A Random Walk down the Options Market[1]

George J. Jiang
Finance Department
Eller College of Management
University of Arizona
Tucson, Arizona 85721-0108
Phone: 520-621-3373
Email: gjiang@eller.arizona.edu

Yisong S. Tian [2]
Finance Area
Schulich School of Business
York University
4700 Keele Street
Toronto, ON M3J 1P3
Phone: 416-736-2100 ext. 77943
Email: ytian@schulich.yorku.ca

May 14, 2009

[1]We would like to thank Peter Christoffersen, Jin-Chuan Duan, Kris Jacobs, Jangkoo Kang, Chris Lamoureux, Moshe Milevsky, Chris Telmer, Jason Wei, Shu Yan, and participants of finance seminars at the *University of Arizona*, *National University of Singapore*, and the *Third International Conference of Financial Markets* for helpful comments and suggestions. The financial support of the *Social Sciences and Humanity Research Council of Canada* is gratefully acknowledged.

[2]Corresponding author.

A Random Walk down the Options Market

Abstract

Forward variance forms a *martingale* and changes in forward variance follow a *random walk*. In this paper, we propose a *model-free* approach to extracting forward variance from option prices and empirically test the implications of the martingale restriction on forward variance. We find evidence of negative autocorrelation in daily changes in forward variance, contradicting the random walk hypothesis. However, this anomalous pattern of negative autocorrelation is fully explained by illiquidity effects and is statistically significant only in the first half of the sample period and when daily trading volume in option contracts is low. In addition, we find no evidence of market misreaction to volatility shocks across the variance term structure, contradicting previous studies that document "overreaction" of long-term volatility relative to changes in short-term volatility. We further show that the so-called "overreaction" anomaly is the result of model misspecification. Overall, our findings support the random walk hypothesis and the informational efficiency of the options market.

JEL Classification: G13, G14

Key Words: Model-free forward variance, random walk, expectations hypothesis, market illiquidity, model misspecification

A Random Walk down the Options Market

Unlike spot variance (e.g., realized variance or the Black-Scholes' (B-S, 1973) implied variance), forward variance forms a martingale in the sense that the current forward variance is an unbiased forecast for the corresponding forward variance in the future. This martingale property implies that any change in forward variance must be driven by the arrival of unexpected news and thus *orthogonal* to the conditional information set. These restrictions on the dynamics of the forward variance have direct implications for the informational efficiency of the options market. In particular, the *orthogonality condition* leads to the hypothesis that changes in forward variance form a random walk. While widely used in testing stock market efficiency and the expectations hypothesis on the term structure of interest rates,[1] such random walk tests have not been carried out in the options market. The main challenge is that such tests cannot be implemented without an appropriate empirical measure for forward variance. Existing studies have generally relied on various option-pricing models to extract the required variance measure (both spot and forward variance), leading to test results that are subject to model specification errors.

In this paper, we propose a *model-free* approach to extract forward variance from option prices and empirically test the implications of the martingale restriction. Drawing insight from Britten-Jones and Neuberger (2000), we construct *model-free forward variance* from option prices without making any model specification. The extracted forward variance reflects the market's aggregate expectations of future variance over the period between two option maturity dates. Holding the two maturity dates fixed, the time series of forward variances forms a martingale. We recognize that this *martingale restriction* may not hold under the objective probability measure due to forward risk premium.[2] To mitigate this common problem in market efficiency tests, we focus on hypothesis testing at the daily frequency since the effect of risk premium is shown to be negligible at the daily frequency. Performing random walk tests on daily changes in forward variance, we find evidence supporting the informational efficiency of the options market. An anomalous

[1] See for example Campbell and Shiller (1984, 1991), Keim and Stambaugh (1986), Fama and French (1988), Lo and MacKinlay (1988, 1990), Bekaert and Hodrick (1992), and Longstaff (2000).

[2] Forward risk premium is the difference between forward variance and lagged forward variance. In comparison, the more commonly used variance risk premium is the difference between realized variance and option-implied variance.

pattern of negative autocorrelation is detected but fully explained by illiquidity effects.

Compared with prior empirical research, our approach has several important advantages. First of all, we perform tests of market efficiency using forward variance instead of spot variance or option prices. Neither spot variance nor option prices (which are functions of spot variance) form a martingale. In contrast, changes in forward variance satisfy the orthogonality condition which directly leads to random walk tests that parallel those in the vast literature on the informational efficiency of the stock and bond markets. As a result, we are able to adopt a similar empirical approach that has proven successful in the study of stock returns and term structure of interest rates.

Secondly, our empirical measure of forward variance is model free. As in Britten-Jones and Neuberger (2000), we make no specific assumptions about the underlying asset price process or variance process and construct the forward variance directly from market prices of options across strike prices and maturities. In contrast, previous studies mostly rely on the B-S implied variance and then make additional assumptions about the variance dynamics or variance term structure in order to develop their empirical tests of market efficiency (e.g., Stein (1989) and Campa and Chang (1995)). This leads to misspecification problems and potentially incorrect inferences in empirical analysis. Indeed, Heynen, Kemna and Vorst (1994) demonstrate that the same empirical test can lead to contradictory findings on market misreaction if different assumptions about the underlying variance process are made.

Finally, we perform both time series tests (e.g., random walk tests) and term structure tests of market efficiency. The model-free forward variance is the key ingredient for developing empirical tests to detect anomalous behavior in both variance time series and cross-sectional relationships between long-term and short-term variances. In particular, it allows us to investigate whether or not changes in forward variance are predictable by lagged measures of variance term structure or whether or not long-term variance misreacts to changes in short-term variance or variance spread. Interestingly, we find no evidence of market misreaction to volatility shocks across the variance term structure, contradicting previous studies that document "over-reactions" of long-term volatility relative to changes in short-term volatility. We provide further evidence that such pattern of market "overreaction" is likely the result of model misspecification.

Empirically, we analyze the S&P 500 (SPX) index options data during the period from June 1, 1988 to December 31, 2007 and find that lagged forward variance is an unbiased forecast for the corresponding forward variance on the following day. This result suggests that the effect of risk premium on forward variance is negligible at the daily frequency, providing empirical support for the expectations hypothesis on forward variance. Nevertheless, further tests on daily changes in forward variance indicate that the orthogonality condition is violated, contradicting the random walk hypothesis. Specifically, we find that daily changes in forward variance exhibit significant negative autocorrelation, with roughly 18% of the daily change being reversed on the following day. The negative autocorrelation is robust across maturity groups, with stronger negative relations at longer option maturities. It does decay fairly quickly, however, and becomes statistically insignificant after two lags (i.e., after two trading days). Note that the negative autocorrelation reported here is not induced by the well-documented mean-reverting behavior of typical variance processes. In fact, further analysis indicates that daily time series of forward variance are not mean reverting. While daily time series of spot variance are mean reverting, it does not lead to negative autocorrelation in daily changes in spot variance.

Is the rejection of random walk evidence against the informational efficiency of the options market? It depends on the source of negative autocorrelation (e.g., Shleifer and Visney (1997)) – market frictions, fundamental risk associated with limits to arbitrage, or investor irrationality.[3] Rational investors may not be able to trade on the anomaly if there is fundamental risk associated with the execution of trading strategies. For example, costly information gathering or delay in dissemination may prevent the timely execution of potentially profitable arbitrage trades. This leads to fundamental risk to arbitrage and may deter even rational investors from engaging in the necessary arbitrage trades that are required to eliminate the anomalous pattern. Until we can rule out alternative explanations, it is inappropriate to attribute the violation of

[3]Another alternative explanation is *Bayesian learning* (e.g., Brav and Heaton (2002) and Lewellen and Shanken (2002)). Market partipants are rational but may not have access to the full information set and cannot fully assess the impact of volatility shocks on option prices. They initially may not be able to distinguish between alternative hypotheses regarding the impact of volatility shocks but do learn and improve their ability to detect them and assess their impact over time. This explanation cannot distinguish between negative and positive correlations, however.

orthogonality to investor irrationality, leading to potentially false rejection of market efficiency.

Indeed, several aspects of the documented negative autocorrelation appear to be linked to market illiquidity. First, negative autocorrelation is identified in daily changes in forward variance and disappears at the weekly or lower frequencies. We note that daily data used in our analysis provide more powerful tests of the orthogonality condition but are also more prone to market microstructure concerns. Secondly, we find negative as opposed to positive autocorrelation, consistent with the effect of bid-ask bounce or stale prices documented in the market microstructure literature (e.g., Scholes and Williams (1977) and Lo and MacKinlay (1988)). Finally, the negative autocorrelation becomes stronger as option maturity increases, consistent with the negative relationship between liquidity and option maturity. For these reasons, we perform additional analysis to further investigate the link between negative autocorrelation and market illiquidity.

We first use trading volume as a proxy for market liquidity and examine the link between autocorrelation and trading volume. We divide the sample period into high-volume days and low-volume days, based on the median daily trading volume after adjusting for trends and seasonality. Comparing autocorrelation between low-volume and high-volume days, we find that negative autocorrelation is present on low-volume days but entirely absent on high-volume days. This finding provides a direct linkage between the documented negative autocorrelation and market illiquidity.

As a robustness check, we also divide the full sample period into two non-overlapping subperiods. The average daily trading volume during the second subperiod (from January 1, 1998 to December 31, 2007) nearly quadruples the average daily trading volume in the first subperiod (from June 1, 1988 to December 31, 1997), suggesting markedly improved liquidity over time. We thus rerun our random walk tests separately for each subperiod and see whether the difference in market liquidity has any impact on autocorrelation. The subperiod results indicate that negative autocorrelation is present in the first subperiod but completely absent in the second subperiod. This finding provides further support that market illiquidity is the source of negative autocorrelation.

Finally, our model-free tests on variance term structure provide further support for the expectations hypothesis on forward variance. Given the orthogonality condition, changes in forward variance must be due to the arrival of unexpected news and thus uncorrelated with any lagged measures of the variance term structure. Regressing daily change in forward variance against lagged short-term variance and lagged variance spread, we find no evidence that lagged term structure measures have any predictive power for changes in forward variance. The evident support for the orthogonality condition here is in sharp contrast with previous research (e.g., Stein (1989)). These studies perform model-based tests on variance term structure behavior and find overreactions of long-term variance relative to changes in short-term variance. While we find similar evidence of "market overreaction" if we replicate their model-based tests using our sample, no such evidence is found if we use the model-free approach to perform the same tests. Further examining the patterns of B-S implied volatilities (skew or smirk) and their effect on model-based tests of variance term structure, we show that the previously documented overreaction anomaly is likely the result of model misspecification. These findings corroborate with previous studies (e.g., Heynen, Kemna and Vorst (1994) and Campa and Chang (1995)) that also report mixed evidence on market overreaction using either different data sets or different model-based approaches. Our model-free approach thus provides a logical explanation for such conflicting findings in previous studies.

The rest of the paper is organized as follows. Section I describes the model-free forward variance, the martingale restriction, and tests of market efficiency. Section II presents the options data and descriptive statistics. Sections III and IV present empirical evidence on the random walk hypothesis in the options market and evaluate alternative explanations of our findings. In Section V, we perform new tests of rational expectations on variance term structure and reconcile conflicting findings in previous studies. The final section concludes.

I. Forward Variance, Martingale Restriction, and the Random Walk Hypothesis

In this section, we describe the theoretical underpinning of forward variance, martingale restriction and the orthogonality condition. Testable hypotheses are then developed for our subsequent empirical investigation of the random walk hypothesis in the options market. A key ingredient of this empirical investigation is the construction of model-free forward variance which removes the need to estimate forward variance using potentially misspecified models.

A. Forward Variance and the Martingale Restriction

Forward interest rate plays a key role in the expectations theory of the term structure of interest rates (e.g., Cox, Ingersoll and Ross (1981) and Campbell and Shiller (1984, 1991)). *Forward variance* is a parallel concept in the options market, representing the market's projection of variance over a certain time period in the future:

$$v(t; T_1, T_2) \equiv \frac{1}{T_2 - T_1} E_t^Q \left[\int_{T_1}^{T_2} \left(\frac{dS_\tau}{S_\tau} \right)^2 \middle| \Omega_t \right]. \tag{1}$$

where t is the current date, S_τ is the price of the underlying asset on day τ, T_1 and T_2 are two future dates (with $t < T_1 < T_2$), and the expectation is taken under the risk-neutral probability measure Q conditional on the information set Ω_t on day t.

A unique feature of forward variance is the *martingale restriction*. Holding the two future dates T_1 and T_2 fixed, forward variance $v(t; T_1, T_2)$ forms a time series as it evolves over time from $t = 0$ to $t = T_1$. Although these forward variances are observed at different points in time, they are all projections of future variance over the same period of $[T_1, T_2]$. Applying iterative expectations, we have the following martingale relationship:

$$v(t; T_1, T_2) = E_t^Q [v(s; T_1, T_2) | \, \Omega_t], \tag{2}$$

for all $t < s \leq T_1$. In other words, forward variance on date t is an unbiased forecast for the corresponding forward variance on any future date s. Neither realized variance nor the B-S implied variance possesses this martingale property.

