

Leverage Effect, Volatility Feedback, and Self-Exciting Market Disruptions

PETER CARR and LIUREN WU*

This version: September 20, 2011; First draft: September 4, 2007

ABSTRACT

The S&P 500 index return interacts negatively with its volatility. This paper traces the negative interaction to three distinct economic channels and proposes to disentangle the relative contribution of each channel using S&P 500 index options. First, equity volatility increases proportionally with the level of financial leverage, the variation of which is dictated by managerial decisions on a company's capital structure based on economic conditions. Second, irrespective of financial leverage, a positive shock to business risk increases the cost of capital and reduces the valuation of future cash flows, generating an instantaneous negative correlation between asset returns and asset volatility. Finally, large, negative market disruptions often generate self-exciting behaviors. The occurrence of one negative disruption induces more disruptions to follow, thus raising market volatility. Model estimation highlights the information in the large cross-section of equity index options in identifying the economic channels underlying the variations of the equity index and its volatility.

JEL Classification: F34, G12, G13.

Keywords: Leverage effect; volatility feedback; self-exciting; market disruptions; jumps; option pricing; implied volatility; constant elasticity of variance; structural models; time-changed Lévy process; fast Fourier transform; unscented Kalman filter.

*Carr is from Courant Institute, New York University; pcarr@nyc.rr.com. Wu is from Zicklin School of Business, Baruch College, The City University of New York; liuren.wu@baruch.cuny.edu. The authors thank Campbell Harvey (the editor), the associate editor, the anonymous referee, Bruno Dupire, Peter Fraenkel, Jingzhi Huang, John Hull, Dilip Madan, Tom McCurdy, George Panyotov, Matthew Richardson, Allen White, participants at AQR, Bloomberg, New York University, Rutgers University, University of Toronto, the 2008 Princeton conference on implied volatility models at Huntington Beach, California, the 2010 Conference on Latest Developments in Heavy-Tailed Distributions in Brussels, Belgium, the 2010 China International Conference in Finance in Beijing, China, the 10th Annual Meeting of the Brazilian Finance Society in Sao Paulo, and the 2011 American Finance Association meetings in Denver for comments. We also thank Sergey Nadtochiy for research assistance and Richard Holowczak for computing support.

Equity index returns interact negatively with return volatilities. This paper traces the negative interaction to three distinct economic channels and proposes to disentangle the relative contribution of each channel by exploiting the information in the large cross-section of equity index options.

First, equity return volatility increases with financial leverage. Financial leverage variation can come from either market price movement or active capital structure management. For a company with a fixed amount of debt, its financial leverage increases when the market price of its stock declines. Black (1976) first proposes this *leverage effect* to explain the negative correlation between equity returns and return volatilities. A company can also proactively vary its financial leverage based on variations in market conditions (Adrian and Shin (2010)). Both types of variation in financial leverage generate variations in equity return volatility. In capital structure models such as Merton (1974), Leland (1994), and Leland and Toft (1996), financial leverage variation is the only source of variation in equity return volatility because these models assume constant volatility for the asset return.

Second, asset return volatility can also vary over time, due to variation in either the risk of a particular business or the composition of different businesses. Furthermore, with fixed future cash flow projections, a forecasted increase in business risk increases the discount rate and reduces the present value of the business, generating the *volatility feedback effect* (Choi and Richardson (2008)). This volatility feedback effect can also show up as a negative correlation between equity return and volatility, irrespective of the level of financial leverage.

Third, financial crises have shown a disconcerting pattern that worries both policy makers and financial managers: A large, negative financial event often increases the chance of more such events to follow. Researchers have labeled this phenomenon as the *self-exciting behavior*,¹ which can show up both cross-sectionally and intertemporally. Cross-sectionally, the default of one company

¹See for example, Azizpour and Giesecke (2008), Ding, Giesecke, and Tomecek (2009), and Ait-Sahalia, Cacho-Diaz, and Laeven (2010).

has been found to increase the likelihood of default on other companies. Over time, one market turmoil often increases the chance of another one to follow, generating intertemporal clustering in negative market events. Both types of self-exciting propagations can have a destabilizing effect on the financial market.

This paper proposes an equity index dynamics that accommodates all three economic channels of return volatility variation, and estimates the model using observations on the S&P 500 index options. Estimation shows that the volatility feedback effect reveals itself mainly in the variation of short-term options, the self-exciting behavior affects both short-term and long-term option variations, whereas the financial leverage variation has its largest impact on long-dated options. Thus, using option observations across a wide range of strikes and maturities, one can effectively disentangle the three sources of volatility variation.

The disentangling of the volatility variation allows us to address important economic questions on the capital structure decision. Model estimation shows that contrary to common wisdom, the aggregate financial leverage in the U.S. market does not always decline with increased business risk. Instead, the market responds differently to different types of business risk. The financial leverage can actually increase with increasing business risk if the business risk is driven by small, diffusive market movements. Only when the perceived risk of self-exciting market disruption increases does the market become truly concerned and start the de-leveraging process.

