On the Relation between S&P 500 Options and VIX Derivatives

Yang-Ho Park

Federal Reserve Board

September 29, 2016

This paper introduces the variance parity that must hold between S&P 500 (SPX) options and VIX derivatives. Exploiting this parity condition, I find that VIX derivatives contribute more to the price discovery of volatility risk than SPX options, even before the 2008 financial crisis when the former were thinly traded. Additionally, I find that consistent with the theory of limited arbitrage, dislocations between SPX option and VIX derivative prices are associated with measures of funding and market liquidity, even after the peak of the crisis is excluded. Specifically, as liquidity conditions deteriorate, investors appear willing to pay premiums for VIX futures relative to SPX options.

JEL Classification: G01; G13; G14

Keywords: VIX derivatives; price discovery; market dislocation; financial crisis; liquidity risk; implied variance

This paper was previously circulated under the title "Price Dislocation and Price Discovery in the S&P 500 Options and VIX Derivatives." I am very grateful for comments and suggestions from Celso Brunetti and seminar participants at the Federal Reserve Board; the 2015 Paris Financial Management Conference; the 2016 Lancaster Financial Econometrics and Empirical Asset Pricing Conference; and Research in Behavioral Finance Conference. This paper also benefits from the excellent research assistance of Juliette Lu and Robertson Wang. The analysis and conclusions set forth are those of the author and do not indicate concurrence by the Board of Governors or other members of its research staff. Send correspondence to Yang-Ho Park, Risk Analysis Section, Board of Governors of the Federal Reserve System, 20th & C Streets, NW, Washington, DC 20551, Tel: (202) 452-3177, e-mail: yang-ho.park@frb.gov.

On the Relation between S&P 500 Options and VIX Derivatives

This paper introduces the variance parity that must hold between S&P 500 (SPX) options and VIX derivatives. Exploiting this parity condition, I find that VIX derivatives contribute more to the price discovery of volatility risk than SPX options, even before the 2008 financial crisis when the former were thinly traded. Additionally, I find that consistent with the theory of limited arbitrage, dislocations between SPX option and VIX derivative prices are associated with measures of funding and market liquidity, even after the peak of the crisis is excluded. Specifically, as liquidity conditions deteriorate, investors appear willing to pay premiums for VIX futures relative to SPX options.

JEL Classification: G01; G13; G14

Keywords: VIX derivatives; price discovery; market dislocation; financial crisis; liquidity risk; implied variance

1 Introduction

Variance constitutes a tradable asset class, even a popular one. The last decade has seen a variety of new variance derivatives. In particular, the Chicago Board Options Exchange (CBOE) has been successful in promoting two kinds of VIX derivatives: the VIX futures and options. The VIX derivatives market did not develop in a vacuum. Its underlying asset is the VIX index, which is in turn derived from the SPX option prices. In short, the VIX derivatives are derivatives of derivatives. The introduction of the VIX derivatives thus raises many interesting questions: (i) Which of the SPX options and the VIX derivatives contribute more to the price discovery process? (ii) Are VIX derivative prices consistent with those of SPX options, especially during a crisis? (iii) If not, what causes the price inconsistency between the two markets? And so on.

I introduce the variance parity relation that must hold between the SPX options and the VIX derivatives absent market friction and asymmetric information. At an intuitive level, the prices of the VIX derivatives should be tied to those of the SPX options through variance measures because they both reflect the market's expectation of return variance. It is well established in the existing studies that a model-free measure of return variance can be extracted from stock options.[1] In this paper, I develop a new approach to backing out a model-free measure of return variance from the VIX derivatives. That said, the SPX-options-implied variance (SIV) data should convey the same information on the term structure of implied variance as the VIX-derivatives-implied variance (VIV) data.

Exploiting the variance parity condition, I find two pieces of evidence that the VIX derivatives contribute more to the price discovery of volatility risk than the SPX options. For the first evidence, I replicate the VIV measure with a maturity of T using a pair of the SIV measures with maturities T and T plus 30 calendar days, and run a vector error correction (VEC) model for the replicated and original VIV measures, which should be consistent under market integration. I then compute the relative contribution of each market to price discovery following the common factor approach of Gonzalo and Granger (1995) and the information share approach of Hasbrouck (1995). The results suggest that information transmits mainly from the VIX derivatives to the SPX options rather than vice versa; colloquially, the tail wags the dog. This result is valid even in the pre-2008 financial crisis period in which the

[1] There is a vast literature on this issue, including Carr and Madan (1998); Britten-Jones and Neuberger (2000); Demeterfi, Derman, Kamal, and Zou (1999); Jiang and Tian (2007); and Carr and Wu (2009), among others.

VIX derivatives were thinly traded.

Second, I further investigate the price discovery process by applying a collection of four-factor stochastic volatility models. My models have two distinctive features. First, each of the SIV and the VIV measures is driven by different two-factor stochastic variance models. Second, I allow for interactions between the SIV and the VIV by introducing error correction mechanisms in the latent volatility dynamics. The error correction terms permit convergence in the short-run and long-run volatility dynamics between the two markets. If there were to exist a convergence in the latent volatility dynamics, I would interpret such a convergence as evidence for gradual information diffusion between the two markets. In particular, I look into speed-of-convergence parameters to measure the relative contribution of each market to the price discovery process, similar to the idea of Gonzalo and Granger (1995). The models are estimated using both the SIV and VIV data, via an unscented Kalman filter.

The joint analysis of the SIV and the VIV confirms that the VIX derivative prices are more informative than the SPX option prices. Several model selection criteria favor the inclusion of error correction mechanisms in the stochastic volatility dynamics, suggesting that the two markets interact with each other in a way that reduces price dislocations. Importantly, it is the SPX option prices—not the VIX derivative prices—that quickly adjust to eliminate dislocations between two short-run volatility factors, even before the 2008 financial crisis. In addition, the joint analysis makes a sizable improvement only in the pricing of the SIV, and I interpret this result as implying that the VIX derivatives contain some information that has not yet been reflected in the SPX option prices. Taken together, the proposed stochastic volatility models reinforce the finding that information transmits mainly from the VIX derivatives to the SPX options.

Three explanations are plausible for the superior role of the VIX derivatives in the price discovery process. First, the VIX futures allow for direct hedging and speculation in variance, as opposed to the SPX options. Second, transaction costs such as bid-ask spreads are far smaller on the VIX futures contracts than the SPX options contracts. Third, the VIX futures require a smaller margin than the SPX options. For these reasons, the VIX futures may be a better investment vehicle for informed investors who may want to lever up their position.

Next, I attempt to identify the drivers of price dislocations between the two derivatives markets. To this end, I introduce two price dislocation indicators. One is based on the logarithmic differences between the replicated and original VIV measures, which I call model-free bases. The other is based on the model-implied bases, which are defined as follows. For each day I obtain a set of the model prices of the SIV and the VIV as if they were all driven by the SPX-relevant variance dynamics, and another set of the model prices as if they were all driven by the VIX-relevant variance dynamics. The model-implied bases are then defined as the logarithmic differences between the two sets of model prices.

A striking feature of the two price dislocation indicators is that they both plunged in the wake of the Lehman Brothers bankruptcy, suggesting that the VIX derivative prices implied far higher levels of return variance at that time than the SPX option prices. It thus appears that investors are willing to pay higher prices for the VIX futures during a crisis because these contracts allow for cost-effective hedging against stock market crashes, as I discussed earlier. Related to this finding, Bardgett, Gourier, and Leippold (2013) argue that the SPX options and the VIX derivatives carry conflicting information during the crisis.

I focus on two liquidity factors as determinants of price dislocations: funding liquidity and market liquidity. The importance of funding liquidity can be deduced from the limits-to-arbitrage theories, developed by Shleifer and Vishny (1997), Gromb and Vayanos (2002, 2010), Liu and Longstaff (2004), and Basak and Croitoru (2006), among others. The theories are based on the simple real-world feature that arbitrage requires capital, which is a scarce, expensive asset in tumultuous periods. In the context of this study, options and futures trading requires posting margin requirements, which often need to be financed out of arbitrageurs' own pockets. The margin accounts are subject to daily marking-to-market, so that arbitrageurs would become reluctant to engage in arbitrage during a crisis when they are fearful of a margin call or an unwanted liquidation of their positions at a loss.

More directly related to this paper, Gârleanu and Pedersen (2011) show that in a model with heterogeneous agents facing margin constraints, price gaps between two identical assets should depend on the shadow cost of capital, which can be captured through an interest rate spread between unsecured and secured loans. Motivated by this theory, I take the Libor-overnight index swaps (OIS) spread as a measure of funding liquidity and examine whether this variable can explain time variation in the price dislocation indicators.

Market illiquidity may be another source of price dislocations for two reasons. First, market illiquidity can cause a price dislocation on its own because it impedes the completion of arbitrage and makes it costly (see, for example, Roll, Schwartz, and Subrahmanyam, 2007; Deville and Riva, 2007; and Oehmke, 2011).[2] These papers suggest that the speed of

[2] Related to the issue, Kamara and Miller (1995) argue that the severity and frequency of arbitrage

convergence in arbitrage is associated with market liquidity. Second, market liquidity can have an indirect influence on price dislocations through its interaction with funding liquidity. In the limits-to-arbitrage theory, an arbitrageur plays a role as a liquidity provider (at least in normal periods). However, if a funding condition were to tighten, an arbitrageur would be less willing to provide liquidity, which in turn makes it more difficult to engage in arbitrage. Brunnermeier and Pedersen (2009) show that funding liquidity and market liquidity can mutually reinforce each other during a crisis. Therefore, taking the relative bid-ask spreads as a measure of market liquidity, I look into whether they can further explain time variation in the price dislocation indicators in addition to funding liquidity.

As is expected, I find that both funding and market liquidity are critical drivers of price dislocations. Although this finding is most apparent in the full sample including the Lehman Brothers crisis, it is still valid even in a subsample period in which the peak of the Lehman Brothers crisis is excluded. Notably, every liquidity measure, either funding or market, is negatively related to each of the price dislocation indicators. These negative relations mean that the price dislocations occurred during the liquidity crises in a way that the VIX derivatives implied higher levels of return variance than the SPX options but not the other way around.

In summary, this paper introduces the variance parity relation that must hold between the SPX options and the VIX derivatives, and documents two empirical findings on the relation between the two derivatives markets utilizing that parity relation. First, most of the price discovery occurs in the VIX derivatives rather than in the SPX options, even before the 2008 financial crisis. Second, as measures of funding and market liquidity deteriorate, the VIX derivatives tend to imply higher levels of variance than the SPX options, indicating that investors are willing to pay premiums for the former relative to the latter.

There is a related literature. Some researchers have been interested in developing a stochastic jump-diffusion model that can consistently price both SPX options and VIX derivatives (see, for example, Cont and Kokholm, 2013; Madan and Pistorius, 2014; and Kokholm and Stisen, 2015). Chung, Tsai, Wang, and Weng (2011) show that the VIX options have predictive information for the stock index returns and volatility, incremental to the SPX options. Recently, Song and Xiu (2016) performed a joint analysis of SPX options and VIX options to study pricing kernels.

The rest of the paper is organized as follows. Section 2 introduces the variance parity

breakdowns are associated with market liquidity.

relation. Section 3 describes the models and estimation technique used in this paper. The data description is provided in Section 4. Sections 5 and 6 discuss price discovery and price dislocation, respectively. Section 7 concludes.

2 Variance parity

This paper aims to study the empirical relation between the SPX options and the VIX derivatives. The difficulty arises because there is no direct lockstep relation between the prices of the SPX options and those of the VIX derivatives, as they are unequally loaded on stochastic risk factors. For example, return jumps may be more important for out-of-the-money (OTM) SPX put options, whereas variance jumps may affect only VIX options. Therefore, a direct comparison of the SPX option prices with the VIX derivative prices is meaningless.

To address this issue, I isolate a summary variable of derivative prices, a model-free measure of expected return variance, that should be identical between the two markets. This section shows that a model-free measure of return variance can be extracted not only from the SPX options but also from the VIX derivatives. If the two markets are fully integrated with no friction and asymmetric information, the two measures of return variance should be consistent with each other. In what follows, I detail the procedure of backing out such a return variance measure from each of the two markets.

2.1 Extracting return variance from the SPX options

It is well recognized that a risk-neutral return variance can be replicated by a static position in the OTM stock options. Early work by Carr and Madan (1998); Britten-Jones and Neuberger (2000); and Demeterfi, Derman, Kamal, and Zou (1999) introduces model-free formulas for return variance under the assumption of continuous stock prices, whereas subsequent researchers such as Jiang and Tian (2007) and Carr and Wu (2009) extend the original idea to the case in which stock prices are driven by both diffusion and jump components.

Let me explain how a risk-neutral measure of return variance is related to stock option prices in a jump-diffusion framework. Assume a risk-neutral probability space $(\Omega, \mathcal{F}, \mathbb{Q})$ and information filtration $\{\mathcal{F}_t\}$. Let the stock price, S_t, take the following stochastic differential equation under the $\mathbb{Q}$ measure:

$$\frac{dS_t}{S_t} = r_t dt + \sigma_t dW_t^{\mathbb{Q}} + \int_{\mathbb{R}} (\exp(x) - 1)[J^{\mathbb{Q}}(dx, dt) - \nu_t^{\mathbb{Q}}(dx)dt], \tag{1}$$

where r_t is the risk-free rate, σ_t is the instantaneous diffusive volatility, $W_t^{\mathbb{Q}}$ is a standard Brownian motion under the $\mathbb{Q}$ measure, $J^{\mathbb{Q}}(dx, dt)$ is a random jump measure, and $\nu_t^{\mathbb{Q}}(dx)$ is a jump compensator for the log price.

Given Equation (1), return variance, or annualized quadratic variation, may be expressed as the sum of integrated variance and jump variation:

$$\mathcal{V}(t,T) = \frac{1}{T-t}\Bigg[\underbrace{\int_t^T \sigma_s^2 ds}_{\text{integrated variance}} + \underbrace{\int_t^T \int_{\mathbb{R}} x^2 J^{\mathbb{Q}}(dx, ds)}_{\text{jump variation}}\Bigg], \tag{2}$$

where $\mathcal{V}(t,T)$ denotes the return variance over the $(t,T]$ horizon. The risk-neutral expectation of the return variance is given by

$$E_t^{\mathbb{Q}}[\mathcal{V}(t,T)] = \frac{1}{T-t} E_t^{\mathbb{Q}}\left[\int_t^T \lambda_s^{\mathbb{Q}} ds\right], \tag{3}$$

where

$$\lambda_s^{\mathbb{Q}} = \sigma_s^2 + \int_{\mathbb{R}} x^2 \nu_s^{\mathbb{Q}}(dx)$$

is referred to as the $\mathbb{Q}$-spot variance process. Jiang and Tian (2005) and Carr and Wu (2009) show that Equation (3) can be approximated using the prices of OTM SPX options up to a high-order error term as follows:

$$E_t^{\mathbb{Q}}[\mathcal{V}(t,T)] \approx \frac{2\exp(r_t(T-t))}{T-t}\left[\int_{F_t(T)}^{\infty} \frac{C_t(T,K)}{K^2} dK + \int_0^{F_t(T)} \frac{P_t(T,K)}{K^2} dK\right], \tag{4}$$

where $F_t(T)$ is the SPX future price and $C_t(T,K)$ and $P_t(T,K)$ are the OTM SPX call and put prices with a maturity of T and a strike price of K, respectively.[3]

[3] In fact, Aït-Sahalia, Karaman, and Mancini (2015) argue that the high-order error term may be non-trivial by comparing the synthetic variance swaps implied by the SPX options and the actual OTC variance swap rates.

2.2 Extracting return variance from the VIX derivatives

I now introduce an alternative model-free approach to extracting a measure of return variance from the VIX derivatives. Before doing so, I begin by explaining why the (squared) VIX future price alone is insufficient in replicating a certain measure of return variance. Let $VF_t(T)$ denote the VIX future price with a maturity of T. By its nature, the VIX future price is the same as the risk-neutral expectation of the time-T VIX index level: $VF_t(T) = E_t^{\mathbb{Q}}[VIX_T]$. By Jensen's inequality, the squared VIX future price must be less than or equal to the risk-neutral expectation of a forward-starting return variance as follows:

$$\begin{aligned} VF_t(T)^2 &= \left(E_t^{\mathbb{Q}}[VIX_T]\right)^2 \\ &\leq E_t^{\mathbb{Q}}[VIX_T^2] \\ &= E_t^{\mathbb{Q}}[E_T^{\mathbb{Q}}[\mathcal{V}(T, T+30d)]] \\ &= E_t^{\mathbb{Q}}[\mathcal{V}(T, T+30d)], \end{aligned} \tag{5}$$

where $30d$ stands for 30 calendar days and $\mathcal{V}(T, T+30d)$ denotes the return variance starting on a future date T with a fixed 30-day window. The equality holds when $T = t$.