Rewrite the martingale restriction in terms of change in forward variance:

$$E_t^Q [v(s; T_1, T_2) - v(t; T_1, T_2) | \, \Omega_t] = 0 \tag{3}$$

This provides an alternative interpretation of the martingale restriction that is more useful for empirical testing: the change in forward variance is white noise and orthogonal to the conditional information set Ω_t. Intuitively, as forward variance $v(t; T_1, T_2)$ has already incorporated all available information on day t (i.e., the conditional information set Ω_t), any subsequent revision to forward variance must be due to the arrival of unexpected news. Consequently, the change in forward variance should be *orthogonal* to the conditional information set Ω_t. Following the literature on random walks (e.g., Campbell, Lo and MacKinlay (1997)), we call this property the *orthogonality condition*.

B. Random Walk and the Expectations Hypothesis

It is important to emphasize that the martingale restriction in Eq. (2) holds under the risk-neutral probability measure. Since empirical tests are performed under the objective probability measure, forward variance may not actually form a martingale due to the presence of forward risk premium. As a result, empirical tests of the martingale restriction are joint tests of the martingale restriction and zero risk premium. This is of course not a unique problem in the options market. Instead, it is a common problem in virtually all tests of market efficiency such as tests of the expectations hypothesis in the bond market (e.g., Longstaff (2000)) or random walk tests in equity market (e.g., Lo and MacKinlay (1988)).

Nevertheless, empirical tests based on the martingale restriction are still meaningful. In particular, the dynamic properties of the forward variance can have important implications for the informational efficiency of the options market. Firstly, the orthogonality condition in Eq. (3) implies that the change in forward variance should be orthogonal to the conditional information set Ω_t. No subset of conditional information should have any predictive power on the change in forward variance. While this is true only under the risk-neutral probability measure, the conditional information set provides useful guidance for the choices of conditional variables in our empirical tests. Secondly, violations of the martingale restriction under the objective probability measure are due to either violations of the martingale restriction or the presence of forward risk premium. To mitigate the risk premium effect, we focus on hypothesis testing at the daily frequency. This is appropriate as we show subsequently that the effect of risk premium is negligible at the daily frequency and thus should have no material impact on the interpretation of our empirical results.

In subsequent empirical tests, we thus focus on daily changes in forward variance:

$$\Delta v(t; T_1, T_2) \equiv v(t; T_1, T_2) - v(t-1; T_1, T_2).$$

The use of daily time series has at least two advantages. As mentioned previously, forward risk premium is expected to be smaller for daily changes in forward variance than for either weekly or monthly changes. It is thus possible that the martingale property is supported by daily time series empirically. In addition, the use of daily changes also takes into account all available information and gain an improved power for empirical tests.

In our first empirical test, we directly examine the martingale restriction and investigate whether or not the lagged forward variance (i.e., $v(t-1; T_1, T_2)$) is an unbiased forecast for the corresponding forward variance on the following day (i.e., $v(t; T_1, T_2)$). The null hypothesis is that the risk premium is zero and the martingale restriction holds. Although a rejection of the null hypothesis is not necessarily a rejection of the martingale restriction, results from this joint test do provide useful guidance for subsequent empirical tests. Similar joint tests have proven quite useful in testing the expectations hypothesis in the equity, interest rate and foreign currency markets (e.g., Leroy (1973), Lucas (1978), Lo and MacKinlay (1988), Fama (1991), and Longstaff (2000)).

The second empirical test is a random walk test based on the orthogonality condition. As discussed previously, daily change in forward variance $\Delta v(t; T_1, T_2)$ is induced by unexpected news arrivals (or variance shocks) on day t. Likewise, the lagged daily change in forward variance $\Delta v(t-1; T_1, T_2)$ is induced by unexpected news arrivals on day $t-1$. Since unexpected news are by nature uncorrelated, the time series of daily changes in forward variance should be serially uncorrelated. Any serial correlation represents predictable patterns which contradicts the orthogonality condition. This leads to a simple random walk test that daily changes in forward variance are serially uncorrelated. .

Our final empirical test focuses on the dynamic behavior of the variance term structure and determines whether future movements of forward variance are predictable by the level and shape of the variance term structure. Changes in forward variance must be due to the arrival of unexpected news and are thus uncorrelated with previously available information including all lagged measures of the variance term structure. Our null hypothesis is that the daily change in forward variance is not predictable by either the lagged short-term variance and or the lagged variance spread (proxies for the level and shape of the variance term structure, respectively). Predictability of daily change in forward variance is a rejection of the orthogonality condition.

C. Model-Free Forward Variance

Empirical tests developed previously require an appropriate empirical measure for forward variance. Previous studies (e.g., Stein (1989), Heynen, Kemna and Vorst (1994), and Campa and Chang (1995)) typically use the B-S implied volatility to approximate the required forward variance using a particular volatility model (e.g., AR(1) or Hull and White (1987)). As both the B-S model and the assumed volatility model are most likely misspecified,[4] it is unclear whether the rejection of the null hypothesis is due to model misspecification or true violations of the expectations hypothesis.

In this study, we deviate from the existing literature and use a model-free approach to extract forward variance directly from option prices without any model specification. This is possible due to the important contribution by Britten-Jones and Neuberger (2000) who show that forward variance is a function of market prices of options across a continuum of strike prices and two maturities:

$$\begin{aligned} v(t;T_1,T_2) &\equiv \frac{1}{T_2-T_1} E_t^Q\left[\int_{T_1}^{T_2}\left(\frac{dS_\tau}{S_\tau}\right)^2 \middle| \Omega_t\right] \\ &= \frac{2}{T_2-T_1}\int_0^\infty \frac{C_t[T_2,K/B(t,T_2)]-C_t[T_1,K/B(t,T_1)]}{K^2}dK \end{aligned} \tag{4}$$

where T_1 and T_2 are option maturity dates (with $T_1 < T_2$), $B(t,T)$ is the market price on day t of a zero-coupon bond that pays \$1 on day T, and $C_t(T,K)$ is the market price on day t of a call option with strike price K maturing on day T. This empirical measure of forward variance will be referred to as the *model-free forward variance* or simply *forward variance* whenever its meaning is clear within context. This model-free

[4]See for example evidence in Bates (1996, 2000, 2004), Bakshi, Cao and Chen (1997, 2000), Andersen, Benzoni and Lund (2002), Pan (2002), Carr and Wu (2003), Chernov, Gallant, Ghysels and Tauchen (2003), and Eraker, Johannes and Polson (2003).

measure of forward variance allows us to implement the empirical tests described previously and examine the expectations hypothesis in the options market.[5]

Note that Britten-Jones and Neuberger's (2000) model-free forward variance is a generalization of the model-free *spot* variance (also known as the model-free implied variance) which originates from the study of variance and volatility swaps (e.g., Dupire (1994), Neuberger (1994), Carr and Madan (1998), and Demeterfi, Derman, Kamal and Zou (1999)). Capturing the expected variance over the remaining life of the options, model-free spot variance is a function of market prices of options with the same maturity but across a continuum of strike prices:

$$\begin{aligned} v(t,T) &\equiv \frac{1}{T-t} E_t^Q \left[\int_t^T \left(\frac{dS_\tau}{S_\tau} \right)^2 \middle| \Omega_t \right] \\ &= \frac{2\exp[r(T-t)]}{T-t} \left[\int_0^{F_0} \frac{P_t(T,K)}{K^2} dK + \int_{F_0}^{\infty} \frac{C_t(T,K)}{K^2} dK \right] \end{aligned} \tag{5}$$

where t is the current date, T is the option maturity date, F_0 is the forward price on date t with maturity coinciding with the option maturity, $C_t(T,K)$ and $P_t(T,K)$ are date t market prices of call and put options, respectively, with strike price K and maturity T, and r is the risk-free rate with continuous compounding. Although there is a fast-growing literature (e.g., Jiang and Tian (2005), Christoffersen, Jacobs, and Vainberg (2008), and Carr and Wu (2009), among others) utilizing this model-free spot variance to study volatility forecasting, variance risk premium, and implied beta, this VIX-type implied variance is not a measure of forward variance and does not possess the martingale property discussed previously. As a result, model-free spot variance is not suitable for testing the informational efficiency of the options market.

II. Data and Descriptive Statistics

We perform our empirical tests using daily SPX options data. Daily data are obviously subject to market microstructure concerns (e.g., Lo and MacKinlay (1988)). Nevertheless, we choose to use daily as opposed to weekly or less frequently observed option prices for several reasons. First of all, more frequently observed

[5]Note that Britten-Jones and Neuberger (2000) derive the model-free forward variance for arbitrary diffusion processes. Jiang and Tian (2005) generalize it to arbitrary martingale processes with jumps.

data contain greater information content and the use of a larger sample provides an improved statistical power for empirical testing. Secondly, the use of daily forward variance is likely to diminish the effect of risk premium on the martingale restriction under the objective probability measure. Our subsequent empirical analysis provides strong support for this argument. Finally, market participants are likely to react swiftly to volatility shocks as they are by nature unexpected and may contain relevant information content. A more frequently observed sample (such as daily data) is more useful for capturing the effect of volatility shock. Of course, we do realize that SPX options are not immune to market microstructure concerns even though they are among the most actively traded option contracts in the marketplace. We do investigate the robustness of our empirical results and ensure that they are not driven by market microstructure concerns.

We pool several data sources together for our subsequent empirical tests. Daily closing data for SPX options are from the Chicago Board Options Exchange (CBOE). Daily Treasury bill yields (our proxy for risk-free rates) are obtained from the Federal Reserve Bulletin. Daily cash dividends are from the Standard and Poors' DRI database. Following Harvey and Whaley (1991, 1992a, 1992b), we use daily cash dividends instead of dividend yield.[6] As the dividend data are only available after May 1988, we choose our sample period to be June 1988 – December 2007. Trading in SPX options is also less active prior to 1988.

To finalize our sample of SPX options, several commonly used data filters are applied. First, option quotes less than \$3/8 are excluded from the sample. These prices may not reflect true option value due to proximity to tick size. Secondly, options with less than a week remaining to maturity are excluded from the sample. These options may have market illiquidity and microstructure concerns. Next, we exclude in-the-money options from the sample. In-the-money options are more expensive and often less liquid than at-the-money or out-of-the-money options. Following Aït-Sahalia and Lo (1998), we use both call and put options in our sample in order to span the available range of strike prices. Put-call parity is used to fill in the missing in-the-money call (put) prices from the corresponding out-of-the-money put (call) prices. Finally, options violating the boundary, monotonicity and convexity conditions are eliminated from the sample. This

[6] For robustness, we also consider dividend yield implied by option prices (using the put-call parity) and find no material change to our empirical results.

is to ensure that we only use a subset of all available option prices that are absent of arbitrage opportunities (e.g., Aït-Sahalia and Duarte (2003) and Carr and Madan (2005)).

We then construct daily series of model-free spot and forward variances using daily closing option prices. Following previous research (e.g., Bakshi, Cao and Chen (1997, 2000)), we use the midpoint of closing bid and ask quotes instead of the last transaction price of the day in order to minimize nonsynchronous trading and bid-ask bounce problems. Both spot and forward variances are calculated for all available option maturity months.

To calculate forward and spot variances as defined in Eqs. (4) and (5), respectively, we need option prices across all strike prices over an infinite range $[0, +\infty]$. Of course, options are only traded over a finite range of strike prices at discrete increments. As shown in Jiang and Tian (2005, 2007), this problem can be overcome effectively by a curve-fitting method with appropriate extrapolation.[7] Between the maximum and minimum available strike prices, natural cubic splines are used to fit a smooth function to the B-S implied volatilities at listed strike prices. Option prices at strike prices that are not listed are then extracted from the fitted function. Beyond the maximum and minimum available strike prices, option prices are extrapolated from the B-S implied volatility at the closest available strike price (either the maximum or minimum available strike price) in order to reduce truncation errors. Once all the needed option prices are extracted, forward and spot variances are calculated using Eqs. (4) and (5), respectively, in conjunction with a numerical integration method. Details of the methodology and its effectiveness can be found in Jiang and Tian (2005, 2007).

Table I about here

Table I provides descriptive statistics of the daily series of spot and forward variances in our sample. We divide the full sample of options into five maturity groups on the basis of expiration months (from T_1 to T_5). These maturity groups will be used in all subsequent empirical tests. The first group includes options

[7] In a related study, Carr and Wu (2009) use a linear interpolation method (also with flat extrapolation) to calculate the model-free implied variance. Their results indicate that this simple method is also quite effective.

with the nearest expiration month, the second group consists of options with the next expiration month, and so on. After the fourth expiration month, all remaining options are aggregated into a single maturity group. We will refer to these groups as the *spot-maturity groups* when it is necessary to distinguish them from forward-maturity groups to be defined next. For each of these five maturity groups, we compile descriptive statistics of spot variance and summarize them in Panel A of Table I. For forward variance, we need two expiration dates (a short one and a long one) to form a single forward period. We thus form four forward-maturity groups from the five spot-maturity groups. We choose the nearest expiration month (T_1) as the short maturity for all four forward-maturity groups and vary the long maturity from T_2 to T_5.[8] For each of the four forward-maturity groups, we compile descriptive statistics of forward variance and summarize them in Panel B of Table I.