The disentangling also reveals economic insights on other empirical observations. For example, an important recent finding in the options market is that variance risk generates a strongly negative risk premium (Bakshi and Kapadia (2003a,b) and Carr and Wu (2009)). Several studies provide interpretations and explore its relations with other financial phenomena.² Our structural decomposition shows that volatility risk premium can come from many different sources, e.g.,

²See, for example, Bakshi and Madan (2006), Bollerslev, Gibson, and Zhou (2011), Bollerslev, Tauchen, and Zhou (2009), Drechsler and Yaron (2010), and Zhou (2010).

the market prices of financial leverage risk, asset volatility risk, and self-exciting market crashes. Estimation shows that the negative risk premium is mainly from the latter two sources.

This paper is the first in the literature that proposes a working solution to disentangling the different economic sources of volatility variation. The reduced-form option pricing literature, e.g., Duffie, Pan, and Singleton (2000), can accommodate multiple sources of stochastic volatilities, but they are specified as purely statistical factors, without any direct linkage to economic sources. The absence of economic linkage prevents these models from addressing economic questions. For example, many option pricing models allow a negative correlation between return and volatility. This negative correlation is often referred to as the “leverage effect” without further distinction on whether it is really coming from financial leverage variation or a purely volatility feedback effect that has nothing to do with leverage. Furthermore, the generic factor structure often poses identification issues that limit most empirical estimation to one-factor volatility specifications, with only a few exceptions on two-factor volatility models. By applying the economic structures, we not only address important economic questions, but also develop a rich and yet parsimonious three-volatility-factor specification that can be well-identified from the option observations.

The capital structure literature pioneered by Merton (1974) is rich in economic structures but remains both too stylized and too complicated to be a working solution for equity option pricing. As such, it remains difficult for these structural models to benefit from the information in the vast equity option observations. Cremers, Driessen, and Maenhout (2008) specify a jump-diffusion stochastic volatility dynamics for the asset value and compute the equity option values as compound options on the asset value. Through this specification, they are able to calibrate the average credit spreads on corporate bonds to the average variance and jump risk premiums estimated from equity index options. Their resolution of the credit spread puzzle highlights the virtue of exploiting information in equity options, but their stylized calibration exercise also highlights the inherent difficulty in making the structural approach a feasible solution for capturing the time variation of

equity options. Our specification retains the flexibility and tractability of reduced-form option pricing models, but incorporates economic structures motivated by the capital structure literature. As a result, the model can extract the rich information in equity index options to show how capital structure decisions vary with different types of economic risks. The model's performance in pricing equity index options is significantly better than all other models estimated in the literature.

The rest of the paper is organized as follows. Section I specifies the equity index dynamics and discusses its linkage to various strands of the literature. Section II summarizes the data set and the estimation methodology. Section III discusses the estimation results. Section IV concludes. The Appendix provides the details on option pricing, data statistics, model estimation methodology, option pricing performance comparison, and estimation results from alternative data sets.

I. Model Specification

We fix a filtered complete probability space $\{\Omega, \mathcal{F}, \mathbb{P}, (\mathcal{F}_t)_{t \geq 0}\}$ and assume no-arbitrage in the economy. Under certain technical conditions, there exists a risk-neutral probability measure \mathbb{Q} , absolutely continuous with respect to \mathbb{P} , such that the gains process associated with any admissible trading strategy deflated by the riskfree rate is a martingale. Let F_t denote the time- t forward level of the equity index over some fixed time horizon and A_t denote the corresponding time- t forward asset value underlying the equity index, with $X_t \equiv F_t/A_t$ measuring the forward equity to asset ratio. Standard textbooks also refer the reciprocal of X_t as the accounting leverage.

We start with specifications on the risk-neutral \mathbb{Q} -dynamics, cognizant of the tractability requirement in pricing options. We derive the statistical \mathbb{P} -dynamics by introducing market prices for various sources of risks, with a focus on the capital structure decisions. We then discuss alternative representations of our model and highlight how the model extends and unifies several strands of the literature.

A. Separating Leverage Effect from Volatility Feedback and Self-Exciting Jumps

We model the asset value dynamics (A_t) and the financial leverage variation (X_t) separately through the multiplicative decomposition,

$$F_t = X_t A_t. \quad (1)$$

This decomposition is merely a tautology, but it allows us to disentangle the impacts from financing decisions and business operation decisions. To capture the financial leverage variation, we model the equity-to-asset ratio X_t as a constant elasticity of variance (CEV) process under the risk-neutral measure \mathbb{Q} ,

$$dX_t/X_t = \delta X^{-p} dW_t, \quad p > 0, \quad (2)$$

where W_t denotes a standard Brownian motion, δ is a positive constant, and p is a power coefficient that determines how the equity index return volatility varies with the level of financial leverage. With $p > 0$, the CEV process captures the *leverage effect*: Conditional on a fixed asset value, a decline in X increases the financial leverage and reduces the equity value by definition, while it also raises the equity volatility via the power term X_t^{-p} in equation (2).