The difference between the risk-neutral expectation of the forward-starting variance (with a fixed 30-day window) and the squared VIX future price is called a convexity adjustment term. The convexity adjustment is associated with the variance of the time-T VIX index, which can be backed out from a cross-section of OTM VIX option prices.

Proposition 1. *Under no arbitrage, the variance of the time-T VIX index, which I denote by $var_t(VIX_T)$, may be expressed in terms of a cross-section of the VIX option prices with different strike prices but with the same maturity T:*

$$var_t(VIX_T) = 2\exp\left(r_t(T-t)\right)\left[\int_{VF_t(T)}^{\infty} VC_t(T,K)dK + \int_0^{VF_t(T)} VP_t(T,K)dK\right], \tag{6}$$

where $VC_t(T,K)$ and $VP_t(T,K)$ are the OTM VIX call and put prices with a maturity of T and a strike price of K, respectively.

See Appendix A for the proof. Roughly speaking, Equation (6) measures the price volatility of the VIX index, as opposed to the return volatility.[4] Combining the convexity adjustment term in Equation (6) with the squared VIX future price, I can infer the risk-neutral expectation of a forward-starting return variance with a fixed 30-day window.

[4] Unlike Equation (6), the formula of Martin (2013) measures the return volatility of an asset price.

Corollary 1. *Under no arbitrage, the risk-neutral expectation of the forward-starting return variance over the horizon* $(T, T+30d]$ *can be expressed in terms of the prices of the VIX futures and options with different strike prices but with the same maturity* T*:*

$$E_t^{\mathbb{Q}}[\mathcal{V}(T, T+30d)] = VF_t(T)^2 + 2\exp\left(r_t(T-t)\right)\left[\int_{VF_t(T)}^{\infty} VC_t(T,K)dK + \int_0^{VF_t(T)} VP_t(T,K)dK\right]. \tag{7}$$

2.3 Comparing two return variance measures

Let $SIV_t(T)$ denote the return variance implied by the SPX options with a maturity of T, and let $VIV_t(T)$ denote the return variance implied by the VIX derivatives with a maturity of T. These two measures of return variance are defined in Equations (4) and (7), respectively. It is important to understand the difference in horizons between $SIV_t(T)$ and $VIV_t(T)$. The former captures the risk-neutral expectation of return variance over the horizon $(t, T]$, whereas the latter does so over the horizon $(T, T+30d]$; that is, $SIV_t(T) = E_t^{\mathbb{Q}}[\mathcal{V}(t, T)]$ and $VIV_t(T) = E_t^{\mathbb{Q}}[\mathcal{V}(T, T+30d)]$.

The two measures of return variance are associated with different over-the-counter (OTC) variance swaps. A variance swap is an OTC swap that exchanges the future realized variance of the underlying asset price for a prespecified strike price at the expiration date. As no money changes hands at the initiation of the contract, a no-arbitrage condition requires that the strike price of the variance swap be equal to the risk-neutral expectation of the return variance over the horizon. That said, $SIV_t(T)$ may be considered a synthetic strike price for a variance swap starting today with a tenor of $T-t$, whereas $VIV_t(T)$ may be considered a synthetic strike price for a forward-starting variance swap with a fixed 30-day tenor.

Although a single SIV observation covers a horizon distinct from that of a single VIV observation, a cross-section of $SIV_t(T)$ with different maturities should convey the same information on the term structure of implied variance as another cross-section of $VIV_t(T)$ with different maturities. This parity relation that must hold between the two markets constitutes the basis for my inquiries.

3 Methodology

This section develops a new family of affine diffusion models for a joint analysis of the SIV and the VIV, which help us understand the dynamic relation between two derivatives markets. In addition, I derive semi-analytic solutions to the valuation of the SIV and the VIV and explain an estimation method based on unscented Kalman filtering.

3.1 Stochastic volatility models with error corrections

Many of the existing studies find that a two-factor variance structure is essential to describe the term structure of return variance (see, for example, Gallant, Hsu, and Tauchen, 1999; Christoffersen, Heston, and Jacobs, 2009; and Aït-Sahalia, Amengual, and Manresa, 2015). Furthermore, modeling the logarithm of variance fits the data better than modeling the variance itself (see, for example, Jones, 2003; Aït-Sahalia and Kimmel, 2007; and Durham, 2013). Thus, I accommodate these two important features of the data in my joint modeling of the SIV and the VIV.

A unique feature of my model is that each of the SIV and the VIV is assumed to be governed by different short-run and long-run volatility dynamics. Let v_{1t} and u_{1t} denote the two variance factors that drive the term structure of the SIV. Similarly, let v_{2t} and u_{2t} denote the two variance factors that drive the term structure of the VIV. In particular, v_{1t} and v_{2t} determine the short-run volatility behavior in each of the two markets, whereas u_{1t} and u_{2t} dictate the long-run volatility behavior. The latent state vector, $X_t = (v_{1t}, v_{2t}, u_{1t}, u_{2t})'$, is assumed to take the following stochastic differential equation forms under the $\mathbb{Q}$ measure:

$$\begin{aligned} dv_{1t} &= -\kappa_{v_1} v_{1t} dt + \gamma_{v_1}(v_{1t} - v_{2t})dt + \sigma_{v_1} dB^{\mathbb{Q}}_{1t} \\ dv_{2t} &= -\kappa_{v_2} v_{2t} dt + \gamma_{v_2}(v_{1t} - v_{2t})dt + \sigma_{v_2} dB^{\mathbb{Q}}_{2t} \\ du_{1t} &= -\kappa_{u_1} u_{1t} dt + \gamma_{u_1}(u_{1t} - u_{2t})dt + \sigma_{u_1} dB^{\mathbb{Q}}_{3t} \\ du_{2t} &= -\kappa_{u_2} u_{2t} dt + \gamma_{u_2}(u_{1t} - u_{2t})dt + \sigma_{u_2} dB^{\mathbb{Q}}_{4t}, \end{aligned} \tag{8}$$

where κ_{v_1}, κ_{v_2}, κ_{u_1}, and κ_{u_2} are the persistence parameters; σ_{v_1}, σ_{v_2}, σ_{u_1}, and σ_{u_2} are the variance-of-variance parameters; γ_{v_1}, γ_{v_2}, γ_{u_1}, and γ_{u_2} are the speeds of adjustment; and $B^{\mathbb{Q}}_{1t}$, $B^{\mathbb{Q}}_{2t}$, $B^{\mathbb{Q}}_{3t}$, and $B^{\mathbb{Q}}_{4t}$ are standard Brownian motions under the $\mathbb{Q}$ measure. Here, I assume zero correlations among all Brownian motions.

Although the SIV and the VIV are driven by different state variables, the state variables are allowed to interact with each other via error correction mechanisms. Specifically, the two terms, $\gamma_{v_1}(v_{1t} - v_{2t})$ and $\gamma_{v_2}(v_{1t} - v_{2t})$, allow for convergence between the short-run volatility factors and thus between short-dated SIV and VIV. The other two error correction terms, $\gamma_{u_1}(u_{1t} - u_{2t})$ and $\gamma_{u_2}(u_{1t} - u_{2t})$, allow for convergence between the long-run volatility factors and thus between long-dated SIV and VIV. If there were to exist such convergences, γ_{v_1} and γ_{u_1} would be negative and γ_{v_2} and γ_{u_2} would be positive.

The virtue of error correction mechanisms is that they are indicative of price informativeness in each of the two equivalent markets. If one market was informationally efficient on its own, it would not respond to price deviations from another market. Conversely, if one market was less informationally efficient than another market, the former would adjust its prices in a way that would reconcile with the latter market. In other words, the market with a large error correction means that it has a lower degree of informational efficiency relative to the competing market. Therefore, the speed-of-adjustment parameters allow us to compare price informativeness between the two markets: the larger the speed of adjustment in magnitude, the less informationally efficient.

The general specification in Equation (8) describes the most extensive of the models used in this paper, which I refer to as the SVEC2 model. I also look at two additional nested models with some restrictions on error correction terms. One is the most simplistic model with no error corrections (with $\gamma_{v_1} = 0$, $\gamma_{v_2} = 0$, $\gamma_{u_1} = 0$, and $\gamma_{u_2} = 0$), which I refer to as the SV model. The SV model would be suitable if the two derivatives markets were fully segmented with no interaction. The other is the model with only short-run error corrections (with $\gamma_{u_1} = 0$ and $\gamma_{u_2} = 0$), which I refer to as the SVEC1 model.

3.2 Moment generating function

Let $M(X_t, t, T; \phi)$ denote the moment-generating function of X_T at time t under the $\mathbb{Q}$ measure:

$$M(X_t, t, T; \phi) = E_t^{\mathbb{Q}}\left[\exp(\phi' X_T)\right]. \tag{9}$$

Given the model in Equation (8), the moment-generating function can be quasi-analytically calculated by solving a system of Riccati ordinary differential equations (ODEs). To be specific, under some regularity conditions, the moment-generating function takes an exponentially affine form of X_t:

$$M(X_t, t, T; \phi) = \exp\left(\alpha(s) + \alpha_{v_1}(s)v_{1t} + \alpha_{v_2}(s)v_{2t} + \alpha_{u_1}(s)u_{1t} + \alpha_{u_2}(s)u_{2t}\right), \tag{10}$$

where $s = T - t$ and the coefficients, $\alpha(s)$, $\alpha_{v_1}(s)$, $\alpha_{v_2}(s)$, $\alpha_{u_1}(s)$, and $\alpha_{u_2}(s)$, satisfy the system of ODEs:

$$\begin{aligned}
\dot{\alpha}(s) &= \frac{1}{2}(\sigma_{v_1}\alpha_{v_1}(s))^2 + \frac{1}{2}(\sigma_{v_2}\alpha_{v_2}(s))^2 + \frac{1}{2}(\sigma_{u_1}\alpha_{u_1}(s))^2 + \frac{1}{2}(\sigma_{u_2}\alpha_{u_2}(s))^2 \\
\dot{\alpha}_{v_1}(s) &= -(\kappa_{v_1} - \gamma_{v_1})\alpha_{v_1}(s) + \gamma_{v_2}\alpha_{v_2}(s) \\
\dot{\alpha}_{v_2}(s) &= -(\kappa_{v_2} + \gamma_{v_2})\alpha_{v_2}(s) - \gamma_{v_1}\alpha_{v_1}(s) \\
\dot{\alpha}_{u_1}(s) &= -(\kappa_{u_1} - \gamma_{u_1})\alpha_{u_1}(s) + \gamma_{u_2}\alpha_{u_2}(s) \\
\dot{\alpha}_{u_2}(s) &= -(\kappa_{u_2} + \gamma_{u_2})\alpha_{u_2}(s) - \gamma_{u_1}\alpha_{u_1}(s),
\end{aligned} \tag{11}$$

where the boundary conditions are given as $\alpha(0) = 0$ and $[\alpha_{v_1}(0), \alpha_{v_2}(0), \alpha_{u_1}(0), \alpha_{u_2}(0)] = \phi'$.

3.3 Valuation formula for return variances

I assume that the SIV and the VIV are governed by different $\mathbb{Q}$-spot variance processes in the form of exponential-affine transformations of the state vector as follows:

$$\begin{aligned}
\lambda_{1t}^{\mathbb{Q}} &= \exp\left(\mu_1 + v_{1t} + u_{1t}\right) \\
\lambda_{2t}^{\mathbb{Q}} &= \exp\left(\mu_2 + v_{2t} + u_{2t}\right),
\end{aligned} \tag{12}$$

where $\lambda_{1t}^{\mathbb{Q}}$ and $\lambda_{2t}^{\mathbb{Q}}$ denote the $\mathbb{Q}$-spot variances governing the SIV and the VIV, respectively, and μ_1 and μ_2 capture long-term levels of the two $\mathbb{Q}$-spot variances.

Given Equations (8) and (12), the model prices of the SIV and the VIV can be derived by integrating out the moment generating functions with different boundary conditions as follows:

$$\begin{aligned}
SIV_t(T;\mathcal{M}) &= \frac{\exp\left(\mu_1\right)}{T-t}\int_t^T M(X_t, t, \tau; \phi = (1,0,1,0)')d\tau \\
VIV_t(T;\mathcal{M}) &= \frac{\exp\left(\mu_2\right)}{30d}\int_T^{T+30d} M(X_t, t, \tau; \phi = (0,1,0,1)')d\tau,
\end{aligned} \tag{13}$$

where $SIV_t(T;\mathcal{M})$ and $VIV_t(T;\mathcal{M})$ denote the model-$\mathcal{M}$ price of the SIV and the VIV, respectively.

I would like to mention two limitations in my modeling framework. First, I use the $\mathbb{Q}$-spot variance as a primitive variable to model the dynamics of return variance. Such a direct modeling of the $\mathbb{Q}$-spot variance was introduced by Filipović, Gourier, and Mancini (2016). As a result, I am unable to disentangle diffusive and jump variances. Second, I do not include the return dynamics in my framework, modeling only the variance dynamics.

Nevertheless, I believe that my modeling framework is well-suited to my research objectives, as I focus on the term structure of return variance.

3.4 Estimation

The models are estimated using a quasi maximum likelihood method via an unscented Kalman filter. Christoffersen, Dorion, Jacobs, and Karoui (2014) show that unscented Kalman filtering is superior to extended Kalman filtering in the application of interest rate derivatives and is comparable to particle filtering, which is computationally far more expensive.

A Kalman filter requires us to recast the models into a state space form that comprises observation equations and state equations. With respect to observation equations, I assume that the SIV and the VIV are observed with measurement errors. Specifically, I assume that logarithmic return variances have constant measurement errors as follows:

$$\begin{aligned} \log(SIV_t(T)) &= \log(SIV_t(T;\mathcal{M})) + \sigma_{e_1}\xi_{1,t} \\ \log(VIV_t(T)) &= \log(VIV_t(T;\mathcal{M})) + \sigma_{e_2}\xi_{2,t}, \end{aligned} \tag{14}$$

where σ_{e_1} and σ_{e_2} capture the sizes of measurement errors and $\xi_{1,t}$ and $\xi_{2,t}$ are independent standard normal random variables.

To define state equations, I specify the dynamics of the state vector under the physical $\mathbb{P}$ measure as:

$$\begin{aligned} dv_{1t} &= -\kappa_{v_1} v_{1t}dt + \eta_{v_1} v_{1t}dt + \gamma_{v_1}(v_{1t} - v_{2t})dt + \sigma_{v_1} dB^{\mathbb{P}}_{1t} \\ dv_{2t} &= -\kappa_{v_2} v_{2t}dt + \eta_{v_2} v_{2t}dt + \gamma_{v_2}(v_{1t} - v_{2t})dt + \sigma_{v_2} dB^{\mathbb{P}}_{2t} \\ du_{1t} &= -\kappa_{u_1} u_{1t}dt + \eta_{u_1} u_{1t}dt + \gamma_{u_1}(u_{1t} - u_{2t})dt + \sigma_{u_1} dB^{\mathbb{P}}_{3t} \\ du_{2t} &= -\kappa_{u_2} u_{2t}dt + \eta_{u_2} u_{2t}dt + \gamma_{u_2}(u_{1t} - u_{2t})dt + \sigma_{u_2} dB^{\mathbb{P}}_{4t}, \end{aligned} \tag{15}$$

where the four additional parameters, η_{v_1}, η_{v_2}, η_{u_1}, and η_{u_2}, capture variance risk premiums. By applying an Euler approximation to Equation (15), I am able to define discrete-time state equations, which will be used when updating the state vector (see Appendix B for further details on the unscented Kalman filter).

4 Data and convexity ratio

4.1 Data

The data set comprises the daily prices of SPX options, VIX futures, and VIX options. The SPX and VIX options data come from OptionMetrics and the futures data are obtained from Thomson Reuters Datastream. The sample period is restricted by the short history of the VIX futures and options. The VIX futures market started on March 26, 2004, and the VIX options market opened about two years later on February 24, 2006. Because the trading of VIX options was illiquid in the very beginning, my sample period spans from July 1, 2006 through August 31, 2014.

The options market has experienced vibrant changes in the trading environment. The CBOE recently introduced weekly and end-of-month options, but I do not include them in my analysis because they were not accounted for in the computation of the VIX index until October 6, 2014.[5] The CBOE introduced overnight trading hours for the SPX options and the VIX derivatives, but I consider only regular trading hours in this paper.[6] Option prices are taken from the bid-ask midpoint at the close of regular trading hours, 3:15 p.m.

