swap rate quotes and the corresponding model-implied values, both in volatility percentage points. The last row reports the maximized log likelihood values for the two models.

	One-factor model					Two-factor model				
Maturity	Mean	Rmse	Auto	Max	R^2	Mean	Rmse	Auto	Max	R^2
2	-0.15	1.77	0.82	7.86	93.19	0.27	0.80	0.62	4.20	98.77
3	-0.20	1.13	0.89	4.49	97.05	-0.00	0.00	0.35	0.01	100.00
6	0.00	0.00	0.44	0.00	100.00	-0.09	0.40	0.75	2.59	99.61
12	0.04	1.08	0.90	3.88	96.82	0.00	0.00	0.18	0.00	100.00
24	-0.60	1.77	0.95	5.05	92.01	-0.05	0.49	0.68	3.39	99.32
Average	-0.18	1.15	0.80	4.25	95.81	0.03	0.34	0.52	2.04	99.54
Likelihood			−5793.4					−3318.4		

Table IV
Parameter estimates of affine stochastic variance models

Entries report the maximum likelihood parameter estimates and t-values (in parentheses) of the one-factor and the two-factor affine stochastic variance models. The estimation employs weekly data on variance swap rates at maturities of two, three, six, 12, and 24 months and ex post realized variances at maturities of seven, 30, 60, 90, and 150 days. The sample is from January 10, 1996, to March 28, 2007, 586 observations for each series.

X_t	κ	θ	σ	γ	$\kappa^{\mathbb{P}}$	$\theta^{\mathbb{P}}$
A: One-factor variance risk model						
v_t	0.1547	0.1220	0.2550	-17.0141	4.4929	0.0042
	(13.81)	(32.55)	(47.47)	(43.06)	(71.42)	(21.52)
B: Two-factor variance risk model						
v_t	4.3730	–	0.4221	-16.3746	11.2851	0.0158
	(38.72)		(44.10)	(37.70)	(51.96)	(3.71)
m_t	0.1022	0.0838	0.1581	-0.6844	0.2104	0.0407
	(9.30)	(24.31)	(47.64)	(1.84)	(3.40)	(3.68)

Table V
Subsample parameter estimates of affine stochastic variance models

Entries report the maximum likelihood parameter estimates and t-values (in parentheses) of the one-factor and the two-factor affine stochastic variance models over two subsample periods. The first subsample is from January 10, 1996 to June 27, 2001, 286 weekly observation for each series. The second subsample is from July 4, 2001 to March 28, 2007, 300 weekly observations for each series. The estimation employs weekly data on variance swap rates at maturities of two, three, six, 12, and 24 months and ex post realized variances at maturities of seven, 30, 60, 90, and 150 days.

X_t	κ	θ	σ	γ	$\kappa^{\mathbb{P}}$	$\theta^{\mathbb{P}}$
A: One-factor variance risk model						
Subsample period: 1996- 2001						
v_t	0.0885	0.2442	0.3056	-17.0344	5.3097	0.0041
	(3.51)	(4.47)	(29.93)	(27.07)	(44.08)	(13.23)
Subsample period: 2001-2007						
v_t	0.7281	0.0495	0.2316	-12.1940	3.5517	0.0101
	(21.82)	(87.41)	(32.66)	(25.54)	(33.96)	(28.19)
B: Two-factor variance risk model						
Subsample period: 1996- 2001						
v_t	3.3945	–	0.4635	-16.2472	10.9250	0.0153
	(16.53)		(27.27)	(25.12)	(29.84)	(3.88)
m_t	0.1857	0.0715	0.2086	-0.4029	0.2697	0.0492
	(7.61)	(15.14)	(26.98)	(1.11)	(3.17)	(3.84)
Subsample period: 2001-2007						
v_t	5.3789	–	0.3476	-15.6418	10.8160	0.0239
	(33.83)		(27.86)	(19.44)	(33.91)	(0.46)
m_t	0.0010	4.2187	0.1058	-0.8444	0.0903	0.0480
	(0.06)	(0.06)	(24.42)	(0.46)	(0.46)	(0.46)

Table VI
Out-of-sample pricing performance on the variance swap rates

We estimate the model during the first subsample from January 10, 2996 to June 27, 2001, and use the model parameters to price variance swap rates out of sample from July 4, 2001 to March 28, 2007. The pricing errors are defined as the difference between the variance swap rate quotes and the corresponding model-implied values, both in volatility percentage points. Entries report the summary statistics of the model pricing errors on the variance swap rates, including the sample average (Mean), root mean squared error (Rmse), weekly autocorrelation (Auto), maximum absolute error (Max), and explained percentage variation (R^2), defined as one minus the variance of the pricing error to the variance of the original swap rate quotes.