Consider first the summary statistics of spot variance reported in Panel A of Table I. The mean spot variance varies in a narrow range from 0.0372 to 0.0403 across the five maturity groups.[9] Spot variance tends to be higher for options with longer maturities, suggesting a generally upward-sloping variance term structure. To further illustrate this term structure relationship, we plot the time series of one-month and three-month spot volatilities in Fig. 1.[10] The three-month volatility is typically higher than the corresponding one-month volatility, confirming that the term structure of spot variances is mostly upward sloping. However, there are also numerous occasions when the three-month volatility is lower than the one-month volatility, indicating the frequent occurrences of downward-sloping term structures. This is consistent with the findings by Xu and Taylor (1994) on the term structure of the B-S implied volatilities.

Fig. 1 about here

[8]We also consider other maturity groupings with more distant maturities as the short maturity, such as (T_2, T_i) for i = 3, 4, 5. All results remain qualitatively unchanged if these additional groupings are included.

[9]This range of mean implied variance translates into a range of mean implied volatility from 19.29% to 20.07%. Although implied volatility is more intuitive, we use implied variance throughout the paper as the martingale restriction applies to implied variance, not implied volatility.

[10]Options with exactly one or three month maturity are usually not available. We thus plot the spot volatility for options with the nearby and the third maturity month instead. The average maturity of these options in our sample period are roughly one month and three months, respectively.

In addition, autocorrelations of the daily series of spot variances suggest strong persistence over time, with the first-order correlation ranging from 0.927 to 0.983. Although higher-order partial autocorrelations drop off substantially, the second- and third-order partial autocorrelations (not reported) are not negligible in general. In particular, the second-order partial autocorrelations are all significant, suggesting that there may be at least two important factors driving the dynamics of spot variance. In contrast, Poterba and Summers (1986), Stein (1989) and Campa and Chang (1995) find that AR(1) is an appropriate representation of the implied volatility series. However, it is important to emphasize that these studies use the B-S implied volatility instead of the model-free implied variance. Although AR(1) may be an appropriate process for the B-S implied volatility, we cannot conclude that it is also an appropriate process for the model-free implied variance.

In Panel B of Table I, the summary statistics of forward variance are presented. These summary statistics are mostly similar to the corresponding summary statistics of spot variance in Panel A with two notable exceptions. One is that forward variance tends to have stronger autocorrelation than spot variance does, especially at shorter maturities. This is likely due to the differences in information content between forward and spot variances. While forward variances on day t and day $t+1$ are both projections of the expected realized variance over the same forward period $[T_1, T_2]$, the corresponding spot variances on the same two days are projections of expected realized variance over two overlapping periods of $[t, T_2]$ and $[t+1, T_2]$. The difference in the overlap may be small when the maturity date T_2 is far in the future, but can be more substantial when it is near. Another notable difference is that forward variance has very similar levels of autocorrelation across maturity groups while spot variance can have very different levels of autocorrelation across maturity groups. For example, the first-order autocorrelation varies in a very narrow range between 0.978 and 0.987 for forward variance across maturity groups but in a much wider range between 0.927 and 0.983 for spot variance. Such levels of differences are also observed at higher-order autocorrelations as well. Again, such differences between spot and forward variances are consistent with the nature of these variance measures.

III. Random Walk Tests in the Options Market

In this section, we perform tests of the martingale restriction and orthogonality condition. We first investigate whether forward variance violates the martingale restriction. Even though the martingale restriction is valid under the risk-neutral measure only, it is still interesting to verify whether it also holds under the objective probability measure. This may help to quantify the effect of risk premium on forward variance. We then examine the orthogonality condition and its implications for the informational efficiency of the options market.

A. Unconditional Tests

Following Longstaff (2000), we begin with an unconditional test that the daily change in forward variance has zero mean. This martingale condition is of course only true under the risk-neutral probability measure. The expected daily change in forward variance may not be zero under the objective probability measure. It is an entirely empirical matter whether the risk premium is sufficiently large for the mean to deviate from zero. This unconditional test is thus not a true model-free test but a joint test of the martingale relationship and zero risk premium.

We perform the unconditional test separately for the four forward-maturity groups defined previously in Table I. Again, each group is defined by a forward period (i.e., (T_1, T_i) for $i = 2$, 3, 4, 5), spanning a future period between two option maturity dates. A time series of forward variance is then constructed from option prices using Eq. (4). Previous studies (e.g., Stein (1989)) typically use only one maturity group, constructed using options with the shortest and second shortest maturities. To ensure robustness, we construct four maturity groups with the first group identical to the one used in previous studies. Taking the first difference of the time series, we obtain daily changes in forward variance which is the subject of the unconditional test.

Results of the unconditional test are reported in Table II. The average daily change in forward variance is close to zero in all four maturity groups. They range from -3.9×10^{-5} to 4.1×10^{-5}, with two positive means and two negative means. As the variance measure is already annualized, these means are quite small. If the average annual variance is 0.04 (equivalent to a standard deviation of 20% a year), these means imply that the average daily change in forward variance represents at most a 0.1% deviation from the average variance. Not surprisingly, none of the means comes close to being significant at any conventional significance level. We thus cannot reject the null hypothesis that the daily change in forward variance has zero mean. The unconditional tests provide empirical support for the martingale restriction.

Table II about here

B. Orthogonality Tests

Depending on the choice of conditioning variables, the orthogonality condition can be examined using a wide variety of random walk tests. Here we concentrate on tests that only require information on past forward variances. In particular, we examine whether or not the daily change in forward variance is orthogonal to (1) lagged forward variance and (2) lagged daily change in forward variance.

We first test whether or not daily change in forward variance is orthogonal to lagged forward variance. This is a common conditional test for the expectations hypothesis (e.g., Stein (1989)). The conditional information set is proxied by the subset of information contained in the lagged forward variance. A univariate linear regression is implemented to perform this test:

$$\Delta v(t; T_1, T_2) = \alpha + \beta v(t-1; T_1, T_2) + \varepsilon_t. \tag{6}$$

The orthogonality condition translates into a null hypothesis that $\beta = 0$ (i.e., daily change in forward variance is orthogonal to lagged forward variance). We also test the joint hypothesis that both the slope coefficient and intercept are zero ($\alpha = \beta = 0$), adding the additional restriction of zero mean. To ensure robustness, we perform regression (6) separately for each maturity group. Market efficiency may differ across maturity groups because long-term options may attract different types of investors than short-term options do and market liquidity also tends to differ across options with different maturities. The results are summarized in Table III, with the standard errors for the estimated coefficients reported in brackets and the

χ^2-test statistic for the joint null hypothesis and its p-value (in brackets) reported in the last column. The standard errors of the estimated coefficient are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (Newey and West (1987)).

Table III about here

As the results in Table III indicate, the intercept (α) is virtually zero and the slope coefficient (β) is negative but very small. None of them is statistically significant at any conventional significance level in any maturity group. The χ^2-test for the joint hypothesis that $\alpha = \beta = 0$ cannot reject the null hypothesis in any of the four maturity groups, with all p-values larger than 10%. This is clear evidence that daily change in forward variance is orthogonal to lagged forward variance, supporting the orthogonality condition. It also suggests that today's forward variance is an unbiased forecast for the corresponding forward variance on the following day and the effect of risk premium is negligible (as α is statistically insignificant) at the daily frequency. These results lend support for our subsequent random walk tests of the expectations hypothesis.

The previous test, though commonly used, is not necessarily the most effective. Lagged forward variance is only one of many variables in the conditional information set and may not contain sufficient information content on the dynamic behavior of forward variance. The regression R^2 in Table III confirm that the lagged forward variance contain virtually no information for the daily change in forward variance. Our next orthogonality test is a random walk test that utilizes a larger number of potentially more informative conditioning variables. In particular, we regress daily change in forward variance against a number of lagged daily changes in forward variance using the following autoregressive regression:

$$\Delta v(t; T_1, T_2) = \alpha + \sum_{i=1}^{k} \phi_i \, \Delta v(t-i; T_1, T_2) + \varepsilon_t, \tag{7}$$

where k is the number of lag terms included. The daily change in forward variance on day t is regressed against lagged daily changes on prior days from $t-1$ to $t-k$. Under the orthogonality condition, the ϕ coefficients in regression (7) should all be zero. In other words, the daily change on day t should be uncorrelated with any of the lagged daily changes before day t if the orthogonality condition holds. The results are summarized in Table IV for the four maturity groups. For each maturity group, we run regression (7) with 25 lags (or a lag period of just over one month). For brevity, we only report coefficients for the first ten lags in Table IV but plot all coefficients in Fig. 2.

Table IV about here

As shown in Table IV, daily changes in forward variance exhibit strong negative autocorrelation, with the first two lag terms (i.e., ϕ_1 and ϕ_2) capturing most of the autocorrelation. Beyond the first two lags, autocorrelation diminishes rapidly in both magnitude and statistical significance. In fact, even the largest ϕ coefficient (in absolute value) beyond the first two lags is smaller than 0.05 and none of them is statistically significant at any conventional significance level. This is true for all four maturity groups. This pattern of rapid decay is also clear from Fig. 2 which visually illustrates the ϕ coefficients over 25 lags, together with the 95% confidence band, for all four forward-maturity groups.

To further analyze the documented negative autocorrelation, we focus on the first two lag terms (i.e., ϕ_1 and ϕ_2). The coefficient of the first lag is consistently negative and varies from -0.08 to -0.27, with an average value of -0.18 across the four maturity groups. It is statistically significant at the 5% level in the first maturity group but at the 1% level in the remaining three groups. These numbers suggest that, on average, about 18% of the daily change in forward variance on day $t-1$ is reversed on day t.[11] The strongest reversal is found in Group IV (27%) while the weakest reversal is in Group I (8%). The negative autocorrelation weakens after the first lag but remains mostly significant at the second lag (except in Group I). This is evidence of predictability in daily changes in forward variance, rejecting the orthogonality condition.

Fig. 2 about here

[11]The estimated ϕ_1 coefficient implies that a one-basis point increase in lagged forward variance leads to an 0.18-basis point decline in the current forward variance.

One potential concern is that negative autocorrelation reported in Table IV could be induced by the well-known mean-reverting property of a typical variance process. To rule out this possibility, we perform further analysis and present evidence that a) daily forward variance is not mean reverting and b) daily spot variance is mean reverting but does not lead to negative autocorrelation in daily changes in spot variance. These results suggest that negative autocorrelation in daily change in forward variance has little to do with variance mean reversion.

First, the forward variance process does not seem to exhibit much mean reversion. In Table III, we present the results of regressing daily change in forward variance against lagged forward variance. If the forward variance process is mean reverting, high (low) levels of forward variance tend to be followed by a downward (upward) revision to forward variance. We thus expect to see a negative slope coefficient (β) in Table III if the forward variance process is mean reverting. Although it is indeed negative, none of them comes close to being significant at any conventional significance level. Mean reversion is thus rejected for the forward variance process.

Secondly, mean reversion in a variance process does not necessarily lead to negative autocorrelation in the first difference of the variance process. Consider the four daily spot variance series with maturities matching the long maturity of each of the four forward-maturity groups (from T_2 to T_5). For each daily spot variance series, we regress daily change in spot variance against lagged spot variance. Untabulated results indicate that all slope coefficients are negative and statistically significant, suggesting mean reversion for the daily spot variance process. Does mean reversion in the daily spot variance process lead to negative autocorrelation in daily changes in spot variance? To find out, we regress the daily change in spot variance against lagged daily changes in spot variance using 25 lags (similar to the analysis for the daily change in forward variance in Table IV). Untabulated results indicate that although some slope coefficients (ϕ's) are indeed negative, none of them is statistically significant for any of the four spot variance series. Mean reversion in the daily spot variance series thus does not lead to negative autocorrelation in daily changes in spot variance.

IV. Illiquidity Effect and the Expectations Hypothesis

The efficient market hypothesis postulates that security prices accurately and instantaneously incorporate all available information. Evidence against orthogonality presented in the previous section seems to suggest otherwise. Negative autocorrelation we document in the previous section suggests that forward variance contains predictable components, with roughly 18% of the daily change in forward variance being reversed on the following day. This level of predictability appears inconsistent with the efficient market hypothesis.

In this section, we perform additional tests in order to investigate the source of the negative autocorrelation we document in Table IV. Simply attributing the anomalous pattern to the effect of risk premium is inappropriate as we have already provided evidence (see Table III) that forward risk premium is negligible at the daily frequency. A much more likely source is market illiquidity since negative autocorrelation seems to strengthen as market liquidity deteriorates. As shown in Table IV, while only the first lag coefficient is significant in Group I, the first two lags are significant in the remaining three groups. Both the magnitude of the coefficient and its statistical significance appear to increase with the forward period (from Group I to Group IV). As long-term options are usually less actively traded than short-term options, the anomaly seems to be stronger when market illiquidity is more severe.