Various data filters are applied to eliminate inaccurate or illiquid options. Specifically, I delete both SPX and VIX options for which the mid price is less than 0.05; the Black-Scholes implied volatility is null; the deviation between the best bid and offer prices is larger than five; or the lower bound constraint is violated. I also exclude the SPX options with fewer than eight days or more than one year to maturity, and the VIX options with fewer than eight days or more than 11 months to maturity.[7] Lastly, the SPX options market closes 15 minutes after the closing of the SPX cash market. To address the issue of nonsynchronous trading hours, I back out the SPX spot price for each of the first three pairs of at-the-money SPX put and call options by using the put-call parity and take an average of the three extracted spot prices. This idea originates from Aït-Sahalia and Lo (1998, 2000).

[5] Including weekly options would put too much weight on fitting the short end of the term structure of return variance.

[6] The trading hours were extended on March 9, 2015 for the SPX options; June 23, 2014 for the VIX futures; and March 2, 2015 for the VIX options.

[7] Note that the VIX options with 11 months to maturity cover the return variance between 11 and 12 months into the future.

4.2 Convexity ratio

In Section 2.2, I show that the squared VIX future price is not a correct measure of return variance because of Jensen's inequality, and that the convexity adjustment term can be implied from a cross-section of the VIX option prices. Here I evaluate the relative magnitude of the convexity adjustment term in the VIV measure. To this end, I define the convexity ratio, $CVRT_t(T)$, as the ratio of the convexity adjustment to the VIV measure:

$$CVRT_t(T) = \frac{\mathrm{var}_t(VIX_T)}{VIV_t(T)}, \tag{16}$$

where the numerator and the denominator are given by Equations (6) and (7), respectively.

Table 1 shows the summary statistics for the convexity ratios that are broken down into short and long maturities, which are defined as less than three months and more than three months to expiration, respectively. Figure 1 shows the histograms of the short- and long-term convexity ratios. The average short-term convexity ratio is 9%, implying that the VIX options contribute to the magnitude of the total VIV measure by that number of percentage points. The long-term convexity ratios (with an average of 16%) are larger than the short-term ones. To sum up, the convexity adjustment component composes a modest fraction of the VIV measures, with a larger effect on long-term contracts.

5 Price discovery

An important goal of this study is to compare price informativeness between the SPX options and the VIX derivatives. I address this issue by looking into whether information on volatility risk diffuses from one market to the other. In the presence of information diffusion, the informationally inefficient market will adjust its prices based on the departure from the informationally efficient market. I examine such a possibility in two econometric settings: the VEC model and the stochastic volatility model as introduced in Section 3.1.

5.1 Model-free analysis

A traditional way of examining price discovery is to define two asset prices that should be identical absent market friction and asymmetric information and run a VEC model for time series of the two asset prices. In this section, I introduce a method of replicating the VIV measure from the SIV data and perform a price discovery analysis based on the estimation result of a VEC model. I refer to this section as a model-free analysis—even though the VEC model is applied—because model-free measures of return variances are used to discuss price discovery, without reliance on any parametric model of return variances.

5.1.1 Replicating the VIV from the SIV

The SIV and the VIV are not directly comparable because they cover different horizons, as explained in Section 2.3. Instead, it is possible to replicate the VIV with a maturity of T using a pair of the SIV measures with maturities of T and $T + 30d$:

$$\begin{aligned}\widetilde{VIV_t(T)} &= E_t^{\mathbb{Q}}[\mathcal{V}(T, T+30d)] \\ &= \left(\frac{T+30d-t}{30d}\right) E_t^{\mathbb{Q}}[\mathcal{V}(t, T+30d)] - \left(\frac{T-t}{30d}\right) E_t^{\mathbb{Q}}[\mathcal{V}(t, T)] \\ &= \left(\frac{T+30d-t}{30d}\right) SIV_t(T+30d) - \left(\frac{T-t}{30d}\right) SIV_t(T),\end{aligned} \tag{17}$$

where $\widetilde{VIV_t(T)}$ denotes a replicated version of the VIV computed from the SIV measures. In practice, the two SIV measures in Equation (17) may not be available, so I interpolate them from the available SIV data using a cubic smoothing spline and replicate the VIV measures from the interpolated SIV measures. If the SPX options and the VIX derivatives markets are fully integrated, the replicated and original VIV measures, $\widetilde{VIV_t(T)}$ and $VIV_t(T)$, should be identical, allowing for an apples-to-apples comparison.

Figure 2 shows the time series plots of the replicated VIV measure (solid line) in comparison with the original VIV measure (dotted line) for a three-month maturity. Panels A and B correspond to the entire sample period and a subsample of the Lehman Brothers crisis, respectively. Although the two VIV measures are quite close to each other most of the time (see Panel A), massive dislocations are observed in the wake of the Lehman Brothers bankruptcy (see Panel B). Furthermore, the replicated VIV measures take some negative values in October 2008. The negative values of replicated VIV measures arise because the near-term SIV measures were too high compared with the far-term SIV measures; in other words, the term structure of the SIV was too steep during the period immediately following the Lehman Brothers bankruptcy.

5.1.2 Testing for a unit root and co-integration

I start by computing the replicated and original constant-maturity VIV measures via a cubic smoothing spline. Table 2 reports the summary statistics of the square root of the VIV measures for three different maturities: one, three, and five months. Here, I consider only two subperiods: the pre-crisis period (July 1, 2006 to August 31, 2008) and the post-crisis period (December 1, 2008 to August 31, 2014). Note that the analysis in this section excludes the three-month period from September 1 through November 30, 2008, which corresponds to the peak of the Lehman Brothers crisis, as the replicated VIV measures sometimes take negative values.

The table shows that there are no apparent differences in the summary statistics between the replicated and original VIV measures, regardless of maturities. This is true of both the pre- and post-crisis periods. Overall, the replicated VIV measure has an unconditional distribution that is very similar to that of the original VIV measure (as long as the peak of the Lehman Brothers crisis is excluded).

Before running a VEC model, I test whether each time series is a stationary or integrated process by implementing the KPSS (Kwiatkowski, Phillips, Schmidt, and Shin, 1992) and ADF (Augmented Dickey Fuller) tests. Table 3 reports the test results for the logarithm of the replicated and original VIV measures, with an optimal lag chosen using Akaike information criteria. The null hypothesis of the KPSS test is that the time series is stationary against the existence of a unit root. As can be seen in the second and third columns of the table, we can reject the null hypothesis of stationarity at 1% levels for both the pre- and post-crisis periods, regardless of the maturities considered. The same conclusion can be reached based on the ADF test (shown in the fourth and fifth columns of the table) as we cannot reject the null hypothesis that the time series has a unit root for every maturity considered. To conclude, both tests suggest that all VIV measures, replicated or original, have a unit root.[8]

To examine the interaction between the two equivalent VIV measures, I apply the Johansen (1988, 1991) co-integration test for the following form of a VEC model, with an order of p:

$$\Delta y_t = \alpha(\beta' y_{t-1} + c_0) + A_j \sum_{j=1}^{p} \Delta y_{t-j} + \varepsilon_t, \tag{18}$$

[8] Stock volatility may be stationary with a long-memory property but it is often difficult to reject stationarity in a relatively short sample period as in my application.

where $y_t = (\log(\widetilde{VIV_t(T)}), \log(VIV_t(T)))'$ is a vector of the logarithm of replicated and original VIV measures with a maturity of T, α is the speed-of-adjustment matrix, β is the co-integrating relation matrix, A_j are the autoregressive coefficient matrices, and ε_t are *i.i.d.* normal shocks. Note that there is an intercept in the co-integrating relation. An optimal order of the VEC model is determined using a likelihood ratio test.

The Johansen maximum eigenvalue test assesses the null hypothesis that the number of co-integrating relations, the rank of $\alpha\beta'$, is equal to r against $r + 1$. The test results are provided in the sixth and seventh columns of Table 3. With a 1% significance level, we can reject the null of $r = 0$ but cannot reject the null of $r = 1$ for every maturity considered in both the pre- and post-crisis periods.[9] These results indicate that the replicated and original VIV measures are likely to have a single co-integrating relation at a 1% level in both the pre- and post-crisis periods. That said, the error correction mechanism is a critical feature of the dynamics of the replicated and original VIV measures.

5.1.3 Measuring relative contributions to price discovery

The VEC model as given in Equation (18) indicates informational efficiency between the two derivatives markets. Given the result that there exists a single co-integrating relation, I let $\alpha = (\alpha_1, \alpha_2)'$, where α_1 and α_2 are scalars. The speed-of-adjustment parameters dictate the direction and magnitude of information diffusion between the two markets. For example, if information transmits from the VIX derivatives to the SPX options, α_1 should be negative. Or, if information diffuses from the SPX options to the VIX derivatives, α_2 should be positive.

The speed-of-adjustment parameters are reported in the eighth and ninth columns of Table 3. As can be seen in the table, α_1 is estimated to be negative for every maturity in both the pre- and post-crisis periods. The negative signs indicate that the SPX option prices adjust to eliminate dislocations from the VIX derivative prices as information diffuses from the latter market to the former market. Similarly, α_2 is estimated to be positive for every maturity in both the pre- and post-crisis periods, suggesting that information transmission also occurs in the other direction from the SPX options to the VIX derivatives.

Two common approaches to measuring the relative contribution of each market to price discovery are Gonzalo and Granger (1995) and Hasbrouck (1995). Following the former ap-

[9] There is only one case in which the Johansen test can reject the null of $r = 1$ against $r = 2$ at a 5% level: the one-month maturity in the post-crisis period. This rejection means that both replicated and original VIV measures with a maturity of one month may be stationary processes in this case.

proach, I define the portion of the price discovery that is attributable to the VIX derivatives, which I denote by GG, as

$$GG = \frac{\alpha_1}{\alpha_1 - \alpha_2}. \tag{19}$$

This GG metric essentially measures the relative contribution of innovations in the VIX derivative prices to innovations in the efficient prices.

Unlike the Gonzalo and Granger (1995) approach, Hasbrouck (1995) defines the relative contribution in price discovery as each market's contribution to the total variance of innovations in the efficient prices, which is called an information share. Using both the speed-of-adjustment parameters and the variance matrix of errors in the VEC model, I define the lower and upper bounds of the VIX derivatives' information share, which I denote by HAS_L and HAS_U, respectively, as

$$\begin{aligned} HAS_L &= \frac{\alpha_1^2(\sigma_2^2 - \sigma_{12}^2/\sigma_1^2)}{\alpha_2^2\sigma_1^2 - 2\alpha_1\alpha_2\sigma_{12} + \alpha_1^2\sigma_2^2} \\ HAS_U &= \frac{(\alpha_1\sigma_2 - \alpha_2\sigma_{12}/\sigma_2)^2}{\alpha_2^2\sigma_1^2 - 2\alpha_1\alpha_2\sigma_{12} + \alpha_1^2\sigma_2^2}, \end{aligned} \tag{20}$$

where σ_1^2, σ_{12}, and σ_2^2 constitute the covariance matrix of ε_t.

The last three columns of Table 3 show the price discovery metrics. Panels A and B correspond to the pre- and post-crisis periods, respectively. In the pre-crisis period, the GG metric ranges from 76% to 94%, depending on the maturities, suggesting that most of the price discovery of volatility risk takes place in the VIX derivatives rather than in the SPX options. A similar result can be found in the post-crisis period as the GG metric ranges from 42% to 96%, depending on the maturities.

In the pre-crisis period, the information share metrics indicate that the relative contribution of the VIX derivatives can be as small as 19% to 47% (based on HAS_L) but as large as 86% to 100% (based on HAS_U). Although the information share metrics have a wide range, their numbers overall suggest that the VIX derivatives contribute more to the price discovery than the SPX options, rather than vice versa. Again, a similar result can be found in the post-crisis period.

In sum, information on future variance gets impounded into the VIX derivative prices more quickly than into the SPX option prices. The SPX option prices then adjust to eliminate the price departure from the VIX derivative prices due to information diffusion. This result holds even in the pre-crisis period in which the VIX derivatives were not actively traded, as well as in the post-crisis period.

5.2 Model-based analysis

This section extends the price discovery analysis by comparing three stochastic volatility models, with or without error correction mechanisms, as introduced in Section 3.1. This model-based analysis complements the previous model-free analysis (Section 5.1) in two respects. First, the model-free analysis excludes the three-month crisis period in which the replicated VIV measures take negative values. In contrast, the analysis here can be applied to the full sample period as the SIV and VIV data are directly fed into the estimation step. Second, whereas the model-free analysis does not account for measurement errors in the replicated and original VIV measures, the analysis here allows for measurement errors in the SIV and VIV data. In what follows, I will look into estimation results and discuss the price discovery process based on estimates of error correction parameters and pricing errors in the SIV and the VIV.

5.2.1 Parameter estimates and model comparison

Table 4 presents parameter estimates across three model specifications: SV, SVEC1 and SVEC2 models. Panels A, B, and C correspond to the full sample period, the pre-crisis period, and the post-crisis period, respectively. Interestingly, the SIV-relevant parameters are somewhat different from the VIV-relevant parameters. For example, the SIV-implied variance dynamics have a lower long-run mean and a higher variance-of-variance than the VIV-implied variance dynamics, regardless of the sample periods. These differences suggest that the two derivatives markets may be segmented to some extent rather than fully integrated.

The table also reports two kinds of model comparison criteria: Akaike and Schwarz information criteria. These criteria are based on some penalty scores, which depend on the number of free parameters, less log likelihood values. The lower the information criteria, the better is the model. Both information criteria suggest that the SVEC1 model is preferred to the SV model. That is, the inclusion of the short-run error correction terms improves likelihood functions substantially in all of the sample periods considered. Compared to the improvement by the short-run error correction terms, the long-run error correction terms make only a very small improvement in likelihood functions on top of the short-run error correction terms, although the SVEC2 model is preferred to the SVEC1 model based on the information criteria.

The main variables of interest are the speed-of-adjustment parameters as they are indicative of information transmission. As is expected, γ_{v_1} turns out to be negative and γ_{v_2} turns out to be positive in all of the three sample periods considered. The results imply that information transmission occurs in a way that eliminates dislocations in the short-run volatility dynamics, or in the short-dated return variance measures.

However, I do not find any clear evidence for convergence in the long-run volatility dynamics as the speed-of-convergence parameters, γ_{u_1} and γ_{u_2}, have unexpected signs. For example, they are both estimated to be positive in the full and the post-crisis sample; and negative in the pre-crisis sample. The unexpected results suggest that the error correction terms in the long-run volatility dynamics may be invalid. This result is not that surprising given that they have a very small effect on likelihood functions when short-run error correction terms are included.

5.2.2 Relative contributions to price discovery

The aforementioned results indicate that price departure is corrected mainly through a convergence in the short-run volatility dynamics. Therefore, I will discuss the relative contribution of the VIX derivatives to the price discovery process by focusing on the short-run speed-of-adjustment parameters.

Similar to the idea of Gonzalo and Granger (1995), my method quantifies the relative contribution of the VIX derivatives to price discovery as

$$PDR = \frac{\gamma_{v_1}}{\gamma_{v_1} - \gamma_{v_2}}, \tag{21}$$

where PDR stands for the model-based price discovery ratio. This ratio is reported in the bottom row of Table 4. For the full sample period, the price discovery ratio turns out to be 73% in both the SVEC1 and SVEC2 models. These numbers suggest that the VIX derivatives contribute to price discovery about three times more than the SPX options, consistent with the result of the model-free analysis.

Figure 3 shows the trading volume of SPX options, VIX futures, and VIX options. Note that the trading volumes of VIX futures and options have experienced dramatic growth in recent years, whereas those of the SPX options have declined to some extent over time.

Because of this observation, one may conjecture that the relative contribution of the VIX derivatives to price discovery may have increased in the post-crisis period compared with the pre-crisis period. However, contrary to expectation, PDR is slightly smaller in the post-crisis period, despite the increased trading volume of VIX derivatives. That said, the VIX derivatives played a larger role in the price discovery of volatility risk even in the pre-crisis period in which they were thinly traded.

5.2.3 Variance pricing errors: Indirect evidence for price discovery

This section aims to provide indirect evidence for the price informativeness of the VIX derivatives by evaluating pricing errors in the SIV and the VIV. The idea is that if one market contains more information than another market, a joint analysis may help improve pricing performance, particularly, in an informationally inefficient market. To see if this is the case, I define root mean squared pricing errors (RMSEs), which are computed separately for the SIV and the VIV:

$$\begin{aligned} RMSE_{i,SPX} &= \sqrt{\frac{1}{N_{SPX}} \sum_{n=1}^{N_{SPX}} \Big(\log(SIV_t(T_n;\mathcal{M}_i)) - \log(SIV_t(T_n)) \Big)^2} \\ RMSE_{i,VIX} &= \sqrt{\frac{1}{N_{VIX}} \sum_{n=1}^{N_{VIX}} \Big(\log(VIV_t(T_n;\mathcal{M}_i)) - \log(VIV_t(T_n)) \Big)^2}, \end{aligned} \tag{22}$$

where $RMSE_{i,SPX}$ and $RMSE_{i,VIX}$ denote the RMSEs of the model i for the SIV and the VIV, respectively, and N_{SPX} and N_{VIX} indicate the total number of SIV and VIV observations included in the sample, respectively. Hereafter, $RMSE_i$ is used to denote either $RMSE_{i,SPX}$ or $RMSE_{i,VIX}$ according to the context of analysis.

