

Variance Risk Premia*

PETER CARR[†]

Bloomberg L.P. and Courant Institute

LIUREN WU[‡]

Zicklin School of Business, Baruch College

First draft: July 26, 2003

This version: May 27, 2004

Filename: VarianceSwap15.tex

*We thank David Hait and OptionMetrics for providing the option data, Alex Mayus for clarifying trading practices, Rui Yao for help on data processing, and Turan Bali, David Bates, Menachem Brenner, Mikhail Chernov, Robert Engle, Dilip Madan, Benjamin Wurzburger, Jin Zhang, Chu Zhang, and seminar participants at Baruch College, Hong Kong University of Science and Technology, New York University, for comments and discussions. We assume full responsibility for any errors. We welcome comments, including references we have inadvertently missed.

[†]499 Park Avenue, New York, NY 10022; tel: (212) 893-5056; fax: (917) 369-5629; pcarr4@bloomberg.com.

[‡]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>.

Variance Risk Premia

1. Introduction

The grant of the 2003 Nobel prize in economics has made available to the general public the well-documented observation that return variances are random over time. Therefore, when investing in a security such as a stock or a stock portfolio, an investor faces at least two sources of uncertainty, namely the uncertainty about the return as captured by the return variance, and the uncertainty about the return variance itself.

It is important to know how investors deal with the uncertainty in return variance to effectively manage risk and allocate assets, to accurately price and hedge derivative securities, and to understand the behavior of financial asset prices in general. We develop a direct and robust method for quantifying the return variance risk premium on an asset using the market prices of options written on this asset. Our method uses the notion of a variance swap, which is an over-the-counter contract that pays the difference between a standard estimate of the realized variance and the fixed swap rate. Since variance swaps cost zero to enter, the variance swap rate represents the risk-neutral expected value of the realized return variance. We theoretically and numerically show that the variance swap rate can be synthesized accurately by a particular linear combination of option prices. Ignoring the small approximation error, the difference between the ex-post realized variance and this synthetic variance swap rate quantifies the variance risk premium. Using a large options data set, we synthesize variance swap rates and analyze the historical behavior of variance risk premia on five stock indexes and 35 individual stocks.

If variance risk is not priced, the time series average of the realized return variance should equal the variance swap rate. Otherwise, the difference between the expected value of the return variance under the statistical probability measure and the variance swap rate reflects the magnitude of the variance risk premium. Therefore, by comparing the variance swap rate to the ex-post realized return variance, we can empirically investigate the behavior of the variance risk premium.

Widespread appreciation of the significance of variance risk by the practitioner community has recently engendered the introduction of a slew of financial products with payoffs that are directly tied

to estimates of realized variance or volatility. Nowadays, variance and volatility swaps trade actively over the counter on major stocks, stock indexes, and currencies. On September 22, 2003, the Chicago Board Options Exchange (CBOE) redefined its well-known volatility index (VIX) in such a way that it approximates the 30-day variance swap rate of the S&P 500 index return. On March 26, 2004, the CBOE launched a new exchange, the CBOE Futures Exchange (CFE) to start trading futures on VIX. These futures contracts represent a simple way to trade variance realized over a future time period. At the time of this writing, options on the VIX are also planned.

Despite the recent surge in liquidity in volatility contracts, high-quality historical time-series data on variance swap rates are not yet available. In this paper, we circumvent this issue by synthesizing return variance swap rates. Working in complete generality, we show how the payoff of a return variance swap can be accurately approximated theoretically by combining the payoff from a static position in a continuum of European options with a dynamic trading strategy in the underlying futures. We show that a sufficient condition for our replication strategy to be exact is that the underlying asset's return dynamics are continuous over time. It is important to appreciate that no restrictive assumptions are necessary on the dynamics followed by the return variance. In particular, the instantaneous variance rate can jump and it need not even be observable. In this sense, the replicating strategy is robust.

When the underlying asset price can jump, the strategy fails to replicate perfectly. We show that the instantaneous approximation error is third order in the size of the jump. When applying the theoretical relation in practice, we also introduce an approximation error due to the interpolation and extrapolation needed to generate the required continuum of option prices from the finite number of available option quotes. We numerically show that both sources of approximation errors are small under realistic price processes and market settings.

Variance swaps are not the only volatility derivatives that can be robustly replicated. Carr and Lee (2003a) develop robust replicating strategies for any contracts with terminal payoffs that are functions of the realized variance and final price. In particular, they develop the replicating strategy for a volatility swap, the payoff of which is linear in the square root of the realized variance. They argue that the Black and Scholes (1973) at-the-money implied volatility is an accurate approximation of the volatility swap rate. We numerically confirm the accuracy of their theoretical arguments. We conclude that variance

swap rates and volatility swap rates can both be accurately approximated using market prices of options and their underlying assets.

Given these conclusions, we synthesize variance and volatility swap rates using options data on five of the most actively traded stock indexes and 35 of the most actively traded individual stocks during the past seven years. We compare the synthetic variance swap rates to the corresponding realized return variance and investigate the historical behavior of the variance risk premia for different assets. We find that the average risk premia on return variances are strongly negative for the S&P 500 and 100 indexes and for the Dow Jones Industrial Average. The variance risk premia for the Nasdaq 100 index and for most individual stocks are also negative, but with a smaller absolute magnitude. The negative sign on the variance risk premia indicates that variance buyers are willing to suffer a negative average excess return to hedge away upward movements in the index return variance.

We investigate whether the classical Capital Asset Pricing Model (CAPM) can explain the negative variance risk premia. We find that the well-documented negative correlation between index returns and volatility generates a strongly negative beta, but this negative beta can only explain a small portion of the negative variance risk premia. The common risk factors identified by Fama and French (1993) cannot explain the strongly negative variance risk premia, either. Therefore, we conclude that either the market for variance risk is highly inefficient or else the majority of the variance risk is generated by an independent risk factor, which the market prices heavily.

We further analyze the dynamics of the variance risk premia by formulating regressions based on various forms of the expectation hypothesis that assume constant or independent variance risk premia. Under the null hypothesis of constant variance risk premia, a regression of the realized variance on the variance swap rate will result in a slope estimate of one. We find that the sample estimates of the regression slope are positive for all stocks and stock indexes, but are significantly lower than the null value of one for over half of the stocks and stock indexes.

The distributions of the return variance and variance risk premia are highly non-normal. The distribution becomes much closer to normal when we represent the variance in log terms and the variance risk premia in log differences. Under the null hypothesis of constant or independent log variance risk premia, a regression of the log realized variance on the log variance swap rate should result in a slope of

one. We find that this hypothesis is supported by the data. At the 95 percent confidence level, the null hypothesis cannot be rejected for any of the five stock indexes and for 24 of the 35 individual stocks.

Since the floating part of the variance swap payoff is just the square of the floating part of the volatility swap payoff, Jensen's inequality dictates that the variance swap rate is greater than the square of the volatility swap rate. The difference between the variance swap rate and the volatility swap rate squared measures the risk-neutral variance of the return volatility. Using the synthesized variance swap rate and the at-the-money implied volatility, we obtain a time series of the risk-neutral variance of the return volatility for each of the five stock indexes and the 35 stocks. Since variance or volatility risk premia compensate for uncertainty in return volatility, we hypothesize that the variance risk premia become more negative when the variance of the return volatility is high. Regressing the negative of the variance risk premia on the variance of volatility, we obtain positive slope estimates for most of the stock indexes and individual stocks, with more than half of them statistically significant.

Finally, we run an expectation hypothesis regression that uses the log variance and controls for the variation in the variance of volatility. The regression slope estimate on the log variance swap rate is no longer significantly different from its null value of one for all but five of the individual stocks.

In the vast literature on stock market volatility, the papers most germane to our study are the recent works by Bakshi and Kapadia (2003a,b). These studies consider the profit and loss (P&L) arising from delta-hedging a long position in a call option. They persuasively argue that this P&L is approximately neutral to the directional movement of the underlying asset return, but is sensitive to the movement in the return volatility. By analyzing the P&L from these delta-hedged positions, Bakshi and Kapadia are able to infer some useful qualitative properties for the variance risk premia without referring to a specific model. Our approach maintains and enhances the robustness of their approach. In addition, our approach provides a quantitative measure of the variance risk premia. As a result, we can analyze not only the sign, but also the quantitative properties of the premia. The quantification enables us to investigate whether the magnitude of the variance risk premia can be fully accounted for by the classical CAPM or by Fama-French factors, and whether the variance risk premia satisfy various forms of the expectation hypothesis.

Chernov (2003), Eraker (2003), Jones (2003), and Pan (2002) analyze the variance risk premia in conjunction with return risk premia by estimating various parametric option pricing models. Their results and interpretations hinge on the accuracy of the specific models that they use in the analysis. Ang, Hodrick, Xing, and Zhang (2003) form stock portfolios ranked by their sensitivity to volatility risk and analyze the difference among these different portfolios. From the analysis, they can infer indirectly the impact of volatility risk on the expected stock return. Also related is the work by Coval and Shumway (2001), who analyze how expected returns on options investment vary with strike choices and whether the classic capital asset pricing theory can explain the expected option returns.

The underlying premise for studying variance risk premia is that return variance is stochastic. Numerous empirical studies support this premise. Prominent empirical evidence based on the time series of asset returns includes Andersen, Benzoni, and Lund (2002), Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Diebold, and Labys (2001, 2003), Ding, Engle, and Granger (1993), Ding and Granger (1996), and Eraker, Johannes, and Polson (2003). Evidence from the options market includes Bakshi, Cao, and Chen (1997, 2000a,b), Bakshi and Kapadia (2003a,b), Bates (1996, 2000), Carr and Wu (2003), Eraker (2001), Huang and Wu (2004), and Pan (2002).

Our analysis of the variance risk premia is based on our theoretical work on synthesizing a variance swap using European options and futures contracts. Carr and Madan (1998), Demeterfi, Derman, Kamal, and Zou (1999a,b), and Britten-Jones and Neuberger (2000) have used the same replicating strategy, but under the assumption of continuity in the underlying asset price. Our derivation is under the most general setting possible. As a result, our theoretical work quantifies the approximation error induced by jumps. In a recent working paper,

Also relevant is the large strand of literature that investigates the information content of Black-Scholes implied volatilities. Although conclusions from this literature have at times contradicted each other, the present consensus is that the at-the-money Black-Scholes implied volatility is an efficient, although biased, forecast of the subsequent realized volatility. Examples of these studies include Latane and Rendleman (1976), Chiras and Manaster (1978), Day and Lewis (1988), Day and Lewis (1992), Lamoureux and Lastrapes (1993), Canina and Figlewski (1993), Day and Lewis (1994), Jorion (1995), Fleming (1998), Christensen and Prabhala (1998), Gwilym and Buckle (1999), Hol and Koopman

(2000), Blair, Poon, and Taylor, (2000a,b), Hansen (2001), Christensen and Hansen (2002), Tabak, Chang, and de Andrade (2002), Shu and Zhang (2003), and Neely (2003).

The remainder of this paper is organized as follows. Section 2 shows the extent to which the payoff to a variance swap can be theoretically replicated by combining the payoff from a static position in European options with the gains from a dynamic position in futures on the underlying asset. We also discuss the relation between volatility swaps and variance swaps in this section. Section 3 uses three standard models of return dynamics to numerically investigate the magnitude of the approximation error due to price jumps and discrete strikes. Section 4 lays down the theoretical foundation for various expectation hypothesis regressions. Section 5 discusses the data and the methodologies used to synthesize variance and volatility swap rates and to calculate realized variance. Section 6 empirically investigates the behavior of the variance risk premia. Section 7 concludes.

2. Synthesizing a Return Variance Swap

A return variance swap has zero net market value at entry. At maturity, the payoff to the long side of the swap is equal to the difference between the realized variance over the life of the contract and a constant fixed at inception called the variance swap rate. If t denotes the entry time and T denotes the payoff time, the terminal payoff to the long side of the swap at T is:

$$[RV_{t,T} - SW_{t,T}]L, \tag{1}$$

where $RV_{t,T}$ denotes the realized annualized return variance between time t and T , and $SW_{t,T}$ denotes the fixed swap rate, which is determined at time t and is paid at time T . The letter L denotes the notional dollar amount that converts the variance difference into a dollar payoff. Since the contract has zero market value at initiation, no-arbitrage dictates that the variance swap rate equals the risk-neutral expected value of the realized variance,

$$SW_{t,T} = \mathbb{E}_t^{\mathbb{Q}} [RV_{t,T}], \tag{2}$$

where $\mathbb{E}_t^{\mathbb{Q}}[\cdot]$ denotes the expectation operator under some risk-neutral measure \mathbb{Q} and conditional on the information up to time t .

In what follows, we show that under relatively weak assumptions on the price process of the underlying, the risk-neutral expected value of the return quadratic variation from time t to T can be approximated from the time- t prices of out-of-the-money European options maturing at time T . Numerical calculations from realistic price processes and strike spacings indicate that the total approximation error is small. Hence, the risk-neutral expected value at t of the increase in the return quadratic variation over $[t, T]$ can be effectively determined at t from an implied volatility smile of maturity T . Thus, assuming continuous monitoring of the underlying asset's price path, we have effectively determined the fixed rate for a variance swap.

2.1. Synthesizing the return quadratic variation by trading options and futures

It is well known that the geometric mean of a set of positive numbers is never more than the arithmetic mean. Furthermore, the larger the variance of the numbers, the greater is the difference between the arithmetic mean and the geometric mean. This section exploits these observations to extract the risk-neutral expected value of realized variance from option prices.