To examine the impact of market illiquidity on violations of orthogonality, we use trading volume as a proxy to market liquidity and separate trading days in our sample period into high-volume days and low-volume days.[12] For each maturity group, we aggregate trading volumes for all options in the group on each trading day and use the total volume as our proxy for liquidity. As trading volume tends to rise over time and exhibit seasonality (such as day of the week effect), we account for trends and seasonality in classifying high-volume vs. low-volume days. Following prior literature (e.g., Chen, Hong and Stein (2001) and Gervais, Kaniel and Mingelgrin (2001)), we first detrend the daily trading volume over time for each maturity group. We then classify a trading day as a high-volume (low-volume) day if the trading volume on that day is higher (lower) than the median daily trading volume over the past month. Using a dummy

[12]We do investigate the robustness of our proxy for market liquidity subsequently and find no material change to our results if other liquidity measures are used instead.

variable for low-volume days, we modify and rerun regression (7) by adding cross terms of the lag variables with the low-volume dummy:

$$\Delta v(t; T_1, T_2) = \alpha + \sum_{i=1}^{k} (\phi_i + \gamma_i \, D_{t-i}) \cdot \Delta v(t-i; T_1, T_2) + \varepsilon_t, \tag{8}$$

where D_t is a dummy variable that takes on a value of 1 if day t is a low-volume day and zero otherwise. In this dummy variable regression, the coefficient ϕ_i is the effect of the lagged daily change in forward variance on high-volume days while the coefficient $(\phi_i + \gamma_i)$ is the corresponding effect on low-volume days. Coefficients of the cross terms (γ_i) capture the difference between high- and low-volume days and thus the impact of market illiquidity. If the documented negative autocorrelation is unrelated to market liquidity, we expect the cross term coefficient (γ_i) to be zero. This dummy variable regression approach is supported by findings in previous studies on the informational efficiency of the options market. Pool (2005) finds that prices of SPX options are more informative on high-volume days relative to low-volume days and market illiquidity is a likely explanation for mispricing on low-volume days. Similarly, Donaldson and Kamstra (2005) find that the B-S implied volatility is a better forecast for future volatility than ARCH forecasts when trading volume is high.

The results of the dummy variable regression in (8) are presented in Table V. Consider first the cross terms of the lag variables with the low-volume dummy (γ's) which capture the difference in autocorrelation between low-volume days and high-volume days. As shown in Table V, the coefficient of the first cross term (γ_1) is all negative and statistically significant (at 5% in Group I and 1% in Groups II-IV). Even the coefficient of the second cross term (γ_2) is negative and statistically significant (at 5%) in two of the four cases. Furthermore, no coefficient of any of the cross terms is ever positive and statistically significant. These results suggest that negative autocorrelation in daily changes of forward variance is indeed stronger on low-volume days than on high-volume days. Market illiquidity is thus likely an important source of the negative autocorrelation in daily changes of forward variance. More importantly, is the improved liquidity on high-volume days sufficient to eliminate negative autocorrelation altogether? If it is, market illiquidity is clearly the source of the documented negative autocorrelation. Indeed, the autocorrelation coefficients on high-volume days (ϕ_i) are all quite small and none of them is statistically significant at any conventional significance level. We thus find no evidence of any significant autocorrelation on high-volume days. The random walk test cannot reject the orthogonality condition on high-volume days. This evidence lends support to a limits-to-arbitrage explanation of rational expectations.

Table V about here

To ensure the robustness of our findings in Table V, we also rerun the dummy variable regression (8) using alternative measures of market illiquidity. We use dollar value of trade (i.e., option volume multiplied by option price) instead of trading volume, define high-volume (low-volume) days as trading days with total trading volume in the first (last) quartile, or replace the dummy variable in regression (8) with the log trading volume. Untabulated results indicate that our findings in Table V are robust to alternative measures of trading volume, alternative definitions of high- vs. low-volume days, and alternative regression specifications. We are thus confident that negative autocorrelation in daily changes in forward variance is driven by market illiquidity.

We next investigate whether or not negative autocorrelation is persistent over time. As market liquidity tends to improve over time, the magnitude of the negative autocorrelation may become weaker over time. If that is indeed the case, we have further support for market illiquidity as the source of negative autocorrelation. Following Lo and MacKinlay (1988) and other related studies, we divide the full sample period into two non-overlapping subperiods. The first subperiod is just short of ten years and covers the period between June 1988 and December 1997, while the second subperiod is exactly ten years and covers the period between January 1998 and December 2007. The average daily trading volume is 57,611 contracts in the first subperiod, compared to 207,569 contracts in the second subperiod. The 260% increase in daily trading volume suggests a dramatic improvement in market liquidity in the second subperiod. We thus expect to find weaker autocorrelation in the second subperiod than in the first subperiod.

We rerun regression (7) separately for each subperiod and report the results in Table VI. It is clear that the same, albeit stronger, pattern of negative autocorrelation is observed in the first subperiod as in the full sample period. The first lag coefficients vary from -0.23 to -0.38, much larger than the corresponding coefficients in the full sample period. All of them are significant at the 1% level. The second lag coefficients are also all negative and significant at the 1% level in three of the four maturity groups. In the last maturity group, even the third lag coefficient is negative and significant at the 5% level. In contrast, no significant autocorrelation is found at all in the second subperiod. The coefficients of the lagged daily changes (ϕ_i) are all quite small and none of them is statistically significant at any conventional significance level. This is a striking result as there is no evidence at all against orthogonality in the second subperiod. Whatever factors at play that have lead to negative autocorrelation in the first subperiod, they are entirely absent in the second subperiod. This finding thus suggests that violations of orthogonality are not persistent over time and our model-free tests find no such violation in the second half of the sample period.

Table VI about here

Results in Tables V and VI establish a strong linkage between market illiquidity and negative autocorrelation in daily changes in forward variance. When the market is sufficiently liquid (e.g., on high-volume days or in the second half of the sample period), no significant autocorrelation is present at all. Negative autocorrelation is detected only when market liquidity is lacking (e.g., on low-volume days or in the first half of the sample period). Such a strong linkage means that we cannot simply attribute the anomalous pattern to investor irrationality. Perhaps, investors are entirely rational in forming their expectations on future volatility and fully aware of the pattern of negative autocorrelations. Market illiquidity prevents investors from executing trades on their expectations, keeping option prices from converging to fundamental values. This interpretation is also consistent with the empirical findings by Cao, Li and Yu (2005) who document that trading strategies based on certain implied variance anomalies are not profitable after transaction costs.

V. Tests on the Variance Term Structure

Option implied variances (either spot or forward variance) across maturities do not move independently of one another. Long-term implied variance should neither move in perfect lockstep with short-term implied variance nor independently of it (e.g., Stein (1989)). The dynamic properties of variance term structure can thus provide insight on rational expectations and the informational efficiency of the options market. Previous attempts to perform such tests are flawed as they all use model-based implied variance to approximate the variance term structure. In this section, we examine movements of the variance term structure using model-free implied variance.

A. Dynamic Behavior of the Variance Term Structure

The martingale restriction in Eq. (3) also has important implications on the dynamic behavior of variance term structure. The daily change in forward variance ($\Delta v(t; T_1, T_2)$) reflects the market consensus revision to projected future variance based on newly arrived information on day t. This revision to variance term structure at the long-term end (i.e., $[T_1, T_2]$) should not be predictable by any previously known information on day $t-1$. In particular, the daily change in forward variance should not be predictable by either lagged spot variance $v(t-1, T_1)$ or lagged variance spread ($v(t-1, T_2) - v(t-1, T_1)$) which contains previously known information on the level or slope of variance term structure. A simple test of rational expectations on variance term structure follows directly from either the lagged spot variance regression or the lagged variance spread regression:

$$\Delta v(t; T_1, T_2) = \alpha + \beta v(t-1, T_1) + \varepsilon_t, \tag{9}$$

$$\Delta v(t; T_1, T_2) = \alpha + \beta [v(t, T_2) - v(t-1, T_1)] + \varepsilon_t, \tag{10}$$

In either case, the null hypothesis is that the slope coefficient is zero ($H_0 : \beta = 0$).

Using our sample of SPX options, we perform regressions (9) and (10) separately for the four forward-maturity groups and summarize the results in Table VII. Results for the two subperiods are similar and thus omitted for brevity. The standard errors of the estimated coefficients are shown in brackets and are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (see Newey and West (1987)). The last column in the table reports the results from the χ^2 test for the joint hypothesis $H_0 : \alpha = \beta = 0$, with the p-value shown in brackets.

As shown in Table VII, the β coefficients are quite small in all maturity groups for both regressions. None of them is statistically significant at any conventional significance level. In fact, the χ^2 test does not reject the joint hypothesis $H_0 : \alpha = \beta = 0$ at the 5% significance level in any of maturity groups. The null hypothesis is marginally rejected at the 10% level in some cases. This is true for both the lagged spot variance regression (Panel A) and the lagged variance spread regression (Panel B). These results thus do not reject the orthogonality condition and provide support for rational expectations of the variance term structure.

Table VII about here

B. Misreaction or Misspecification? A re-examination of previous empirical tests

Existing studies have also examined rational expectations on the variance term structure. Some of these studies present evidence against rational expectations (e.g., Stein (1989)) while others do not (e.g., Campa and Change (1995)). Our model-free tests in the previous subsection (see Table VII) show that changes in forward variance are not predictable by either the short variance or the variance spread, supporting rational expectations on the dynamic behavior of the variance term structure. How do we reconcile the conflicting findings and what are the sources of discrepancies? We argue and present evidence subsequently that model misspecification in previous studies is the reason behind the conflicting findings.

Unlike our model-free approach, empirical tests in previous studies suffer from at least two types of model specification errors. The first type of specification errors results from the use of implied volatility based on a specific option-pricing model (the B-S model in most cases). If the assumed option-pricing model is misspecified, systematic biases may be impounded into implied volatilities leading to spurious patterns in empirical tests of the orthogonality condition. Any test based on the term structure of the B-S implied volatilities is thus potentially flawed, leading to erroneous and conflicting findings (e.g., Stein

(1989) and Campa and Chang (1995)).

The second type of specification errors arises from assumptions about the variance term structure relationship. In addition to using the B-S implied volatility, these studies also make further assumptions about the underlying variance process (e.g., AR(1)) which then lead to a particular relationship between long-term and short-term variances. This model-specific term structure relationship is then tested to draw inferences on market efficiency. Problems with this type of specification errors are obvious. As demonstrated by Heynen, Kemna and Vorst (1994), the same empirical test can lead to contradictory results if different assumptions about the underlying variance process are made. We thus focus on the first type of specification errors (i.e., those due to the use of model-based implied variance/volatility).

We first show that the most commonly used empirical test in previous studies can detect a seemingly anomalous pattern of "market overreaction" due to model specification errors. These tests originate from Stein (1989) who finds that long-term variance overreacts to shocks to the short-term variance. In particular, Stein (1989) investigates the correlation between the short-term variance and the error of using the "forward variance" to forecast the future spot variance. He finds that the variance forecast error is negatively correlated with the short-term variance and interprets it as evidence of market overreaction. However, his measure of forward variance is misspecified as it is based on the assumption of a linear variance term structure and constructed from the B-S implied variances at the short and long maturities. The potential specification errors render such tests invalid as they are joint tests of market efficiency and the assumed models. We now correct this misspecification problem by re-examining the same relationship using the model-free forward variance and compare the results to those from the corresponding model-based test in Stein (1989).

Recall that the regression test in Stein (1989) is as follows:

$$\Delta\widehat{v}_t(T_1, T_2) = \alpha + \beta v(t, T_1) + \varepsilon_t, \tag{11}$$

where

$$\Delta\widehat{v}_t(T_1, T_2) = v(T_1, T_2) - \widehat{v}(t; T_1, T_2), \tag{12}$$

$$\widehat{v}(t; T_1, T_2) = v(t, T_1) + \left(\frac{T_2 - t}{T_2 - T_1}\right)[v(t, T_2) - v(t, T_1)]$$

and $v(t,T)$ is spot variance approximated using the B-S implied variance (from the date t market price of at-the-money option maturing on date T), $\widehat{v}(t;T_1,T_2)$ is a proxy for forward variance constructed from the corresponding long-term and short-term spot variance, and $\Delta\widehat{v}_t(T_1,T_2)$ is the error of using the forward variance proxy to forecast the future spot variance. A negative β coefficient suggests negative correlation between variance forecast error and the short-term variance.

Note that the proxy for forward variance constructed this way is precisely correct only if the two spot variances are measured using the model-free spot variance. This is because both model-free spot and forward variances are integrated variance and it is straightforward to show that model-free forward variance is related to model-free spot variance as follows:

$$v(t;T_1,T_2) = v(t,T_1) + \left(\frac{T_2-t}{T_2-T_1}\right)[v(t,T_2)-v(t,T_1)]$$

The proxy for forward variance is thus appropriate if model-free spot variance is used to construct it.

In contrast, the proxy for forward variance is not correct when the B-S implied variance is used to construct it. Regression (11) is thus misspecified if the B-S implied variance is used as input, which is what previous studies (e.g., Stein (1989)) have done. Even though a negative relationship (i.e., negative β) is found in these studies, it is unclear whether it is due to misspecification or market misreaction.

To provide new insight on this relationship, we replicate Stein's (1989) test using our sample of SPX options and compare with the corresponding model-free test when the B-S implied variance is replaced with the model-free implied variance. In other words, we implement regression (11) using both the model-free implied variance and the B-S implied variance. The results are summarized in Table VIII, with the model-free results in Panel A and model-based results in Panel B.

We first examine the results from the model-based test reported in Panel B of Table VIII. The β coefficient is negative in all four forward-maturity groups and significant at the 5% level in two of the four cases, including the first forward-maturity group used in Stein (1989). The χ^2 test strongly rejects the joint hypothesis $H_0 : \alpha = \beta = 0$, with the p-value virtually zero in all four maturity groups. These results are consistent with the findings in Stein (1989) who interprets it as evidence of market overreaction. In both studies, the statistical significance of β is not particularly strong, never significant at the 1% level.