I next compute the pricing differences of a model against the SV model, denoted by $\Delta RMSE_i$, which are defined as the logarithmic difference between the $RMSE$ of a particular model i and that of the SV model. That said, $\Delta RMSE_i$ is indicative of the pricing performance of error correction mechanisms relative to the model without them. In particular, a negative (positive) value of $\Delta RMSE_i$ means that the model i yields lower (higher) pricing errors than the SV model.

To assess the statistical significance of pricing differences, I apply the Diebold and Mariano (2002) test. To this end, I compute a time series of RMSEs for each model, $\{\text{RMSE}_{i,t}\}_{t=1}^{D}$, where D is the total number of dates and $\text{RMSE}_{i,t}$ is the date-t RMSE implied by a model

i. Next, I obtain a series of the differences in RMSEs between every pair of the models, $d_{i,j,t} = \text{RMSE}_{i,t} - \text{RMSE}_{j,t}$, $t = 1, \ldots, D$, and define a test statistic between models i and j as

$$z_{i,j} = \frac{\overline{d_{i,j,t}}}{\sigma(d_{i,j,t})}, \tag{23}$$

where $\overline{d_{i,j,t}}$ is a sample mean of the differences and $\sigma(d_{i,j,t})$ is a heteroskedasticity- and autocorrelation-consistent estimate of the sample standard deviation of the differences. According to Diebold and Mariano (2002), the test statistic in Equation (23) follows a standard normal distribution under the null hypothesis that there is no statistically significant difference between models i and j.

Tables 5 and 6 present the variance pricing performance and the Diebold and Mariano (2002) test statistics across different model specifications. In each table, Panels A, B, and C correspond to the full sample, the pre-crisis sample, and the post-crisis sample, respectively. Note that the models with the short-run error correction terms (the SVEC1 and SVEC2 models) substantially improve the pricing of the SIV by 11% to 16% over the SV model, with the largest improvement found in the full sample period. These results are statistically significant at 1% levels, with z-statistics between -3.50 and -9.18 in the SVEC1 model and between -3.48 and -10.09 in the SVEC2 model. In fact, most of the improvements are achieved in the short-term contracts with less than three-month maturities; the improvements are far smaller in the long-term contracts with greater than three-month maturities.

However, the benefit of the inclusion of short-run error correction terms comes at a cost. It turns out that the SVEC1 and SVEC2 models deteriorate the pricing of the VIV by 5% to 10% relative to the SV model. These results are statistically significant at 1% levels, with z-statistics between 6.41 and 8.54 in the SVEC1 model and between 5.63 and 8.52 in the SVEC2 model.

In sum, a joint analysis (with a short-run error correction mechanism) improves the pricing of the SIV notably while impairing the pricing of the VIV. This asymmetric result may suggest that the VIX derivative prices contain some information that has not yet been reflected in the SPX option prices, but not likely the other way around. This is true regardless of whether a long-run error correction mechanism is included. Thus, I consider this result to be indirect evidence that the VIX derivative prices are more informative than the SPX option prices.

5.3 Why are VIX derivatives more informative than SPX options?

Both model-free and model-based analyses point to the superior role of the VIX derivatives in price discovery, regardless of the sample periods. Why is this the case? It is important to note that the VIX futures have three nice characteristics for variance trading compared with the SPX options. First, the VIX futures allow for pure speculation and hedging in variance, whereas the SPX options are affected by both variance and underlying price risks. Second, transaction costs are far lower in the VIX futures than the SPX options. For example, the relative bid-ask spreads are, on average, 0.01 in the VIX futures market and 0.29 in the SPX options market (see Table 7).[10] Third, the VIX futures require a smaller margin than options contracts; in particular, shorting options requires margins that are as large as the full prices of the options. Owing to these differences, the VIX futures would be more attractive for informed investors who may want to lever up their position.[11] Thus, for further research, it might be interesting to scrutinize whether time-varying margins or time-varying bid-ask spreads can cause a time-varying pattern in the price discovery process.

6 Price dislocation

Another important objective of this paper is to understand the drivers of price dislocations between the SPX options and the VIX derivatives. To serve this purpose, this section introduces two indicators of price dislocations between the two derivatives markets and investigates whether measures of funding liquidity and market liquidity can explain time-variation in the price dislocation indicators.

6.1 Introducing two indicators of price dislocation

This section introduces two dislocation indicators. The first dislocation indicator is based on the logarithmic difference between the replicated and original VIV measures:

$$\Delta VIV_t(T) \equiv \log\left(\widetilde{VIV_t(T)}\right) - \log\left(VIV_t(T)\right), \tag{24}$$

[10] See Section 6.2.2 for a further discussion on relative bid-ask spreads.

[11] Black (1975) argues that the market that allows for leverage should contribute to the price formation more than the market that does not, with all else being equal.

which I call model-free bases in the VIV. Here, I exclude the observations for which the replicated VIV measures are negative. To see the time-evolution of price dislocations, I compute an average of the model-free bases with different maturities for a given day, which I denote by BSS_t^{MF}:

$$BSS_t^{MF} = \frac{1}{N_{t,VIV}} \sum_{n=1}^{N_{t,VIV}} \Delta VIV_t(T_n), \tag{25}$$

where $N_{t,VIV}$ denotes the number of model-free bases with different maturities at time t. If the two derivatives markets are fully integrated, BSS_t^{MF} should be equal to zero. In contrast, a positive (negative) value of BSS_t^{MF} means that the two markets may be dislocated in the way that the SPX options imply higher (lower) levels of return variance than the VIX derivatives.

The second dislocation indicator is based on the model prices of the SIV and the VIV. Recall that $SIV_t(T;\mathcal{M})$ denotes the model price of the SIV as viewed by the SPX market players. Here, I compute a counterfactual model price of the SIV as if it were traded in the VIX derivatives market. That is, the cross-model price, which I denote by $\widehat{SIV_t(T;\mathcal{M})}$, represents the price of $SIV_t(T)$ that would be fair in the VIX derivatives market. Similarly, I compute a counterfactual model price of the VIV as if it were traded in the SPX options market, which I denote by $\widehat{VIV_t(T;\mathcal{M})}$. I then define the model-implied bases in the SIV and the VIV using their true and counterfactual model prices:

$$\begin{aligned} \Delta SIV_t(T;\mathcal{M}) &\equiv \log\left(SIV_t(T;\mathcal{M})\right) - \log\left(\widehat{SIV_t(T;\mathcal{M})}\right) \\ \Delta VIV_t(T;\mathcal{M}) &\equiv \log\left(\widehat{VIV_t(T;\mathcal{M})}\right) - \log\left(VIV_t(T;\mathcal{M})\right), \end{aligned} \tag{26}$$

where $\Delta SIV_t(T;\mathcal{M})$ and $\Delta VIV_t(T;\mathcal{M})$ denote the model-implied bases in the SIV and the VIV, respectively. I then calculate an average of the model-implied bases for a given day, which I denote by BSS_t^{SV}:

$$BSS_t^{SV} = \frac{1}{K_{t,SIV} + K_{t,VIV}} \left(\sum_{n=1}^{K_{t,SIV}} \Delta SIV_t(T_n;\mathcal{M}) + \sum_{n=1}^{K_{t,VIV}} \Delta VIV_t(T_n;\mathcal{M}) \right), \tag{27}$$

where $K_{t,SIV}$ and $K_{t,VIV}$ denote the number of model-implied bases in the SIV and the VIV, respectively. As in BSS_t^{MF}, a positive (negative) value of BSS_t^{SV} means that the SPX options imply higher levels of return variance than the VIX derivatives.

Figure 4 shows the evolution of the two dislocation indicators, with $BSS^{MF}$and BSS^{SV} shown in Panels A and B, respectively. Here, the SV model is used to compute BSS^{SV}

because the model effectively assumes that the two derivatives markets are fully segregated. A striking feature of the figure is that the dislocation indicators both plummeted abysmally down to about −100% in the wake of the Lehman Brothers bankruptcy. This number means that the VIX derivatives-implied return variances were about 100% larger than their SPX options-implied counterparts.

The massive dislocation is consistent with a large literature reporting no-arbitrage relation breakdowns during the 2008 financial crisis. The list includes the interest rate parity violation (Coffey, Hrung, and Sarkar, 2009; Baba and Packer, 2009; Fong, Valente, and Fung, 2010; and Mancini Griffoli and Ranaldo, 2012); the CDS-bond parity violation (Gârleanu and Pedersen, 2011; and Bai and Collin-Dufresne, 2013); and the TIPS-Treasury parity violation (Fleckenstein, Longstaff, and Lustig, 2014). Recently, Pasquariello (2014) introduced an aggregate market dislocation index, which might be useful for monitoring financial market stability.

Why do the VIX derivatives tend to be more expensive relative to the SPX options, especially during the Lehman Brothers crisis? Variance trading is primarily used as a hedging tool against stock market downturns. As I explained earlier, the VIX futures make it easy to hedge against stock market crashes and require smaller transaction costs than the SPX options. It thus appears that during crisis periods investors are willing to pay premiums for VIX futures relative to SPX options.

6.2 Drivers of price dislocation

In this section, I will investigate the economic sources of price dislocations. In particular, I attempt to link the dislocation indicators to measures of funding and market liquidity using ordinary linear regressions. The summary statistics and correlation matrix for the dependent and explanatory variables are provided in Table 7.

6.2.1 Role of funding liquidity

There is an extensive list of papers showing that limited capital obstructs arbitrage activity, leaving parity violations sustained over an extended period of time. This effect becomes most severe during a crisis when capital becomes a scarce, expensive asset. Limits-to-arbitrage theories have been developed by many researchers, including Shleifer and Vishny (1997), Gromb and Vayanos (2002), Liu and Longstaff (2004), and Basak and Croitoru (2006), among others.

In particular, Gârleanu and Pedersen (2011) show that when agents are constrained by margins, a parity violation relies on the shadow cost of capital, which can be captured through an interest rate spread between uncollateralized and collateralized loans. Following the theory, I take the Libor-OIS spread ($LIBOIS$) as a measure of funding liquidity. Figure 5 shows the time series plot of $LIBOIS$. Before August 9, 2007, $LIBOIS$ averaged only about 10 basis points. The next day, $LIBOIS$ rose to about 40 basis points and continued to climb to 365 basis points until the peak of the crisis. Starting from the fourth quarter of 2009, $LIBOIS$ returned to near the pre-crisis range.

Case 1 of Table 8 shows the regression results of the two dislocation indicators against $LIBOIS$ alone. Panels A and B correspond to the full sample and the ex-crisis sample, respectively, where the ex-crisis sample is the one excluding the three-month window corresponding to the peak of the Lehman Brothers crisis. All explanatory variables are standardized in the regression analysis, so each coefficient can be interpreted as the change in the dependent variable in response to a one standard deviation change in the explanatory variable. In the full sample, the measure of funding liquidity alone can explain 16% and 30% of the time variation in $BSS^{MF}$and BSS^{SV}, respectively. However, in the ex-crisis sample, the explanatory power significantly is reduced to 5% for $BSS^{MF}$and 4% for BSS^{SV}. Nevertheless, all of the results are statistically significant at 1% levels, even after the peak of the Lehman Brothers crisis is excluded from the regression analysis.

Note that $LIBOIS$ is negatively associated with each price dislocation indicator in both sample periods. The negative relations indicate that, as funding liquidity worsens, the VIX derivatives imply higher levels of return variance than the SPX options. For example, a one standard deviation increase in $LIBOIS$ would lead to a 2% increase in the original VIV measure relative to the replicated VIV measure (in the full sample).

I also consider other interest rate spreads, such as the Libor-GC (general collateral) repo spread and the TED spread, but the results (not shown in this paper) are similar to or slightly weaker than those with $LIBOIS$. Overall, regardless of the choice of an interest rate spread, measures of funding liquidity can explain a fraction of the time variation in price dislocations.

6.2.2 Role of market liquidity

Market liquidity may be another important driver of price dislocations because it facilitates arbitrage. Authors such as Roll, Schwartz, and Subrahmanyam (2007); Deville and Riva (2007); and Oehmke (2011) provide empirical and theoretical evidence that market illiquidity may disrupt arbitrage, particularly delaying convergence between asset prices.

As a proxy for market liquidity, I consider the relative bid-ask spread, which is the bid-ask spread divided by the mid price. In particular, I obtain three relative bid-ask spreads, $SPRDS$, $SPRDF$, and $SPRDV$, which denote weekly averages of the relative bid-ask spreads for SPX options, VIX futures, and VIX options, respectively. When taking a weekly average of the relative bid-ask spreads for each derivatives market, I include only the contracts with more than 10 trading volumes on that date.

Figure 6 shows the time series plots of the three relative bid-ask spreads, with Panels A, B, and C corresponding to the SPX options, the VIX futures, and the VIX options, respectively. Interestingly, all three bid-ask spreads increased in the wake of the Lehman Brothers bankruptcy, although the increases were not as dramatic as those seen in $LIBOIS$.

Case 2 of Table 8 shows the regression results of $BSS^{MF}$and BSS^{SV}onto $SPRDS$, $SPRDF$, and $SPRDV$. In the full sample, the three market liquidity proxies together can explain 15% to 31% of the time variation in the dislocation indicators. In the ex-crisis sample, the explanatory power is reduced to 5% for $BSS^{MF}$and 4% for BSS^{SV}. Statistical significances are obtained at different levels. First, $SPRDS$ is statistically significant at a 1% level in the full sample and a 5% level in the ex-crisis sample. Second, $SPRDF$ is statistically significant at a 1% level in both the full and ex-crisis samples. $SPRDV$ is statistically significant at a 1% to 5% level only in the full sample.

Note that all of the coefficients on relative bid-ask spreads have negative signs. These negative signs mean that, as market liquidity deteriorates, the VIX derivatives imply higher levels of return variance than the SPX options. To sum up, market liquidity is a critical driver of the time-variation in price dislocations.

6.2.3 Joint role of funding liquidity and market liquidity

Here I examine the combined effect of funding liquidity and market liquidity on price dislocations. In the limits-to-arbitrage theory and, in particular, in the framework of Brunnermeier and Pedersen (2009), funding liquidity and market liquidity can mutually reinforce each other. Indeed, the funding liquidity measure is positively associated with each of the three bid-ask spread measures, with a correlation between 0.39 and 0.47 (see Table 7). As a consequence, a multivariate regression analysis helps us illuminate the true source of price dislocations. It could be the case that either funding liquidity or market liquidity can wash out the effect of the other.

Case 3 of Table 8 shows the results of regressing the dislocation indicators against $LIBOIS$, $SPRDS$, $SPRDF$, and $SPRDV$. The combined explanatory power increases up to 21% for $BSS^{MF}$and 39% for BSS^{SV}in the full sample. Again, the explanatory power decreases significantly if the peak of the crisis is excluded from the regression. Nonetheless, $LIBOIS$ is still statistically significant at 1% levels in both samples. $SPRDS$ is still statistically significant at a 1% level in the full sample and a 5% level in the ex-crisis sample, and $SPRDF$ is still statistically significant at a 5% level in both samples. Overall, the results suggest that funding liquidity and market liquidity are both critical drivers of the time-variation in price dislocations.

6.2.4 Robustness

The SIV and the VIV depart from each other because of some risk factors. That said, I add to the regression the CBOE VIX index as a proxy for variance risk and the CBOE SKEW index as a proxy for return jump risk. The results are provided in Case 4 of Table 8. Note that even with the VIX and SKEW indexes included, the statistical significances on the measures of funding and market liquidity are similarly maintained in both the full and ex-crisis samples.

6.3 Further discussion

The preceding results show that the departure of SPX option prices from the VIX derivative prices is associated with both funding and market liquidity measures. Aside from the liquidity factors, price dislocations could be caused by demand or hedging pressures in futures and options markets as well. This speculation is related to a large literature suggesting that demand or hedging pressure can distort the prices of futures and options, affecting their expected returns. Hirshleifer (1989, 1990), Bessembinder (1992), and de Roon, Nijman, and Veld (2000) are good examples for futures markets. This particular story goes as far back as Keynes (1930) and Hicks (1939). Similar work has been done for options markets as well. Bollen and Whaley (2004) and Gârleanu, Pedersen, and Poteshman (2009) show that a higher demand pressure is associated with a higher option price, a potential resolution for the option pricing puzzle.