To fix notation, we let S_t denote the spot price of an asset at time $t \in [0, \mathcal{T}]$, where \mathcal{T} is some arbitrarily distant horizon. We let F_t denote the time- t futures price of maturity $T > t$. For simplicity, we assume that the futures contract marks to market continuously. We also assume that the futures price is always positive, although it can get arbitrarily close to zero. No arbitrage implies that there exists a risk-neutral probability measure \mathbb{Q} defined on a probability space $(\Omega, \mathcal{F}, \mathbb{Q})$ such that the futures price F_t solves the following stochastic differential equation:

$$dF_t = F_{t-} \sigma_{t-} dW_t + \int_{\mathbb{R}^0} F_{t-} (e^x - 1) [\mu(dx, dt) - \mathbf{v}_t(x) dx dt], \quad t \in [0, \mathcal{T}], \quad (3)$$

starting at some fixed and known value $F_0 > 0$. In equation (3), W_t is a \mathbb{Q} standard Brownian motion, \mathbb{R}^0 denotes the real line excluding zero, F_{t-} denotes the futures price at time t just prior to any jump at t , and the random counting measure $\mu(dx, dt)$ realizes to a nonzero value for a given x if and only if the futures price jumps from F_{t-} to $F_t = F_{t-} e^x$ at time t . The process $\{\mathbf{v}_t(x), x \in \mathbb{R}^0, t \in [0, \mathcal{T}]\}$

compensates the jump process $J_t \equiv \int_0^t \int_{\mathbb{R}^0} (e^x - 1) \mu(dx, ds)$, so that the last term in equation (3) is the increment of a \mathbb{Q} -pure jump martingale. This compensating process $v_t(x)$ must satisfy (Prokhorov and Shiryaev (1998)):

$$\int_{\mathbb{R}^0} (|x|^2 \wedge 1) v_t(x) dx < \infty, \quad t \in [0, T].$$

In words, the compensator must integrate the square of the small jumps ($|x| < 1$) to have a well-defined quadratic variation. Furthermore, large jumps ($|x| > 1$) must not be so frequent as to have infinite aggregate arrival rate. Thus, equation (3) models the futures price change as the sum of the increments of two orthogonal martingale components, a purely continuous martingale and a purely discontinuous (jump) martingale. This decomposition is generic for any continuous-time martingale (Jacod and Shiryaev (1987)).

To avoid notational complexity, we assume that the jump component of the returns process exhibits finite variation:

$$\int_{\mathbb{R}^0} (|x| \wedge 1) v_t(x) dx < \infty, \quad t \in [0, T].$$

The time subscripts on σ_{t-} and $v_t(x)$ indicate that both are stochastic and predictable with respect to the filtration \mathcal{F}_t . We further restrict σ_{t-} and $v_t(x)$ so that the futures price F_t is always positive. Finally, we assume deterministic interest rates so that the futures price and the forward price are identical.¹ So long as futures contracts trade, we need no assumptions on dividends.

Under the specification in equation (3), the quadratic variation on the futures return from time t to T is

$$V_{t,T} = \int_t^T \sigma_{s-}^2 ds + \int_t^T \int_{\mathbb{R}^0} x^2 \mu(dx, ds). \quad (4)$$

The annualized quadratic variation is $RV_{t,T} = \frac{1}{T-t} V_{t,T}$. We show that this return quadratic variation can be replicated up to a higher-order error term by a static position in a portfolio of options of the same horizon T and a dynamic position in futures. As futures trading is costless, the risk-neutral expected value of the quadratic variation can be approximated by the forward value of the portfolio of European options. The approximation is exact when the futures price process is purely continuous. When the futures price can jump, the instantaneous approximation error at time t is of order $O\left(\left(\frac{dF_t}{F_{t-}}\right)^3\right)$.

¹We can alternatively assume the weaker condition of zero quadratic covariation between the futures price and the price of a pure discount bond of the same maturity.

Theorem 1 *Under no arbitrage, the time- t risk-neutral expected value of the return quadratic variation of an asset over horizon $T - t$ defined in (4) can be approximated by the continuum of European out-of-the-money option prices across all strikes $K > 0$ and with maturity T :*

$$\mathbb{E}_t^{\mathbb{Q}} [RV_{t,T}] = \frac{2}{T-t} \int_0^{\infty} \frac{\Theta_t(K, T)}{B_t(T)K^2} dK + \varepsilon, \quad (5)$$

where ε denotes the approximation error, $B_t(T)$ denotes the time- t price of a bond paying one dollar at T , and $\Theta_t(K, T)$ denotes the time- t value of an out-of-the-money option with strike price $K > 0$ and maturity $T \geq t$ (a call option when $K > F_t$ and a put option when $K \leq F_t$). The approximation error ε is zero when the futures price process is purely continuous. When the futures price can jump, the approximation error ε is of order $O\left(\left(\frac{dF_t}{F_t}\right)^3\right)$ and is determined by the compensator of the discontinuous component,

$$\varepsilon = \frac{-2}{T-t} \mathbb{E}_t^{\mathbb{Q}} \int_t^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] v_s(x) dx ds. \quad (6)$$

Proof. Let $f(F)$ be a twice differentiable function of F . By Itô's lemma for semi-martingales:

$$\begin{aligned} f(F_T) &= f(F_t) + \int_t^T f'(F_{s-}) dF_s + \frac{1}{2} \int_t^T f''(F_{s-}) \sigma_{s-}^2 ds \\ &\quad + \int_t^T \int_{\mathbb{R}^0} [f(F_{s-}e^x) - f(F_{s-}) - f'(F_{s-})F_{s-}(e^x - 1)] \mu(dx, ds), \end{aligned} \quad (7)$$

Applying equation (7) to the function $f(F) = \ln F$, we have:

$$\ln(F_T) = \ln(F_t) + \int_t^T \frac{1}{F_{s-}} dF_s - \frac{1}{2} \int_t^T \sigma_{s-}^2 ds + \int_t^T \int_{\mathbb{R}^0} [x - e^x + 1] \mu(dx, ds). \quad (8)$$

Adding and subtracting $2\left[\frac{F_T}{F_t} - 1\right] + \int_t^T x^2 \mu(dx, ds)$ and re-arranging, we obtain the following representation for the quadratic variation of returns:

$$\begin{aligned} V_{t,T} &\equiv \int_t^T \sigma_{s-}^2 ds + \int_t^T x^2 \mu(dx, ds) = 2 \left[\frac{F_T}{F_t} - 1 - \ln \left(\frac{F_T}{F_t} \right) \right] + 2 \int_t^T \left[\frac{1}{F_{s-}} - \frac{1}{F_t} \right] dF_s \\ &\quad - 2 \int_t^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] \mu(dx, ds). \end{aligned} \quad (9)$$

A Taylor expansion with remainder of $\ln F_T$ about the point F_t implies:

$$\ln F_T = \ln F_t + \frac{1}{F_t}(F_T - F_t) - \int_0^{F_t} \frac{1}{K^2}(K - F_T)^+ dK - \int_{F_t}^{\infty} \frac{1}{K^2}(F_T - K)^+ dK. \quad (10)$$

Combining equations (9) and (10) and noting that $F_T = S_T$, we have:

$$\begin{aligned} V_{t,T} &= 2 \left[\int_0^{F_t} \frac{1}{K^2}(K - S_T)^+ dK + \int_{F_t}^{\infty} \frac{1}{K^2}(S_T - K)^+ dK \right] \\ &\quad + 2 \int_t^T \left[\frac{1}{F_{s-}} - \frac{1}{F_t} \right] dF_s \\ &\quad - 2 \int_t^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] \mu(dx, ds). \end{aligned} \quad (11)$$

Thus, we can replicate the return quadratic variation up to time T by the sum of (i) the payoff from a static position in $\frac{2dK}{K^2}$ European options on the underlying spot at strike K and maturity T (first line), (ii) the payoff from a dynamic trading strategy holding $2B_s(T) \left[\frac{1}{F_{s-}} - \frac{1}{F_t} \right]$ futures at time s (second line), and (iii) a higher-order error term induced by the discontinuity in the futures price dynamics (third line). The options are all out-of-the money forward, i.e., call options when $F_t > K$ and put options when $K \leq F_t$.

Taking expectations under measure \mathbb{Q} on both sides, we obtain the risk-neutral expected value of the quadratic variation on the left hand side. We also obtain the forward value of the sum of the startup cost of the replicating strategy and the replication error on the right hand side:

$$\mathbb{E}_t^{\mathbb{Q}} [V_{t,T}] = \int_0^{\infty} \frac{2\Theta_t(K, T)}{B_t(T)K^2} dK - 2\mathbb{E}_t^{\mathbb{Q}} \int_t^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] v_s(x) dx ds.$$

By the martingale property, the expected value of the gains from dynamic futures trading is zero under the risk-neutral measure. Dividing by $(T - t)$ on both sides, we obtain the result on the annualized return quadratic variation. ■

Equation (5) forms the theoretical basis for our empirical study. We will numerically illustrate that the approximation error is small. Then we use the first term on the right hand side to determine the synthetic variance swap rate on stocks and stock indexes. The relevant return variance underlying the variance swap is that of the futures, which is equal to that of the forward under our assumption of

deterministic interest rates. Comparing the synthetic variance swap rate to the realized return variance, we will investigate the behavior of the variance risk premia on different stocks and stock indexes.

2.2. Volatility swaps

In many markets especially currencies, an analogous volatility swap contract also exists that pays the difference between the realized volatility and a fixed volatility swap rate,

$$[\sqrt{RV_{i,T}} - VS_{i,T}]L, \quad (12)$$

where $VS_{i,T}$ denotes the fixed volatility swap rate. Since the contract has zero value at inception, no-arbitrage dictates that the volatility swap rate equals the risk-neutral expected value of the square root of the realized variance,

$$VS_{i,T} = \mathbb{E}_t^{\mathbb{Q}} [\sqrt{RV_{i,T}}]. \quad (13)$$

Volatility swaps and variance swaps serve similar purposes in hedging against uncertainty in return volatility. Carr and Lee (2003b) show that there is a robust replicating portfolio for a volatility swap under the sufficient conditions of continuous futures prices and a stochastic volatility process whose coefficients and increments are independent of returns. The replicating portfolio requires dynamic trading in both futures and options, rendering the replication much more difficult in practice than the replication of a variance swap. However, it is actually much easier to robustly approximate the initial price of a volatility swap than a variance swap. Carr and Lee (2003a) show that the volatility swap rate is well approximated by the Black and Scholes (1973) implied volatility for the at-the-money forward ($K = F$) option of the same maturity, $ATMV$,

$$VS_{i,T} \doteq ATMV_{i,T}. \quad (14)$$

This approximation is accurate up to the third order $O(\sigma^3)$ when the underlying futures price is purely continuous and the volatility process is uncorrelated with the return innovation. The at-the-money implied volatility remains a good first-order approximation in the presence of jumps and return-volatility correlations. Appendix A provides more details on the derivation.

Comparing the definitions of the variance swap rate in equation (2) and the volatility swap rate in equation (13), we observe the following relation between the two:

$$\text{Var}_t(\sqrt{RV_{i,T}}) = SW_{i,T} - VS_{i,T}^2, \quad (15)$$

where $\text{Var}_t^{\mathbb{Q}}(\cdot)$ denotes the conditional variance operator under the risk-neutral measure. The standard quotation convention for variance swaps and volatility swaps is to quote both in volatility terms. Using this convention, the variance swap rate should always be higher than the volatility swap rate by virtue of Jensen's inequality. When the variance swap rate and the volatility swap rate are both represented in terms of variance, the difference between the two is just the risk-neutral variance of realized volatility. The two swap rates coincide with each other when return volatility is constant.

Remark 1 *The difference between the variance swap rate and the volatility swap rate squared measures the degree of randomness in return volatility.*

The remark is an important observation. The existence of risk premia for return variance or volatility hinges on the premise that the return variance or volatility is stochastic in the first place. The remark provides a direct measure of the perceived riskiness in return volatility based on observations from the options market. Using the market prices of options of the same maturity but different strikes, we can approximate the variance swap rate according to equation (5). We can also approximate the volatility swap rate using the Black-Scholes implied volatility from the at-the-money option. The difference between the two swap rates reveals the (risk-neutral) variance of the return volatility and hence provides a direct measure of the perceived riskiness in return volatility.

3. Numerical Illustration of Standard Models

The attempted replication of the payoff to a variance swap in equation (5) has an instantaneous error of order $O\left(\left(\frac{dF_t}{F_t}\right)^3\right)$. We refer to this error as *jump error* as it vanishes under continuous path monitoring if there are no jumps. Even if we ignore this jump error, the pricing of a variance swap still requires a continuum of option prices at all strikes. Unfortunately, option price quotes are only available in

practice at a discrete number of strike levels. Clearly, some form of interpolation and extrapolation is necessary to determine the variance swap rate from the available quotes. The interpolation and extrapolation introduce a second source of error, which we term *discretization error*. The discretization error would disappear if option price quotes were available at all strikes.

To gauge the magnitude of these two sources of approximation error, we numerically illustrate three standard option pricing models: (1) the Black-Scholes model (BS), (2) the Merton (1976) jump-diffusion model (MJD), and (3) a combination of the MJD model with Heston (1993) stochastic volatility (MJDSV). The MJDSV model is due to Bates (1996), who estimates it on currency options. Bakshi, Cao, and Chen (1997) estimate the models on S&P 500 index options.

The risk-neutral dynamics of the underlying futures price process under these three models are:

$$\begin{aligned}
\text{BS:} \quad dF_t/F_t &= \sigma dW_t, \\
\text{MJD:} \quad dF_t/F_{t-} &= \sigma dW_t + dJ(\lambda) - \lambda g dt, \\
\text{MJDSV:} \quad dF_t/F_{t-} &= \sqrt{v_t} dW_t + dJ(\lambda) - \lambda g dt,
\end{aligned} \tag{16}$$

where W denotes a standard Brownian motion and $J(\lambda)$ denotes a compound Poisson jump process with constant intensity λ . Conditional on a jump occurring, the MJD model assumes that the size of the jump in the log price is normally distributed with mean μ_j and variance σ_j^2 , with the mean percentage price change induced by a jump given by $g = e^{\mu_j + \frac{1}{2}\sigma_j^2} - 1$. In the MJDSV model, the diffusion variance rate v_t is stochastic and follows a mean-reverting square-root process:

$$dv_t = \kappa(\theta - v_t) dt + \sigma_v \sqrt{v_t} dZ_t, \tag{17}$$

where Z_t is another standard Brownian motion, correlated with W_t by $\mathbb{E}^{\mathbb{Q}}[dZ_t dW_t] = \rho dt$.

The MJDSV model nests the MJD model, which in turn nests the BS model. We regard the progression from BS to MJD and then from MJD to MJDSV as one of increasing complexity. All three models are analytically tractable, allowing us to numerically calculate risk-neutral expected values of variance and volatility, without resorting to Monte Carlo simulation. The difference in the BS model between the synthetic variance swap rate and the constant variance rate are purely due to the discretization error, since there are no jumps. The increase in the error due to the use of the MJD model instead of BS

allows us to numerically gauge the magnitude of the jump error in the presence of discrete strikes. The change in the approximation error from the MJD model to the MJDSV model allows us to numerically gauge the impact of stochastic volatility in the presence of discrete strikes and jumps. In theory, the addition of stochastic diffusion volatility does not increase the approximation error in the presence of a continuum of strikes. However, the reality of discrete strikes forces us to numerically assess the magnitude of the interaction effect.