We model the asset value dynamics under the risk-neutral measure \mathbb{Q} as,

$$\frac{dA_t}{A_t} = \sqrt{v_t^Z} dZ_t + \int_{-\infty}^0 (e^x - 1) (\mu(dx, dt) - \pi(x) dx v_t^J dt), \quad (3)$$

$$dv_t^Z = \kappa_Z (\theta_Z - v_t^Z) dt + \sigma_Z \sqrt{v_t^Z} dZ_t^v, \quad \mathbb{E}[dZ_t^v dZ_t] = \rho dt < 0, \quad (4)$$

$$dv_t^J = \kappa_J (\theta_J - v_t^J) dt - \sigma_J \int_{-\infty}^0 x (\mu(dx, dt) - \pi(x) dx v_t^J dt), \quad (5)$$

where Z_t and Z_t^v denote two standard Brownian motions, $\mu(dx, dt)$ denotes a counting measure for

jumps, and $\pi(x)v_t^J$ denotes the time- t arrival rate of jumps of size x in log asset value $\ln A_t$, with

$$\pi(x) = e^{-|x|/v_J}|x|^{-1}. \quad (6)$$

Equation (3) decomposes the asset value variation into a diffusion component with stochastic variance v_t^Z and a discontinuous component with stochastic jump intensity v_t^J . A negative correlation between Z_t and Z_t^J in equation (4) generates the *volatility feedback* effect: A positive shock to the business risk increases the cost of capital and reduces the asset value.

The negative jump component in equation (3) captures the impact of market turmoils. The intensity of market turmoils is stochastic and follows the dynamics specified in equation (5), where a downside jump in the asset value is associated with an upside jump in the jump intensity. The coefficient $\sigma_J > 0$ captures the proportional scale on the jump size in the intensity per each jump in the log asset value, and the negative sign in front of σ_J highlights the opposite effect of the jump on the asset value and the jump intensity. The specification captures the *self-exciting behavior*: The occurring of a downside jump event increases the intensity of future downside jump events.

Although we can accommodate both positive and negative jumps, equation (6) only incorporates negative jumps for parsimony.³ We further assume independence between the Brownian innovation (dW_t) in financial leverage and the Brownian innovations in the asset value (dZ_t) and the asset return volatility (dZ_t^J). The innovation independence assumption allows us to model A_t and X_t as separate martingales, and facilitates the pricing of equity index options. Nevertheless, this assumption does not exclude dynamic interactions between business risks and the financial leverage decisions under the real world probability measure \mathbb{P} . Indeed, the dynamic interactions are one focal point of our analysis, as discussed in the next subsection.

³When we incorporate positive jumps in the specification, estimation on equity index options often finds that the positive jump size is not significantly different from zero under the risk-neutral measure.

B. Active Capital Structure Decisions and Market Prices of Economic Risks

Through the market pricing specification for the financial leverage risk W_t , we allow the variation of the equity-to-asset ratio X_t to depend on all three sources of economic risk under the real world probability measure \mathbb{P} .

$$dX_t/X_t = X_t^{-p} (a_X - \kappa_{XX}X_t - \kappa_{XZ}v_t^Z - \kappa_{XJ}v_t^J) dt + X_t^{-p} dW_t^{\mathbb{P}}, \quad (7)$$

Under this specification, the market adjusts the capital structure based on the current level of the financial leverage (X_t), the current level of the business diffusion risk v_t^Z , and the current level of the business jump risk v_t^J . We use $\kappa_L \equiv [\kappa_{XX}, \kappa_{XZ}, \kappa_{XJ}]^\top$ to denote the leverage loading coefficient vector on the three state variables (X_t, v_t^Z, v_t^J). The constant term a_X allows the market to set a long-run target on the equity-to-asset ratio. Through model estimation, we can infer how the market's capital structure responds to different types of risks in the market.

The statistical dynamics in equation (7) can be reconciled with the risk-neutral dynamics in equation (2) via the following market price specification for the financial leverage risk W_t ,

$$\gamma_t^X = (a_X - \kappa_{XX}X_t - \kappa_{XZ}v_t^Z - \kappa_{XJ}v_t^J). \quad (8)$$

Furthermore, we perform the following decomposition on the diffusion asset value risk,

$$Z_t = \rho Z_t^v + \sqrt{1 - \rho^2} \tilde{Z}_t, \quad (9)$$

and assume proportional market prices (γ^Z, γ^v) on the independent diffusion return risk \tilde{Z}_t and the diffusion variance risk Z^v . The proportional market price (γ^Z) of the independent diffusion return risk (\tilde{Z}_t) generates an instantaneous risk premium on the asset return of $\gamma^Z \sqrt{1 - \rho^2} v_t^Z$. The proportional market price (γ^v) of the diffusion variance risk (Z_t^v) generates a drift adjustment term

$\gamma^v \sigma_Z v_t^Z$ for v_t^Z . It also generates an instantaneous asset return risk premium $\rho \gamma^v v_t^Z$. Finally, we assume constant market price (γ^J) on the jump return risk, which generates an exponential tilting on the jump arrival rate under the statistical measure \mathbb{P} ,

$$\pi^{\mathbb{P}}(x) = e^{\gamma^J x} e^{-|x|/v_J} |x|^{-1} = e^{-|x|/v_J^{\mathbb{P}}} |x|^{-1}, \quad (10)$$

with $v_J^{\mathbb{P}} = v_J / (1 + \gamma^J v_J)$.