	One-factor model					Two-factor model				
Maturity	Mean	Rmse	Auto	Max	R^2	Mean	Rmse	Auto	Max	R^2
2	0.525	1.591	0.840	8.097	95.29	0.20	0.64	0.56	2.68	99.23
3	0.272	0.863	0.896	3.963	98.33	-0.00	0.00	0.18	0.00	100.00
6	0.000	0.000	0.368	0.000	100.00	-0.03	0.20	0.73	0.65	99.87
12	-0.623	0.969	0.910	3.956	97.57	0.00	0.00	0.49	0.00	100.00
24	-2.138	2.327	0.900	5.871	95.60	-0.21	0.60	0.89	2.39	98.33
Average	-0.393	1.150	0.783	4.377	97.36	-0.01	0.29	0.57	1.14	99.49

Table VII
Summary statistics of investment strategies

Entries report the mean excess return, excess standard deviation (Std), skewness (Skew), excess kurtosis (Kurt), and annualized Sharpe ratios on three different portfolio strategies. $\mathcal{S}_1$ is the optimal strategy with stock, bond, and two variance swaps. $\mathcal{S}_2$ only invests in the stock and bond. $\mathcal{S}_3$ only invests in the two variance swaps. Data are weekly (every Wednesday) from January 10, 1996, to March 30, 2007, 586 observations for each full sample data series.

Strategy	Mean	Std	Skew	Kurt	Sharpe
A. Whole Period: 1996-2007					
$\mathcal{S}_1$	7.4303	4.9255	1.1945	7.0545	1.4984
$\mathcal{S}_2$	3.3106	10.7539	-0.1838	0.4997	0.3077
$\mathcal{S}_3$	7.9814	5.6856	1.0733	7.3502	1.3966
B. In-sample period: 1996-2001					
$\mathcal{S}_1$	9.0779	5.8340	1.3357	4.7407	1.6003
$\mathcal{S}_2$	6.1913	11.5308	-0.1158	0.2267	0.5396
$\mathcal{S}_3$	10.1087	6.6544	1.3312	5.1225	1.5699
C. Out-of-sample period: 2001-2007					
$\mathcal{S}_1$	5.9342	3.8114	-0.0381	3.8359	1.6522
$\mathcal{S}_2$	0.7085	10.0286	-0.4028	0.5058	0.0732
$\mathcal{S}_3$	6.0522	4.4910	-0.2057	4.8652	1.4083

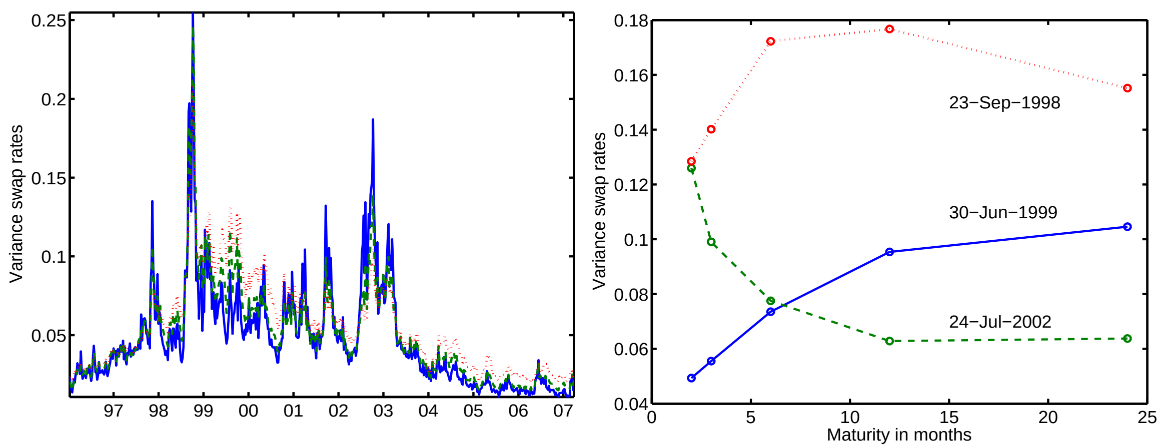

Figure 1. Time series and term structure of the return variance swap rates.
The left panel plots the time series of the variance swap rate quotes (in volatility percentage points) at three selected time to maturities: two months (solid line), six months (dashed line), and 24 months (dotted line). The right panel plots representative term structure shapes at different dates.