In contrast, the β coefficient is consistently positive in the model-free test as reported in Panel A of Table VIII. In all four forward-maturity groups, the null hypothesis that the β coefficient is zero cannot be rejected at any conventional significance level. In addition, the χ^2 test does not come close to rejecting the joint hypothesis $H_0 : \alpha = \beta = 0$ in three of the four maturity groups. The null hypothesis is marginally rejected at the 10% level in the first maturity group. The model-free tests thus do not detect any evidence of market misreaction at all.

Combining results from both model-free and model-based tests, it is clear that the model-based test can erroneously reject the null hypothesis when the model-free test does not. This finding casts doubt on the validity of the market overreaction anomaly identified in previous studies. When a model-based test finds evidence of negative correlation, it does not necessarily suggest market overreaction. Model-free tests should be used to corroborate the findings. One must be cautious when drawing conclusions from the results of model-based tests.

Table VIII about here

An interesting question is why the use of the B-S implied variance leads to negative correlation in regression (11). To answer this question, we recall a well-known empirical stylized fact that the B-S implied volatility exhibits a skew or smirk pattern for stock index options. This pattern can be much more pronounced for short-term options than for long-term options. A direct implication of the skew or smirk is that at-the-money implied volatility underestimates the true volatility, as corroborated by recent empirical research (e.g., Jiang and Tian (2005) and Carr and Wu (2009)). In addition, B-S implied volatilities from short-maturity options are also more sensitive to volatility shocks than their counterparts from long-maturity options. When the underlying volatility is high, a more pronounced skew of short-maturity options leads to a larger underestimation of the true volatility by short-term at-the-money B-S implied volatility than long-term at-the-money B-S implied volatility. In other words, long-term implied volatility is "too high" when short-term implied volatility is high. This is precisely the type of "market overreaction" documented by Stein (1989). As a result, the misspecified B-S model alone can lead to an incorrect diagnosis of market overreaction.

To quantify the impact of specification errors induced by the use of the B-S implied volatility, we examine the relationship between long-term and short-term B-S implied volatilities in the presence of model specification errors. In particular, we consider a scenario where the B-S model is assumed to be the true model but the underlying asset is in fact described by the *Stochastic Volatility with Jumps* (SVJ) model:

$$\begin{aligned} dS_t/S_t &= (r - \lambda_0 \mu_J)dt + \sqrt{V_t} dW_t^S + J_t dQ_t(\lambda_0) \\ dV_t &= \kappa^*(\theta^* - V_t)dt + \sigma \sqrt{V_t} dW_t^V \\ dW_t^S dW_t^V &= \rho dt \end{aligned} \tag{13}$$

where W_t^S and W_t^V are correlated standard Brownian motions, V_t is the instantaneous variance, $Q_t(\lambda_0)$ is a Poisson process with jump intensity λ_0, and the jump size follows an i.i.d. log-normal distribution $\ln(1+J_t) \sim N\left(\ln(1+\mu_J) - \frac{1}{2}\sigma_J^2, \sigma_J^2\right)$. The SVJ model is used for three reasons. First, previous empirical research (e.g. Bakshi, Cao and Chen (1997) and Pan (2002)) suggests that the SVJ model fits SPX options data very well. Secondly, the SVJ model is a Markov process and consistent with rational expectations. In other words, it does not internally generate any misreaction of long-term variance relative to changes in short-term variance. Finally, the SVJ model contains the B-S model as a special case and thus provides a consistent framework for analyzing model misspecification.

We calibrate the SVJ model to SPX option prices observed on two representative dates, May 14, 1991 and December 28, 1990, in our sample. On these two dates, the SPX options market exhibits relatively low (18%) and high (25%) volatility levels as suggested by the one-month model-free implied volatility on those days. These volatility levels are well within the 10%–40% range indicated by the volatility plot in Fig. 1. To be consistent with Stein (1989), we consider options with two-week (short-term) and six-week (long-term) maturities which are roughly the average maturities of the nearby and next maturity contracts. We then simulate option prices using the calibrated model with either an 18% or 25% instantaneous volatility.

The B-S implied volatilities are then calculated from the simulated SVJ option prices and plotted in Panels A and C of Fig. 3. From these plots, it is clear that at-the-money (moneyness = 1) B-S implied volatility underestimates the true volatility and the underestimation is more severe when volatility is high. More importantly, when volatility is high (at 25%) long-term at-the-money B-S implied volatility is much higher relative to short-term at-the-money B-S implied volatility than when volatility is low (at 18%). This is clearly illustrated in Panels B and D of Fig. 3 which plot the difference in B-S implied volatilities between long-term and short-term options at the same moneyness levels. For at-the-money options, the difference is positive but virtually zero when volatility is low (at 18% in Panel B). In contrast, the difference can be quite large (and positive) when volatility is high (at 25% in Panel D). This is precisely the pattern reported in Stein (1989) that long-term implied volatility is "too high" when the underlying volatility is high.

Fig. 3 about here

VI. Conclusions

Though an important tool in previous research on stock and bond market efficiency, random walk tests have never been applied in the options market before. In this paper, we perform such tests in the SPX options market and find support for informational efficiency and the expectations hypothesis. Although we do detect an anomalous pattern of negative autocorrelation in daily changes of forward variance, it is fully explained by an illiquidity effect. In particular, negative autocorrelation is completely absent when the market is sufficiently liquid (e.g., on high-volume days or in the second half of the sample period). We also rule out mean reversion in daily forward variance as the explanation for the documented negative autocorrelation since there is no evidence that daily forward variance is mean reverting. In contrast, daily spot variance is mean reverting but we find no significant negative autocorrelation in daily changes in spot variance.

Our study makes several contributions to the literature. First of all, we draw insight from Britten-Jones and Neuberger (2000) and implement a model-free approach to the construction of forward variance from market prices of options across two maturities. This model-free forward variance allows us to perform random walk tests of the expectations hypothesis in the options market. Testable hypotheses are developed to examine the orthogonality condition that daily changes in forward variance are not serially correlated. Such tests are not possible with variance measures commonly used in prior research such as realized variance, the B-S implied variance, or the VIX-type model-free spot variance. These are spot variance instead of forward variance and do not have to satisfy the martingale restriction or orthogonality condition.

Secondly, we perform new model-free tests on the dynamic behavior of variance term structure and examine how long-term variance reacts to changes in short-term variance. We find no evidence of market misreaction, supporting rational expectations of the variance term structure. This is in sharp contrast to previous studies (e.g., Stein (1989)) that perform model-based tests and find overreactions of long-term variance relative to changes in short-term variance. While we find similar evidence of market overreaction if we replicate their model-based tests, no such evidence is found if we use the model-free approach to perform the same tests. We thus argue that model misspecification is behind the false evidence of market overreaction and has lead to conflicting findings in previous studies (e.g., Stein (1989), Heynen, Kemna and Vorst (1994) and Campa and Chang (1995)).

Finally, model-free forward variance may provide an opportunity for further research on volatility modeling, forecasting and other related issues. For example, does forward variance contain useful information for realized variance or implied spot variance? If it does, can it be used to forecast future realized variance or future implied spot variance? Are forward variance, realized variance and implied spot variance driven by the same underlying factors? Although we already find evidence that daily forward variances are not mean reverting while daily spot variances are, are there any other differences in the dynamic properties of forward variance, implied spot variance and realized variance? The potential applications are likely to expand as interest in model-free forward variance grows in the future.

References

Ackert, Lucy, and Yisong S. Tian, 2001, Efficiency in index options markets and trading in stock baskets, *Journal of Banking and Finance* 25, 1607-1634.

Aït-Sahalia, Yacine, and Jefferson Duarte, 2003, Nonparametric option pricing under shape restrictions, *Journal of Econometrics* 116, 9-47.

Aït-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.

Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2002, An empirical investigation of continuous-time equity return models, *Journal of Finance* 57, 1239-1284.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003-2049.

Bakshi, Gurdip., Charles Cao, and Zhiwu Chen, 2000, Pricing and hedging long-term options, *Journal of Econometrics* 94, 277-318.

Bates, David S., 1996, Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche mark options, *Review of Financial Studies* 9, 69-107.

Bates, David S., 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181-238.

Bates, David S., 2006, Maximum likelihood estimation of latent affine processes, *Review of Financial Studies* 19, 909-965.

Bekaert, Geert, and Robert J. Hodrick, 1992, Characterizing predictable components in excess returns on equity and foreign exchange markets, *Journal of Finance* 47, 467-509.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637-659.

Brav, A., and J.B. Heaton, 2002, Competing theories of financial anomalies, *Review of Financial Studies* 15, 575-606.

Britten-Jones, Mark, and Anthony Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *Journal of Finance* 55, 839-866.

Campa, José M., and P.H. Kevin Chang, 1995, Testing the expectations hypothesis on the term structure of volatilities in foreign exchange options, *Journal of Finance* 50, 529-547.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, The econometrics of financial markets, *Princeton University Press*, Princeton, New Jersey.

Campbell, John Y., and Robert Shiller, 1984, A simple account of the behavior of long-term interest rates, *American Economic Review* Papers and Proceedings 74, 44-48.

Campbell, John Y., and Robert Shiller, 1991, Yield spreads and interest rate movements: a bird's eye view, *Review of Economic Studies* 58, 495-514.

Cao, Charles, Haitao Li, and Fan Yu, 2005, Is investor misreaction economically significant? Evidence from short- and long-term S&P 500 index options, *Journal of Futures Markets* 25, 717-752.

Carr, Peter, and Dilip B. Madan, 1998, Towards a theory of volatility trading. In: Jarrow, R. (eds), *Risk Book on Volatility Risk*, New York.

Carr, Peter, and Dilip B. Madan, 2005, A note on sufficient conditions for no arbitrage, *Finance Research Letters* 2, 125-130.

Carr, Peter, and Liuren Wu, 2003, What type of process underlies options? A simple robust test, *Journal of Finance* 58, 2581–2610.

Carr, Peter, and Liuren Wu, 2009, Variance risk premia, *Review of Financial Studies* 22, 1311-1341.

Chance, Don M., 1987, Parity tests of index options, *Advances in Futures and Options Research* 2, 47-64.

Chen, Joseph, Harrison Hong, and Jeremy C. Stein, 2001, Forecasting crashes: trading volume, past returns, and conditional skewness in stock prices, *Journal of Financial Economics* 61, 345-381.

Chernov, Mikhail, A. Ronald Gallant, Eric Ghysels, and George Tauchen, 2003, Alternative models for stock price dynamics, *Journal of Econometrics* 116, 225–257.

Christensen, B.J., C.S. Hansen, and N.R. Prabhala, 2001, The telescoping overlap problem in options data, Working paper, *University of Aarhus* and *University of Maryland*.

Christensen, B.J., and N.R. Prabhala, 1998, The relation between implied and realized volatility, *Journal of Financial Economics* 50,125-150.

Christoffersen, Peter, Kris Jacobs, and Gregory Vainberg, 2008, Forward-looking betas, Working paper, *McGill University*.

Cox, John, Jonathan Ingersoll, and Stephen Ross, 1981, A reexamination of traditional hypotheses about the term structure of interest rates, *Journal of Finance* 36, 321-346.

Demeterfi, Kresimir, Emanuel Derman, Michael Kamal, and Joseph Zou, 1999, A guide to volatility and variance swaps, *Journal of Derivatives* 6, 9-32.

Donaldson, R. Glen, and Mark J. Kamstra, 2005, Volatility forecasts, trading volume and the ARCH vs. option-implied volatility tradeoff, *Journal of Financial Research* 27, 519-538.

Duffie, Darrell, Jun Pan, Kenneth J. Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

Dupire, Bruno, 1994, Pricing with a smile, *Risk* 7, 18-20.

Eraker, Bjorn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in equity index volatility and returns, *Journal of Finance* 58, 1269–1300.

Fama, Eugene, and Kenneth French, 1988, Permanent and temporary components of stock prices, *Journal of Political Economy* 96, 246-273.

Gervais, Siman, Ron Kaniel, and Dan H. Mingelgrin, 2001, The High-volume return premium, *Journal of Finance* 56, 877-919.

Harvey, Campbell R., and Robert E. Whaley, 1991, S&P 100 index option volatility, *Journal of Finance* 46, 1551-1561.

Harvey, Campbell R., and Robert E. Whaley, 1992a, Market volatility prediction and the efficiency of the S&P 100 index option market, *Journal of Financial Economics* 31, 43-73.

Harvey, Campbell R., and Robert E. Whaley, 1992b, Dividends and S&P 100 index option valuation, *Journal of Futures Markets* 12, 123-137.

Heynen, Ronald, Angelien Kemna, and Ton Vorst, 1994, Analysis of the term structure of implied volatilities, *Journal of Financial and Quantitative Analysis* 29, 31-56.

Jiang, George J., and Yisong S. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305-1342.

Jiang, George J., and Yisong S. Tian, 2007, Extracting model-free volatility from option prices: An examination of the VIX index, *Journal of Derivatives* 14, 35-60.

Kamara, Avraham, and Thomas W. Miller, Jr., 1995, Daily and intraday tests of European put-call parity, *Journal of Financial and Quantitative Analysis* 30, 519-539.

Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357-390.

Lewellen, J., and J. Shanken 2002, Learning, asset-pricing tests, and market efficiency, *Journal of Finance* 57, 1113-1145.