In the numerical illustrations, we normalize the current futures price to \$100 and assume a constant riskfree rate at $r = 5.6$ percent. We consider the replication of a return variance swap rate over a one-month horizon. The option prices under the Black-Scholes model can be computed analytically. Under the MJD model, they can be computed using a weighted average of the Black-Scholes formula. For the MJDSV model, we rely on the analytical form of the characteristic function of the log return, and compute the option prices based on the fast Fourier inversion method of Carr and Madan (1999). Table 1 summarizes the model parameter values used in the numerical illustrations. These parameters reflect approximately those estimated from S&P 500 index option prices, e.g., in Bakshi, Cao, and Chen (1997).

3.1. Variance swap rate

Under the BS model, the annualized return variance rate is constant at σ^2 . Under the MJD model, this variance rate is also constant at $\sigma^2 + \lambda(\mu_j^2 + \sigma_j^2)$. Under the MJDSV model, the realized return variance rate is stochastic. The risk-neutral expected value of the annualized variance rate, hence the variance swap rate, depends on the current level of the instantaneous variance rate v_t ,

$$\mathbb{E}_t^{\mathbb{Q}} [RV_{t,T}] = \sigma_t^2 + \lambda(\mu_j^2 + \sigma_j^2), \quad (18)$$

where σ_t^2 is given by

$$\sigma_t^2 \equiv \frac{1}{T-t} \mathbb{E}_t^{\mathbb{Q}} \int_t^T v_s ds = \theta + \frac{1 - e^{-\kappa(T-t)}}{\kappa(T-t)} (v_t - \theta). \quad (19)$$

Our replicating strategy implicit in equation (5) is exact when the underlying dynamics are purely continuous, but has a higher order approximation error in the presence of jumps. Thus, under the BS model, the theoretical approximation error is zero: $\varepsilon = 0$. Under the other two jump models MJD and MJDSV, the compound Poisson jump component has the following compensator:

$$v(x) = \lambda \frac{1}{\sqrt{2\pi\sigma_j^2}} e^{-\frac{(x-\mu_j)^2}{2\sigma_j^2}}. \quad (20)$$

We can compute the approximation error ε from equation (6):

$$\varepsilon = 2\lambda (g - \mu_j - \sigma_j^2/2). \quad (21)$$

Thus, the approximation error depends on the jump parameters $(\lambda, \mu_j, \sigma_j)$.

The other obvious source of error is from the interpolation and extrapolation needed to obtain a continuum of option prices from the finite number of available option quotes. To numerically gauge the impact of this discretization error, we assume that we have only five option quotes at strike prices of \$80, \$90, \$100, \$110, and \$120, based on a normalized futures price level of \$100. All the stock indexes and individual stocks in our sample average no less than five strikes at each chosen maturity. Hence, the choice of just five strike prices is conservative.

To gauge the magnitude of the total approximation error, we first compute the option prices under the model parameters in Table 1 and compute the option implied volatility at the five strikes. Then, we linearly interpolate the implied volatility across the five strikes to obtain a finer grid of implied volatilities. For strikes below \$80, we use the implied volatility level at the strike of \$80. Similarly, for strikes above \$120, we use the implied volatility level at the strike of \$120. This interpolation and extrapolation scheme is simple and conservative. There might exist more accurate schemes, but we defer the exploration of such schemes for future research.

With the interpolated and extrapolated implied volatility quotes at all strikes, we apply the Black-Scholes formula to compute the out-of-the-money option prices at each strike level. Then, we approximate the integral in equation (5) with a sum over a fine grid of strikes. We set the lower and upper bounds of the sum at ± 8 standard deviations away from at the money, where the standard deviation

is based on the return variance calculation given in equation (18). The fine grid used to compute the sum employs 2,000 strike points within the above bounds. We perform this analysis based on a one-month horizon ($T - t = 1/12$). Following this numerical approximation procedure, we compute the synthesized annualized variance swap rate over this horizon, $\widehat{SW}_{t,T}$, where the hat stresses the approximations involved. The difference between this approximate variance swap rate \widehat{SW} and the analytically computed annualized variance $\mathbb{E}_t^{\mathbb{Q}}[RV_{t,T}]$ represents the aggregate approximation error.

Table 2 summarizes our numerical results on the approximation error of the variance swap rates under the title “Variance Swap.” Under the BS model, the analytical approximation error is zero. Furthermore, since the implied volatility is constant and equal to σ at all strikes, there is no interpolation or extrapolation error on the implied volatility. The only potential error can come from the numerical integration. Table 2 shows that this error is not distinguishable from zero up to the fourth reported decimal point.

Under the MJD model, the analytical error due to jumps is 0.0021, about 1.51 percent of the total variance (0.1387). The aggregate error via numerical approximation is also 0.0021. Hence again, numerical approximation via five strike levels does not induce noticeable additional errors.

Under the MJDSV model, we consider different instantaneous variance levels, represented as its log difference from the mean, $\ln(v_t/\theta)$. As the current instantaneous variance level v_t varies, the analytical error due to the jump component is fixed at 0.0021, because the arrival rate of the jump component does not change. But as the aggregate variance level varies from 0.0272 to 2.3782, the percentage error due to jumps varies accordingly from 7.72 percent to 0.09 percent. The aggregate numerical error also varies at different volatility levels, but the variation and the magnitude are both fairly small. The interpolation across the five option strikes does not add much additional approximation error, indicating that our simple interpolation and extrapolation strategy works well.

Our numerical results show that the jump error is small under commonly used option pricing models and reasonable parameter values. The additional numerical error due to discretization is also negligible. Hence, we conclude that the synthetic variance swap rate matches closely the analytical risk-neutral expected value of the return variance.

3.2. Volatility swap rate

Under the BS model, the volatility swap rate and the variance swap rate coincide with each other and with the realized return variance σ^2 when they are represented in the same units. Under the MJD model, the return quadratic variation $V_{t,T}$ as defined in (4) is random due to the random arrival of jumps of random size. Under the MJDSV model, $V_{t,T}$ has another source of randomness due to stochastic volatility. The randomness in $V_{t,T}$ under these two models generates a difference between the variance swap rate and the volatility swap rate due to Jensen's inequality, as captured by equation (15).

To compute the analytical volatility swap rate under the MJD and MJDSV models, we use the following mathematical equality for any positive number q ,

$$\sqrt{q} = \frac{1}{2\sqrt{\pi}} \int_0^\infty \frac{1 - e^{-sq}}{s^{3/2}} ds. \quad (22)$$

Appendix B provides the proof for this equality. Then, by replacing q with $V_{t,T}$ and taking expectations on both sides, we can represent the volatility swap rate as a function of the Laplace transform of the quadratic variation,

$$\mathbb{E}_t^{\mathbb{Q}} [\sqrt{V_{t,T}}] = \frac{1}{2\sqrt{\pi}} \int_0^\infty \frac{1 - \mathbb{E}_t^{\mathbb{Q}} [e^{-sV_{t,T}}]}{s^{3/2}} ds. \quad (23)$$

Under the MJD model, this Laplace transform can be represented as an infinite sum:

$$\mathbb{E}_t^{\mathbb{Q}} [e^{-sV_{t,T}}] = e^{-s\sigma^2(T-t)} \sum_{n=0}^{\infty} e^{\lambda(T-t)} \frac{(\lambda(T-t))^n}{n!} (1 + 2\sigma_j^2 s)^{-\frac{n}{2}} e^{-\frac{n\mu_j^2 s}{1+2\sigma_j^2 s}}, \quad (24)$$

where the first term is due to the constant diffusion component. Under the MJDSV model, this first term changes due to stochastic volatility,

$$\mathbb{E}_t^{\mathbb{Q}} [e^{-sV_{t,T}}] = e^{-b(T-t)v_t - c(T-t)} \sum_{n=0}^{\infty} e^{\lambda(T-t)} \frac{(\lambda(T-t))^n}{n!} (1 + 2\sigma_j^2 s)^{-\frac{n}{2}} e^{-\frac{n\mu_j^2 s}{1+2\sigma_j^2 s}}, \quad (25)$$

where

$$\begin{aligned} b(t) &= \frac{2s(1 - e^{-\eta t})}{2\eta - (\eta - \kappa)(1 - e^{-\eta t})}, \\ c(t) &= \frac{\kappa\theta}{\sigma_v^2} \left[2 \ln \left(1 - \frac{\eta - \kappa}{2\eta} (1 - e^{-\eta t}) \right) + (\eta - \kappa)t \right], \end{aligned} \quad (26)$$

and

$$\eta = \sqrt{\kappa^2 + 2\sigma_v^2 s}.$$

Given these two Laplace transforms, we can solve for the volatility swap rate for the two models via numerical integration of equation (23). We use an adaptive Lobatto quadrature method to evaluate this integral numerically with a tolerance level of 10^{-9} . We then compare how the volatility swap rates match the at-the-money implied volatility from each model.

Under the title ‘‘Volatility Swap,’’ Table 2 reports the accuracy of using the at-the-money implied volatility to approximate the volatility swap rate. For ease of comparison to the variance swap rate, we report the squares of the volatility. Under the Black-Scholes model, the volatility swap rate and the implied volatility coincide because $\sigma = 0.37$ is constant.

Under the MJD and MJDSV models, we observe some differences between the at-the-money implied volatility and the analytical volatility swap rate. But in all cases, the differences are fairly small, with the magnitudes similar to the approximation errors for the variance swap rates.

Historically, many studies have used at-the-money implied volatilities as proxies for the true volatility series to study its time series property and forecasting capabilities. Our numerical results, together with the theoretical results in Carr and Lee (2003a), show that these studies have indeed chosen a good proxy. Although it is calculated using the Black-Scholes formula, the at-the-money implied volatility represents an accurate approximation for the risk-neutral expected value of the return volatility under much more general settings.

4. Expectation Hypotheses

If we use \mathbb{P} to denote the statistical probability measure, we can link the variance swap rate and the annualized realized variance as follows:

$$SW_{t,T} = \frac{\mathbb{E}_t^{\mathbb{P}} [M_{t,T} RV_{t,T}]}{\mathbb{E}_t^{\mathbb{P}} [M_{t,T}]} = \mathbb{E}_t^{\mathbb{P}} [m_{t,T} RV_{t,T}], \quad (27)$$

where $M_{t,T}$ denotes a pricing kernel and $m_{t,T}$ represents its normalized version, which is a \mathbb{P} -martingale, $\mathbb{E}_t^{\mathbb{P}}[m_{t,T}] = 1$. Assuming a constant interest rate, we have:

$$\mathbb{E}_t^{\mathbb{P}}[M_{t,T}] = B_t(T) = e^{-r(T-t)}. \quad (28)$$

For traded assets, no-arbitrage guarantees the existence of at least one such pricing kernel (Duffie (1992)).

We decompose equation (27) into two terms:

$$SW_{t,T} = \mathbb{E}_t^{\mathbb{P}}[m_{t,T}RV_{t,T}] = \mathbb{E}_t^{\mathbb{P}}[RV_{t,T}] + \text{Cov}_t^{\mathbb{P}}(m_{t,T}, RV_{t,T}). \quad (29)$$

The first term $\mathbb{E}_t^{\mathbb{P}}[RV_{t,T}]$ represents the time-series conditional mean of the realized variance. The second term captures the conditional covariance between the normalized pricing kernel and the realized variance. The negative of this covariance defines the return variance risk premium.

Dividing both sides of (29) by $SW_{t,T}$, we can also represent the decomposition in excess returns:

$$1 = \mathbb{E}_t^{\mathbb{P}}\left[m_{t,T} \frac{RV_{t,T}}{SW_{t,T}}\right] = \mathbb{E}_t^{\mathbb{P}}\left[\frac{RV_{t,T}}{SW_{t,T}}\right] + \text{Cov}_t^{\mathbb{P}}\left(m_{t,T}, \frac{RV_{t,T}}{SW_{t,T}}\right). \quad (30)$$

If we regard $SW_{t,T}$ as the forward cost of our investment, $(RV_{t,T}/SW_{t,T} - 1)$ captures the excess return from going long the variance swap. The negative of the covariance term in equation (30) represents the variance risk premium in terms of the excess return. Based on the decompositions in equations (29) and (30), we analyze the behavior of the variance risk premia. We also test several forms of the expectation hypothesis on the variance risk premia.

Using the volatility swap rate, we can analogously define the volatility risk premium and analyze its empirical properties. We have done so. The results are qualitatively similar to the results on the variance risk premia. We only report the results on the variance risk premia in this paper to avoid repetition.

4.1. The average magnitude of the variance risk premia

From equation (29), a direct estimate of the average variance risk premium is the sample average of the difference between the variance swap rate and the realized variance, $RP_{i,T} \equiv RV_{i,T} - SW_{i,T}$. This difference also measures the terminal capital gain from going long on a variance swap contract. From equation (30), we can also compute an average risk premia in excess return form by computing the average excess return of a long swap position. To make the distribution closer to normality, we represent the excess return in continuously compounded form and label it as the log variance risk premium, $LRP_i \equiv \ln(RV_{i,T}/SW_{i,T})$.

The most basic form of the expectation hypothesis is to assume zero variance risk premium. Therefore, the null hypothesis is: $RP_{i,T} = 0$ and $LRP_{i,T} = 0$. We empirically investigate whether the average (log) variance risk premium is significantly different from zero.

4.2. Expectation hypothesis on constant variance risk premia

A weaker version of the expectation hypothesis is to assume that the variance risk premium is constant or independent of the variance swap rate. Then, we can run the following regressions to test the hypothesis:

$$RV_{i,T} = a + bSW_{i,T} + e_{i,T}, \quad (31)$$

$$\ln RV_{i,T} = a + b \ln SW_{i,T} + e_{i,T}. \quad (32)$$

The null hypothesis underlying equation (31) is that $RP_{i,T}$ is constant or independent of the variance swap rate. Under this null hypothesis, the slope estimate b should be one. The null hypothesis underlying equation (32) is that the log variance risk premia $LRP_{i,T}$ is constant or independent of the log variance swap rate. The null value of the slope estimate is also one. Under the null hypothesis of zero risk premia, the intercepts of the two regressions should be zero. Therefore, tests of these expectation hypotheses amount to tests of the null hypotheses: $a = 0$ and $b = 1$ for the two regressions.