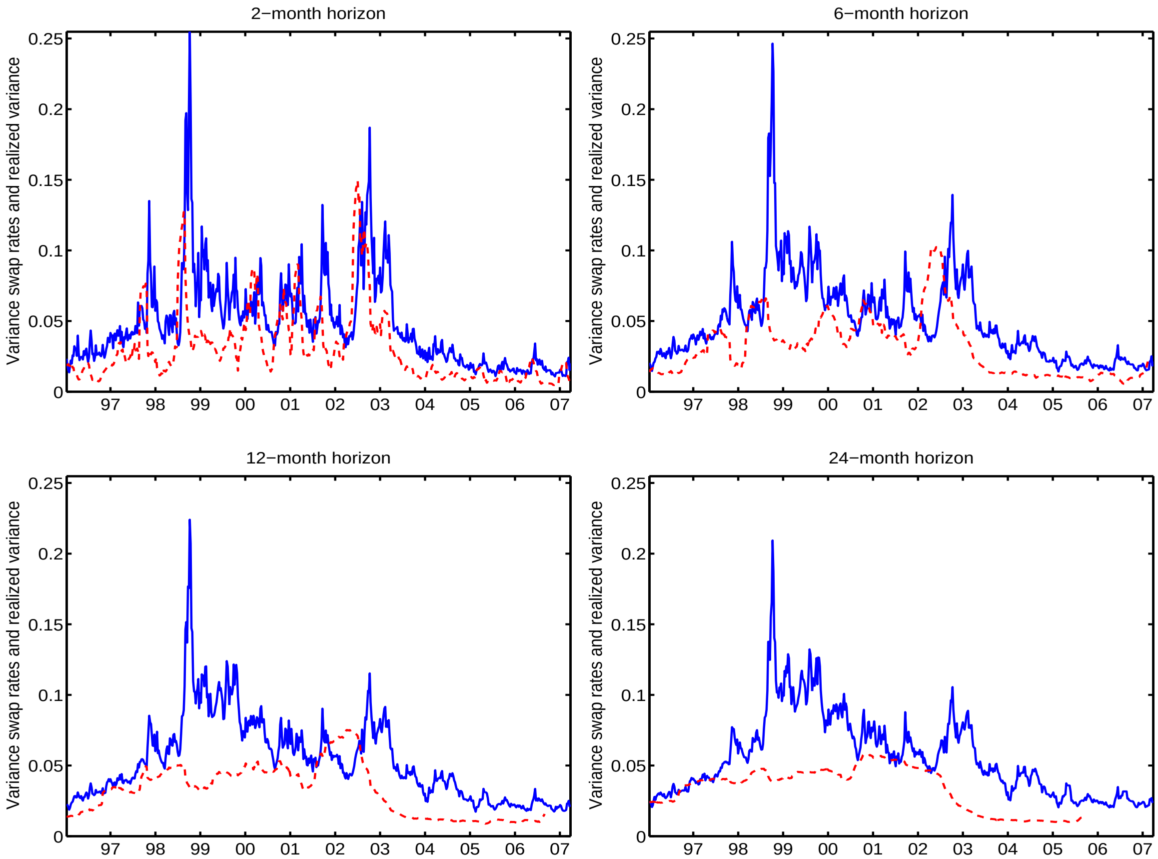

Figure 2. Comparing variance swap rates to the ex post realized variance.
The graph contrasts the variance swap rate time series (solid lines) with the corresponding ex post realized variance (dashed lines) at selected maturities.