Lo, Andrew W., and A. Craig Mackinlay, 1988, Stock market prices do not follow random walks: evidence from a simple specification test, *Review of Financial Studies* 1, 41-66.

Lo, Andrew W., and A. Craig Mackinlay, 1990, When are contrarian profits due to stock market overreaction? *Review of Financial Studies* 3, 175-206.

Longstaff, Francis A., 1995, Option pricing and the martingale restriction, *Review of Financial Studies* 8, 1091-1124.

Longstaff, Francis A., 2000, The term structure of very short-term rates: new evidence for the expectations hypothesis, *Journal of Financial Economics* 58, 397-415.

Neuberger, Anthony, 1994, The log contract, *Journal of Portfolio Management* 20, 74-80.

Newey, W.K., and K.D. West, 1987, A simple positive definite heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Pool, Veronika, 2005, Is heavy trading good or bad for price discovery? Working paper, *Vanderbilt University*.

Poterba, James M., and Lawrence H. Summers, 1986, The persistence of volatility and stock market fluctuations, *American Economic Review* 76, 1142-1151.

Scholes, Myron, and Joseph T. Williams, 1977, Estimating betas from nonsynchronous data, *Journal of Financial Economics* 5, 309-328.

Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35-55.

Stein, Jeremy, 1989, Overreaction in the options market, *Journal of Finance* 44, 1011-1023.

Xu, Xinzhong, and Stephen J. Taylor, 1994, The term structure of volatility implied by foreign exchange options, *Journal of Financial and Quantitative Analysis* 29, 57-74.

Table I

Summary Statistics of the Model-Free Implied Variance

Panels A and B report the descriptive statistics of the daily model-free implied spot and forward variances for different maturity groups. The first maturity group of the spot variance consists of options with the nearest maturity month, the second maturity group consists of options with the next maturity month, and so on. Mean and median number of days to maturity ($T_i, i = 1, \cdots, 5$) are also reported for each maturity group. The table also reports the first five autocorrelations of the daily model-free implied variance.

Panel A. Spot Variance

	Maturity group				
	[0, T_1]	[0, T_2]	[0, T_3]	[0, T_4]	[0, T_5]
Mean, Median (T_i)	21, 21	53, 51	97, 85	171, 165	255, 254
N	4930	4930	4930	4925	4705
A.1: Moments					
Mean	0.0372	0.0379	0.0386	0.0391	0.0403
Median	0.0289	0.0301	0.0326	0.0332	0.0341
St.dev.	0.0278	0.0253	0.0234	0.0213	0.0199
Skewness	2.3287	1.8866	1.5879	1.4731	1.3779
Kurtosis	9.2825	6.0580	3.6573	3.2570	2.9020
A.2: Autocorrelations					
1	0.9266	0.9591	0.9833	0.9832	0.9743
2	0.9014	0.9425	0.9713	0.9734	0.9665
3	0.8817	0.9315	0.9611	0.9642	0.9566
4	0.8632	0.9209	0.9506	0.9552	0.9499
5	0.8480	0.9095	0.9413	0.9473	0.9394

Panel B. Forward Variance

	Maturity group			
	I	II	III	IV
	[T_1, T_2]	[T_1, T_3]	[T_1, T_4]	[T_1, T_5]
N	4927	4927	4921	4653
B.1: Moments				
Mean	0.0387	0.0392	0.0395	0.0406
Median	0.0293	0.0312	0.0330	0.0345
St.dev.	0.0244	0.0227	0.0208	0.0196
Skewness	1.5628	1.5089	1.4198	1.3564
Kurtosis	3.2451	3.3964	2.9813	2.7757
B.2: Autocorrelations				
1	0.9867	0.9842	0.9855	0.9782
2	0.9655	0.9702	0.9701	0.9645
3	0.9549	0.9619	0.9615	0.9530
4	0.9453	0.9514	0.9527	0.9466
5	0.9361	0.9433	0.9455	0.9367

Table II
Summary Statistics of Daily Change in Forward Variance

This table reports the mean and standard error of daily change in forward variance ($\Delta v(t; T_1, T_2) = v(t; T_1, T_2) - v(t-1; T_1, T_2)$ where $t < T_1 < T_2$, T_1 and T_2 are option maturity dates). The standard errors are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (see Newey and West (1987)) and reported in the brackets next to the coefficient estimates. N denotes the number of observations.

Maturity group	N	Mean of Daily Change in Forward Variance (10^{-5})	t-Statistic
I	4927	-3.8538 (6.2501)	-0.6165
II	4927	4.0831 (5.2987)	0.7706
III	4921	-2.6902 (4.8463)	-0.5551
IV	4653	1.3147 (5.1219)	0.2567

Table III

Regression of Daily Change in Forward Variance against Lagged Forward Variance

This table reports the OLS results of the following regression:

$$\Delta v(t;T_1,T_2) = \alpha + \beta v(t-1;T_1,T_2) + \varepsilon_t,$$

where $t < T_1 < T_2$, T_1 and T_2 are option maturity dates, and $\Delta v(t;T_1,T_2) = v(t;T_1,T_2) - v(t-1;T_1,T_2)$ is the daily change in forward variance. The standard errors of the coefficient estimates are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (see Newey and West (1987)) and reported in the brackets next to the coefficient estimates. The χ^2 test is for the joint hypothesis $H_0 : \alpha = \beta = 0$ with p-value in the brackets. The ** and * indicate α or β significantly different from zero at the 1% and 5% level, respectively. D-W denotes the Durbin-Watson statistic.

Maturity group	N	α	β	Adj. R^2	D-W	χ^2 test
I	4927	0.000 (0.000)	-0.013 (0.011)	0.002	2.120	3.307 (0.193)
II	4927	0.000 (0.000)	-0.016 (0.012)	0.001	2.216	3.954 (0.138)
III	4921	0.000 (0.000)	-0.014 (0.009)	0.002	2.135	3.730 (0.155)
IV	4653	0.000 (0.000)	-0.022 (0.013)	0.003	2.019	4.110 (0.128)

Table IV

Auto-Regression of Daily Change in Forward Variance

This table reports the OLS results of the following regression for the full sample period (June 1988 - December 2007):

$$\Delta v(t;T_1,T_2) = \alpha + \sum_{i=1}^{k} \phi_i \, \Delta v(t-i;T_1,T_2) + \varepsilon_t,$$

where $t < T_1 < T_2,\ T_1$ and T_2 are option maturity dates, and $\Delta v(t-i;T_1,T_2) = v(t-i;T_1,T_2) - v(t-1-i;T_1,T_2)$ $(0 \leq i \leq k)$ are the current and lagged daily changes in forward variance. The regressions are run with 25 lags with the first ten ϕ coefficients reported in the table and all 25 ϕ coefficients plotted in Figure 2, together with a 95% confidence band. The standard errors of the coefficient estimates are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (see Newey and West (1987)) and reported in the brackets next to the coefficient estimates. The ** and * indicate α or ϕ_i $(1 \leq i \leq 10)$ significantly different from zero at the 1% and 5% level, respectively. D-W denotes the Durbin-Watson statistic.

	Maturity group							
	I		II		III		IV	
α	-0.0006	(0.6028)	0.0001	(0..0001)	-0.0000	(0..0000)	0.0000	(0.0001)
ϕ_1	-0.0781*	(0.0319)	-0.1974**	(0..0421)	-0.1906**	(0..0370)	-0.2691**	(0.0406)
ϕ_2	-0.0191	(0.0218)	-0.0725*	(0..0325)	-0.0730*	(0..0320)	-0.1349**	(0.0350)
ϕ_3	-0.0268	(0.0249)	0.0061	(0..0252)	-0.0024	(0..0234)	-0.0153	(0.0304)
ϕ_4	-0.0230	(0.0245)	-0.0135	(0..0282)	-0.0241	(0..0247)	0.0025	(0.0328)
ϕ_5	-0.0099	(0.0214)	0.0118	(0..0292)	0.0369	(0..0237)	-0.0103	(0.0312)
ϕ_6	-0.0146	(0.0208)	-0.0170	(0..0270)	0.0019	(0..0224)	-0.0423	(0.0279)
ϕ_7	-0.0280	(0.0194)	-0.0257	(0..0247)	-0.0101	(0..0227)	-0.0040	(0.0274)
ϕ_8	0.0062	(0.0228)	0.0236	(0..0264)	-0.0318	(0..0266)	-0.0171	(0.0317)
ϕ_9	0.0070	(0.0243)	0.0171	(0..0260)	0.0303	(0..0233)	0.0424	(0.0285)
ϕ_{10}	0.0321	(0.0212)	0.0157	(0..0211)	-0.0277	(0..0208)	0.0243	(0.0295)
N	4927		4927		4921		4653	
Adj. R^2	0.011		0.044		0.046		0.154	
D-W	2.000		2.001		2.001		1.996	

Table V

The Effect of Illiquidity

The table reports the OLS results of the following regressions:

$$\Delta v(t;T_1,T_2) = \alpha + \sum_{i=1}^{k} (\phi_i + \gamma_i\, D_{t-i}) \cdot \Delta v(t-i;T_1,T_2) + \varepsilon_t,$$

where $t < T_1 < T_2$, T_1 and T_2 are option maturity dates, and $\Delta v(t-i;T_1,T_2) = v(t-i;T_1,T_2) - v(t-1-i;T_1,T_2)$ $(0 \le i \le k)$ are the current and lagged daily changes in forward variance, D_t is a dummy variable that takes on a value of 1 if day t is a low-volume day and zero otherwise. The regressions are run with 25 lags with the first ten ϕ and γ coefficients reported in the table. The standard errors of the coefficient estimates are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (see Newey and West (1987)) and reported in the brackets next to the coefficient estimates. The ** and * indicate $\phi_i,\ \gamma_i$ $(1 \le i \le 10)$ significantly different from zero at the 1% and 5% level, respectively. D-W denotes the Durbin-Watson statistic.

	Maturity group							
	I		II		III		IV	
α	-0.0001	(0.0001)	0.0001	(0.0001)	-0.0000	(0.0000)	0.0000	(0.0001)
ϕ_1	-0.0253	(0.0637)	-0.0535	(0.0835)	-0.1179	(0.0651)	-0.1173	(0.0608)
γ_1	-0.0834*	(0.0424)	-0.1907**	(0.0565)	-0.1455**	(0.0429)	-0.2800**	(0.0564)
ϕ_2	0.0016	(0.0301)	-0.0242	(0.0334)	0.0477	(0.0562)	-0.0502	(0.0474)
γ_2	-0.0564	(0.0458)	-0.0902	(0.0557)	-0.0878*	(0.0370)	-0.1007*	(0.0469)
ϕ_3	0.0278	(0.0482)	0.0400	(0.0436)	0.0181	(0.0444)	0.0413	(0.0632)
γ_3	-0.0379	(0.0328)	-0.0206	(0.0350)	-0.0252	(0.0348)	-0.0467	(0.0478)
ϕ_4	-0.0104	(0.0443)	0.0018	(0.0519)	0.0694	(0.0452)	0.0046	(0.0638)
γ_4	-0.0123	(0.0318)	-0.0140	(0.0355)	-0.0527	(0.0366)	0.0089	(0.0562)
ϕ_5	0.0534	(0.0457)	-0.0245	(0.0591)	0.0383	(0.0344)	-0.0447	(0.0579)
γ_5	-0.0193	(0.0316)	0.0258	(0.0432)	0.0492	(0.0358)	0.0194	(0.0423)
ϕ_6	0.0334	(0.0430)	-0.0407	(0.0553)	-0.0174	(0.0450)	-0.0195	(0.0546)
γ_6	-0.0331	(0.0278)	0.0098	(0.0368)	-0.0609	(0.0445)	-0.0473	(0.0395)
ϕ_7	-0.0142	(0.0393)	-0.0389	(0.0548)	-0.0097	(0.0507)	0.0033	(0.0551)
γ_7	-0.0217	(0.0277)	-0.0083	(0.0425)	-0.0178	(0.0341)	-0.0212	(0.0391)
ϕ_8	0.0360	(0.0442)	0.0057	(0.0486)	0.0040	(0.0457)	-0.0368	(0.0530)
γ_8	-0.0181	(0.0315)	0.0183	(0.0293)	-0.0256	(0.0319)	-0.0010	(0.0374)
ϕ_9	0.0157	(0.0404)	-0.0145	(0.0440)	0.0032	(0.0277)	-0.0156	(0.0533)
γ_9	0.0072	(0.0275)	0.0277	(0.0286)	0.0471	(0.0431)	0.0472	(0.0378)
ϕ_{10}	0.0408	(0.0365)	0.0111	(0.0364)	0.0182	(0.0305)	0.0337	(0.0562)
γ_{10}	0.0110	(0.0279)	0.0180	(0.0251)	-0.0342	(0.0365)	-0.0008	(0.0486)
N	4927		4927		4921		4653	
Adj. R^2	0.009		0.043		0.051		0.143	
D-W	1.998		1.997		1.997		2.000	

Table VI

Auto-Regression of Daily Change in Forward Variance for Two Subperiods

The table reports the OLS results of the following regression over two subperiods:

$$\Delta v(t;T_1,T_2) = \alpha + \sum_{i=1}^{k} \phi_i \, \Delta v(t-i;T_1,T_2) + \varepsilon_t,$$

where $t < T_1 < T_2$, T_1 and T_2 are option maturity dates, and $\Delta v(t-i;T_1,T_2) = v(t-i;T_1,T_2) - v(t-1-i;T_1,T_2)$ $(0 \le i \le k)$ are the current and lagged daily changes in forward variance. The regressions are run with 25 lags. The standard errors of the coefficient estimates are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (see Newey and West (1987)) and reported in the brackets next to the coefficient estimates. The ** and * indicate α or ϕ_i $(1 \le i \le 10)$ significantly different from zero at the 1% and 5% level, respectively. D-W denotes the Durbin-Watson statistic.