4.3. Hypothesis on the link between the variance risk premia and variance of volatility

The existence of nonzero variance risk premia hinges on the existence of randomness in volatility. In a world where return variances are constant, no risk and hence no premium would exist on volatility. We hypothesize that the magnitude of the variance risk premium is positively correlated with the magnitude of the uncertainty in the return volatility.

Remark 1 proposes an observable measure for the uncertainty in return volatility. The difference between the variance swap rate and the volatility swap rate squared measures the variance of the return volatility under the risk-neutral measure. Therefore, we can run the following regression:

$$\ln(RV_{t,T}/SW_{t,T}) = a + b(SW_{t,T} - VS_{t,T}^2) + e, \quad (33)$$

and test whether the slope coefficient differs from zero.

4.4. Expectation hypothesis under the Heston model

To illustrate the economic intuition behind the average variance risk premia and the expectation hypothesis regression slope estimates, we go through a simple example based on the stochastic volatility model of Heston (1993). This model assumes that the instantaneous return variance, v_t , follows a square-root process under the risk-neutral measure \mathbb{Q} :

$$dv_t = \kappa(\theta - v_t) dt + \sigma_v \sqrt{v_t} dZ_t, \quad (34)$$

where Z_t denotes a standard Brownian motion, θ is the long-run mean instantaneous variance rate, κ is the mean-reversion speed, and σ_v is a parameter governing the instantaneous volatility of variance.

A common assumption for the square-root model is that the market price of risk due to shocks in the Brownian motion Z is proportional to the diffusion component of the instantaneous variance process:²

$$\gamma(v_t) = \gamma\sigma_v\sqrt{v_t}. \quad (35)$$

In words, a zero cost portfolio with unit exposure to the increment dZ_t would be expected to change in value as compensation for uncertainty in the realization of Z . Under the statistical measure \mathbb{P} , the assumed absolute appreciation rate for this portfolio is $\gamma\sigma_v\sqrt{v_t}$ per unit time, where γ is real and possibly negative.

Under assumption (35), Girsanov's theorem implies that the diffusion of the v_t process remains the same under the statistical measure \mathbb{P} , but the drift of v_t changes to the following,

$$\mu(v_t) = \kappa(\theta - v_t) + \gamma\sigma_v^2 v_t = \kappa^P (\theta^P - v_t), \quad (36)$$

which remains affine in the instantaneous variance rate v_t . The \mathbb{P} -long-run mean and the mean-reversion speed are

$$\theta^P = \frac{\kappa}{\kappa - \gamma\sigma_v^2} \theta, \quad \kappa^P = \kappa - \gamma\sigma_v^2. \quad (37)$$

When the market price of Z risk is positive ($\gamma > 0$), the long-run mean of the variance rate under the statistical measure \mathbb{P} , θ^P , becomes larger than the long-run mean θ under the risk-neutral measure \mathbb{Q} . The mean-reversion speed κ^P under measure \mathbb{P} becomes smaller (slower). The opposite is true when the market price of Z risk is negative.

²Examples of square-root stochastic volatility models with proportional market price of risk include Pan (2002) and Eraker (2003). Many term structure models also assume proportional market price of risk on square-root factors. Examples include Cox, Ingersoll, and Ross (1985), Duffie and Singleton (1997), Roberds and Whiteman (1999), Backus, Foresi, Mozumdar, and Wu (2001), and Dai and Singleton, (2000, 2002).

Assuming the square-root process in (34) and the proportional market price of Z risk in (35), we can derive the conditional expected value of the realized aggregate variance under the two measures:

$$SW_{t,T} \equiv \mathbb{E}_t^{\mathbb{Q}} [RV_{t,T}] = \theta + \frac{1 - e^{-\kappa(T-t)}}{\kappa(T-t)} (v_t - \theta), \quad (38)$$

$$\mathbb{E}_t^{\mathbb{P}} [RV_{t,T}] = \theta^P + \frac{1 - e^{-\kappa^P(T-t)}}{\kappa^P(T-t)} (v_t - \theta^P). \quad (39)$$

Both are affine in the current level of the instantaneous variance rate v_t . Therefore, the conditional variance risk premium as measured by the difference between the two expected values, $RP_t = \mathbb{E}_t^{\mathbb{P}} [RV_{t,T}] - \mathbb{E}_t^{\mathbb{Q}} [RV_{t,T}]$, is also affine in v_t and is hence also given by a stochastic process.

The long-run mean of v_t is θ^P and θ under measures \mathbb{P} and \mathbb{Q} , respectively. The unconditional mean of the variance risk premium under measure \mathbb{P} is equal to:

$$\begin{aligned} \mathbb{E}^{\mathbb{P}} [RP_t] &= \mathbb{E}^{\mathbb{P}} [RV_{t,T} - SW_{t,T}] = \theta^P - \left[\theta + \frac{1 - e^{-\kappa(T-t)}}{\kappa(T-t)} (\theta^P - \theta) \right] \\ &= \left[1 - \frac{1 - e^{-\kappa(T-t)}}{\kappa(T-t)} \right] \frac{\gamma \sigma_v^2}{\kappa - \gamma \sigma_v^2} \theta. \end{aligned} \quad (40)$$

Therefore, the average variance risk premium is positive when the market price of Z risk γ is positive and negative when the market price of Z risk γ is negative. The average risk premium becomes zero when $\gamma = 0$.

Now we consider the expectation hypothesis regression:

$$RV_{t,T} = a + bSW_{t,T} + e. \quad (41)$$

The missing variable in the expectation regression is the variance risk premium, RP_t , which is affine in v_t . Since the swap rate $SW_{t,T}$ is also affine in v_t , the missing risk premium in the regression is correlated with the regressor. Thus, the slope estimate for b will deviate from its null value of one.

From equations (38) and (39), we can derive the population value for the regression slope:

$$b = \frac{\text{Cov}^{\mathbb{P}} (\mathbb{E}_t^{\mathbb{P}} [RV_{t,T}], SW_{t,T})}{\text{Var}^{\mathbb{P}} (SW_{t,T})} = \frac{\kappa (1 - e^{-\kappa^P(T-t)})}{\kappa^P (1 - e^{-\kappa(T-t)})}, \quad (42)$$

where $Var^{\mathbb{P}}(\cdot)$ and $Cov^{\mathbb{P}}(\cdot, \cdot)$ denote variance and covariance under measure \mathbb{P} , respectively. The slope is equal to the null value of one only when $\kappa = \kappa^P$. To see exactly how the slope deviates from the null value, we Taylor expand the two exponential functions up to second order and obtain:

$$b = \frac{\kappa \left(1 - e^{-\kappa^P(T-t)}\right)}{\kappa^P \left(1 - e^{-\kappa(T-t)}\right)} \doteq \frac{1 - \frac{1}{2}\kappa^P(T-t)}{1 - \frac{1}{2}\kappa(T-t)}. \quad (43)$$

Therefore, the slope is less than one when $\kappa^P > \kappa$, or when $\gamma < 0$. The slope is greater than one when $\kappa^P < \kappa$, or $\gamma > 0$.

The relation becomes complicated when the regression is on log variance. Taylor expanding the logarithms of $SW_{t,T}$ and $\mathbb{E}_t^{\mathbb{P}}[RV_{t,T}]$ around their respective long-run means generates the following first-order approximations:

$$\ln SW_{t,T} \doteq \ln \theta + \frac{1 - e^{-\kappa(T-t)}}{\theta \kappa (T-t)} (v_t - \theta), \quad (44)$$

$$\ln \mathbb{E}_t^{\mathbb{P}}[RV_{t,T}] \doteq \ln \theta^P + \frac{1 - e^{-\kappa^P(T-t)}}{\theta^P \kappa^P (T-t)} (v_t - \theta^P). \quad (45)$$

The regression slope on the log variances is approximately,

$$b = \frac{Cov^{\mathbb{P}}(\ln \mathbb{E}_t^{\mathbb{P}}[RV_{t,T}], \ln SW_{t,T})}{Var^{\mathbb{P}}(\ln SW_{t,T})} \doteq \frac{\theta \left(1 - \frac{1}{2}\kappa^P(T-t)\right)}{\theta^P \left(1 - \frac{1}{2}\kappa(T-t)\right)}. \quad (46)$$

Whether this slope is greater or less than the null value of one becomes ambiguous. For example, when $\gamma > 0$, we have $\theta < \theta^P$, but $\left(1 - \frac{1}{2}\kappa^P(T-t)\right) > \left(1 - \frac{1}{2}\kappa(T-t)\right)$. The two conflicting impacts generate ambiguous regression slopes that will depend on the exact value of the model parameters.

Finally, under the Heston model with proportional market price of Z risk, the variance risk premium is proportional to the instantaneous variance rate. Therefore, any other variable that is related (and ideally proportional) to the instantaneous variance rate would also have explanatory power for the risk premium. Equation (33) proposes to use the risk-neutral variance of return volatility, $Var_t^{\mathbb{Q}}(\sqrt{RV_{t,T}})$ as the explanatory variable. Under the Heston model and the proportional market price of risk assumption, this conditional variance of volatility is indeed related to v_t , but in a complicated nonlinear way. Thus,

we expect the variable to have some explanatory power for the variance risk premium at least under the Heston example.

5. Data and Methodologies

Our options data are from OptionMetrics, a financial research and consulting firm specializing in econometric analysis of the options markets. The “Ivy DB” data set from OptionMetrics is the first widely-available, up-to-date, and comprehensive source of high-quality historical price and implied volatility data for the U.S. stock and stock index options markets. The Ivy DB database contains accurate historical prices of options and their associated underlying instruments, correctly calculated implied volatilities and option sensitivities, based on closing quotes at the Chicago Board of Options Exchange (CBOE). Our data sample starts from January 1996 and ends in February 2003.

From the data set, we filter out market prices of options on five stock indexes and 35 individual stocks. We choose these stocks and stock indexes mainly based on the quote availability, which approximates the stocks’ trading activity. Table 3 provides the list of the five stock indexes and 35 individual stocks in our sample, as well as the starting and ending dates, the sample length (N), and the average number of strikes (NK) at the chosen maturities for each stock (index). The list includes options on the S&P 500 index (SPX), the S&P 100 index (OEX), the Dow Jones Industrial Index (DJX), and the Nasdaq-100 index (NDX). The index options on SPX, DJX, and NDX are European options on the spot indexes. The OEX options and options on the other 35 individual stocks and the QQQ (the Nasdaq-100 tracking stock) are all American options on the underlying spot.

Index options are more active than the individual stock options. On average, more than 20 strikes are available at the chosen maturity for the S&P index options, but the number of available strikes at the chosen maturity for individual stock options is mostly within single digits. Therefore, inferences drawn from the index options data could be more accurate than those drawn from the individual stock options.

The data set includes closing quotes for each option contract (bid and ask) along with Black-Scholes implied volatilities based on the mid quote. For the European options, implied volatilities are directly

inferred from the Black-Scholes option pricing formula. For the American options, OptionMetrics employs a binomial tree approach that takes account of the early exercise premium. The data set also includes the interest rate curve and the projected dividend yield.

In parallel with our numerical studies in the previous section, we choose a monthly horizon for the synthesis of variance swap rates. At each date for each stock or stock index, we choose to the two nearest maturities, except when the shortest maturity is within eight days, under which scenario we switch the next two maturities to avoid the potential microstructure effects of the very short-dated options. We only retain options that have strictly positive bid quotes and where the bid price is strictly smaller than the ask price.

Analogous to the numerical illustrations, at each maturity, we first linearly interpolate implied volatilities at different moneyness levels, defined as $k \equiv \ln(K/F)$, to obtain a fine grid of implied volatilities. For moneyness levels k below the lowest available moneyness level in the market, we use the implied volatility at the lowest strike price. For k above the highest available moneyness, we use the implied volatility at the highest strike. Using this interpolation and extrapolation procedure, we generate a fine grid of 2,000 implied volatility points with a strike range of ± 8 standard deviations from at-the-money. The standard deviation is approximated by the average implied volatility.

Given the fine grid of implied volatility quotes, IV , we compute the forward price of a European option of strike K and maturity T using the Black (1976) formula,

$$\frac{\Theta_t(K, T)}{B_t(T)} = \begin{cases} F_t N(d_1) - KN(d_2) & K > F_t \\ -F_t N(-d_1) + KN(-d_2) & K \leq F_t \end{cases}, \quad (47)$$

with

$$d_1 = \frac{\ln(F_t/K) + IV^2(T-t)/2}{IV\sqrt{T-t}}, \quad d_2 = d_1 - IV\sqrt{T-t}. \quad (48)$$

We can rewrite the initial cost of the approximate replicating portfolio in equation (5) as

$$\mathbb{E}_t^{\mathbb{Q}}[RV_{t,T}] \doteq \frac{2}{T-t} \left[\int_{-\infty}^0 \left(-e^{-k} N(-d_1(k)) + N(-d_2(k)) \right) dk + \int_0^{\infty} \left(e^{-k} N(d_1(k)) - N(-d_2(k)) \right) dk \right], \quad (49)$$

with

$$d_1(k) = \frac{-k + IV^2(k)(T-t)/2}{IV(k)\sqrt{T-t}}, \quad d_2(k) = d_1(k) - IV(k)\sqrt{T-t}. \quad (50)$$

Therefore, the value of this portfolio does not depend directly on the spot or forward price of the underlying, but only on the moneyness level k and the implied volatility at each moneyness level k .

Based on the implied volatilities at the two nearest maturities that are no shorter than eight days, we compute the synthetic variance swap rates at these two maturities. Then, we linearly interpolate to obtain the variance swap rate at a 30-day horizon. We also linearly interpolate to obtain the at-the-money implied volatility over a 30-day horizon as an approximation for the volatility swap rate. We do not extrapolate. When the shortest maturity is over 30 days, we use the variance swap rate and at-the-money implied volatility at the shortest maturity.

At each day, we also compute the relevant forward price F of each stock based on the interest rates, dividend yields, and the spot price level. Then, we match the variance swap rate with an ex-post annualized realized variance estimate over the next 30 calendar days,

$$RV_{t,t+30} = \frac{365}{30} \sum_{i=1}^{30} \left(\frac{F_{t+i,t+30} - F_{t+i-1,t+30}}{F_{t+i-1,t+30}} \right)^2, \quad (51)$$

where $F_{t,T}$ denotes the time- t forward price with expiry at time T . The estimation of the ex-post realized variance defined in equation (51) is similar to the way that the floating component of the payoff to a variance swap contract is calculated in practice. A small difference exists between the return variance defined in equation (51) and the quadratic variation in (4) due to the difference between daily monitoring and continuous monitoring. The forward price has a fixed maturity date and hence a shrinking time-to-maturity as calendar time rolls forward. Since the stock prices in the OptionMetrics data set are not adjusted for stock splits, we manually adjust the stock splits for each stock in calculating the realized variance. We have also downloaded stock prices from Bloomberg to check for robustness. Furthermore, we have also computed alternative realized variances based on spot prices, and based on demeaned returns. These variations in the definition of the realized variance do not alter our conclusions. We report our results based on the realized variance definition in equation (51).