More directly related to this work, Cheng (2014) and Mixon and Onur (2015) show that the demand pressures broken down into several groups—dealers, asset managers, and hedge funds—affect the difference between the VIX future price and the VIX index. Similarly, it would be reasonable to speculate that my results may be partially explained by demand or hedging pressures from distinct groups of investors. For example, it might be the case that the demand pressure from asset managers, who are net buyers in the VIX futures market, may distort the VIX future prices upward during the crisis, as these contracts are primarily used as a hedging tool against stock market downturns. To test this speculation, I obtain demand pressure measures across distinct investor groups based on the disaggregated Commitments of Traders report from the Commodity Futures Trading Commission. Specifically, the demand pressure measures are defined as the net position of each group divided by the overall open interest, as in Cheng (2014). However, contrary to the conjecture, I find little evidence that the demand pressures (at least in the VIX futures market) can explain time variation in the price dislocation indicators, although the results are not reported in this paper.

7 Conclusion

Motivated by the recent growing interest in variance markets, my goal in this paper is to perform an analysis on the relation between the SPX options and the VIX derivatives. To this end, I introduce a new model-free approach to backing out a measure of return variance from the VIX futures and options. This approach leads to the variance parity relation between the SPX options and the VIX derivatives, along with the existing approaches to extracting a measure of return variance from stock options. That is, absent friction and asymmetric information, the term structure of variance implied by the SPX options should be consistent with that implied by the VIX derivatives.

The variance parity relation helps us shed light on the price discovery process and the price dislocation in the two derivatives markets. My chief findings are twofold. First, I find that the VIX derivatives contribute more to the price discovery of volatility risk than the

SPX options, even before the 2008 financial crisis when the former were thinly traded. In short, the tail wags the dog. This result may be attributable to the fact that, compared with the SPX options, the VIX futures allow for pure speculation and hedging in variance, incur lower transaction costs, and require lower margin requirements.

Second, I find that violations of the parity condition are associated with measures of funding and market liquidity, consistent with the theory of limited arbitrage. Specifically, as liquidity conditions deteriorate, the VIX derivative prices tend to imply higher levels of return variance than the SPX option prices, but not the other way around. Thus, during a liquidity crisis, investors seem willing to pay premiums for the VIX futures relative to the SPX options as the former allow for cost-effective hedging against a stock market crash.

References

Aït-Sahalia, Y., D. Amengual, and E. Manresa (2015): "Market-based estimation of stochastic volatility models," *Journal of Econometrics*, 187(2), 418–435.

Aït-Sahalia, Y., M. Karaman, and L. Mancini (2015): "The term structure of variance swaps and risk premia," Working paper, Princeton University.

Aït-Sahalia, Y., and R. Kimmel (2007): "Maximum likelihood estimation of stochastic volatility models," *Journal of Financial Economics*, 83, 413–452.

Aït-Sahalia, Y., and A. W. Lo (1998): "Nonparametric estimation of state-price densities implicit in financial asset prices," *Journal of Finance*, 53, 499–547.

——— (2000): "Nonparametric risk management and implied risk aversion," *Journal of Econometrics*, 94, 9–51.

Baba, N., and F. Packer (2009): "From turmoil to crisis: Dislocations in the FX swap market before and after the failure of Lehman Brothers," *Journal of International Money and Finance*, 28(8), 1350–1374.

Bai, J., and P. Collin-Dufresne (2013): "The CDS-Bond basis," Working paper, Georgetown University and Swiss Finance Institute.

Bardgett, C., E. Gourier, and M. Leippold (2013): "Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets," Working paper, Swiss Finance Institute.

Basak, S., and B. Croitoru (2006): "On the role of arbitrageurs in rational markets," *Journal of Financial Economics*, 81(1), 143–173.

Berndt, E., B. Hall, R. Hall, and J. Hausman (1974): "Estimation and inference in nonlinear structural models," *Annals of Economic and Social Measurement*, 3, 653–665.

Bessembinder, H. (1992): "Systematic risk, hedging pressure, and risk premiums in futures markets," *Review of Financial Studies*, 5(4), 637–667.

Black, F. (1975): "Fact and fantasy in the use of options," *Financial Analysts Journal*, 31(4), 36–41.

Bollen, N., and R. Whaley (2004): "Does net buying pressure affect the shape of implied volatility functions?," *Journal of Finance*, 59, 711–753.

Britten-Jones, M., and A. Neuberger (2000): "Option prices, implied price processes, and stochastic volatility," *Journal of Finance*, 55, 839–866.

Brunnermeier, M. K., and L. H. Pedersen (2009): "Market liquidity and funding liquidity," *Review of Financial Studies*, 22(6), 2201–2238.

Carr, P., and D. Madan (1998): "Towards a theory of volatility trading," in *Volatility: New Estimation Techniques for Pricing Derivatives*. RISK Publications.

——— (2001): "Optimal positioning in derivative securities," *Quantitative Finance*, 1, 19–37.

Carr, P., and L. Wu (2009): "Variance risk premiums," *Review of Financial Studies*, 22, 1311–1341.

Cheng, I.-H. (2014): "The expected return of fear," Working paper, Dartmouth College.

Christoffersen, P., C. Dorion, K. Jacobs, and L. Karoui (2014): "Nonlinear Kalman filtering in affine term structure models," *Management Science*, 60, 2248–2268.

Christoffersen, P., S. Heston, and K. Jacobs (2009): "The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well," *Management Science*, 55, 1914–1932.

Chung, S.-L., W.-C. Tsai, Y.-H. Wang, and P.-S. Weng (2011): "The information content of the S&P 500 index and VIX options on the dynamics of the S&P 500 index," *Journal of Futures Markets*, 31(12), 1170–1201.

Coffey, N., W. B. Hrung, and A. Sarkar (2009): "Capital constraints, counterparty risk, and deviations from covered interest rate parity," Staff report, Federal Reserve Bank of New York.

Cont, R., and T. Kokholm (2013): "A consistent pricing model for index options and volatility derivatives," *Mathematical Finance*, 23(2), 248–274.

de Roon, F. A., T. E. Nijman, and C. Veld (2000): "Hedging pressure effects in futures markets," *Journal of Finance*, 55(3), 1437–1456.

DEMETERFI, K., E. DERMAN, M. KAMAL, AND J. ZOU (1999): "More than you ever wanted to know about volatility swaps," Research Notes, Goldman Sachs Quantitative Strategies.

DEVILLE, L., AND F. RIVA (2007): "Liquidity and arbitrage in options markets: A survival analysis approach," *Review of Finance*, 11(3), 497–525.

DIEBOLD, F. X., AND R. S. MARIANO (2002): "Comparing predictive accuracy," *Journal of Business and Economic Statistics*, 20(1), 134–144.

DURHAM, G. B. (2013): "Risk-neutral modelling with affine and nonaffine models," *Journal of Financial Econometrics*, 11, 650–681.

FILIPOVIĆ, D., E. GOURIER, AND L. MANCINI (2016): "Quadratic variance swap models," *Journal of Financial Economics*, 119(1), 44–68.

FLECKENSTEIN, M., F. A. LONGSTAFF, AND H. LUSTIG (2014): "The TIPS-Treasury bond puzzle," *Journal of Finance*, 69(5), 2151–2197.

FONG, W.-M., G. VALENTE, AND J. K. FUNG (2010): "Covered interest arbitrage profits: The role of liquidity and credit risk," *Journal of Banking and Finance*, 34(5), 1098–1107.

GALLANT, A. R., C. HSU, AND G. TAUCHEN (1999): "Using daily range data to calibrate volatility diffusions and extract the forward integrated variance," *Review of Economics and Statistics*, 81, 617–631.

GÂRLEANU, N., AND L. PEDERSEN (2011): "Margin-based asset pricing and deviations from the law of one price," *Review of Financial Studies*, 24(6), 1980–2022.

GÂRLEANU, N., L. PEDERSEN, AND A. POTESHMAN (2009): "Demand-based option pricing," *Review of Financial Studies*, 22(10), 4259–4299.

GONZALO, J., AND C. GRANGER (1995): "Estimation of common long-memory components in cointegrated systems," *Journal of Business and Economic Statistics*, 13(1), 27–35.

GROMB, D., AND D. VAYANOS (2002): "Equilibrium and welfare in markets with financially constrained arbitrageurs," *Journal of Financial Economics*, 66(2), 361–407.

——— (2010): "Limits of arbitrage: The state of the theory," Working paper, National Bureau of Economic Research.

Hasbrouck, J. (1995): "One security, many markets: Determining the contributions to price discovery," *Journal of Finance*, 50(4), 1175–1199.

Hicks, J. R. (1939): *Value and capital: An inquiry into some fundamental principles of economic theory*. Oxford: The Clarendon Press.

Hirshleifer, D. (1989): "Determinants of hedging and risk premia in commodity futures markets," *Journal of Financial and Quantitative Analysis*, 24(03), 313–331.

——— (1990): "Hedging pressure and futures price movements in a general equilibrium model," *Econometrica*, 58(2), 411–428.

Jiang, G., and Y. Tian (2005): "The model-free implied volatility and its information content," *Review of Financial Studies*, 18, 1305–1342.

——— (2007): "Extracting model-free volatility from option prices: An examination of the VIX index," *Journal of Derivatives*, 14, 35–60.

Johansen, S. (1988): "Statistical analysis of cointegration vectors," *Journal of Economic Dynamics and Control*, 12(2), 231–254.

——— (1991): "Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models," *Econometrica*, 59(6), 1551–1580.

Jones, C. (2003): "The dynamics of stochastic volatility: Evidence from underlying and options markets," *Journal of Econometrics*, 116, 181–224.

Julier, S. J. (2002): "The scaled unscented transformation," in *American Control Conference. Proceedings of the 2002*, vol. 6, pp. 4555–4559.

Kamara, A., and T. W. Miller (1995): "Daily and intradaily tests of European put-call parity," *Journal of Financial and Quantitative Analysis*, 30(4), 519–539.

Keynes, J. M. (1930): *A treatise on money: The applied theory of money*. London: Macmillan.

Kokholm, T., and M. Stisen (2015): "Joint pricing of VIX and SPX options with stochastic volatility and jump models," *Journal of Risk Finance*, 16(1), 27–48.

Kwiatkowski, D., P. C. Phillips, P. Schmidt, and Y. Shin (1992): "Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?," *Journal of Econometrics*, 54, 159–178.

Liu, J., and F. A. Longstaff (2004): "Losing money on arbitrage: Optimal dynamic portfolio choice in markets with arbitrage opportunities," *Review of Financial Studies*, 17(3), 611–641.

Madan, D. B., and M. Pistorius (2014): "Joint calibration of SPX and VIX option surfaces: With applications to pricing and hedging equity and volatility linked hybrid notes," Working paper, University of Maryland.

Mancini Griffoli, T., and A. Ranaldo (2012): "Limits to arbitrage during the crisis: Finding liquidity constraints and covered interest parity," Working paper, University of St. Gallen.

Martin, I. (2013): "Simple variance swaps," Working paper, Stanford University and National Bureau of Economic Research.

Mixon, S., and E. Onur (2015): "Volatility derivatives in practice: Activity and impact," Staff report, Commodity Futures Trading Commission.

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

Oehmke, M. (2011): "Gradual arbitrage," Working paper, Columbia Business School.

Pasquariello, P. (2014): "Financial market dislocations," *Review of Financial Studies*, 27(6), 1868–1914.

Roll, R., E. Schwartz, and A. Subrahmanyam (2007): "Liquidity and the law of one price: The case of the futures-cash basis," *Journal of Finance*, 62(5), 2201–2234.

Shleifer, A., and R. W. Vishny (1997): "The limits of arbitrage," *Journal of Finance*, 52(1), 35–55.

Song, Z., and D. Xiu (2016): "A tale of two option markets: Pricing kernels and volatility risk," *Journal of Econometrics*, 190(1), 176–196.

Appendix A Proof for Proposition 1

Let me start with the basic mathematical result (see, for example, Carr and Madan, 2001):

$$f(S) = f(S_0) + f'(S_0)(S - S_0) + \int_{S_0}^{\infty} f''(K) \max(S - K, 0)dK + \int_{0}^{S_0} f''(K) \max(K - S, 0)dK, \tag{A.1}$$

where $f(S)$ is a twice differentiable function. Assuming $f(S) = S^2$ and replacing S and S_0 with VIX_T and $VF_t(T)$, respectively, I obtain

$$\begin{aligned} VIX_T^2 &= VF_t(T)^2 + 2VF_t(T)(VIX_T - VF_t(T)) \\ &\quad + 2\int_{VF_t(T)}^{\infty} \max(VIX_T - K, 0)dK + 2\int_{0}^{VF_t(T)} \max(K - VIX_T, 0)dK. \end{aligned} \tag{A.2}$$

Taking a risk-neutral expectation on Equation (A.2) leads to

$$\begin{aligned} E_t^{\mathbb{Q}}[VIX_T^2] &= VF_t(T)^2 + 2VF_t(T)(E_t^{\mathbb{Q}}(VIX_T) - VF_t(T)) \\ &\quad + 2\int_{VF_t(T)}^{\infty} E_t^{\mathbb{Q}}[\max(VIX_T - K, 0)]dK + 2\int_{0}^{VF_t(T)} E_t^{\mathbb{Q}}[\max(K - VIX_T, 0)]dK. \end{aligned} \tag{A.3}$$

Substituting the relations,

$$\begin{aligned} VF_t(T) &= E_t^{\mathbb{Q}}(VIX_T), \\ VC_t(T,K) &= \exp(-r_t(T-t))E_t^{\mathbb{Q}}[\max(VIX_T - K, 0)], \text{ and} \\ VP_t(T,K) &= \exp(-r_t(T-t))E_t^{\mathbb{Q}}[\max(K - VIX_T, 0)], \end{aligned}$$

into Equation (A.3), I obtain

$$\begin{aligned} \mathrm{var}_t(VIX_T) &= E_t^{\mathbb{Q}}[VIX_T^2] - VF_t(T)^2 \\ &= 2\exp(r_t(T-t)) \left[\int_{VF_t(T)}^{\infty} VC_t(T,K)dK + \int_{0}^{VF_t(T)} VP_t(T,K)dK \right]. \end{aligned} \tag{A.4}$$

This equation means that the variance of VIX_T can be replicated by an equally-weighted continuum of the OTM VIX options. Q.E.D.

Appendix B Unscented Kalman filtering

Suppose that the latent state vector, $X_t = (v_{1t}, v_{2t}, u_{1t}, u_{2t})'$, is of the form:

$$X_{t+1} = \Phi_0 + \Phi_x X_t + \epsilon_{t+1}, \tag{B.1}$$

where $\text{var}_t[\epsilon_{t+1}] = Q(X_t)$. Kalman filtering comprises two stages: the prediction and the update stages. The prediction stage involves computing the one-step-ahead predictive mean and variance of X_t given $\hat{X}_t = E_t[X_t]$ and $P_t = \text{var}_t[X_t]$:

$$\begin{aligned} \hat{X}_{t+1|t} &= \Phi_0 + \Phi_x \hat{X}_t \\ P_{t+1|t} &= \Phi_x P_t \Phi_x' + Q(X_t). \end{aligned} \tag{B.2}$$

The update stage starts by defining a set of sigma points that match the predictive mean and variance of the latent vector. To be specific, a set of $2L+1$ sigma points and weights are selected based on the scaled unscented transformation of Julier (2002), where L is the dimension of the latent vector:

$$\begin{aligned} \mathscr{X}_0 &= \hat{X}_{t+1|t} \\ \mathscr{X}_i &= \hat{X}_{t+1|t} + \left(\sqrt{(L+\xi)P_{t+1|t}}\right)_i, \quad i = 1, \ldots, L \\ \mathscr{X}_i &= \hat{X}_{t+1|t} - \left(\sqrt{(L+\xi)P_{t+1|t}}\right)_{i-L}, \quad i = L+1, \ldots, 2L, \end{aligned} \tag{B.3}$$

whose weights are given by

$$\begin{aligned} W_0^m &= \frac{\xi}{L+\xi} \\ W_0^c &= \frac{\xi}{L+\xi} + (1 - \alpha^2 + \beta) \\ W_i^m &= \frac{1}{2(L+\xi)}, \quad i = 1, \ldots, 2L \\ W_i^c &= \frac{1}{2(L+\xi)}, \quad i = 1, \ldots, 2L, \end{aligned} \tag{B.4}$$

where $\xi = \alpha^2(L+\kappa) - L$ and $\left(\sqrt{(L+\xi)P_{t+1|t}}\right)_i$ is the ith column of the matrix square root. In my applications, I set $\alpha = 0.01$, $\beta = 2$, and $\kappa = 0$.