C. Alternative Representations

Combining the specifications in equations (1)-(6), we can write the risk-neutral dynamics for the forward equity index as,

$$dF_t/F_t = \delta \left(\frac{F_t}{A_t} \right)^{-p} dW_t + \sqrt{v_t^Z} dZ_t + \int_{-\infty}^0 (e^x - 1) (\mu(dx, dt) - \pi(x) dx v_t^J dt). \quad (11)$$

By performing a change of variable $v_t^X = \delta^2 X_t^{-2p}$, we can also rewrite the risk-neutral equity index dynamics in the form of a three-factor stochastic volatility model,

$$dF_t/F_t = \sqrt{v_t^X} dW_t + \sqrt{v_t^Z} dZ_t + \int_{-\infty}^0 (e^x - 1) (\mu(dx, dt) - \pi(x) dx v_t^J dt), \quad (12)$$

where the stochastic variance from the leverage effect v_t^X follows a 3/2-process,

$$dv_t^X = \kappa_X (v_t^X)^2 dt - \sigma_X (v_t^X)^{3/2} dW_t, \quad (13)$$

with $\kappa_X = p(2p + 1)$ and $\sigma_X = 2p$ and its innovation is perfectly negatively correlated with the corresponding return innovation component. Under this specification, the index return is driven by two Brownian motion components and a jump component. The instantaneous variances on the two

Brownian motions v_t^X and v_t^Z and the jump intensity process v_t^J are all stochastic and are driven by three separate dynamic processes (13), (4), and (5), respectively.

The alternative representations reveal several new insights. First, via change of variables, equation (13) makes it explicit that the financial leverage variation can be one of the three contributors to the stochastic volatility in equity return. Second, the particular 3/2-volatility of volatility dependence due to the leverage effect in (13) is interesting. The behaviors of 3/2-processes have been studied by several authors, e.g., Heston (1997), Lewis (2000), and Carr and Sun (2007). Within the one-factor diffusion context, several empirical studies find that a 3/2 specification on the variance rate dynamics performs better than the square-root specification.⁴ Thus, by including a 3/2-volatility component in addition to the square-root dynamics for the asset diffusion variance rate v_t^Z and the self-exciting jump dynamics for the jump intensity v_t^J , our model has the potential of generating better pricing performance than existing affine specifications in the literature.

Third, equations (12) and (13) reveal that from the equity index forward level F_t and index options, one cannot identify the scale parameter δ for the equity-to-asset ratio X_t . Henceforth, we set $\delta = 1$ and regard X_t as a proportionally scaled version of the equity-to-asset ratio.

Finally, the literature often makes a dichotomous distinction between the local volatility models of Dupire (1994) and the stochastic volatility models such as Heston (1993). The local volatility models are popular in the industry, but the academic option pricing literature focuses almost exclusively on scale-free stochastic volatility specifications. The two representations of our model show that the gap between the two strands of literature is not as big as generally perceived. We can represent our model as either having a local-volatility component in (11) or as a pure three-factor stochastic volatility model in (12).

⁴Favorable evidence from time-series returns includes Chacko and Viceira (2003), Ishida and Engle (2002), Javaheri (2005), and Jones (2003). Supporting evidence from equity index options include Jones (2003), Medvedev and Scaillet (2003), and Bakshi, Ju, and Ou-Yang (2006).

D. Relations to Literature

Our model extends and unifies several strands of the literature. First, several papers specify the equity index as following a pure CEV process, e.g., Beckers (1980), Cox (1996), Emanuel and MacBeth (1982), and Schroder (1989). The local volatility model of Dupire (1994) also specifies the equity return volatility as a function of the equity level. These specifications, however, are inherently unstable because equity return volatility is a unitless quantity whereas the equity index level has the unit of the domestic currency (e.g., dollar). By specifying a unitless quantity as a function of the equity level, the volatility coefficient must absorb the unit mismatch and as a result must vary with the scaling of the equity level. A rescaling of the equity level leads to a rescaling of the volatility coefficient even if the equity return dynamics remain the same. By contrast, our specification scales the equity index level in the volatility function by the asset value, thus making the instantaneous volatility a unitless quantity, circumventing the stability concern on scale-dependent dynamic specifications.

Second, by specifying the equity return volatility as a function of financial leverage, our model is more economically sound and moves closer to structural models on capital structure. Hurd and Li (2008) show that under certain parametric conditions, the Leland (1994) model, where equity is modeled as a barrier option on the asset value, implies that the equity return volatility is a power function of the equity-to-asset ratio, with the power being $p = 1/2$. Our specification can be regarded as a generalization of the Leland structural model by (i) allowing the power dependence on the leverage ratio to be a free parameter and (ii) accommodating two additional sources of stochastic volatility on the asset return.

Third, our diffusion asset return component with stochastic volatility shares the same specification as the stochastic volatility model of Heston (1993), but our specification models the leverage effect separately and uses the Heston dynamics to capture purely the volatility feedback effect.