At each day, we have computed a 30-day variance swap rate, a 30-day volatility swap rate, and a 30-day ex-post realized variance (the realized variance from that day to 30 days later). In our analysis,

we apply the following filters to delete inactive days that occur mainly for individual stock options: (1) The nearest available maturity must be within 90 days. (2) The actual stock price level must be greater than one dollar. (3) The number of strikes is at least three at each of the two nearest maturities. For a stock with active options trading, the most active options are usually the ones that mature in the current or next month. Hence, an absence of quotes for these short-term options is an indication of inactivity. Furthermore, since a stock will be delisted from the stock exchange if the stock price stays below one dollar for a period of time, options trading on such penny stocks are normally very inactive. The last filter on the number of strikes at each maturity is needed to accurately estimate the variance swap rate. None of these filters are binding for the S&P 500 and 100 index options.

Table 4 reports the summary statistics for the realized variance (RV), the synthetic variance swap rate (SW), and the synthetic volatility swap rate (VS). For ease of comparison, we represent all three series in percentage volatility units. Of the three series, the average value of the realized variance is the lowest, and the variance swap rate is the highest, with the volatility swap rate in the middle. All three rates exhibit positive skewness and positive excess kurtosis for most stocks and stock indexes.

6. The Behavior of Variance Risk Premia

In this section, we empirically investigate the behavior of the variance risk premia. First, we establish the existence, sign, and average magnitude of the variance risk premia. Then, we investigate whether the classical capital asset pricing theory (CAPM) and Fama-French market factors can fully account for the premia. Finally, we analyze the dynamic properties of the risk premia using the various expectation hypotheses formulated in Section 4.

6.1. Do investors price variance risk?

If investors price the variance risk, we expect to see a difference between the sample averages of the realized variance and the variance swap rate. Table 5 reports the summary statistics of the difference between the realized variance and the variance swap rate, $RP = 100 \times (RV_{t,T} - SW_{t,T})$, in the left panel and the log difference $LRP = \ln(RV_{t,T}/SW_{t,T})$ in the right panel. We label RP as the variance risk

premia and LRP the log variance risk premia. The variance risk premia RP show large kurtosis and sometimes also large skewness. The skewness and kurtosis are much smaller for the log variance risk premia LRP .

The mean (log) variance risk premia are negative for all of the stock indexes and for most of the individual stocks. To test its statistical significance, we construct a t -statistic for the risk premia,

$$t\text{-stat} = \sqrt{N}\mu_j/\sigma_j, \quad j = RP, LRP, \quad (52)$$

where N denotes the sample length, μ denotes the sample average, and σ denotes the Newey and West (1987) serial-dependence adjusted standard error, computed with a lag of 30 days. We report the estimated t -values in Table 5. The largest t -statistics come from the S&P 500 and S&P 100 indexes and the Dow Jones Industrial Average, which are strongly significant for both variance risk premia and log variance risk premia. The Nasdaq-100 index and its tracking stock generate t -statistics that are much lower. The t -statistics on the two Nasdaq indexes are not statistically significant for the variance risk premia RP , albeit significant for the log variance risk premia LRP .

The t -statistics on the log variance risk premia are also negative for most of the individual stocks, but the magnitudes are smaller than that for the S&P indexes. The mean log variance risk premia are significantly negative for 21 of the 35 individual stocks. However, the mean variance risk premia (RP) are insignificant for all but three of the 35 individual stocks.

If an investor creates the fixed part of the variance swap payoff by purchasing at time t the proper portfolio of options with expiry date T and then dynamically trading futures, the initial cost of this trading strategy is given by $B_t(T)SW_t$ and the terminal payoff of this strategy at time T is the realized variance $RV_{t,T}$. Therefore, the log risk premium $LRP = \ln(RV_{t,T}/SW_{t,T})$ captures the continuously compounded excess return to such a trading strategy. The mean values of LRP in Table 5 show that on average, the investors are willing to accept a negative excess return for this investment strategy, especially on the S&P and Dow indexes. This excess return is over -50 percent per month for the two S&P 500 indexes and for Dow Jones. Therefore, we conclude that investors price heavily the uncertainty in the variance of the S&P and Dow indexes.

However, the average variance risk premia on the Nasdaq-100 index and the individual stocks are much smaller. The average capital gains from going long the variance swap contract (RP) are mostly insignificant for Nasdaq-100 index and individual stocks. Thus, we conjecture that the market does not price all return variance variation in each single stock, but only prices the variance risk in the stock market portfolio. Based on this hypothesis, the average variance risk premium on each stock is not proportional to the total variation of the return variance, but proportional to the covariation of the return variance with the market portfolio return variance. To test this hypothesis, we use the realized variance on S&P 500 index return as the market portfolio variance, and estimate the “variance beta” as

$$\beta_j^V = Cov(RV_j, RV_{SPX}) / Var(RV_{SPX}), j = 1, \dots, 40, \quad (53)$$

where the variance and covariance are measured using the common sample of the two realized variance series. Then, we expect the average variance risk premium on each stock j (\overline{LRP}_j) is positive related to its variance beta. The regression estimates are as follows,

$$\overline{LRP}_j = 0.0201 + 0.2675 \beta_j^V + e, \quad R^2 = 15.9\%, \quad (54)$$

(0.34) (2.72)

with t -statistics reported in the parentheses below the estimates. The slope estimate is statistically significant and positive at 95 confidence level. Here, we estimate both the variance risk premia and the variance beta using log variance. Figure 1 plots the scatter plot of this regression, from which we also observe an apparent positive relation. Thus, the market charges premium not on the total variance risk for each stock, but on its covariance with a common variance risk factor.

Given the large magnitudes of the variance risk premia on S&P and Dow indexes, it is natural to investigate whether shorting variance swaps on these indexes constitutes an attractive investment strategy. To answer this question, we measure the annualized information ratio for a short position in a variance swap. Figure 2 plots the information ratio estimates. The left panel plots the raw information ratio, defined as the mean excess log return over its standard deviation, scaled by $\sqrt{12}$ for annualization. The standard deviation is the simple sample estimate on the overlapping daily data. In the right panel, we adjust the standard deviation calculation for serial dependence following Newey and West (1987) with 30 lags.

By going short the variance swap contracts on the S&P and Dow indexes, we obtain very high raw information ratios (over three). After adjusting for serial dependence, the Sharpe ratios are still higher than an average stock portfolio investment. Nevertheless, given the nonlinear payoff structure, caution should be applied when interpreting Sharpe ratios on derivative trading strategies (Goetzmann, Ingersoll Jr., Spiegel, and Welch (2002)).

Overall, we find that the market prices heavily the uncertainties in the return variance of the S&P and Dow indexes. The variance risk premia on the Nasdaq index and on individual stocks are smaller. The negative sign of the variance risk premia implies that investors are willing to pay a premium, or receive a return lower than the riskfree rate, to hedge away upward movements in the return variance of the stock indexes. In other words, investors regard market volatility increases as extremely unfavorable shocks to the investment opportunity and demand a heavy premium for bearing such shocks.

6.2. Can we explain the variance risk premia with classical risk factors?

The variance risk premia are strongly negative for S&P and Dow indexes. The classical capital asset pricing theory (CAPM) argues that the expected excess return on an asset is proportional to the beta of the asset, or the covariance of the asset return with the market portfolio return. Qualitatively, the negative excess return on the variance swap contract on the stock indexes is consistent with the CAPM, given the well-documented negative correlation between the index returns and index volatility.³ If investors go long stocks on average and if realized variance is negatively correlated with index returns, the payoff to the long side of a variance swap is attractive as it acts as insurance against an index decline. Therefore, investors are willing to receive a negative excess return for this insurance property.

Can this negative correlation fully account for the negative variance risk premia? To answer this question, we estimate the following regressions,

$$\ln RV_{t,T}/SW_{t,T} = \alpha + \beta_j ER_{t,T}^m + e, \quad (55)$$

³Black (1976) first documented this phenomenon and attributed it to the “leverage effect.” Various other explanations have also been proposed in the literature, e.g., Haugen, Talmor, and Torous (1991), Campbell and Hentschel (1992), Campbell and Kyle (1993), and Bekaert and Wu (2000).

for the five stock indexes and 35 individual stocks. In equation (55), ER^m denotes the excess return on the market portfolio. Given the negative correlation between the index return and index return volatility, we expect that the beta estimates are negative for at least the stock indexes. Furthermore, if the CAPM fully accounts for the variance risk premia, the intercept of the regression α should be zero. This intercept represents the average excess return of a market-neutral investment strategy that goes long one unit of the variance swap and short β units of the market portfolio. Under CAPM, all market-neutral investment strategies should generate zero expected excess returns.

To estimate the relation in equation (55), we consider two proxies for the excess return to the market portfolio. First, we use the S&P 500 index to proxy for the market portfolio and compute the excess return based on the forward price on the index,

$$ER_{t,T}^m = \ln S_T^m / F_{t,T}^m. \quad (56)$$

Since we have already constructed the forward price on S&P 500 index when we construct the time series on the realized variance, we can readily obtain a daily series of the excess returns (ER^m) that match the variance data series.

Our second proxy is the value-weighted return on all NYSE, AMEX, and NASDAQ stocks (from CRSP) minus the one-month Treasury bill rate (from Ibbotson Associates). This excess return is publicly available at Kenneth French's data library on the web.⁴ The data are monthly. The sample period that matches our options data is from January 1996 to December 2002.

We estimate the regressions using the generalized methods of moments (GMM), with the weighting matrix computed according to Newey and West (1987) with 30 lags for the overlapping daily series and six lags for the non-overlapping monthly series.

Table 6 reports the estimates (and t -statistics in parentheses) on the CAPM relation. The results using the daily series on S&P 500 index and the monthly series on the valued-weighted market portfolio are similar. The β estimates are strongly negative for all the stock indexes and most of the individual stocks. The β estimates are the most negative for S&P and the Dow indexes. These negative estimates

⁴The web address is: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

are consistent with the vast empirical literature that documents a negative correlation between stock index returns and return volatility. The negative beta estimates are also consistent with the average negative variance risk premia observed the most strongly on S&P and Dow indexes.

Nevertheless, the intercept α estimates remain strongly negative, especially for the S&P and Dow indexes, implying that the negative beta cannot fully account for the observed negative variance risk premia. Indeed, the estimates for α are not much smaller than the mean variance risk premia reported in Table 5, indicating that the β risk does not tell the full story of the variance risk premia. The results call for additional risk factors.

Fama and French (1993) identify two additional risk factors in the stock market that are related to the firm size (*SMB*) and book-to-market value (*HML*), respectively. We investigate whether these additional common risk factors help explain the variance risk premia. We estimate the following relations on the five stock indexes and 35 individual stocks,

$$\ln RV_{t,T}/SW_{t,T} = \alpha + \beta ER_{t,T}^m + sSMB_{t,T} + hHML_{t,T} + e. \quad (57)$$

Data on all three risk factors are available on Kenneth French's data library. We refer the interested readers to Fama and French (1993) for details on the definition and construction of these common risk factors. The sample period that overlaps with our options data is monthly from January 1996 to December 2002. Again, ER^m denotes the excess return to the market portfolio. Furthermore, both *SMB* and *HML* are in terms of excess returns on zero-cost portfolios. Therefore, the intercept α represents the expected excess return on an investment that goes long one unit of the variance swap contract, short β of the market portfolio, s of the size portfolio, and h of the book-to-market portfolio. This investment strategy is neutral to all three common risk factors.

We use GMM to estimate the relation in (57), with the weighting matrix constructed following Newey and West (1987) with six lags. Table 7 reports the parameter estimates and t -statistics. The intercept estimates for the indexes remain strongly negative, the magnitudes only slightly smaller than the average variance risk premia reported in Table 5. Therefore, the Fama-French risk factors can only explain a small portion of the variance risk premia.

In the joint regression, both the market portfolio ER^m and the size portfolio SMB generate significantly negative loadings, indicating that the return variance is not only negatively correlated with the index returns but also negatively correlated with the SMB factor. Hence, going long the variance swap contract also serves as an insurance against the SMB factor going up. The loading estimates on the HML factor are mostly insignificant.

Fama and French (1993) also consider two bond-market factors, related to maturity ($TERM$) and default (DEF) risks. Furthermore, Jegadeesh and Titman (1993) identify a momentum phenomenon that past winner often continue to outperform past losers. Later studies, e.g., Rouwenhorst (1998, 1999) and Jegadeesh and Titman (2001), have confirmed the robustness of the results. We construct the $TERM$ and DEF factors using Treasury and corporate yield data from the Federal Reserve Statistical Release. Kenneth French's data library also provides a momentum factor (UMD) similar to that from Carhart (1997). However, single-factor marginal regressions on these three factors show that none of these three factors have a significant loading on the variance risk premia. Therefore, they cannot explain the variance risk premia, either.

The bottom line story here is that neither the original capital asset pricing model nor the Fama-French factors can fully account for the negative variance risk premia on the stock indexes. Therefore, either there exist a large inefficiency in the market for index variance or else the majority of the variance risk is generated by an independent risk factor that the market prices heavily. Investors are willing to receive a negative excess return to hedge against market volatility going up, not only because market volatility movement is negatively correlated with stock market portfolio return, but also because investors regard market volatility hikes by themselves as unfavorable shocks and demand a high compensation for bearing such shocks.

We leave the study of economic foundations for the negative variance risk premia for future research. Here, we propose several potential reasons for the negative premia. We consider the holding of the market portfolio of stocks. With the same expected return, the increase in return variance implies an decline in performance in terms of the information ratio. Hence, one way to guarantee a minimum performance is to buy options to hedge against return variance increases. The fact that shorting the variance swap contract generates high information ratios indicates that the high negative premia are not justified based purely on the information ratio measure. Nevertheless, going long the variance swap

contrast is also an effective strategy to hedge against the risks associated with the random arrival of discontinuous price movements. These risks are not well measured by the information ratio. Furthermore, considerations on meeting value-at-risk requirements and preventing shortfalls and draw-downs also make long variance swap an attractive strategy that could generate negative variance risk premia.