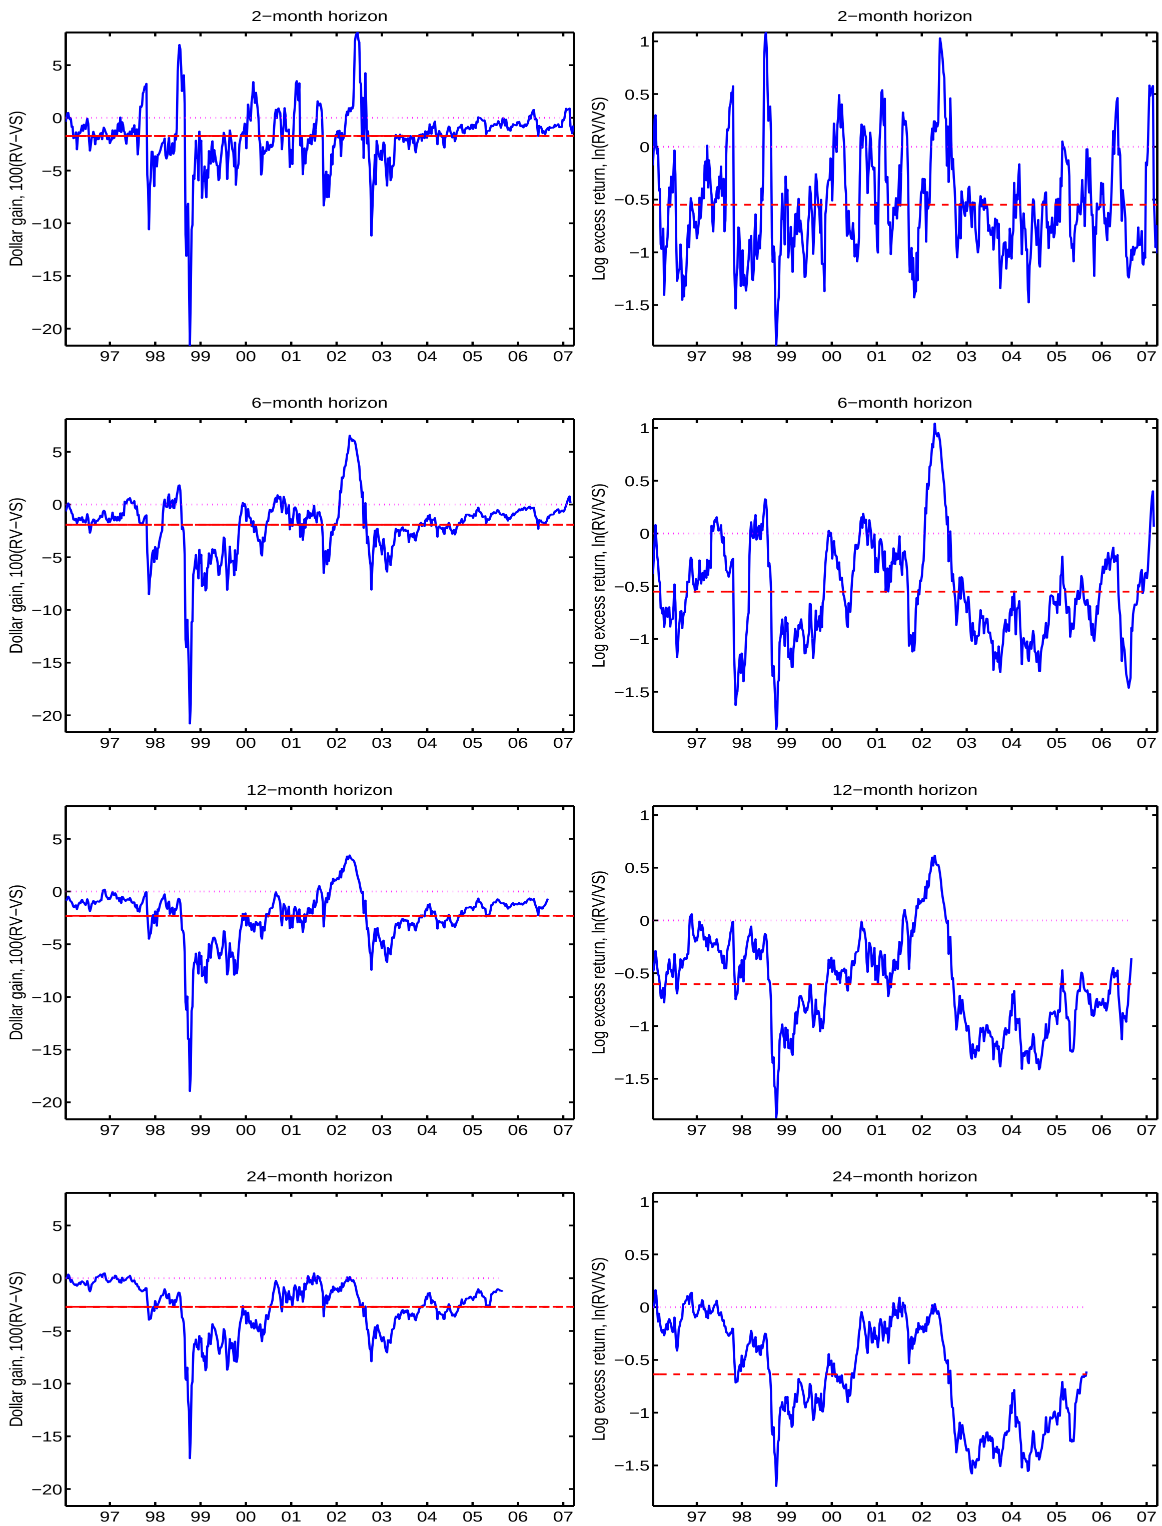

Figure 3. Profit and loss from long variance swaps and holding the contract to maturity.
The solid lines in the left panels plot the time series of the capital gain from long \$100 par variance swap contracts each week and holding them to maturity, measured by $100 \times (RV_{t,T} - VS_{t,T})$. The solid lines in the right panels plot the time series of the log excess return from the investment, $\ln(RV_{t,T}/VS_{t,T})$ by regarding the swap rate as the forward cost. In both panels, the dashed lines represent the sample average and the dotted lines represent the zero-profit benchmark.

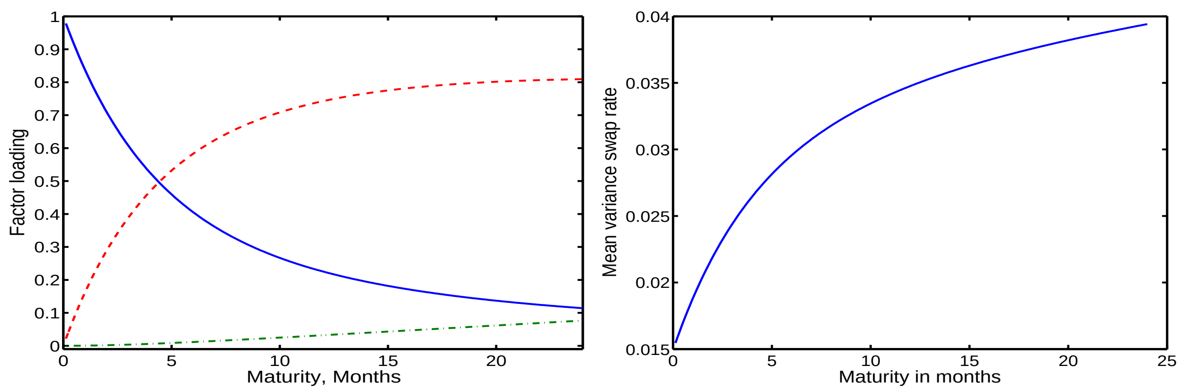

Figure 4. Factor loadings and the mean term structure of variance swap rates.
The left panel plots the contemporaneous response of the variance swap term structure to unit shocks on the variance rate v_t (solid line) and the central tendency factor m_t (dashed line). The