Suppose that the observation equations are of the form:

$$Y_t = H(X_t) + \Omega, \tag{B.5}$$

where $H(X_t)$ represents the pricing formula for the SIV and the VIV and Ω denotes the measurement errors. The predictive mean and variance of Y_t are computed using the sigma points and weights:

$$\begin{aligned} \hat{Y}_{t+1|t} &= \sum_{i=0}^{2L} W_i^m H(\mathscr{X}_i) \\ P^y_{t+1|t} &= \sum_{i=0}^{2L} W_i^c (H(\mathscr{X}_i) - \hat{Y}_{t+1|t})(H(\mathscr{X}_i) - \hat{Y}_{t+1|t})' + \Omega, \end{aligned} \tag{B.6}$$

where $\hat{Y}_{t+1|t} = E_t[Y_{t+1}]$ and $P^y_{t+1|t} = \mathrm{var}_t[Y_{t+1}]$.

Finally, I am able to update the conditional mean and variance of X_t:

$$\begin{aligned} \hat{X}_{t+1} &= \hat{X}_{t+1|t} + K_{t+1}(Y_{t+1} - \hat{Y}_{t+1|t}) \\ P_{t+1} &= P_{t+1|t} - K_{t+1} P^y_{t+1|t} K'_{t+1}, \end{aligned} \tag{B.7}$$

where

$$K_{t+1} = \sum_{i=0}^{2L} W_i^c (\mathscr{X}_i - \hat{X}_{t+1|t})(H(\mathscr{X}_i) - \hat{Y}_{t+1|t})'(P^y_{t+1|t})^{-1}. \tag{B.8}$$

The log likelihood function is then given by

$$\log(L) = -\frac{1}{2}\sum_{t=1}^{T}\left[\log(2\pi)N_t + \log\left(|\det(P^y_{t+1|t})|\right) + (Y_{t+1} - \hat{Y}_{t+1|t})'(P^y_{t+1|t})^{-1}(Y_{t+1} - \hat{Y}_{t+1|t})\right], \tag{B.9}$$

where N_t is the dimension of the observation equations at time t. The log likelihood function is optimized via the Berndt, Hall, Hall, and Hausman (1974) algorithm. The standard errors are computed by inverting the Hessian matrix at the optimal solution.

Table 1: **Summary statistics of convexity ratios**

The table presents summary statistics of the convexity ratio, $CVRT_t(T)$, which is defined as the ratio of the convexity adjustment to the VIV measure:

$$CVRT_t(T) = \frac{\text{var}_t(VIX_T)}{VIV_t(T)},$$

where the numerator and the denominator are computed as in Equations (6) and (7), respectively. The short and long terms are defined as less than three months and more than three months, respectively.

Maturity	1%	5%	Median	95%	99%	Mean	Std.	Skew.	Kurt.
Short term	0.020	0.033	0.09	0.16	0.18	0.09	0.038	0.11	2.33
Long term	0.082	0.097	0.16	0.20	0.22	0.16	0.032	-0.31	2.83
Total	0.023	0.044	0.13	0.19	0.21	0.13	0.047	-0.25	2.30

Table 2: **Summary statistics of the replicated and original VIV measures**

The table presents the summary statistics of the square root of the replicated and original VIV measures with one-, three-, and five-month maturities, which I denote by $\widetilde{VIV_t}(T)$ and $VIV_t(T)$, respectively. The replicated VIV measure is computed from a pair of SIV measures, which are extracted from the SPX option prices. In contrast, the original VIV measure is computed directly from the prices of VIX futures and options. The pre-crisis period (in Panel A) and the post-crisis period (in Panel B) are defined as July 1, 2006 to August 31, 2008; and December 1, 2008 to August 31, 2014, respectively.

Variable	Mean	Median	Min.	Max.	Std.	Skew.	Kurt.	AR(1)
Panel A: Pre-crisis period								
Replicated VIV measures:								
$\sqrt{\widetilde{VIV_t}(1)}$	0.19	0.19	0.11	0.29	0.05	0.11	1.53	0.98
$\sqrt{\widetilde{VIV_t}(3)}$	0.20	0.20	0.13	0.28	0.05	0.13	1.41	0.98
$\sqrt{\widetilde{VIV_t}(5)}$	0.20	0.19	0.14	0.39	0.04	0.27	1.87	0.97
Original VIV measures:								
$\sqrt{VIV_t(1)}$	0.20	0.21	0.12	0.29	0.05	0.00	1.47	0.99
$\sqrt{VIV_t(3)}$	0.21	0.22	0.14	0.29	0.05	-0.01	1.37	0.99
$\sqrt{VIV_t(5)}$	0.21	0.22	0.15	0.29	0.04	0.00	1.38	0.99
Panel B: Post-crisis period								
Replicated VIV measures:								
$\sqrt{\widetilde{VIV_t}(1)}$	0.23	0.21	0.13	0.60	0.09	1.42	5.03	0.98
$\sqrt{\widetilde{VIV_t}(3)}$	0.26	0.24	0.15	0.55	0.08	0.93	3.68	0.99
$\sqrt{\widetilde{VIV_t}(5)}$	0.27	0.26	0.17	0.52	0.07	0.70	3.05	0.98
Original VIV measures:								
$\sqrt{VIV_t(1)}$	0.24	0.21	0.13	0.61	0.09	1.35	4.81	0.99
$\sqrt{VIV_t(3)}$	0.26	0.25	0.15	0.54	0.08	0.84	3.48	0.99
$\sqrt{VIV_t(5)}$	0.28	0.27	0.17	0.49	0.07	0.53	2.76	0.99

Table 3: **A model-free analysis of price discovery**

The table presents the unit-root test results, the co-integration test results, and the price discovery results for the logarithm of replicated and original VIV measures, which I denote by $\widetilde{VIV_t(T)}$ and $VIV_t(T)$, respectively. The replicated VIV measure is computed from a pair of SIV measures, which are extracted from the SPX option prices. In contrast, the original VIV measure is computed directly from the prices of VIX futures and options. The KPSS (Kwiatkowski, Phillips, Schmidt, and Shin, 1992) test evaluates the null hypothesis that the time series is stationary against the existence of a unit root. In contrast, the ADF (Augmented Dickey Fuller) test assesses the null of the existence of a unit root. The Johansen maximum eigenvalue test assesses the null hypothesis that the number of co-integrating relations is equal to r against $r+1$. GG (Gonzalo and Granger, 1995) measures the relative contribution of innovations in the VIX derivative prices to innovations in the efficient prices. HAS_L and HAS_U denote the lower and upper bounds of the VIX derivatives' information share, respectively, as introduced by Hasbrouck (1995). The pre-crisis period (in Panel A) and the post-crisis period (in Panel B) are defined as July 1, 2006 to August 31, 2008; and December 1, 2008 to August 31, 2014, respectively. The symbols *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	KPSS test		ADF test		Johansen test		Speeds-of-adjustment		Price discovery ratios		
Maturity	$\log \widetilde{VIV}$	$\log VIV$	$\log \widetilde{VIV}$	$\log VIV$	$H_0: r=0$	$H_0: r=1$	α_1	α_2	GG	HAS_L	HAS_U
Panel A: Pre-crisis period											
1 month	4.56***	4.81***	−0.73	−0.75	92.36***	2.01	−0.37	0.02	0.94	0.19	1.00
3 month	4.61***	5.11***	−0.83	−0.88	77.44***	1.61	−0.28	0.06	0.83	0.47	0.94
5 month	4.96***	5.34***	−0.99	−0.94	41.15***	1.65	−0.16	0.05	0.76	0.43	0.86
Panel B: Post-crisis period											
1 month	4.59***	4.58***	0.30	0.26	204.38***	9.77**	−0.15	0.21	0.42	0.05	0.92
3 month	5.61***	5.95***	0.62	0.67	155.42***	7.36	−0.19	0.09	0.67	0.21	0.93
5 month	6.82***	8.35***	0.85	0.87	154.16***	5.69	−0.26	0.01	0.96	0.59	1.00

Table 4: **Parameter estimates and information criteria**

This table shows the parameter estimates and the information criteria. Panels A, B, and C correspond to the full sample period (July 1, 2006 to August 31, 2014); the pre-crisis period (July 1, 2006 to August 31, 2008); and the post-crisis sample (December 1, 2008 to August 31, 2014), respectively. Standard errors are in parentheses. AIC and SIC stand for Akaike and Schwarz information criteria, respectively. PDR measures the relative contribution of the VIX derivatives to the price discovery of volatility risk, which is defined in Equation (21).

	Panel A: Full sample period			**Panel B: Pre-crisis period**			**Panel C: Post-crisis period**		
	SV	SVEC1	SVEC2	SV	SVEC1	SVEC2	SV	SVEC1	SVEC2
μ_1	-3.46	-4.29	-7.60	-3.77	-5.14	-2.29	-2.95	-3.21	-3.95
	(0.03)	(0.06)	(0.19)	(0.06)	(0.48)	(0.19)	(0.02)	(0.02)	(0.08)
μ_2	-3.17	-2.97	-5.37	-3.14	-2.61	-1.34	-2.74	-2.79	-3.50
	(0.03)	(0.03)	(0.12)	(0.05)	(0.10)	(0.25)	(0.02)	(0.02)	(0.06)
κ_{v_1}	5.94	4.33	4.57	10.04	8.48	7.42	6.27	5.01	5.91
	(0.05)	(0.21)	(0.23)	(0.18)	(0.46)	(0.40)	(0.08)	(0.27)	(0.25)
κ_{v_2}	4.46	4.15	4.08	5.30	3.67	4.13	5.05	3.63	2.74
	(0.02)	(0.08)	(0.09)	(0.06)	(0.19)	(0.16)	(0.05)	(0.16)	(0.13)
σ_{v_1}	2.30	4.70	4.70	2.58	5.13	4.97	2.24	4.43	4.45
	(0.04)	(0.11)	(0.10)	(0.09)	(0.22)	(0.20)	(0.05)	(0.11)	(0.11)
σ_{v_2}	1.76	1.81	1.84	1.81	2.07	1.94	1.82	1.74	1.77
	(0.04)	(0.04)	(0.04)	(0.09)	(0.13)	(0.12)	(0.06)	(0.06)	(0.05)
η_{v_1}	-0.75	7.43	7.89	-16.76	-3.55	-8.03	-5.32	-6.01	-1.00
	(0.62)	(1.62)	(1.78)	(5.13)	(7.95)	(7.79)	(1.96)	(2.91)	(2.78)
η_{v_2}	1.58	4.30	3.80	-1.21	2.63	3.58	1.45	7.41	4.36
	(0.89)	(1.06)	(1.17)	(3.82)	(4.23)	(3.72)	(1.58)	(1.32)	(1.32)
κ_{u_1}	0.52	0.24	0.28	0.53	0.07	0.03	0.80	0.59	0.50
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.02)
κ_{u_2}	0.24	0.22	0.24	0.20	0.14	0.14	0.63	0.53	0.41
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
σ_{u_1}	1.07	0.74	0.70	0.92	0.58	0.53	0.94	0.60	0.50
	(0.02)	(0.01)	(0.01)	(0.04)	(0.03)	(0.03)	(0.02)	(0.01)	(0.02)
σ_{u_2}	0.58	0.48	0.51	0.55	0.42	0.42	0.73	0.51	0.49
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.03)
η_{u_1}	-0.66	0.10	0.06	-0.44	0.19	-0.40	-0.21	-0.21	-0.22
	(0.43)	(0.20)	(0.06)	(1.39)	(0.30)	(0.38)	(0.65)	(0.42)	(0.21)
η_{u_2}	-0.19	-0.18	-0.03	-0.71	-0.27	-0.16	-0.02	0.08	-0.42
	(0.36)	(0.34)	(0.08)	(1.32)	(0.50)	(0.22)	(0.65)	(0.47)	(0.25)
γ_{v_1}		-33.78	-34.60		-29.88	-29.04		-29.11	-30.28
		(1.42)	(1.45)		(1.82)	(1.62)		(1.54)	(1.54)
γ_{v_2}		12.33	12.55		13.11	12.84		16.53	16.15
		(0.32)	(0.35)		(0.78)	(0.71)		(0.63)	(0.61)
γ_{u_1}			0.49			-0.21			1.24
			(0.03)			(0.07)			(0.08)
γ_{u_2}			0.32			-0.32			1.30
			(0.03)			(0.06)			(0.10)
σ_{e_1}	0.061	0.054	0.053	0.036	0.031	0.030	0.051	0.045	0.044
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
σ_{e_2}	0.029	0.029	0.029	0.022	0.024	0.023	0.028	0.029	0.029
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
log(L)	34,938	36,728	36,790	10,527	11,055	11,065	27,084	28,799	28,898
AIC	-69,843	-73,420	-73,541	-21,021	-22,074	-22,090	-54,137	-57,561	-57,757
SIC	-69,753	-73,319	-73,428	-20,953	-21,996	-22,004	-54,052	-57,466	-57,651
PDR		0.73	0.73		0.70	0.69		0.64	0.65

Table 5: **Variance pricing errors**

This table shows the variance pricing performance across different models. Panels A, B, and C correspond to the full sample period (July 1, 2006 to August 31, 2014); the pre-crisis period (July 1, 2006 to August 31, 2008); and the post-crisis sample (December 1, 2008 to August 31, 2014), respectively. The root mean squared error is separately defined for the SIV and the VIV as follows:

$$RMSE_{i,SPX} = \sqrt{\frac{1}{N_{SPX}} \sum_{n=1}^{N_{SPX}} \Big(\log(SIV_t(T_n; \mathcal{M}_i)) - \log(SIV_t(T_n)) \Big)^2}$$

$$RMSE_{i,VIX} = \sqrt{\frac{1}{N_{VIX}} \sum_{n=1}^{N_{VIX}} \Big(\log(VIV_t(T_n; \mathcal{M}_i)) - \log(VIV_t(T_n)) \Big)^2},$$

where $RMSE_{i,SPX}$ and $RMSE_{i,VIX}$ denote the RMSEs of the model i for the SIV and the VIV, respectively, and N_{SPX} and N_{VIX} indicate the total number of SIV and VIV observations included in the sample, respectively. The table also shows a measure of the pricing improvement of a model against the SV model, denoted by $\Delta RMSE_i$, which is defined as the logarithmic difference between the $RMSE$ of a particular model i and that of the SV model. A negative (positive) value of $\Delta RMSE_i$ means that the corresponding model yields lower (higher) pricing errors than the SV model. The pricing errors are broken down into short terms (less than three months to maturity) and long terms (more than three months to maturity).

	Valuation of the SIV						Valuation of the VIV					
	$RMSE_{i,SPX}$			$\Delta RMSE_{i,SPX}$			$RMSE_{i,VIX}$			$\Delta RMSE_{i,VIX}$		
	Short term	Long term	Total	Short term	Long term	Total	Short term	Long term	Total	Short term	Long term	Total
Panel A: Full sample period												
SV	0.048	0.057	0.053	0.00	0.00	0.00	0.027	0.021	0.024	0.00	0.00	0.00
SVEC1	0.034	0.056	0.048	-0.34	-0.00	-0.11	0.028	0.023	0.026	0.05	0.07	0.05
SVEC2	0.034	0.055	0.047	-0.35	-0.02	-0.12	0.028	0.023	0.026	0.04	0.09	0.06
Panel B: Pre-crisis period												
SV	0.034	0.025	0.030	0.00	0.00	0.00	0.020	0.016	0.019	0.00	0.00	0.00
SVEC1	0.026	0.025	0.025	-0.26	-0.01	-0.16	0.021	0.019	0.020	0.06	0.13	0.09
SVEC2	0.026	0.025	0.026	-0.26	-0.01	-0.15	0.021	0.018	0.020	0.07	0.11	0.08
Panel C: Post-crisis period												
SV	0.040	0.047	0.044	0.00	0.00	0.00	0.026	0.022	0.024	0.00	0.00	0.00
SVEC1	0.024	0.047	0.039	-0.48	-0.01	-0.13	0.028	0.023	0.026	0.07	0.08	0.07
SVEC2	0.025	0.045	0.038	-0.46	-0.05	-0.16	0.028	0.025	0.026	0.08	0.13	0.10

Table 6: **Pairwise model comparison tests**

This table shows the Diebold and Mariano (2002) test statistics for the performance differences in the pricing of the SIV and the VIV. A negative (positive) statistic in a cell (i, j) indicates that the model i outperforms (underperforms) the model j. Panels A, B, and C correspond to the full sample period (July 1, 2006 to August 31, 2014); the pre-crisis period (July 1, 2006 to August 31, 2008); and the post-crisis sample (December 1, 2008 to August 31, 2014), respectively. The symbols *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Valuation of the SIV			Valuation of the VIV		
	SV	SVEC1	SVEC2	SV	SVEC1	SVEC2
Panel A: Full sample						
SV	0.00	9.18***	10.09***	0.00	−8.54***	−8.52***
SVEC1	−9.18***	0.00	5.57***	8.54***	0.00	−3.19***
SVEC2	−10.09***	−5.57***	0.00	8.52***	3.19***	0.00
Panel B: Pre-crisis period						
SV	0.00	3.50***	3.48***	0.00	−6.41***	−5.63***
SVEC1	−3.50***	0.00	−0.49	6.41***	0.00	0.56
SVEC2	−3.48***	0.49	0.00	5.63***	−0.56	0.00
Panel C: Post-crisis period						
SV	0.00	7.06***	7.80***	0.00	−7.50***	−8.35***
SVEC1	−7.06***	0.00	4.39***	7.50***	0.00	−6.70***
SVEC2	−7.80***	−4.39***	0.00	8.35***	6.70***	0.00

Table 7: **Summary statistics and correlation matrix**

BSS^{MF} denotes the price dislocation indicator based on model-free bases, BSS^{SV} denotes the price dislocation indicator based on model-implied bases, $LIBOIS$ denotes the Libor-OIS spread, $SPRDS$ denotes the relative bid-ask spread of the SPX options, $SPRDF$ denotes the relative bid-ask spread of the VIX futures, $SPRDV$ denotes the relative bid-ask spread of the VIX options, VIX_t denotes the CBOE VIX index, and $SKEW$ denotes the CBOE SKEW index.