Fourth, the jump specification (6) in the asset value dynamics has its origin in the variance-gamma model of Madan and Seneta (1990) and Madan, Carr, and Chang (1998). Compared to compound Poisson-type jump specifications used in Merton (1976), Kou (2002), and Duffie, Pan, and Singleton (2000), the variance-gamma specification generates an infinite number of jumps within any finite interval and has been shown to perform better in capturing both the time-series behavior (Carr, Geman, and Madan (2001)) and the option price behavior (Huang and Wu (2004) and Carr and Wu (2007)).

Fifth, although many option pricing models add jumps to the equity dynamics, our specification is unique and makes better economic sense. Bakshi, Cao, and Chen (1997) allow a jump in the index process but with constant intensity. Pan (2002) specifies the jump intensity as a function of the diffusion variance rate, which follows a square-root process. Eraker, Johannes, and Polson (2003), Eraker (2004), and Broadie, Chernov, and Johannes (2007) use synchronized finite-activity jumps to model the equity index return and volatility. Still, the diffusion variance and jump intensity are governed by one process. As a result, these models do not capture the observation that small market movements and large market turmoils can be driven by completely different forces. Huang and Wu (2004) allow diffusion and jumps to generate separate stochastic volatilities, but they do not accommodate the self-exciting behavior. Our specification is the first in the literature that allows small market movements with volatility feedback effect and large market turmoils with self-exciting behavior as separate sources of stochastic volatility.

Taken together, our model offers several advantages over the current option pricing literature. First, our model allows separate channels for leverage effect and volatility feedback, whereas most reduced-form specifications in the literature do not make such a distinction. As a result, although these models can be used to price equity options, they cannot distinguish the contributions from different economic sources. Second, although several models allow jumps in the return and/or volatility process, our model is the first to accommodate the self-exciting behavior as a separate

source of volatility variation. Third, our model has its roots in structural models on capital structure and can thus address capital structure questions that are normally reserved for structural models.

Meanwhile, although structural models are rich in economic structures, most of these models remain highly stylized and cannot be used to capture the equity option variation. Computing equity option values as compound options on the asset value can also be a daunting numerical task. Thus, these models lack both the flexibility and the tractability to price equity options. By contrast, our specification is much more tractable in pricing equity index options and at the same time much more flexible in accommodating the observed equity index option behaviors. Appendix A explains how equity index options can be priced efficiently under our model.

II. Data and Estimation

We estimate the model using S&P 500 index (SPX) options. The options are both listed on the Chicago Board of Options Exchange (CBOE) and traded actively over the counter (OTC). Options transactions on the listed market are concentrated at short maturities. The expiries for listed options are mostly within three years and most transactions are on options with maturities less than one year. By contrast, activities on the OTC market are more concentrated on long-dated contracts. These contracts are often used as hedges for long-dated corporate structural deals and variance swap transactions.

We have data from both markets, but we focus our analysis on the OTC data, mostly because the OTC data cover a much wider span of maturities from one month to five years and can thus help enhance the identification of the different sources of volatility variations. At one particular date and maturity, many different mechanisms can generate an implied volatility skew along the strike dimension; yet, these different mechanisms have different implications on how the implied volatility skew evolve with the option maturity. By using the OTC data that span a wider range

of maturities, we can achieve a better disentanglement of the different mechanisms that our model incorporates. We use the listed options data from CBOE for robustness analysis. The estimation results from this alternative data set are provided in Appendix E.

The OTC SPX options data are from a major bank. The data are in the form of Black and Scholes (1973) implied volatility quotes from January 8, 1997 to March 5, 2008. The data are available daily, but we sample the data weekly every Wednesday to avoid weekday effects. At each date, the quotes are on a matrix of eight fixed time to maturities at one, three, six, 12, 24, 36, 48, and 60 months and five relative strikes at 80, 90, 100, 110, and 120 percent of the spot index level. The OTC quotes are constructed to match the listed option prices at short maturities and to match the over-the-counter transactions at long maturities. All together, the data set contains 40 implied volatility series over 583 weeks, a total of 23,320 observations. Appendix B summarizes the statistical behaviors of the SPX implied volatility surface and explains how our model can accommodate these behaviors.

The model uses three state variables (X_t, v_t^Z, v_t^J) to capture the variation of the implied volatility surface over time. The objective of the estimation is to identify the values of the structural parameters that govern the financial leverage and asset value dynamics, and to extract the levels of the three state variables at different time periods. The estimates for the structural parameters facilitate our understanding of the dynamics and interactions between the different risk sources, whereas the levels of the three state variables at different sample periods shed light on the relative contribution of each risk source at different historical time periods.

To estimate the model parameters, we cast the model into a state-space form by treating the three state variables as the hidden states, and the implied volatility quotes as observations with errors. We employ a nonlinear filtering technique to extract the levels of the states at each date from the implied volatility observations. The model parameters are estimated by maximizing the likelihood defined on the model forecasting errors on the options. Appendix C provides the technical

details on the estimation procedure.