dash-dotted line represents the remaining loading on the common unconditional risk-neutral mean of the variance rate and the central tendency θ_m. The right panel plots the mean term structure of the variance swap rate implied by the estimated two-factor variance risk model.

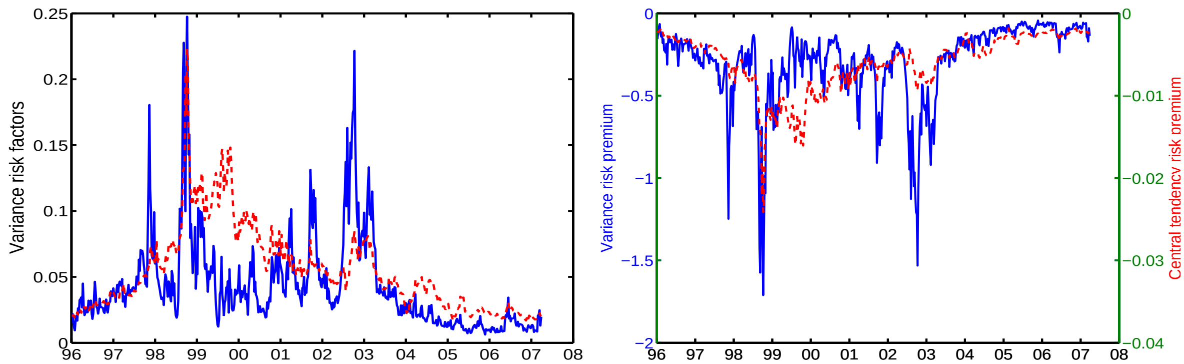

Figure 5. Variance risk dynamics and variance risk premia.
The left panel plots the time series of the instantaneous variance rate v_t (solid line) and the central tendency factor m_t (dashed line) extracted from the two-factor variance risk model using Kalman filter. The right panel plots the instantaneous risk premium on the risk factors. The solid line denotes the risk premium for the instantaneous variance rate ($\gamma^v \sigma_v v_t$) with the scale on the left hand side. The dashed line denotes the instantaneous risk premium for the central tendency factor ($\gamma^m \sigma_m m_t$), with the scale on the right hand side.