6.3. Are variance risk premia constant?

To understand the dynamic behavior of the variance risk premia, we run the following two expectation-hypothesis regressions,

$$RV_{t,T} = a + bSW_{t,T} + e, \quad (58)$$

$$\ln RV_{t,T} = a + b \ln SW_{t,T} + e. \quad (59)$$

Under the null hypothesis of constant variance risk premia, the slope estimate for equation (58) should be one. Under the null hypothesis of constant log variance risk premia, the slope estimate for equation (59) should be one. We estimate the regressions using the generalized method of moments (GMM), with the weighting matrix computed according to Newey and West (1987) with 30 lags to account for the overlapping sample.

Table 8 reports the estimates and t -statistics under the null hypothesis of $a = 0, b = 1$. The columns on the left side summarize the estimation results on equation (58). All of the slope estimates are positive, but many of them are lower than one. The t -statistics show that over half of the stock indexes and individual stocks generate regression slopes that are significantly lower than the null value of one. Our previous analysis shows that under the Heston (1993) stochastic volatility model, these lower-than-one slope estimates lend support to negative market price of variance risk (equation (43)).

The columns on the right side of Table 8 report the estimation results on equation (59) based on log variances. For all the stock indexes and 24 of the 35 individual stocks, the slope estimates are no longer significantly different from one at the 95 confidence level.

The difference between the slope estimates of the two regressions indicates that the risk premia defined in log returns is closer to a constant or independent series than the risk premia defined in level

differences. The Heston (1993) stochastic volatility model with a negatively proportional market price of risk can qualitatively match the results on these two regressions.

6.4. Do variance risk premia increase with variance risk?

Variance risk premia arise as compensation for bearing the uncertainty in return variance. When going shorting a variance swap contract, the investor receives a positive average premium as a compensation for bearing the risk of facing market volatility going up.

We hypothesize that the absolute magnitude of the variance risk premia increases with the riskiness in the return volatility. To test this hypothesis, we use GMM to estimate the following relations on the five stock indexes and 35 stocks,

$$\ln(SW_{i,T}/RV_{i,T}) = a + b(SW_{i,T} - VS_{i,T}^2) + e, \quad (60)$$

$$\ln(SW_{i,T}/RV_{i,T}) = a + b \ln(SW_{i,T}/VS_{i,T}^2) + e. \quad (61)$$

The left hand side of the two equations represents the log excess returns to the investor who goes short a variance swap. The right hand side measures the difference between the variance swap rate and the volatility swap rate, which captures the variance of return volatility. We use this difference as a measure of the riskiness in return volatility. In equation (61), we use the log difference to replace the level difference to obtain better distributional properties.

Table 9 reports the estimation results. Consistent with our hypothesis, the slope estimates are predominantly positive for most stocks and stock indexes. The t -statistics of the slope estimates show that the S&P and Dow indexes, as well as many individual stocks, generate slopes that are significantly positive. Therefore, the absolute magnitude of the variance risk premia increases with the riskiness in return volatility, as measured by the risk-neutral conditional variance of the return volatility.

Given the explanatory power of the variance of volatility on the variance risk premia, we further hypothesize that if we control for the variance of volatility in the expectation-hypothesis regression as follows,

$$\ln RV_{i,T} = a + b \ln SW_{i,T} + c \ln(SW_{i,T}/VS_{i,T}^2) + e, \quad (62)$$

the slope estimate for b would become closer to its null hypothesis value of one. We choose to use the log difference $\ln(SW_{t,T}/VS_{t,T}^2)$ instead of the level difference for better distributional properties.

Table 10 reports the GMM estimation results. Consistent with our hypothesis, the estimates for b become closer to one than in the case without controlling for variance of volatility. The t -statistics suggest that for all the stock indexes and all but five of the individual stocks, the estimates for b are not significantly different from its null value of one.

7. Conclusion

In this paper, we propose a direct and robust method to quantify the variance risk premia on financial assets underlying options. Our method uses the notion of a variance swap, which is an over-the-counter contract that pays the difference between a standard estimate of the realized variance and the fixed swap rate. Since variance swaps cost zero to enter, the variance swap rate represents the risk-neutral expected value of the realized return variance. We theoretically and numerically show that the variance swap rate is well approximated by a particular linear combination of option prices. Hence, the difference between the ex-post realized variance and this synthetic variance swap rate quantifies the variance risk premium. Using a large options data set, we synthesize variance swap rates and analyze variance risk premia on five stock indexes and 35 individual stocks.

We find that the variance risk premia are strongly negative for the S&P 500 and 100 indexes and for the Dow Jones Industrial Average. The magnitude of the premia are smaller for the Nasdaq 100 index and for individual stocks. Investors are willing to pay a large premium, or receive a negative excess return, to take the long option position implicit in a long variance swap. The net effect of doing so is to hedge against market volatility going up. The negative risk premia imply that investors regard market volatility going up as unfavorable shocks.

We investigate whether the classical capital asset pricing theory can explain the negative variance risk premia. We find that the well-documented negative correlation between index returns and volatility generates a strongly negative beta, but this negative beta can only explain a small portion of the negative variance risk premia. The Fama-French factors cannot account for the strongly negative variance risk

premiums, either. Therefore, we conclude that either there is a large inefficiency in the market for index variance or else the majority of the variance risk is generated by an independent risk factor that the market prices heavily.

To analyze the dynamic properties of the variance risk premium, we formulate various forms of expectation-hypothesis regressions. When we regress the realized variance on the variance swap rate, we obtain slope estimates that are all positive, but mostly significantly lower than one, the null value under the hypothesis of constant or independent variance risk premium. The slope estimates become closer to one when the regression is on the logarithm of variance. These regression results indicate that although the log variance risk premium are strongly negative, they are not that strongly correlated with the expected log variance.

Like variance swaps, volatility swaps also trade over the counter and may be synthesized by trading in options. The difference between the variance swap rate and the square of the volatility swap rate measures the risk-neutral variance of volatility. Since we can readily synthesize both variance swap rates and volatility swap rates, this risk-neutral variance of volatility is easily and robustly determined from option prices. We regress the negative of the variance risk premium against this estimate of variance of volatility and find that the slope estimates are mostly positive. This result confirms our hypothesis that the variance risk premium become more negative when the variance of volatility is high.

When we use the log of variance and control for the variance of volatility in the expectation hypothesis regression, the regression slope estimates on the variance swap rate are no longer significantly different from the null value of one for all the five stock indexes and for all but five of the individual stocks. Hence, an observed relative increase in the variance swap rate is on average associated with a subsequent relative increase in realized variance of the same size, once we control for the variance of volatility.

The simple and robust method that we propose to measure the risk-neutral expected value of return variance and variance risk premium opens fertile ground for future research. On top of our research agenda is to understand the dynamics of return variance and the economic meanings of the variance risk premium. In particular, given the predominant evidence on stochastic variance and strongly negative variance risk premium, it is important to understand the pricing kernel behavior as a function of both the

market portfolio returns and return variances. Recent studies, e.g., Jackwerth (2000) and Engle and Rosenberg (2002) have found some puzzling behaviors on the pricing kernel projected on the equity index return alone. Accurately estimating the pricing kernel as a joint function of the index return and return variance represents a challenging task, but accomplishing this task can prove to be very fruitful not only for understanding the behavior of the variance risk premia, but also for resolving the puzzling behaviors observed on the pricing kernels projected on the index return alone.

The empirical analysis in this paper focuses on the variance swap rate over a fixed 30-day horizon. Since we observe option prices at many different maturities, we can construct variance swap rates at these different maturities and construct a term structure of variance swap rates at each day. An important line for future research is to design and estimate stochastic return variance models to capture the dynamics of the term structure. The key advantage of doing so is that we can gain a better understanding of the return variance dynamics without the interference from the return innovation specification.

Appendix A. Approximating the Volatility Swap Rate

Most of the results in this appendix are from Carr and Lee (2003a). We provide them here for completeness. Carr and Lee assume the following \mathbb{Q} -dynamics for the futures price F_t ,

$$dF_t/F_t = \sigma_t dW_t. \quad (\text{A1})$$

Compared to equation (3), they make the extra assumption of no jumps. They further restrict the diffusion volatility σ_t to be independent of the Brownian motion W_t .

Under these assumptions, Hull and White (1987) show that the value of a call option equals the risk-neutral expected value of the Black-Scholes formula value, considered as a function of the random realized volatility. In the special case when the call is at-the-money ($K = F$), the time- t value of a European call option maturing at time T becomes,

$$ATMC_{t,T} = B_t(T) \mathbb{E}_t^{\mathbb{Q}} \left\{ F_t \left[N \left(\frac{\sqrt{RV_{t,T}(T-t)}}{2} \right) - N \left(-\frac{\sqrt{RV_{t,T}(T-t)}}{2} \right) \right] \right\}, \quad (\text{A2})$$

where $RV_{t,T}$ is the random annualized realized return variance over the time period $[t, T]$,

$$RV_{t,T} \equiv \frac{1}{T-t} \int_t^T \sigma_s^2 ds. \quad (\text{A3})$$

As first shown in Brenner and Subrahmanyam (1988), a Taylor series expansion of each normal distribution function in (A2) around zero generates,

$$N \left(\frac{\sqrt{RV_{t,T}(T-t)}}{2} \right) - N \left(-\frac{\sqrt{RV_{t,T}(T-t)}}{2} \right) = \frac{\sqrt{RV_{t,T}(T-t)}}{\sqrt{2\pi}} + O((T-t)^{\frac{3}{2}}). \quad (\text{A4})$$

Substituting (A4) in (A2) implies that:

$$ATMC_{t,T} = B_t(T) \mathbb{E}_t^{\mathbb{Q}} \left[\frac{F_t}{\sqrt{2\pi}} \sqrt{RV_{t,T}(T-t)} + O((T-t)^{\frac{3}{2}}) \right], \quad (\text{A5})$$

and hence the volatility swap rate is given by:

$$VS_{t,T} = \mathbb{E}_t^{\mathbb{Q}} \sqrt{RV_{t,T}} = \frac{\sqrt{2\pi}}{B_t(T) F_t \sqrt{T-t}} ATMC_{t,T} + O((T-t)^{\frac{3}{2}}). \quad (\text{A6})$$

Since an at-the-money call is concave in volatility, $\frac{\sqrt{2\pi}}{B_t(T)F_t\sqrt{T-t}}ATMC_{t,T}$ is a slightly downward biased approximation of the volatility swap rate. As a result, the coefficient on $(T-t)^{\frac{3}{2}}$ is positive. However, Brenner and Subrahmanyam show that the at-the-money implied volatility (ATMV) is also given by:

$$ATMV_{t,T} = \frac{\sqrt{2\pi}}{B_t(T)F_t\sqrt{T-t}}ATMC_{t,T} + O((T-t)^{\frac{3}{2}}). \quad (A7)$$

Once again, $\frac{\sqrt{2\pi}}{B_t(T)F_t\sqrt{T-t}}ATMC_{t,T}$ is a slightly downward biased approximation of the at-the-money implied volatility and hence the coefficient on $(T-t)^{\frac{3}{2}}$ is positive. Subtracting equation (A7) from (A7) shows that the volatility swap rate can be approximated by the at-the-money implied volatility,

$$VS_{t,T} = ATMV_{t,T} + O((T-t)^{\frac{3}{2}}). \quad (A8)$$

In fact, the leading source of error in (A6) is partially cancelled by the leading source of error in (A7). As a result, this approximation has been found to be extremely accurate. The shorter the time to maturity, the better the approximation.

Appendix B. Approximating the Volatility Swap Rate

This appendix follows from an appendix in Carr and Lee (2003b). Let $\Gamma(\alpha) \equiv \int_0^\infty t^{\alpha-1}e^{-t}dt$ be the gamma function with α a positive real number. Then, it is well known that

$$\sqrt{\pi} = \Gamma\left(\frac{1}{2}\right) = \int_0^\infty \frac{e^{-t}}{\sqrt{t}}dt. \quad (B9)$$

Consider the change of variables $s = t/q$ for $q > 0$, we have $t = sq, dt = qds$, and

$$\sqrt{\pi} = \sqrt{q} \int_0^\infty \frac{e^{-sq}}{\sqrt{s}}ds, \quad (B10)$$

from which we obtain one representation for \sqrt{q} :

$$\sqrt{q} = \frac{\sqrt{\pi}}{\int_0^\infty \frac{e^{-sq}}{\sqrt{s}}ds}. \quad (B11)$$

Another representation is obtained by integrating (B10) by parts. Let

$$\begin{aligned} u &= \frac{1}{\sqrt{t}}, & dv &= e^{-t} dt \\ du &= -\frac{1}{2t^{3/2}} dt, & v &= 1 - e^{-t}. \end{aligned}$$

Hence,

$$\sqrt{\pi} = \frac{1 - e^{-t}}{\sqrt{t}} \Big|_{t=0}^{t=\infty} + \frac{1}{2} \int_0^{\infty} \frac{1 - e^{-t}}{t^{3/2}} dt = \frac{1}{2} \int_0^{\infty} \frac{1 - e^{-t}}{t^{3/2}} dt. \quad (\text{B12})$$

Again, consider the change of variables $s = t/q$ for $q > 0$,

$$\sqrt{\pi} = \frac{1}{2} \int_0^{\infty} \frac{1 - e^{-sq}}{(sq)^{3/2}} q ds = \frac{1}{2\sqrt{q}} \int_0^{\infty} \frac{1 - e^{-sq}}{(s)^{3/2}} dt, \quad (\text{B13})$$

from which we can solve for \sqrt{q} as

$$\sqrt{q} = \frac{1}{2\sqrt{\pi}} \int_0^{\infty} \frac{1 - e^{-sq}}{s^{3/2}} ds. \quad (\text{B14})$$

Table 6

Explaining variance risk premia with CAPM beta

Entries report the GMM estimates (and t -statistics in parentheses) of the following relation,

$$\ln RV_{t,T}/SW_{t,T} = \alpha + \beta_j ER_{t,T}^m + e,$$

where ER^m denotes the excess return on the market portfolio, which is proxied by the return on the S&P 500 index forward in the left panel and the excess return on the CRSP valued-weighted stock portfolio in the right panel. The t -statistics are computed according to Newey and West (1987) with 30 lags for the overlapping daily series in the left panel and six lags for the non-overlapping monthly series in the right panel. The columns under “ R^2 ” report the unadjusted R-squared of the regression.