	Maturity group							
	I		II		III		IV	
Panel A: First subperiod (Jun 1988 - Dec 1997)								
α	-0.0000	(0.0001)	0.0001	(0.0001)	-0.0000	(0.0001)	0.0001	(0.0001)
ϕ_1	-0.2291**	(0.0453)	-0.3150**	(0.0470)	-0.3334**	(0.0500)	-0.3759**	(0.0446)
ϕ_2	-0.0442	(0.0313)	-0.1641**	(0.0429)	-0.1582**	(0.0397)	-0.2454**	(0.0475)
ϕ_3	-0.0466	(0.0281)	-0.0347	(0.0348)	-0.0499	(0.0336)	-0.0907*	(0.0414)
ϕ_4	-0.0120	(0.0255)	-0.0472	(0.0333)	-0.0299	(0.0303)	-0.0270	(0.0477)
ϕ_5	0.0042	(0.0274)	-0.0043	(0.0443)	0.0321	(0.0301)	-0.0499	(0.0418)
ϕ_6	-0.0526	(0.0314)	-0.0105	(0.0376)	0.0050	(0.0315)	-0.0700	(0.0387)
ϕ_7	-0.0314	(0.0297)	-0.0404	(0.0284)	0.0023	(0.0329)	-0.0548	(0.0407)
ϕ_8	0.0267	(0.0337)	0.0184	(0.0249)	-0.0348	(0.0323)	-0.0499	(0.0443)
ϕ_9	0.0411	(0.0365)	0.0157	(0.0265)	0.0301	(0.0272)	0.0240	(0.0369)
ϕ_{10}	0.0338	(0.0281)	0.0191	(0.0313)	-0.0405	(0.0231)	0.0299	(0.0392)
N	2415		2415		2409		2141	
R^2	0.058		0.148		0.116		0.217	
D-W	2.006		2.000		2.001		1.998	
Panel B: Second subperiod (Jan 1998 - Dec 2007)								
α	-0.0001	(0.0001)	0.0000	(0.0001)	0.0000	(0.0001)	-0.0000	(0.0000)
ϕ_1	0.0404	(0.0391)	0.0662	(0.0467)	0.0272	(0.0462)	0.0562	(0.0460)
ϕ_2	-0.0351	(0.0290)	-0.0567	(0.0476)	-0.0232	(0.0544)	-0.0216	(0.0497)
ϕ_3	-0.0036	(0.0380)	0.0160	(0.0356)	0.0324	(0.0373)	0.0220	(0.0412)
ϕ_4	-0.0383	(0.0367)	0.0172	(0.0419)	0.0033	(0.0406)	0.0202	(0.0472)
ϕ_5	-0.0273	(0.0306)	-0.0046	(0.0355)	-0.0010	(0.0379)	0.0335	(0.0395)
ϕ_6	0.0224	(0.0283)	-0.0328	(0.0343)	0.0059	(0.0275)	-0.0175	(0.0283)
ϕ_7	-0.0381	(0.0271)	0.0117	(0.0421)	-0.0258	(0.0313)	0.0016	(0.0380)
ϕ_8	0.0032	(0.0313)	-0.0012	(0.0382)	-0.0135	(0.0459)	-0.0248	(0.0383)
ϕ_9	-0.0204	(0.0319)	0.0304	(0.0388)	0.0437	(0.0340)	0.0560	(0.0369)
ϕ_{10}	0.0466	(0.0306)	0.0122	(0.0270)	0.0010	(0.0315)	-0.0103	(0.0319)
N	2512		2512		2512		2512	
Adj. R^2	0.015		0.025		0.019		0.025	
D-W	1.996		1.999		2.000		2.000	

Table VII

Dynamic Behavior of the Variance Term Structure

This table reports the results of the OLS regression of the daily change in forward variance ($\Delta v(t; T_1, T_2)$) against either the lagged spot variance ($v(t-1, T_1)$) or the lagged variance spread ($v(t-1, T_2) - v(t-1, T_1)$):

$$\Delta v(t; T_1, T_2) = \alpha + \beta v(t-1, T_1) + \varepsilon_t$$

$$\Delta v(t; T_1, T_2) = \alpha + \beta(v(t-1, T_2) - v(t-1, T_1)) + \varepsilon_t$$

where $t < T_1 < T_2$, and T_1 and T_2 are option maturity dates. The standard errors of the coefficient estimates are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (see Newey and West (1987)) and reported in the brackets next to the coefficient estimates. The χ^2 test is for the joint hypothesis $H_0 : \alpha = \beta = 0$ with p-value in the brackets. The ** and * indicate α or β significantly different from zero at the 1% and 5% level, respectively. D-W denotes the Durbin-Watson statistic.

Maturity group	N	α	β	Adj. R^2	D-W	χ^2 test
Panel A: The Lagged Spot Variance Regression						
I	4927	-0.000 (0.000)	0.005 (0.004)	0.001	2.150	5.242 (0.073)
II	4927	-0.000 (0.000)	0.006 (0.004)	0.001	2.376	2.002 (0.367)
III	4921	-0.001* (0.000)	0.005 (0.003)	0.001	2.416	2.495 (0.287)
IV	4653	-0.000 (0.000)	0.000 (0.002)	0.000	2.711	0.016 (0.991)
Panel B: The Lagged Variance Spread Regression						
I	4927	0.000 (0.000)	-0.018 (0.029)	0.005	2.183	5.461 (0.065)
II	4927	0.000 (0.000)	-0.010 (0.014)	0.003	2.433	3.661 (0.160)
III	4921	0.000 (0.000)	-0.006 (0.009)	0.002	2.431	4.923 (0.085)
IV	4653	0.000 (0.000)	-0.008 (0.010)	0.001	2.727	1.928 (0.381)

Table VIII

Market Misreaction Tests: Model-Based vs. Model-Free Implied Volatilities

This table reports the OLS estimation results of the following regression:

$$\Delta\hat{v}_t(T_1, T_2) = \alpha + \beta v(t, T_1) + \varepsilon_t$$

where $t < T_1 < T_2$, T_1 and T_2 are option maturity dates, $v(t, T_1)$ is the short-term spot variance, $\Delta\hat{v}_t(T_1, T_2) = v(T_1, T_2) - \hat{v}(t; T_1, T_2)$, and $\hat{v}(t; T_1, T_2) = v(t, T_1) + (T_2 - t)/(T_2 - T_1)[v(t, T_2) - v(t, T_1)]$ is a proxy for forward variance constructed from two spot variances. This proxy is precisely correct if it is constructed using the model-free spot variance, but not so if it is constructed using the Black-Scholes implied variance. Results of the model-free version of the regression are reported in Panel A while the corresponding results for the model-based version are reported in Panel B. The standard errors of the coefficient estimates are computed following a robust procedure taking into account of the heteroscedastic and autocorrelated error structure (see Newey and West (1987)) and reported in the brackets next to the coefficient estimates. The χ^2 test is for the joint hypothesis $H_0 : \alpha = \beta = 0$ with p-value in the brackets. The ** and * indicate α or β significantly different from zero at the 1% and 5% level, respectively. D-W denotes the Durbin-Watson statistic.

Maturity group	N	α	β	Adj. R^2	D-W	χ^2 test
Panel A: Misreaction Tests based on the Model-Free Implied Variance						
I	4927	-0.000 (0.000)	0.005 (0.004)	0.001	2.150	5.242 (0.073)
II	4927	-0.000 (0.000)	0.006 (0.004)	0.001	2.376	2.002 (0.367)
III	4921	-0.001* (0.000)	0.005 (0.003)	0.001	2.416	2.495 (0.287)
IV	4653	-0.000 (0.000)	0.000 (0.002)	0.000	2.711	0.016 (0.991)
Panel B: Misreaction Tests based on the Black-Scholes Implied Variance						
I	4910	-0.001 (0.001)	-0.076* (0.030)	0.011	1.388	19.06 (0.000)
II	4895	-0.003** (0.001)	-0.048* (0.024)	0.005	1.640	54.86 (0.000)
III	4828	-0.004** (0.001)	-0.027 (0.018)	0.002	1.622	98.40 (0.000)
IV	4472	-0.007** (0.001)	-0.038 (0.026)	0.002	1.543	106.6 (0.000)

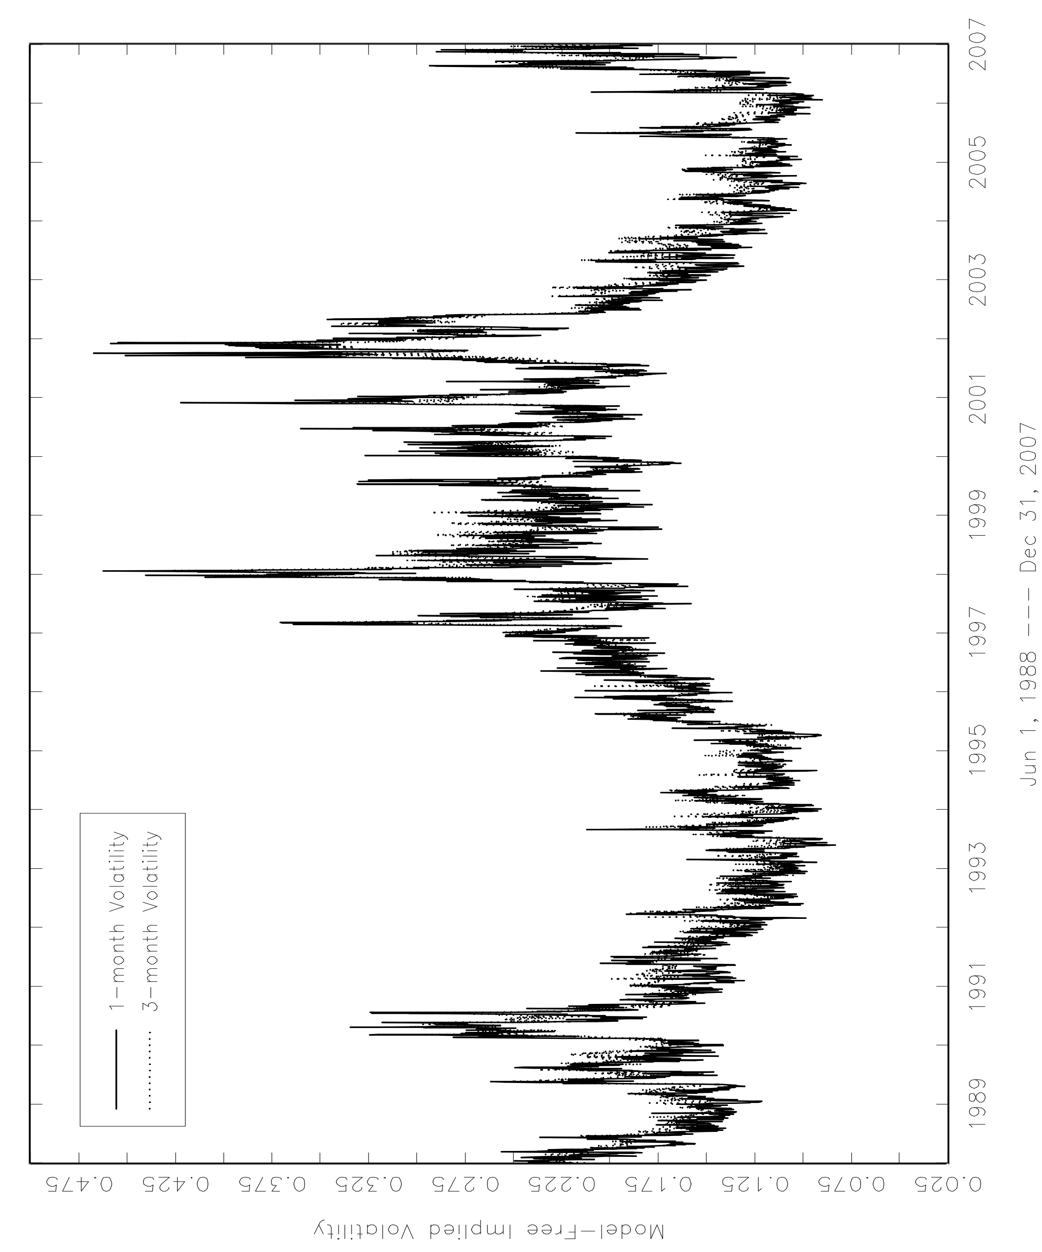

Figure 1. Time series plot of the one-month and three-month model-free implied volatilities