	BSS^{MF}	BSS^{SV}	$LIBOIS$	$SPRDS$	$SPRDF$	$SPRDV$	VIX	$SKEW$
Panel A: Summary statistics								
Mean	-0.07	-0.06	0.35	0.29	0.01	0.22	0.22	120.48
Median	-0.06	-0.06	0.16	0.28	0.01	0.21	0.19	119.92
Min.	-1.29	-1.03	0.06	0.16	0.00	0.13	0.10	0.00
Max.	0.10	0.22	3.64	0.81	0.02	0.38	0.81	143.26
Std.	0.07	0.08	0.43	0.06	0.00	0.04	0.10	6.47
Skew.	-7.09	-4.51	3.47	2.32	1.52	0.60	2.11	-2.63
Kurt.	89.41	43.96	19.24	16.84	6.26	3.22	8.66	61.32
AR(1)	0.32	0.80	1.00	0.96	0.96	0.97	0.98	0.76
Panel B: Correlation matrix								
BSS^{MF}	1.00	0.44	-0.11	-0.13	-0.18	-0.03	-0.07	0.09
BSS^{SV}		1.00	-0.55	-0.46	-0.29	-0.24	-0.36	0.17
$LIBOIS$			1.00	0.39	0.43	0.47	0.78	-0.34
$SPRDS$				1.00	0.03	-0.04	0.41	-0.00
$SPRDF$					1.00	0.43	0.34	-0.29
$SPRDV$						1.00	0.35	-0.26
VIX							1.00	-0.31
$SKEW$								1.00

Table 8: **Drivers of price dislocations**

This table shows the regression results of the price dislocation indicators, BSS^{MF} and BSS^{SV}, on the potential determinants of price dislocations. The full sample (in Panel A) spans from July 1, 2006 through August 31, 2014. The ex-crisis sample (in Panel B) is the one excluding the three-month window (September through November 2008), which corresponds to the peak of the Lehman Brothers crisis. BSS^{MF} denotes the price dislocation indicator based on model-free bases, BSS^{SV} denotes the price dislocation indicator based on model-implied bases, $LIBOIS$ denotes the Libor-OIS spread, $SPRDS$ denotes the relative bid-ask spread of the SPX options, $SPRDF$ denotes the relative bid-ask spread of the VIX futures, $SPRDV$ denotes the relative bid-ask spread of the VIX options, VIX_t denotes the CBOE VIX index, and $SKEW$ denotes the CBOE SKEW index. Newey and West (1987) robust t-statistics with an optimal lag are shown in parentheses. The symbols *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent variable	const.	$LIBOIS$	$SPRDS$	$SPRDF$	$SPRDV$	VIX	$SKEW$	adj. R^2	Nobs
Panel A: Full sample									
Case 1: Funding liquidity									
BSS^{MF}	−0.04***	−0.02***						0.16	1992
	(-7.44)	(-4.39)							
BSS^{SV}	−0.03***	−0.04***						0.30	2050
	(-4.15)	(-4.30)							
Case 2: Market liquidity									
BSS^{MF}	0.12**		−0.02***	−0.01***	−0.01**			0.15	1992
	(2.43)		(-3.22)	(-3.58)	(-2.30)				
BSS^{SV}	0.22***		−0.04***	−0.02***	−0.01***			0.31	2050
	(3.43)		(-3.80)	(-3.60)	(-3.51)				
Case 3: Funding and market liquidity									
BSS^{MF}	0.06*	−0.02***	−0.01***	−0.01**	−0.00			0.21	1992
	(1.88)	(-3.79)	(-3.43)	(-2.49)	(-0.51)				
BSS^{SV}	0.12***	−0.03***	−0.02***	−0.01**	−0.00			0.39	2050
	(2.85)	(-4.73)	(-3.94)	(-2.38)	(-0.92)				
Case 4: All explanatory variables									
BSS^{MF}	−0.01	−0.03***	−0.02***	−0.01**	−0.00	0.01**	0.00	0.23	1992
	(-0.11)	(-4.02)	(-4.03)	(-2.29)	(-0.63)	(2.34)	(1.25)		
BSS^{SV}	0.06	−0.04***	−0.03***	−0.01**	−0.00	0.02***	0.00	0.42	2050
	(0.80)	(-5.38)	(-4.70)	(-2.34)	(-1.08)	(3.34)	(0.79)		
Panel B: Ex-crisis sample									
Case 1: Funding liquidity									
BSS^{MF}	−0.04***	−0.01***						0.05	1944
	(-11.17)	(-4.23)							
BSS^{SV}	−0.05***	−0.01***						0.04	1987
	(-11.61)	(-3.96)							
Case 2: Market liquidity									
BSS^{MF}	0.02		−0.01**	−0.01***	−0.00			0.05	1944
	(0.69)		(-2.49)	(-3.13)	(-0.63)				
BSS^{SV}	0.02		−0.01**	−0.01***	−0.00			0.04	1987
	(0.58)		(-2.39)	(-3.00)	(-0.67)				
Case 3: Funding and market liquidity									
BSS^{MF}	0.01	−0.01***	−0.01**	−0.01**	0.00			0.08	1944
	(0.36)	(-3.12)	(-2.56)	(-2.17)	(0.12)				
BSS^{SV}	0.01	−0.01***	−0.01**	−0.01**	−0.00			0.07	1987
	(0.30)	(-2.84)	(-2.43)	(-2.12)	(-0.03)				
Case 4: All explanatory variables									
BSS^{MF}	−0.12*	−0.01***	−0.01***	−0.01*	0.00	0.01**	0.01**	0.11	1944
	(-1.86)	(-3.54)	(-3.54)	(-1.78)	(0.08)	(2.30)	(2.25)		
BSS^{SV}	−0.11*	−0.01***	−0.01***	−0.01*	−0.00	0.01**	0.01**	0.09	1987
	(-1.73)	(-3.45)	(-3.42)	(-1.76)	(-0.06)	(2.42)	(2.06)		

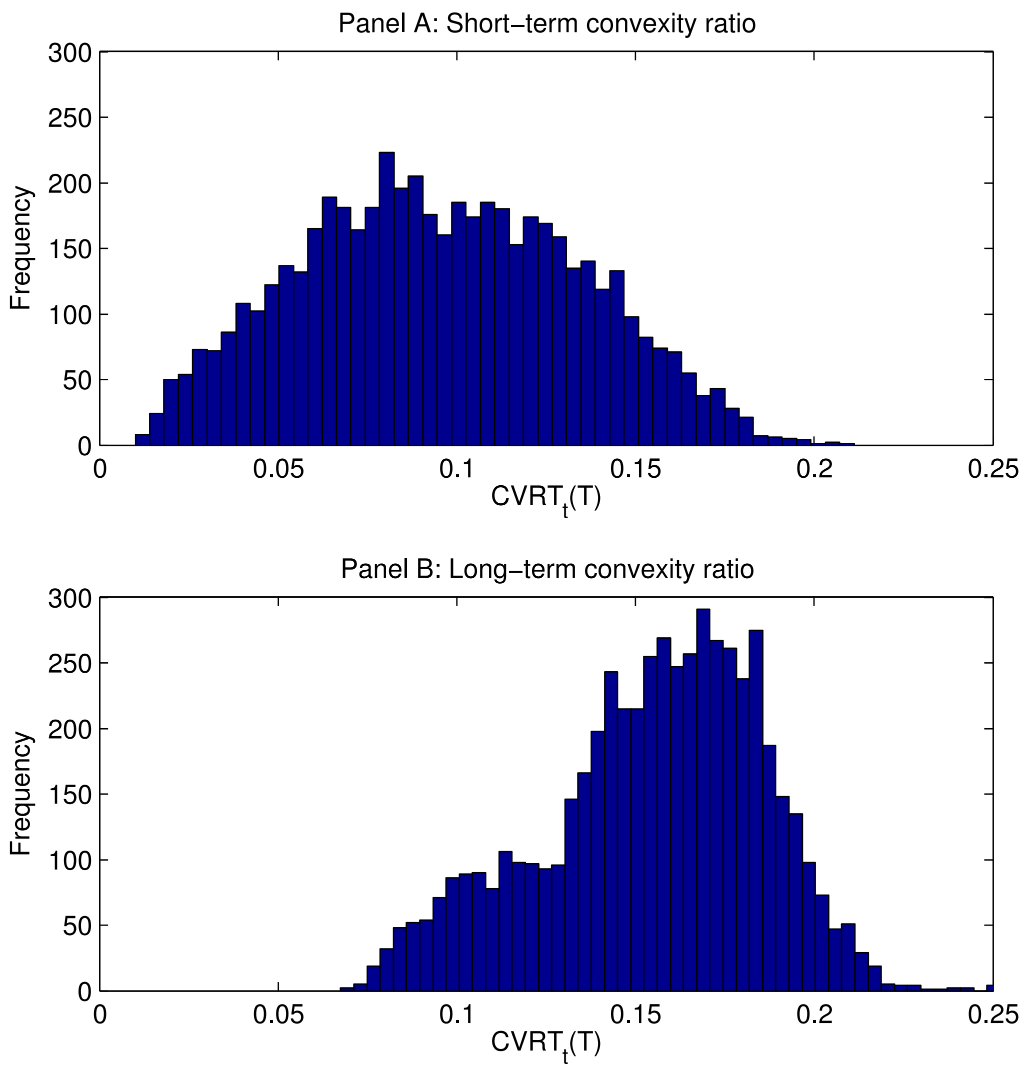

Figure 1: Histograms of the convexity ratios. Panels A and B correspond to the short- and long-term maturities, which are defined as less than three months and more than three months, respectively.

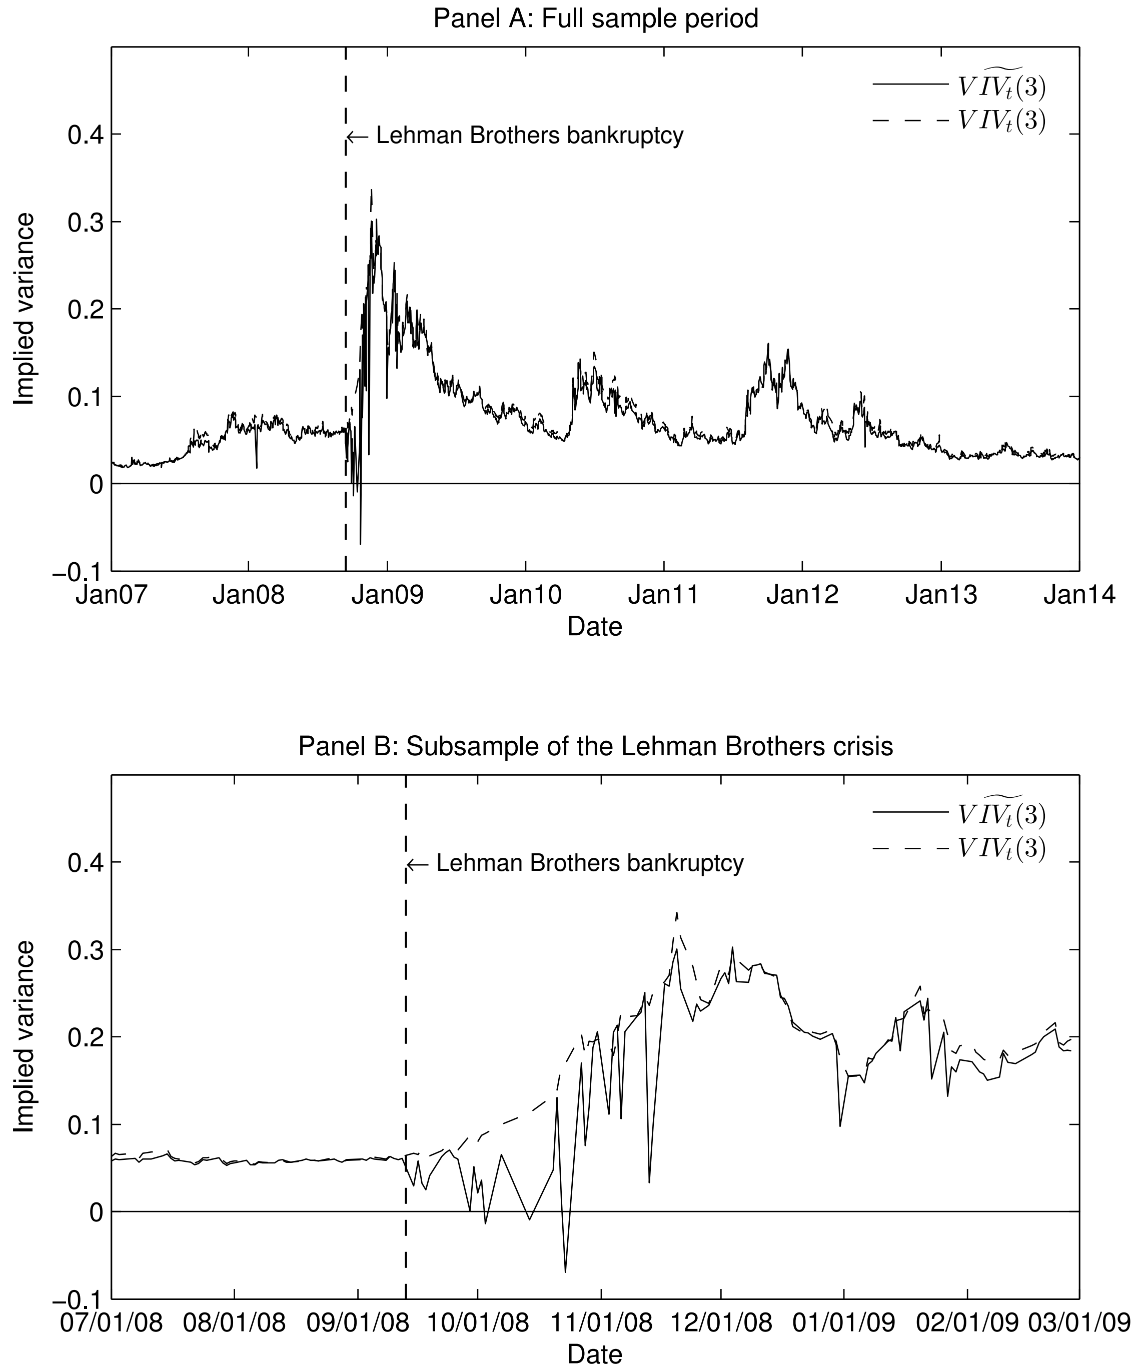

Figure 2: Time series plots of the replicated and original VIV measures with a three-month maturity. The replicated VIV measure (solid line) is computed from a pair of SIV measures, which are computed from the SPX option prices. In contrast, the original VIV measure (dotted line) is computed directly from the VIX derivative prices.