Proxy	S&P 500 Index			Valued-Weighted Market Portfolio		
	α	β	R^2	α	β	R^2
SPX	-0.577 (-12.302)	-4.589 (-5.884)	0.183	-0.568 (-9.378)	-5.299 (-4.623)	0.234
OEX	-0.492 (-10.293)	-4.569 (-5.916)	0.186	-0.498 (-8.007)	-5.335 (-4.880)	0.233
DJX	-0.532 (-9.569)	-4.734 (-5.407)	0.216	-0.531 (-8.492)	-4.513 (-3.761)	0.198
NDX	-0.198 (-4.944)	-2.563 (-4.065)	0.090	-0.151 (-4.113)	-3.526 (-3.242)	0.183
QQQ	-0.267 (-4.754)	-1.226 (-2.015)	0.024	-0.238 (-4.482)	-2.709 (-1.568)	0.107
MSFT	-0.269 (-6.511)	-2.255 (-4.217)	0.058	-0.263 (-4.635)	-2.375 (-2.844)	0.063
INTC	-0.015 (-0.325)	-2.298 (-2.871)	0.059	0.016 (0.336)	-3.669 (-3.084)	0.143
IBM	-0.223 (-4.134)	-2.310 (-2.876)	0.044	-0.183 (-3.057)	-2.040 (-1.665)	0.037
AMER	-0.162 (-3.579)	-2.216 (-3.269)	0.043	-0.167 (-3.530)	-1.521 (-1.459)	0.022
DELL	-0.196 (-3.881)	-2.678 (-3.613)	0.073	-0.189 (-2.715)	-3.224 (-3.408)	0.110
CSCO	-0.266 (-3.577)	-0.957 (-0.599)	0.004	-0.217 (-3.195)	-1.927 (-0.855)	0.040
GE	-0.230 (-5.731)	-2.593 (-3.798)	0.092	-0.227 (-4.284)	-1.621 (-1.512)	0.046
CPQ	-0.110 (-1.975)	-2.398 (-2.312)	0.039	0.047 (0.909)	-3.318 (-2.762)	0.101
YHOO	-0.094 (-1.702)	-0.593 (-0.813)	0.004	-0.109 (-1.745)	0.831 (0.776)	0.008
SUNW	-0.089 (-2.080)	-2.380 (-3.371)	0.062	-0.049 (-0.899)	-3.951 (-3.187)	0.179
MU	-0.092 (-2.498)	-1.350 (-2.234)	0.025	-0.049 (-1.095)	-2.512 (-3.606)	0.094
MO	-0.244 (-3.906)	0.482 (0.540)	0.001	-0.180 (-2.942)	0.702 (0.604)	0.003
AMZN	-0.218 (-3.289)	0.120 (0.122)	0.000	-0.054 (-0.671)	0.284 (0.306)	0.001
ORCL	-0.141 (-2.041)	-2.206 (-2.597)	0.032	-0.124 (-1.697)	-3.674 (-2.567)	0.116
LU	-0.068 (-1.376)	-1.436 (-1.732)	0.019	0.030 (0.634)	-2.736 (-2.849)	0.103
TRV	-0.126 (-2.047)	-1.993 (-2.567)	0.035	-0.097 (-0.980)	-1.022 (-0.647)	0.011
WCOM	-0.101 (-1.405)	-3.430 (-3.307)	0.075	-0.006 (-0.068)	-4.137 (-2.691)	0.129
TYC	-0.353 (-4.066)	-1.724 (-1.490)	0.018	-0.321 (-3.137)	0.923 (0.308)	0.003
AMAT	-0.102 (-2.598)	-1.080 (-1.968)	0.015	-0.054 (-1.175)	-2.736 (-3.498)	0.104
QCOM	-0.154 (-2.889)	-1.305 (-1.646)	0.015	-0.089 (-1.578)	-2.578 (-2.202)	0.062
TXN	-0.028 (-0.677)	-0.724 (-1.346)	0.007	-0.032 (-0.621)	-0.714 (-0.887)	0.007
PFE	-0.200 (-3.342)	-1.957 (-1.878)	0.036	-0.149 (-2.455)	-1.909 (-1.246)	0.037
MOT	-0.031 (-0.581)	-1.954 (-1.861)	0.031	0.032 (0.441)	-3.523 (-2.044)	0.106
EMC	-0.031 (-0.682)	-2.611 (-3.398)	0.081	-0.028 (-0.389)	-3.146 (-3.890)	0.131
HWP	-0.076 (-1.530)	-1.661 (-1.956)	0.025	0.010 (0.185)	-1.956 (-1.536)	0.049
AMGN	-0.162 (-3.050)	-1.129 (-1.281)	0.014	-0.045 (-0.556)	-0.127 (-0.091)	0.000
BRCM	-0.029 (-0.533)	0.878 (1.172)	0.011	-0.013 (-0.146)	-0.613 (-0.304)	0.004
MER	-0.109 (-2.493)	-1.363 (-1.739)	0.024	-0.081 (-1.333)	-1.271 (-1.162)	0.022
NOK	-0.046 (-0.792)	-1.715 (-1.933)	0.030	0.028 (0.442)	-1.928 (-1.490)	0.059
CHL	-0.140 (-2.692)	-1.609 (-1.805)	0.029	-0.053 (-1.269)	-1.953 (-2.115)	0.052
UNPH	-0.005 (-0.085)	-1.444 (-1.258)	0.018	0.011 (0.297)	-2.982 (-1.201)	0.073
EBAY	-0.252 (-3.216)	0.173 (0.171)	0.000	-0.206 (-1.633)	0.018 (0.014)	0.000
JNPR	-0.147 (-2.154)	-0.490 (-0.575)	0.003	-0.133 (-1.720)	-1.912 (-0.924)	0.034
CIEN	-0.027 (-0.361)	-2.422 (-1.827)	0.046	-0.011 (-0.122)	-4.399 (-1.924)	0.154
BRCD	0.008 (0.110)	0.104 (0.111)	0.000	0.078 (0.811)	-2.497 (-1.941)	0.086

Table 7

Explaining variance risk premia with Fama-French risk factors

Entries report the GMM estimates (and t -statistics in parentheses) of the following relation,

$$\ln RV_{i,T}/SW_{i,T} = \alpha + \beta ER_{i,T}^m + sSMB_{i,T} + hHML_{i,T} + e,$$

where the regressors are the three stock-market risk factors defined by Fama and French (1993): the excess return on the market portfolio (ER^m), the size factor (SMB), and the book-to-market factor (HML). The data are monthly from January 1996 to December 2002. The t -statistics are computed according to Newey and West (1987) with six lags. The columns under “ R^2 ” report the unadjusted R-squared of the regression.

Ticker	α	ER^m	SMB	HML	R^2
SPX	-0.561 (-8.365)	-5.038 (-3.765)	-2.831 (-2.132)	-0.287 (-0.342)	0.276
OEX	-0.489 (-7.311)	-5.090 (-3.992)	-3.344 (-2.483)	-0.509 (-0.570)	0.289
DJX	-0.518 (-7.447)	-4.434 (-3.201)	-3.637 (-3.150)	-1.327 (-1.692)	0.273
NDX	-0.150 (-4.032)	-2.777 (-2.635)	-1.948 (-2.472)	1.351 (1.836)	0.272
QQQ	-0.221 (-4.184)	-1.932 (-1.309)	-1.851 (-1.784)	1.504 (2.651)	0.235
MSFT	-0.247 (-5.054)	-2.469 (-2.865)	-4.939 (-5.170)	-1.976 (-2.289)	0.222
INTC	0.023 (0.567)	-3.770 (-3.207)	-2.823 (-2.746)	-1.156 (-1.519)	0.194
IBM	-0.174 (-3.371)	-1.934 (-1.535)	-3.053 (-1.931)	-0.782 (-0.510)	0.085
AMER	-0.153 (-3.765)	-1.576 (-1.173)	-3.291 (-2.377)	-1.125 (-1.124)	0.084
DELL	-0.187 (-2.733)	-2.673 (-3.190)	-3.118 (-2.544)	0.401 (0.338)	0.190
CSCO	-0.227 (-3.598)	-1.009 (-0.444)	1.288 (1.026)	2.082 (2.389)	0.076
GE	-0.208 (-4.997)	-1.512 (-1.294)	-2.617 (-2.689)	-0.738 (-0.884)	0.121
CPQ	0.046 (0.959)	-3.024 (-2.112)	1.069 (0.927)	0.847 (0.809)	0.108
YHOO	-0.107 (-1.738)	0.144 (0.109)	0.574 (0.363)	-0.975 (-0.894)	0.029
SUNW	-0.056 (-1.079)	-3.113 (-2.087)	-1.509 (-1.472)	0.997 (0.959)	0.224
MU	-0.046 (-1.072)	-2.704 (-3.663)	-0.346 (-0.403)	-0.617 (-0.950)	0.099
MO	-0.187 (-3.145)	0.939 (0.688)	-0.306 (-0.169)	0.883 (0.622)	0.008
AMZN	-0.063 (-0.911)	-0.367 (-0.259)	-1.682 (-0.726)	-1.887 (-1.180)	0.043
ORCL	-0.119 (-1.570)	-3.893 (-3.237)	-0.264 (-0.172)	-0.377 (-0.413)	0.117
LU	0.031 (0.652)	-3.475 (-2.738)	-0.859 (-0.544)	-1.438 (-1.364)	0.133
TRV	-0.059 (-0.763)	-0.463 (-0.277)	-5.841 (-5.029)	-1.069 (-1.016)	0.218
WCOM	-0.014 (-0.157)	-4.793 (-2.656)	-2.345 (-1.403)	-1.784 (-1.558)	0.157
TYC	-0.282 (-2.585)	0.404 (0.173)	-4.381 (-1.565)	-2.715 (-1.850)	0.087
AMAT	-0.036 (-1.064)	-2.508 (-2.905)	-4.070 (-4.315)	-1.068 (-1.780)	0.247
QCOM	-0.090 (-1.600)	-1.744 (-1.446)	-3.711 (-3.344)	0.604 (0.581)	0.135
TXN	-0.020 (-0.400)	-0.949 (-1.439)	-4.230 (-4.987)	-1.954 (-2.125)	0.155
PFE	-0.129 (-2.541)	-1.528 (-0.865)	-3.535 (-1.956)	-1.073 (-0.551)	0.103
MOT	0.028 (0.353)	-3.038 (-1.893)	0.784 (0.541)	0.817 (0.737)	0.114
EMC	-0.030 (-0.491)	-1.651 (-1.707)	-1.965 (-2.401)	1.761 (2.539)	0.259
HWP	0.014 (0.250)	-2.384 (-1.600)	-0.826 (-0.690)	-1.046 (-0.946)	0.061
AMGN	-0.040 (-0.550)	-0.426 (-0.287)	-1.142 (-0.844)	-1.194 (-1.485)	0.018
BRCM	0.023 (0.290)	-0.094 (-0.042)	-3.281 (-2.507)	-0.535 (-0.444)	0.086
MER	-0.075 (-1.393)	-0.899 (-0.762)	-1.881 (-1.324)	0.297 (0.332)	0.060
NOK	0.043 (0.651)	-1.874 (-1.517)	-2.407 (-2.105)	-0.815 (-0.947)	0.115
CHL	-0.050 (-1.300)	-1.816 (-1.952)	-2.762 (-2.195)	-0.531 (-0.593)	0.111
UNPH	0.028 (0.418)	-1.791 (-0.796)	-1.815 (-1.015)	1.128 (0.798)	0.148
EBAY	-0.177 (-1.373)	0.590 (0.536)	-2.783 (-1.993)	0.300 (0.268)	0.072
JNPR	-0.049 (-0.686)	-3.209 (-2.277)	-3.556 (-2.444)	-3.600 (-3.620)	0.173
CIEN	0.011 (0.119)	-5.258 (-2.452)	-2.997 (-2.526)	-2.506 (-1.584)	0.211
BRCD	0.138 (1.155)	-1.766 (-1.485)	-3.769 (-2.599)	-0.966 (-0.782)	0.225

Table 8

Expectation hypothesis regressions on constant variance risk premia

Entries report the GMM estimates (and t -statistics in parentheses) of the following relations,

$$\begin{aligned} \text{Left panel: } RV_{i,T} &= a + bSW_{i,T} + e, \\ \text{Right panel: } \ln RV_{i,T} &= a + b \ln SW_{i,T} + e. \end{aligned}$$

The t -statistics are calculated according to Newey and West (1987) with 30 lags, under the null hypothesis of $a = 0, b = 1$. The columns under “ R^2 ” report the unadjusted R-squared of the regression.