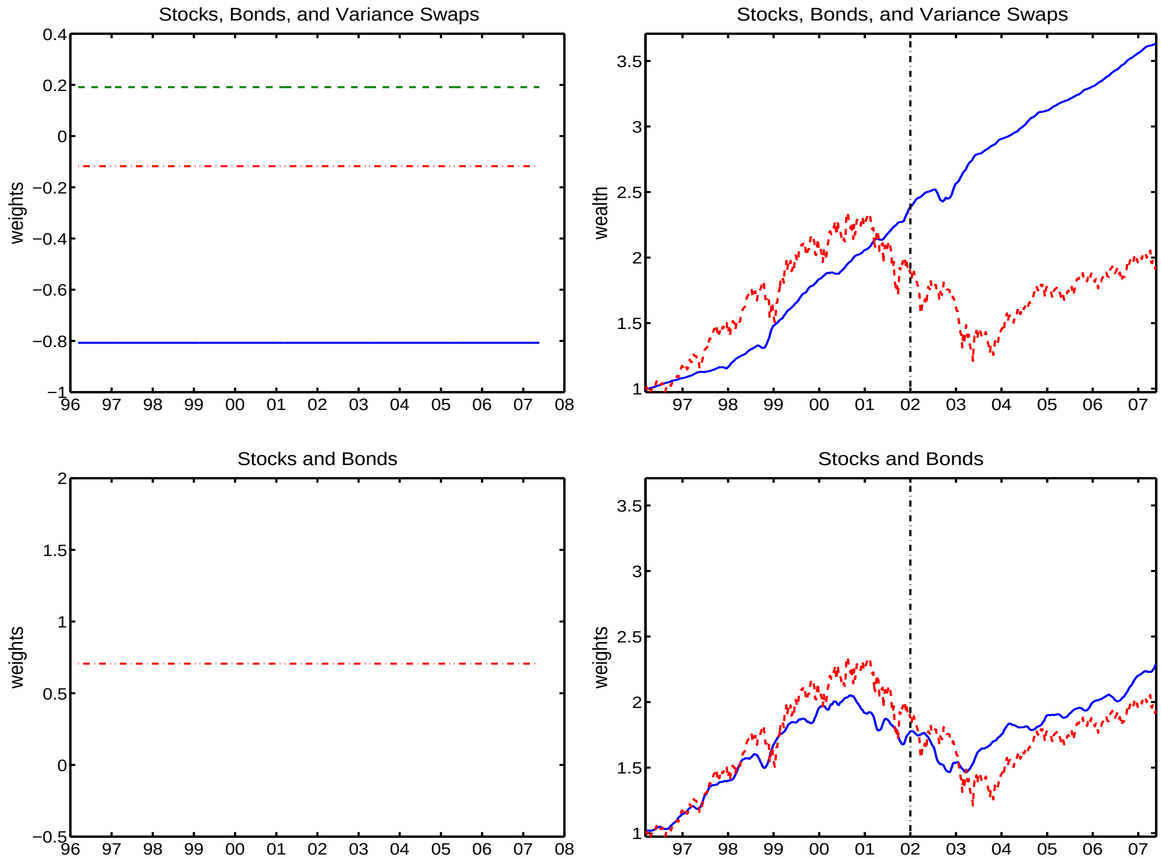

Figure 6. Portfolio fraction and cumulative wealth path of different investment strategies.
The left panels plots the optimal portfolio weights for the stock and variance swaps (upper panel), and the stock only (lower panel). The solid line corresponds to the notional (per dollar wealth) invested in the variance swap with maturity $\tau_1 = 2/12$. The dashed line corresponds to the variance swap with maturity $\tau_2 = 2$. The dash-dotted line corresponds to the stock investment. In the right panels, we plot the cumulative wealth path for the optimal variance swap strategy (upper panel) and the stock-bond strategy (lower panel), together with the S&P 500 index level. The dash-dotted vertical line indicates the beginning of the out-of-sample period.