The model has nine parameters $(p, \kappa_Z, \theta_Z, \sigma_Z, \rho, \kappa_J, \theta_J, \sigma_J, v_J)$ and three state variables (X_t, v_t^Z, v_t^J) to price 40 options each week. In addition, the model has six parameters $(a, \kappa_{XX}, \kappa_{XZ}, \kappa_{XJ}, \gamma^v, \gamma^J)$ to control the market prices of risks and hence the statistical dynamics.

Despite its richness in economic structures and flexibility in allowing three separate sources of volatility variation, the model is very parsimonious in terms of the number of free parameters. The model has nine parameters that govern the risk-neutral dynamics of the equity index and thus the cross-sectional behavior of the equity index options. By contrast, the models considered in Eraker (2004) and Broadie, Chernov, and Johannes (2007) have only one source of stochastic volatility, but can have as many as eight parameters. The two-factor volatility structure considered in Bates (2000) have 11 parameters. Our model's higher dimension in the states dictates that it can match the observed option behavior over a wide range of strikes and maturities, much better than any existing models in the literature can do. On the other hand, the parsimony of our specification, together with the large data set that we use for model estimation, allows us to identify all model parameters with strong statistical significance. Appendix D illustrates the statistical significance of the superior performance of our specification by comparing its option pricing performance to a two-factor reduced-form benchmark.

III. Results and Discussion

Table I reports the model parameter estimates and their standard errors in parentheses. For ease of discussion, we group the parameters into three panels, each describing the dynamics of one source of volatility risk. Given the parsimony of the specification and the large amount of data used for the model estimation, all parameters are estimated with strong statistical significance.

[Table I about here.]

A. Disentangling Different Sources of Stochastic Volatility and Volatility Skew

The model allows three distinct channels of volatility variation: (i) the variation of the equity-to-asset ratio, (ii) the variation of diffusion business risk, and (iii) the variation of discontinuous business risk. The instantaneous return variance contributions from the three sources are, X_t^{-2p} , v_t^Z , $v_J^2 v_t^J$, respectively. Based on the sample averages of the filtered state variable values and the parameter estimates on (p, v_J) , we compute the average return variance from each source, at 0.0117, 0.0231, and 0.0136, respectively, or at 10.84%, 15.19%, and 11.67%, respectively, in volatility terms. Therefore, all three sources contribute to a significant portion of the index return variance.

The three sources of stochastic volatility interact dynamically with the index return and contribute to the implied volatility skew at different maturity ranges. The financial leverage effect is captured by the power coefficient p . A coefficient of zero would imply zero dependence of return variance on the financial leverage. The larger the estimate, the stronger the leverage effect. The estimate for the power coefficient is 2.8427, suggesting a strong dependence of the index return variance on the financial leverage level.

When using a pure CEV model to fit the implied volatility skew of the SPX options, Jackwerth and Rubinstein (1996) find that the fitted values for the power coefficient can be as high as $p = 5$. Our estimate for the power coefficient is much smaller because our model has several other mechanisms that also contribute to the implied volatility skew. The coefficient estimates are also smaller when the model is estimated with the index return time series, e.g., Beckers (1980) and Christie (1982).

The volatility feedback effect is captured through the instantaneous correlation (ρ) between the diffusion movement in the asset value and its variance rate. The estimate for the correlation is highly negative at -0.8388 , suggesting that the volatility feedback effect is also very strong. This feedback effect can generate a negative implied volatility skew at intermediate maturities. The

variance rate v_t^Z is modeled as a mean-reverting square-root process, with κ_Z measuring the risk-neutral mean-reverting speed, θ_Z the risk-neutral long-run mean, and ω_Z the volatility of volatility coefficient. The estimate of $\kappa_Z = 3.0112$ indicates that this variance rate process is highly mean-reverting under the risk-neutral measure. As a result, shocks on the variance rate v_t^Z dissipates quickly as the option maturity increases. The mean estimate of $\theta_Z = 0.0244$ implies a return volatility contribution of $\sqrt{\theta_Z} = 15.62\%$ under the risk-neutral measure from this particular variance rate, slightly higher than the sample averages of the v_t^Z estimates over our sample at 15.19%. The estimate for the volatility of volatility coefficient is at $\sigma_Z = 0.5987$, which contributes to the curvature of the implied volatility smile at intermediate maturities.

Parameter estimates for the self-exciting downside jump dynamics are summarized in Panel C of Table I. The average size of the downside jump under the risk-neutral measure is governed by v_J , which is estimated to be at 19.26%. The negative jump directly induces a negative skewness in the risk-neutral return distribution and accordingly a negative implied volatility skew at short maturities. The self-exciting behavior of the downside jump extends its impact to longer option maturities. In particular, the jump intensity v_t^J shows little mean reversion under the risk-neutral measure as the mean-reversion speed estimate is very small at $\kappa_J = 0.0009$. The slow mean reversion helps sustain the jump effect to long option maturities. The degree of self-excitement is measured by σ_J , which measures the magnitude of the intensity upside jump per downside jump in the asset return. The large estimate of $\sigma_J = 5.6356$ indicates that a downside jump in the asset return evokes a very large upside jump in the jump intensity itself, highlighting the significance of the self-exciting behavior.