Ticker	$RV_{i,T} = a + bSW_{i,T} + e$			$\ln RV_{i,T} = a + b \ln SW_{i,T} + e$		
	a	b	R^2	a	b	R^2
SPX	0.729 (1.148)	0.526 (-3.981)	0.293	-0.492 (-2.385)	0.941 (-0.505)	0.391
OEX	0.466 (0.740)	0.633 (-3.137)	0.315	-0.531 (-2.534)	1.012 (0.104)	0.418
DJX	1.007 (1.243)	0.527 (-3.361)	0.223	-0.183 (-0.696)	0.806 (-1.319)	0.269
NDX	-2.547 (-1.508)	1.085 (0.648)	0.585	-0.421 (-2.272)	1.080 (1.158)	0.690
QQQ	-4.281 (-1.435)	1.052 (0.364)	0.473	-0.655 (-1.778)	1.126 (1.070)	0.523
MSFT	3.042 (1.358)	0.701 (-2.209)	0.333	-0.254 (-0.968)	0.992 (-0.087)	0.453
INTC	2.351 (0.656)	1.007 (0.039)	0.358	0.325 (1.236)	0.888 (-1.276)	0.430
IBM	4.737 (1.885)	0.625 (-2.239)	0.244	0.143 (0.502)	0.862 (-1.320)	0.290
AMER	12.920 (3.494)	0.636 (-4.607)	0.263	0.562 (1.439)	0.801 (-1.969)	0.296
DELL	8.643 (2.092)	0.668 (-2.565)	0.224	0.507 (1.192)	0.797 (-1.666)	0.257
CSCO	-2.402 (-0.871)	1.006 (0.055)	0.534	-0.811 (-1.705)	1.161 (1.239)	0.367
GE	1.309 (1.240)	0.768 (-2.661)	0.407	0.047 (0.272)	0.886 (-1.612)	0.488
CPQ	13.279 (3.027)	0.551 (-2.977)	0.182	0.846 (1.742)	0.708 (-2.083)	0.251
YHOO	27.022 (3.714)	0.615 (-4.217)	0.313	0.696 (1.538)	0.811 (-1.860)	0.344
SUNW	5.287 (1.163)	0.843 (-1.105)	0.483	0.089 (0.245)	0.944 (-0.555)	0.481
MU	17.606 (3.873)	0.663 (-4.220)	0.273	0.971 (2.806)	0.733 (-3.121)	0.299
MO	6.654 (4.740)	0.459 (-6.258)	0.146	0.539 (1.978)	0.697 (-2.956)	0.238
AMZN	53.943 (4.739)	0.336 (-7.649)	0.140	1.013 (1.421)	0.728 (-1.819)	0.276
ORCL	30.942 (3.241)	0.267 (-3.026)	0.108	1.046 (1.254)	0.674 (-1.366)	0.288
LU	18.249 (2.730)	0.418 (-2.631)	0.220	0.338 (1.216)	0.870 (-1.505)	0.492
TRV	4.393 (1.623)	0.799 (-1.694)	0.237	0.414 (1.171)	0.808 (-1.578)	0.266
WCOM	9.398 (3.237)	0.634 (-2.599)	0.304	0.472 (1.308)	0.807 (-1.681)	0.377
TYC	18.120 (3.902)	0.355 (-18.024)	0.246	0.059 (0.208)	0.878 (-1.629)	0.480
AMAT	12.450 (3.235)	0.691 (-3.314)	0.267	0.837 (2.634)	0.748 (-2.962)	0.302
QCOM	8.678 (1.893)	0.786 (-2.259)	0.327	0.469 (1.130)	0.835 (-1.539)	0.292
TXN	4.316 (1.069)	0.929 (-0.506)	0.429	0.332 (1.113)	0.895 (-1.214)	0.466
PFE	4.546 (2.892)	0.584 (-4.106)	0.150	0.501 (1.664)	0.724 (-2.552)	0.195
MOT	6.500 (1.977)	0.820 (-1.157)	0.394	0.498 (1.725)	0.827 (-1.773)	0.449
EMC	-2.871 (-0.559)	1.153 (0.852)	0.493	0.116 (0.334)	0.954 (-0.454)	0.502
HWP	6.807 (2.461)	0.737 (-2.375)	0.332	0.340 (1.193)	0.862 (-1.549)	0.377
AMGN	5.241 (2.190)	0.728 (-2.354)	0.389	0.196 (0.700)	0.884 (-1.330)	0.437
BRCM	23.907 (1.813)	0.761 (-1.396)	0.320	1.045 (2.194)	0.754 (-2.249)	0.343
MER	4.187 (1.369)	0.799 (-1.399)	0.333	0.376 (1.188)	0.841 (-1.563)	0.373
NOK	12.659 (2.964)	0.646 (-2.579)	0.261	0.699 (1.776)	0.778 (-1.995)	0.338
CHL	4.234 (2.136)	0.756 (-1.912)	0.352	0.005 (0.021)	0.946 (-0.632)	0.518
UNPH	18.161 (1.653)	0.868 (-0.696)	0.349	0.514 (1.014)	0.876 (-1.014)	0.389
EBAY	2.708 (0.282)	0.926 (-0.516)	0.415	-1.256 (-3.506)	1.245 (2.857)	0.663
JNPR	32.793 (2.113)	0.632 (-2.301)	0.274	1.506 (2.419)	0.646 (-2.553)	0.235
CIEN	38.400 (2.601)	0.636 (-2.280)	0.204	1.113 (1.848)	0.738 (-1.971)	0.337
BRCD	24.059 (1.628)	0.861 (-0.846)	0.277	1.026 (1.622)	0.774 (-1.615)	0.299

Table 9

Variance risk premia and variance of return volatility

Entries report the estimates (and t -statistics in parentheses) of the following two regressions in the left and right panels, respectively,

$$\ln SW_{i,T}/RV_{i,T} = a + b (SW_{i,T} - VS_i^2) + e,$$

$$\ln SW_{i,T}/RV_{i,T} = a + b \ln (SW_{i,T}/VS_i^2) + e.$$

The t -statistics are calculated according to Newey and West (1987) with 30 lags, under the null hypothesis of $a = 0, b = 0$. The columns under “ R^2 ” report the unadjusted R-squared of the regression.

Ticker	$\ln SW_{i,T}/RV_{i,T} = a + b (SW_{i,T} - VS_i^2) + e$			$\ln SW_{i,T}/RV_{i,T} = a + b \ln (SW_{i,T}/VS_i^2) + e$		
	a	b	R^2	a	b	R^2
SPX	0.462 (7.001)	0.087 (3.117)	0.048	0.283 (4.052)	1.142 (6.233)	0.086
OEX	0.412 (4.598)	0.085 (1.355)	0.013	0.295 (3.218)	1.055 (2.929)	0.028
DJX	0.379 (4.258)	0.130 (2.386)	0.049	0.302 (3.007)	1.211 (3.241)	0.048
NDX	0.204 (4.049)	0.002 (0.092)	0.000	0.154 (2.701)	0.505 (1.306)	0.006
QQQ	0.232 (3.641)	0.010 (0.859)	0.003	0.189 (2.563)	0.671 (1.621)	0.011
MSFT	0.234 (5.065)	0.014 (3.037)	0.023	0.194 (3.255)	0.547 (2.137)	0.019
INTC	-0.012 (-0.189)	0.016 (0.715)	0.004	-0.002 (-0.032)	0.299 (0.444)	0.001
IBM	0.168 (2.839)	0.027 (4.342)	0.035	0.049 (0.723)	1.294 (4.945)	0.050
AMER	0.107 (1.834)	0.012 (2.397)	0.015	0.114 (1.975)	0.455 (1.639)	0.007
DELL	0.106 (1.690)	0.028 (2.527)	0.028	0.087 (1.182)	1.137 (2.150)	0.022
CSCO	0.293 (3.376)	-0.005 (-0.530)	0.001	0.210 (3.000)	0.469 (1.172)	0.003
GE	0.156 (3.410)	0.043 (2.877)	0.039	0.126 (2.725)	0.813 (4.186)	0.045
CPQ	0.055 (0.965)	0.020 (5.554)	0.071	0.004 (0.070)	1.017 (6.947)	0.059
YHOO	0.052 (0.858)	0.005 (4.437)	0.021	0.040 (0.652)	0.545 (2.071)	0.013
SUNW	0.062 (1.300)	0.012 (6.099)	0.038	0.035 (0.669)	0.685 (3.308)	0.029
MU	0.050 (1.340)	0.008 (3.342)	0.026	0.044 (1.161)	0.548 (2.760)	0.016
MO	0.165 (2.547)	0.037 (2.381)	0.050	0.133 (2.058)	0.792 (2.414)	0.040
AMZN	0.047 (0.535)	0.013 (4.195)	0.084	0.023 (0.264)	1.545 (4.041)	0.050
ORCL	0.300 (3.285)	-0.045 (-2.982)	0.175	0.310 (1.790)	-1.761 (-1.335)	0.036
LU	0.053 (1.034)	0.008 (9.332)	0.053	0.045 (0.744)	0.396 (1.720)	0.009
TRV	0.060 (0.796)	0.027 (1.617)	0.015	-0.046 (-0.531)	1.347 (2.944)	0.040
WCOM	-0.017 (-0.210)	0.045 (5.860)	0.108	-0.046 (-0.569)	1.396 (7.922)	0.095
TYC	0.270 (2.818)	0.011 (4.095)	0.039	0.245 (2.472)	0.734 (1.700)	0.012
AMAT	0.003 (0.072)	0.024 (3.123)	0.047	0.014 (0.272)	0.944 (2.550)	0.026
QCOM	0.098 (1.461)	0.012 (1.470)	0.010	0.097 (1.284)	0.588 (1.118)	0.005
TXN	-0.047 (-0.987)	0.020 (3.337)	0.040	-0.083 (-1.552)	1.027 (4.236)	0.040
PFE	0.109 (1.799)	0.064 (4.318)	0.047	0.068 (1.245)	1.231 (3.997)	0.059
MOT	0.004 (0.069)	0.018 (3.217)	0.040	-0.053 (-0.940)	1.245 (7.698)	0.058
EMC	0.033 (0.684)	0.003 (0.613)	0.002	0.017 (0.342)	0.299 (1.352)	0.006
HWP	0.046 (0.890)	0.020 (4.596)	0.022	0.036 (0.661)	0.605 (2.273)	0.011
AMGN	0.084 (1.490)	0.030 (2.701)	0.032	0.030 (0.535)	1.187 (3.188)	0.036
BRCM	-0.039 (-0.674)	0.010 (2.475)	0.039	-0.079 (-1.342)	1.521 (3.027)	0.049
MER	-0.026 (-0.422)	0.067 (3.737)	0.070	-0.025 (-0.393)	1.547 (3.603)	0.053
NOK	0.019 (0.306)	0.014 (4.730)	0.031	0.012 (0.184)	0.692 (2.930)	0.023
CHL	0.079 (1.335)	0.031 (4.601)	0.042	0.038 (0.607)	1.061 (5.001)	0.042
UNPH	-0.024 (-0.355)	0.005 (0.850)	0.004	0.006 (0.101)	-0.008 (-0.016)	0.000
EBAY	0.156 (1.494)	0.017 (1.554)	0.020	-0.143 (-1.336)	4.018 (5.897)	0.217
JNPR	0.067 (0.696)	0.007 (1.104)	0.022	0.118 (1.056)	0.260 (0.285)	0.001
CIEN	-0.044 (-0.522)	0.013 (2.144)	0.027	0.032 (0.388)	0.005 (0.005)	0.000
BRCD	-0.075 (-0.762)	0.009 (1.095)	0.018	-0.055 (-0.547)	0.668 (0.766)	0.007

Table 10

Expectation hypothesis regression on variance risk premia controlling for variance of volatility
 Entries report the GMM estimates (and t -statistics in parentheses) of the following relation,

$$\ln RV_{i,T} = a + b \ln SW_{i,T} + c \ln (SW_{i,T}/VS_i^2) + e.$$

The t -statistics are calculated according to Newey and West (1987) with 30 lags, under the null hypothesis of $a = 0, b = 1, c = 0$. The columns under “ R^2 ” report the unadjusted R-squared of the regression.

Ticker	a	b	c	R^2
SPX	-0.293 (-1.517)	1.007 (0.058)	-1.146 (-6.090)	0.442
OEX	-0.296 (-1.328)	1.001 (0.005)	-1.055 (-2.935)	0.434
DJX	-0.113 (-0.436)	0.878 (-0.874)	-1.072 (-3.106)	0.295
NDX	-0.363 (-1.833)	1.076 (1.106)	-0.459 (-1.188)	0.692
QQQ	-0.602 (-1.672)	1.132 (1.120)	-0.708 (-1.674)	0.529
MSFT	-0.371 (-1.431)	1.068 (0.704)	-0.651 (-2.424)	0.465
INTC	0.325 (1.241)	0.886 (-1.319)	0.047 (0.074)	0.430
IBM	0.033 (0.125)	0.967 (-0.331)	-1.243 (-4.732)	0.319
AMER	0.554 (1.412)	0.814 (-1.786)	-0.308 (-1.017)	0.298
DELL	0.569 (1.406)	0.811 (-1.588)	-1.056 (-1.974)	0.271
CSCO	-0.760 (-1.642)	1.166 (1.271)	-0.527 (-1.277)	0.369
GE	0.031 (0.180)	0.933 (-0.960)	-0.739 (-3.990)	0.506
CPQ	0.805 (1.721)	0.754 (-1.782)	-0.874 (-4.576)	0.285
YHOO	0.626 (1.277)	0.835 (-1.431)	-0.307 (-0.819)	0.346
SUNW	0.048 (0.134)	0.976 (-0.239)	-0.666 (-2.950)	0.494
MU	0.929 (2.602)	0.752 (-2.756)	-0.355 (-1.657)	0.303
MO	0.509 (1.891)	0.741 (-2.538)	-0.612 (-1.928)	0.256
AMZN	0.861 (1.198)	0.794 (-1.312)	-1.165 (-2.717)	0.295
ORCL	0.833 (1.216)	0.698 (-1.434)	1.402 (1.640)	0.305
LU	0.311 (1.149)	0.885 (-1.331)	-0.244 (-1.002)	0.493
TRV	0.341 (1.021)	0.887 (-1.010)	-1.160 (-2.914)	0.285
WCOM	0.530 (1.524)	0.842 (-1.402)	-1.319 (-6.742)	0.431
TYC	0.028 (0.098)	0.904 (-1.248)	-0.409 (-0.968)	0.481
AMAT	0.851 (2.782)	0.765 (-2.854)	-0.814 (-2.259)	0.316
QCOM	0.507 (1.213)	0.839 (-1.496)	-0.535 (-0.996)	0.296
TXN	0.318 (1.079)	0.930 (-0.819)	-0.955 (-4.014)	0.484
PFE	0.448 (1.487)	0.792 (-1.804)	-1.083 (-3.143)	0.232
MOT	0.482 (1.718)	0.860 (-1.427)	-1.118 (-5.689)	0.475
EMC	0.084 (0.244)	0.970 (-0.293)	-0.269 (-1.300)	0.505
HWP	0.344 (1.219)	0.875 (-1.418)	-0.524 (-2.093)	0.382
AMGN	0.285 (1.016)	0.897 (-1.190)	-1.143 (-3.162)	0.455
BRCM	0.784 (1.547)	0.830 (-1.429)	-0.975 (-1.996)	0.353
MER	0.361 (1.189)	0.887 (-1.166)	-1.412 (-3.219)	0.400
NOK	0.598 (1.433)	0.815 (-1.541)	-0.423 (-1.487)	0.343
CHL	-0.066 (-0.283)	1.011 (0.126)	-1.081 (-4.402)	0.536
UNPH	0.600 (1.131)	0.847 (-1.167)	0.460 (0.806)	0.391
EBAY	-0.133 (-0.308)	1.061 (0.646)	-3.755 (-5.153)	0.717
JNPR	1.664 (2.546)	0.598 (-2.653)	0.648 (0.702)	0.241
CIEN	1.232 (2.063)	0.697 (-2.246)	1.073 (1.100)	0.344
BRCB	1.015 (1.642)	0.778 (-1.591)	-0.050 (-0.055)	0.299