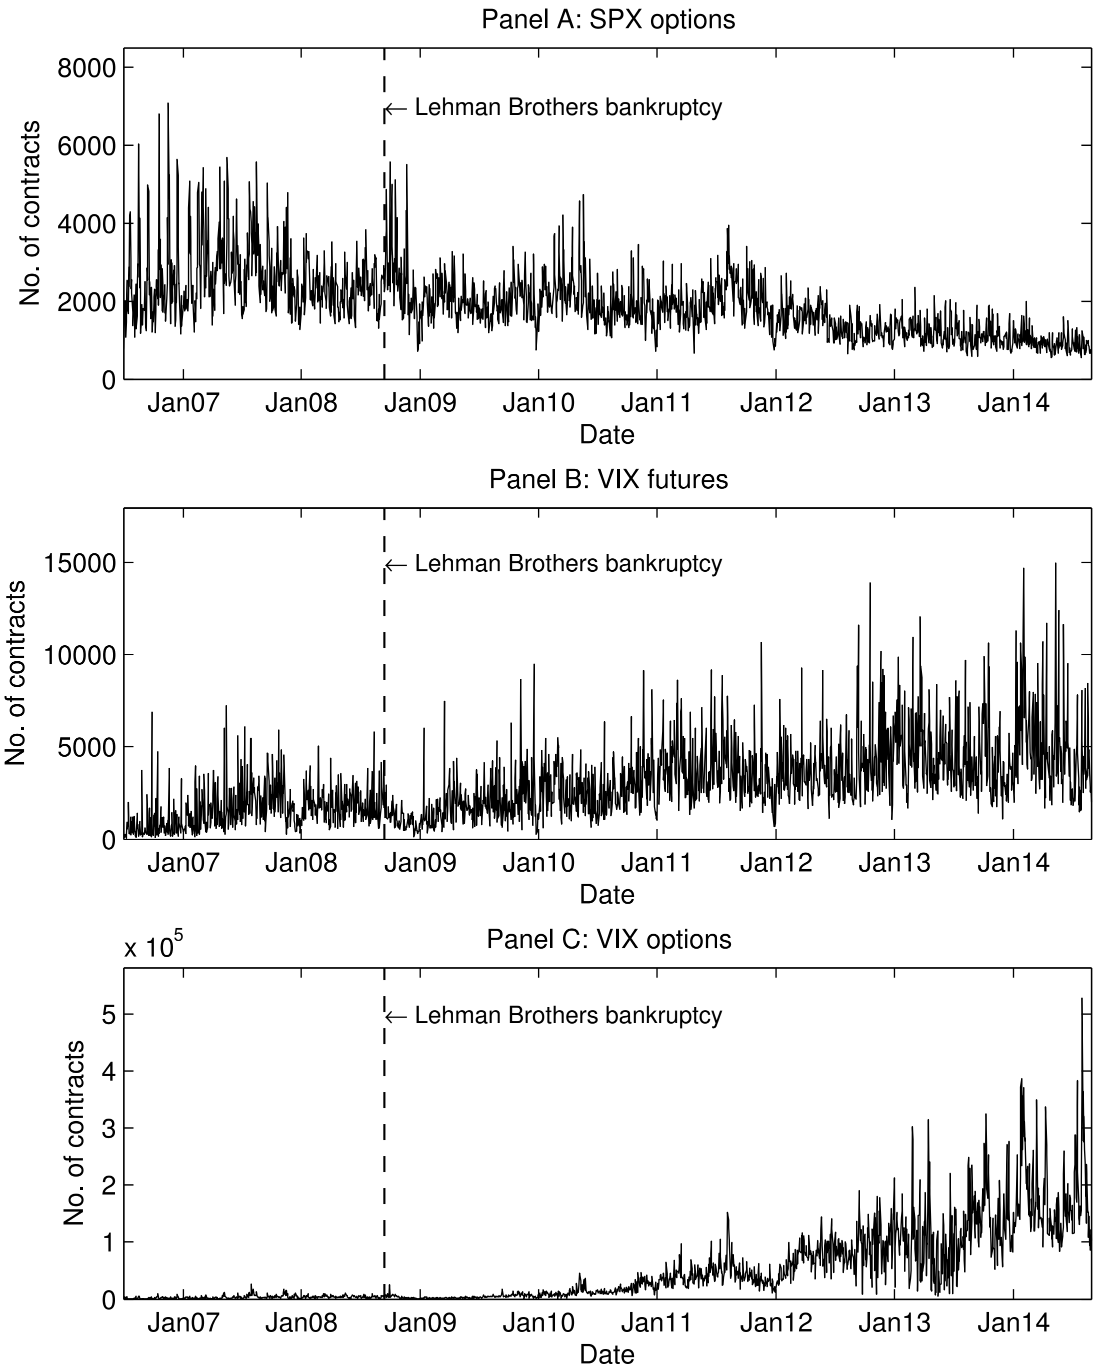

Figure 3: Time series plots of the trading volume for different derivative contracts. Panels A, B, and C correspond to SPX options, VIX futures, and VIX options, respectively.

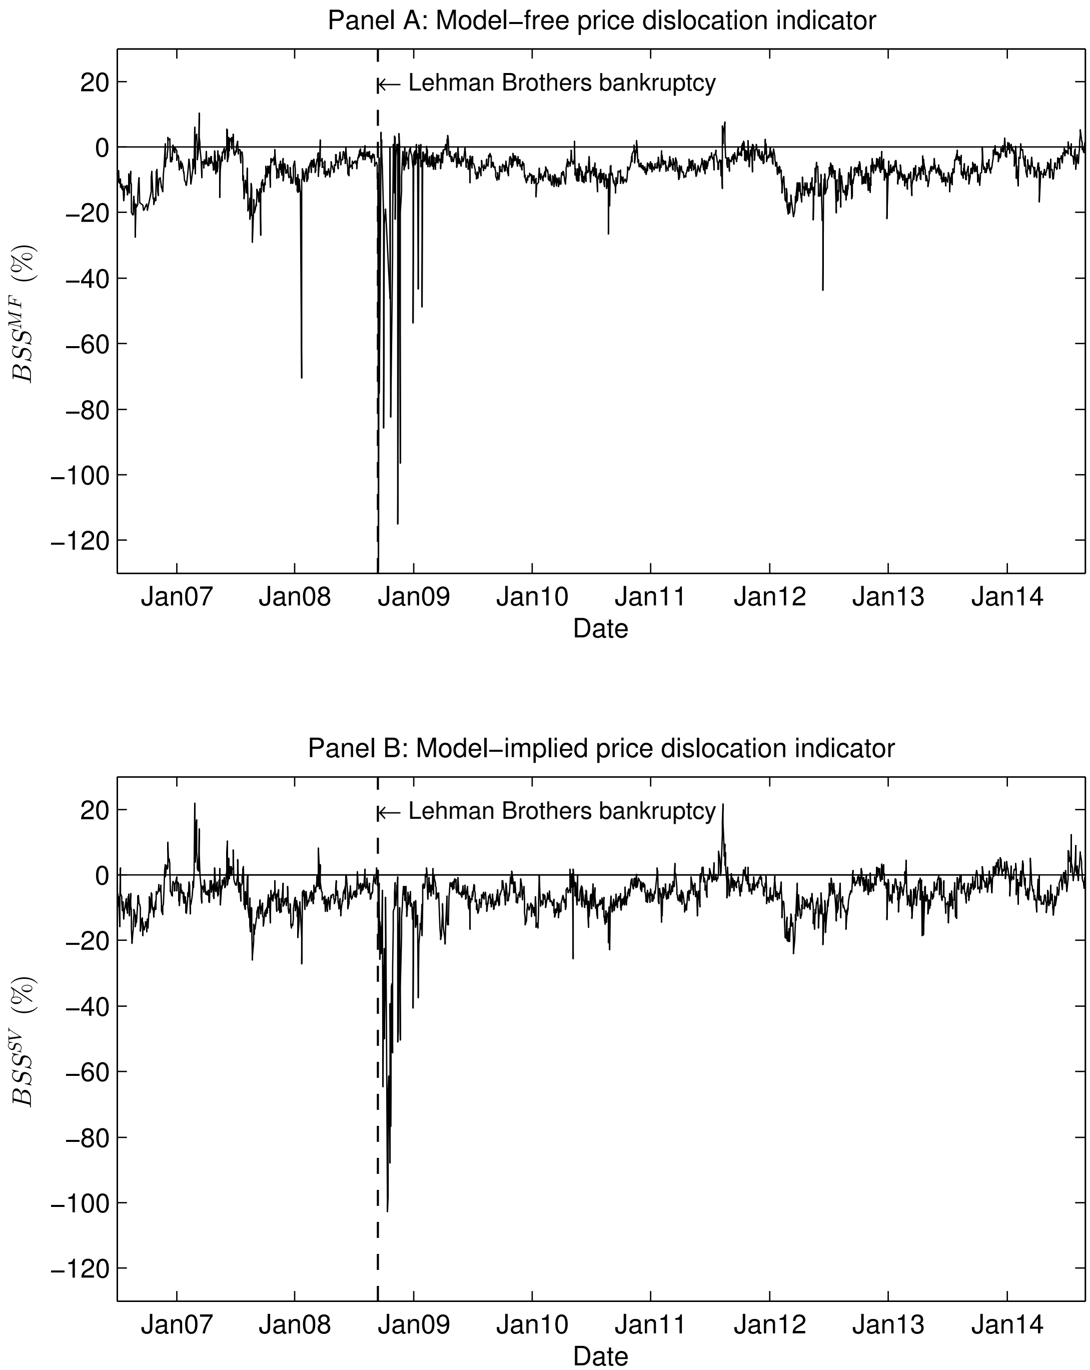

Figure 4: Time series plots of the pricing dislocation indicators. Panel A shows the price dislocation indicator based on model-free bases, whereas Panel B shows the price dislocation indicator based on model-implied bases.

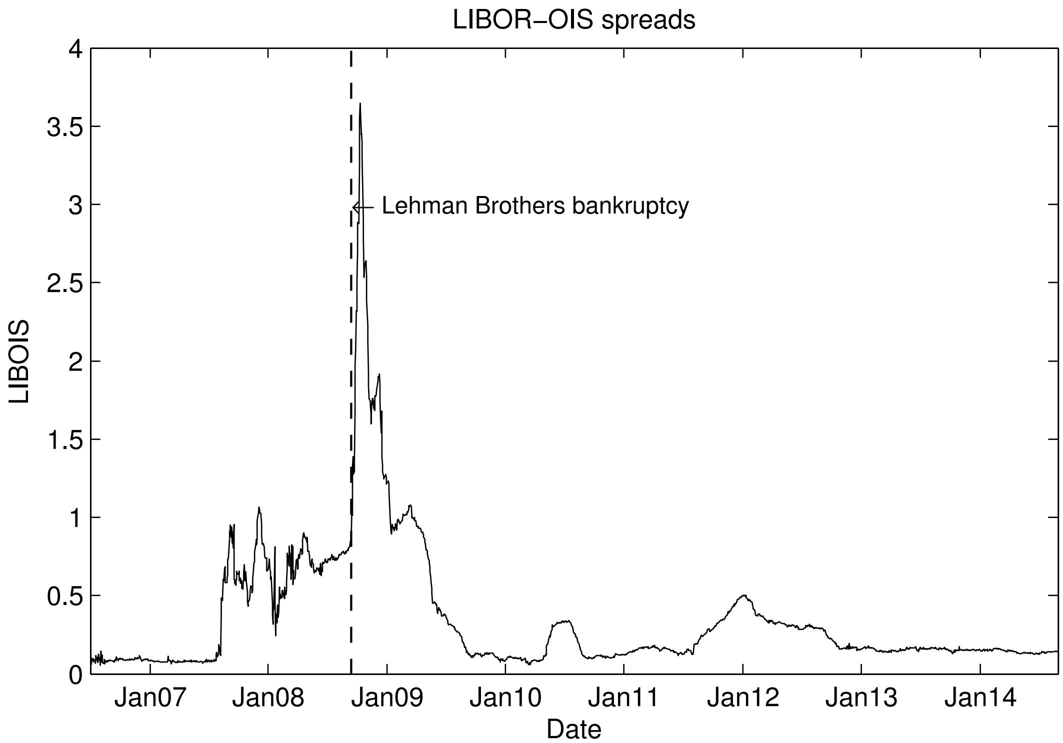

Figure 5: Time series plot of LIBOR-OIS spreads.

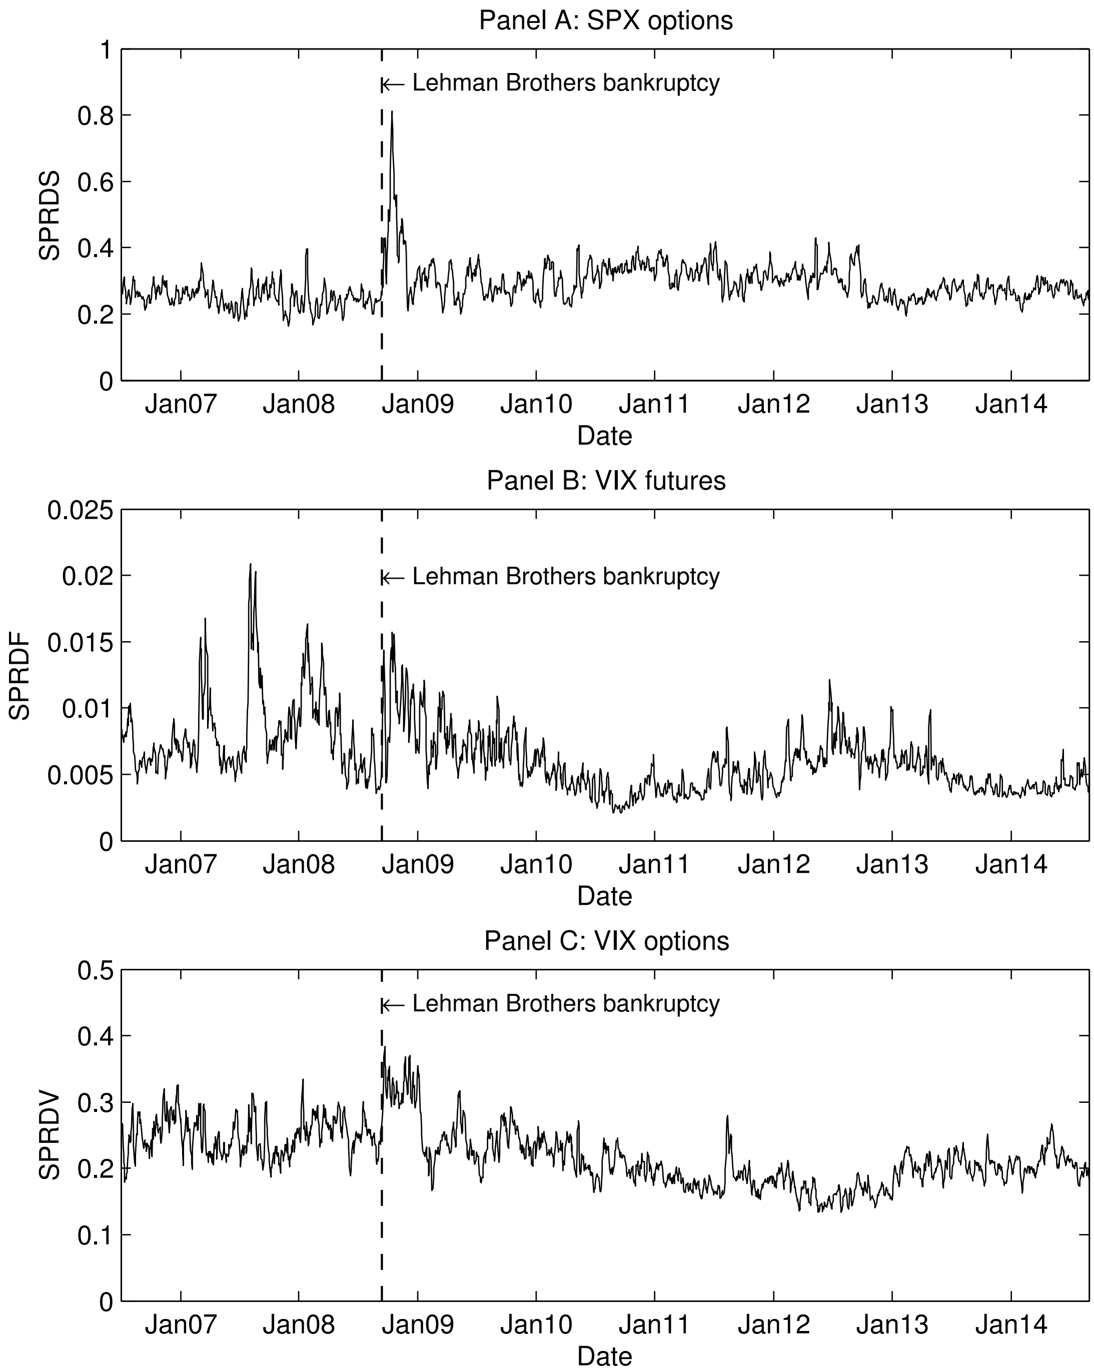

Figure 6: Time series plots of the relative bid-ask spreads for different derivatives markets. Panels A, B, and C correspond to the SPX options, the VIX futures, and the VIX options, respectively.