Taken together, the estimation results show that the variations from all three sources, i.e., the financial leverage, the diffusion business risk, and the jump business risk, all contribute significantly to the variation of the SPX return. The three sources of variation all show strong dynamic interactions with the index return via the leverage effect, the volatility feedback effect, and the

self-exciting behavior, respectively. These interactions contribute to the observed implied volatility skew at different maturities.

B. Differentiating Volatility Responses to Different Types of Shocks

To understand how the volatility surface responds differently to shocks from the three risk sources, we shock each of the three state variables (v_t^X, v_t^Z, v_t^J) from its mean level to the 10th percentile and 90th percentile, respectively, while holding the other two state variables at their respective mean levels. We then compute the implied volatility surface before and after each shock to show the responses.

Figure 1 plots the responses of the volatility skew and term structure to shocks from the three risk sources (X_t, v_t^Z, v_t^J) . The three rows correspond to shocks from the three risk sources. In each row, the left panel plots the response of the one-month implied volatility skew and the right panel plots the response of the at-the-money (100% relative strike) implied volatility term structure. In each panel, the solid line represents the model-generated implied volatility values when evaluated at the sample averages of the state variables; the dashed line represents the model-generated values when evaluated at the 90th percentile of the state variable in question while holding the other two state variables at their sample averages; and the dash-dotted line represents responses to a shift to the 10th percentile for the state variable in question while holding the other two state variables at their sample averages.

[Figure 1 about here.]

When we shock the equity-to-asset ratio X_t from its sample average to its 90th-percentile value, the financial leverage is reduced and its contribution to the return variance (X_t^{-2p}) is reduced accordingly. Thus, as shown in the first row of Figure 1, the dashed lines are below the correspond-

ing solid lines. When the equity-to-asset ratio is reduced to its 10th-percentile value, the financial leverage increases, and the option implied volatilities (dash-dotted lines) increase and hence move above the solid line. The left panel shows that shocks to X_t affect the implied volatilities more at-the-money than out-of-the-money. The right panel shows that the responses of the at-the-money implied volatilities to financial leverage variation are relatively uniform across maturities. A shock in X_t induces a near parallel shift in the at-the-money implied volatility term structure.

When the diffusion variance rate v_t^Z experiences a positive shock, the return variance contribution from the diffusion component (Z_t) increases. In response, the option implied volatilities move up, but the implied volatility skew declines due to the increase of the diffusion component relative to the jump component. This effect shows up clearly in the left panel of the second row in Figure 1, where the dashed line is above the solid line but becomes slightly less skewed. A negative shock, on the other hand, reduces the volatility level but increases the steepness of the skew. The term structure plot on the right side highlights the transient nature of the v_t^Z risk factor due to its high risk-neutral mean reversion speed ($\kappa_Z = 3.0112$). Shocks to v_t^Z induce large responses at short maturities, but the response declines quickly as the option maturity increases.

When the arrival rate of the downside jump v_t^J experiences a positive shock, the return variance increases from the contribution of the jump component. Furthermore, as v_t^J governs the intensity of the negatively skewed return jump component, a positive shock to v_t^J also increases the negative skewness of the instantaneous return distribution. Thus, as shown in the left panel of the third row in Figure 1, a positive shock to v_t^J not only raises the implied volatility level, but also steepens the negative skew by raising the low-strike implied volatility more than it raises the high-strike implied volatility. The right panel shows the term structure effect. Consistent with the low mean-reversion speed estimate ($\kappa_J = 0.0009$), shocks to v_t^J have a persistent impact on the implied volatilities across maturities.

The response analysis shows that all three sources of variation contribute to the return vari-

ance, but their contributions are quite different across different strikes and maturities. Along the strike dimension, leverage (X_t) variation affects at-the-money volatility more than out-of-the-money volatility, diffusion variance (v_t^Z) variation affects volatility at high strikes more than at low strikes, but the jump intensity (v_t^J) variation affects volatility at low strikes more than at high strikes. Across maturities, shocks from the diffusion variance rate are the most transient and thus their impacts are mostly on short-term options; by contrast, both the self-exciting downside jumps and the leverage effect have long-lasting impacts at both short and long maturities. These different response patterns allow us to disentangle the three sources of risks by using options across a wide span of strikes and maturities.

C. Detecting Capital Structure Decisions from Equity Index Options

The model uses four parameters to capture how the equity-to-asset ratio responds to the levels of financial leverage and the two types of business risks, the diffusion risk v_t^Z and the downside jump risk v_t^J . Table I reports the estimates for these parameters in Panel A as part of the statistical dynamics for the equity-to-asset ratio.

The parameter κ_{XX} measures the dependence of the capture structure decision on past leverage levels. The small estimate at $\kappa_{XX} = 0.0001$ suggests that changes in the financial leverage are largely independent of past leverage levels.

The parameter κ_{XZ} captures how the capital structure decision responds to the diffusion business risk v_t^Z . Interestingly, the estimate is strongly positive at $\kappa_{XZ} = 17.5341$, suggesting that when the diffusion-type fluctuations in the underlying asset value increase, X_t declines and hence the financial leverage actually increases. Conventional wisdom holds that one wants to reduce the financial leverage to mitigate the risk when the business risk increases; however, our estimation suggests that the financial leverage does not decrease, but rather increase when diffusion-type

business risk increases.