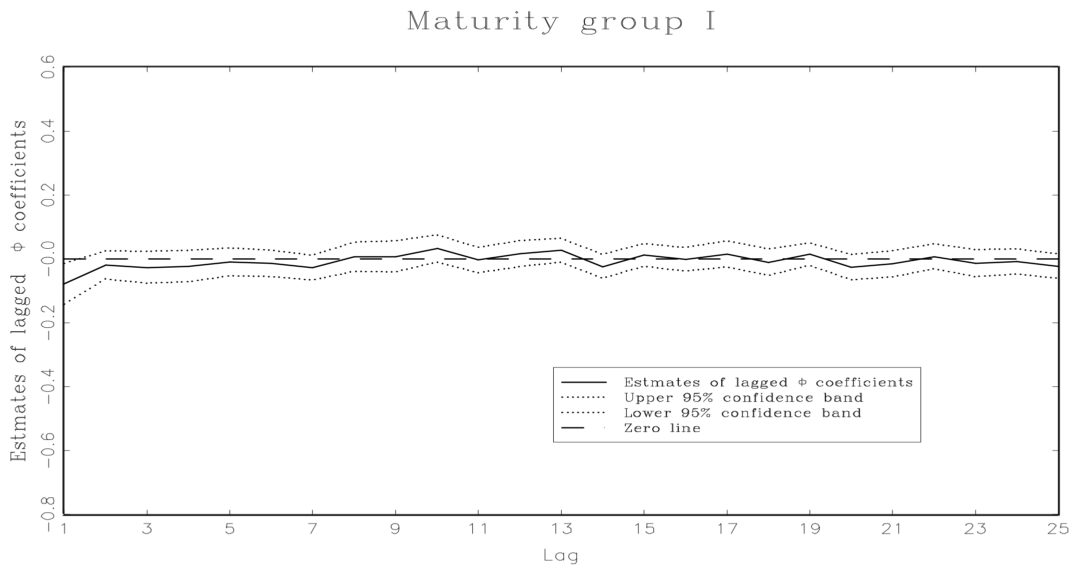
Maturity group I
Estmates of lagged Φ coefficients
Lag
Estmates of lagged Φ coefficients
Upper 95% confidence band
Lower 95% confidence band
Zero line

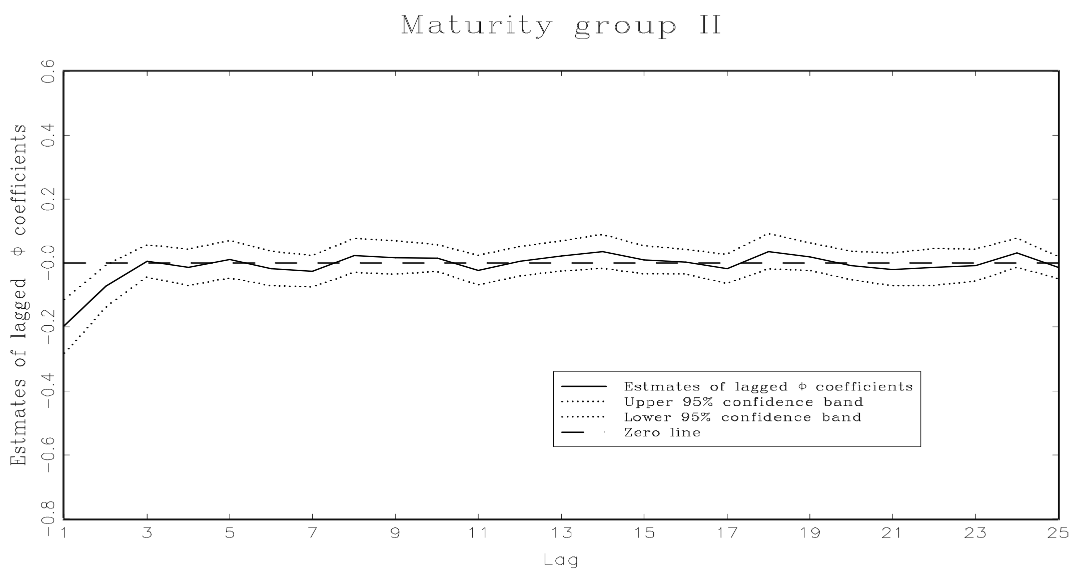
Maturity group II
Estmates of lagged Φ coefficients
Lag
Estmates of lagged Φ coefficients
Upper 95% confidence band
Lower 95% confidence band
Zero line

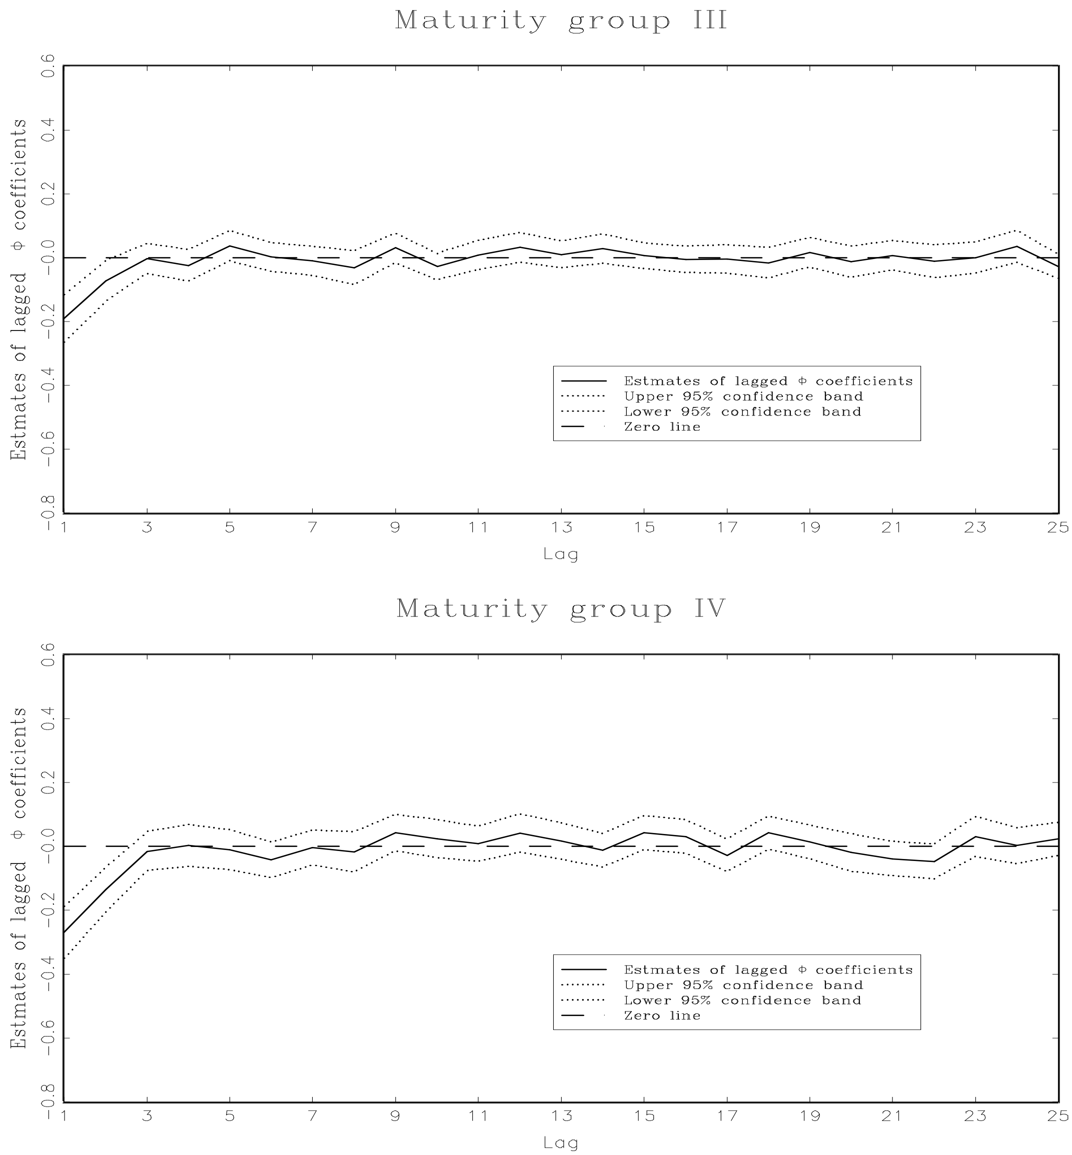

Figure 2. Plots of the lagged ϕ coefficient estimates. The figure illustrates the ϕ coefficients, together with the 95% confidence band, from the following regression:

$$\Delta v(t; T_1, T_2) = \alpha + \sum_{i=1}^{k} \phi_i \, \Delta v(t-i; T_1, T_2) + \varepsilon_t,$$

where $t < T_1 < T_2$, T_1 and T_2 are option maturity dates, and $\Delta v(t-i; T_1, T_2) = v(t-i; T_1, T_2) - v(t-1-i; T_1, T_2)$ $(0 \leq i \leq k)$ are the current and lagged daily changes in forward variance.

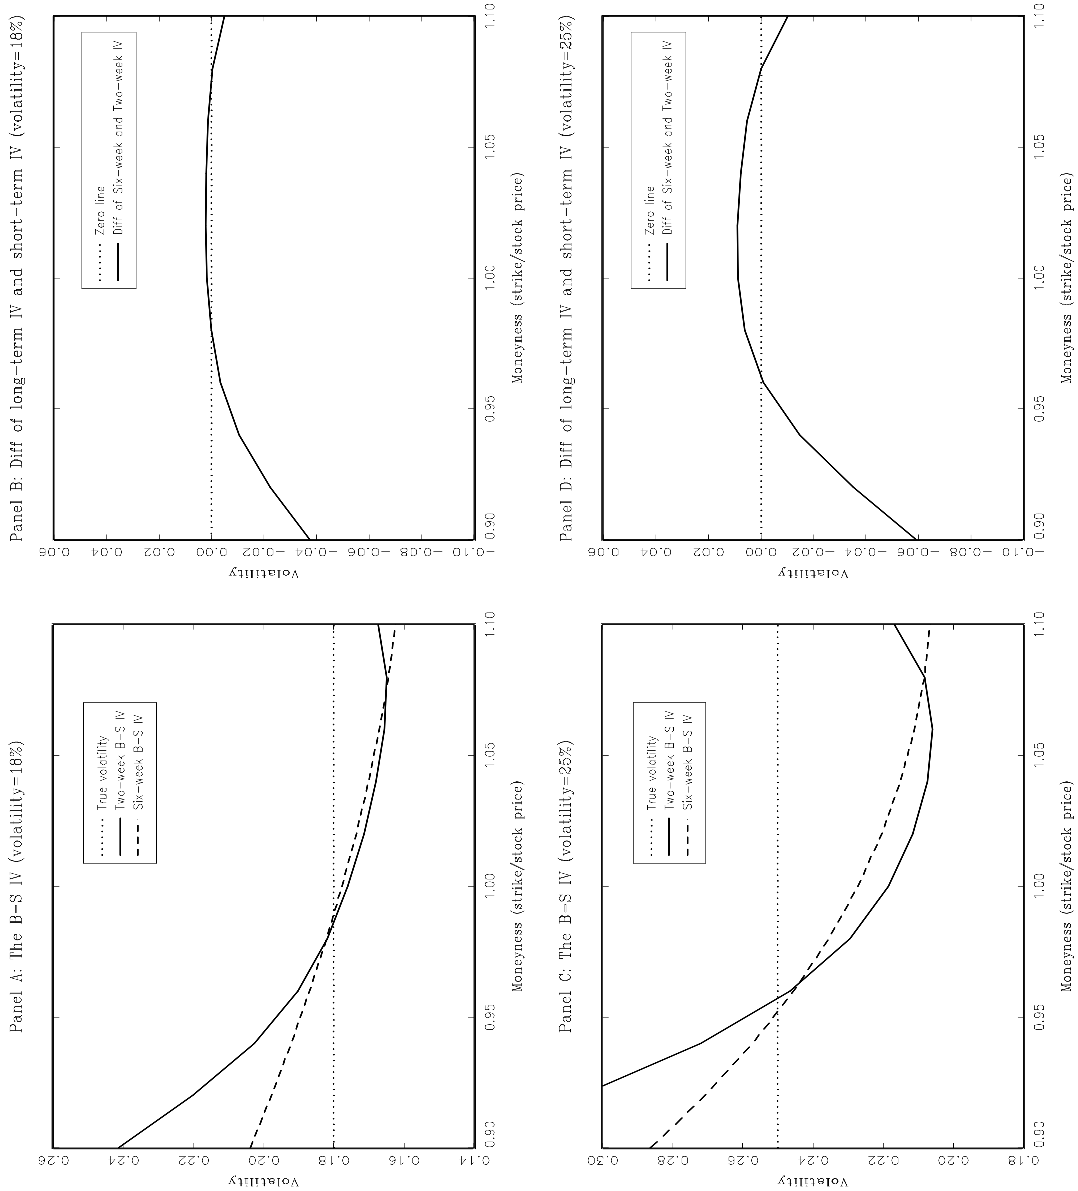

Figure 3. Plots of the B-S implied volatility from the SVJ model. This figure plots the B-S implied volatility from the SVJ model. In Panels A and B, the model parameters are calibrated using option prices on May 14, 1991 and the instantaneous volatility is set equal to 18%. In Panels C and D, the model parameters are calibrated using option prices on December 28, 1990 and the instantaneous volatility is set equal to 25%.