On the other hand, the capital structure response to the jump business risk is estimated to be negative at $\kappa_{XJ} = -0.0774$. The market reduces financial leverage when the expected downside jump risk increases but increases financial leverage when the diffusion risk increases. Therefore, the main concern with using financial leverage is not the normal daily business fluctuations, but rather the unexpected, large downside jumps that can lead to a self-exciting spiral.

Traditional corporate finance often links the level of financial leverage to the riskiness of the underlying business. Our model decomposes the business risk into two types and allows the financial leverage decision to respond differently to variations in the two types of business risk. Model estimation shows that the aggregate financial leverage in the U.S. market does not always decline with increased business risk. The financial leverage can actually increase with increasing business risk if the risk is driven by small, diffusive market movements. Only when the perceived risk of self-exciting market disruptions increases, does the market become truly concerned by starting the de-leveraging process.

D. Decomposing Different Sources of Variance Risk Premiums

Several studies, e.g., Bakshi and Kapadia (2003a,b) and Carr and Wu (2009), have documented strongly negative variance risk premiums on stock indexes. Several ensuing studies propose explanations for the negative variance risk premium (e.g., Bakshi and Madan (2006) and Drechsler and Yaron (2010)), or explore the variance risk premium as a predictor for other financial behaviors (e.g., Bollerslev, Gibson, and Zhou (2011), Bollerslev, Tauchen, and Zhou (2009), and Zhou (2010)). Missing from these studies is the realization that variance risk and hence variance risk premium can come from several distinct sources. Before one can develop a valid economic theory, it is important to understand the relative contribution of each risk source to the risk premium. If

the purpose is to use the variance risk premium to predict other financial behaviors, it can be useful to separate the risk premium from different sources and investigate the predictability of each risk premium component.

Our model decomposes the index return into three risk sources (W_t, Z_t, J_t) , all contributing to the stochastic return variance and thus return variance risk premium. Based on the parameter estimates and the extracted state variables, we calculate the average contribution of each risk source to the variance risk premium.

The financial leverage factor X_t generates an instantaneous index return variance of $v_t^X = X_t^{-2p}$. The instantaneous risk-neutral drift of v_t^X is $\mu(v_t^X)^{\mathbb{Q}} = p(2p+1)(v_t^X)^2$, and the corresponding statistical drift is $\mu(v_t^X)^{\mathbb{P}} = p(2p+1)(v_t^X)^2 - 2p(v_t^X)^{3/2}(a_X - \kappa_{XX}X_t - \kappa_{XZ}v_t^Z - \kappa_{XJ}v_t^J)$. The sample averages of the two drifts are 0.0040 and 0.0069, respectively, generating a slightly positive average risk premium of 0.0029. Thus, our estimation suggests that the observed negative variance risk premium does not come from the variance risk induced by financial leverage variations.

The diffusive component of the asset return contributes to the instantaneous index return variance by v_t^Z , which follows a square-root diffusion dynamics, with the instantaneous risk premium on v_t^Z given by $\gamma^v \sigma_Z v_t^Z$. With a negative estimate on the market price $\gamma^v = -17.4512$, the instantaneous risk premium averages at -0.2412 .

The jump component of the asset return contributes to the instantaneous index return variance by $(v_J)^2 v_t^J$ under the risk-neutral measure and by $(v_J^{\mathbb{P}})^2 v_t^J$ under the statistical measure. The difference is induced by the market price of jump risk γ^J , which is estimated at $\gamma^J = 0.4468$. The sample averages of the two variance series are 0.0136 (11.67% in volatility term) and 0.0116 (10.75% in volatility term), respectively, thus generating a negative variance risk premium of -0.0021 .

Finally, the market price of the jump risk also induces a difference between the statistical and risk-neutral drifts of the jump intensity process v_t^J at $\sigma_J(v_J^{\mathbb{P}} - v_J)v_t^J$, which averages at -0.0316 .

Our model estimation confirms the existence of a negative variance risk premium on the equity index, but it also shows the complex composition of the variance risk premium. The analysis shows that the instantaneous variance risk premium is a combined result of four different types of risk premiums. When one investigates the variance risk premium over different investment horizons and using different variance instruments (e.g., one month investment in one-month variance swap as in Carr and Wu (2009) or two-month investment in two- to 12-month variance swaps as in Egloff, Leippold, and Wu (2010)), these four sources of instantaneous variance risk premium will further interact with the corresponding variance dynamics to determine the expected excess returns. Since the diffusion variance rate v_t^Z is more transient than the jump arrival rate v_t^J , we expect the risk premium from the diffusion variance rate to dominate investments on short-term variance swap contracts. By contrast, excess returns on long-term variance swap contracts can be dominated by risk premiums from the persistent jump intensity v_t^J and the mean-repelling leverage variation X_t . This decomposition is consistent with the empirical findings in Egloff, Leippold, and Wu (2010) that short-term variance swap contracts generate more negative risk premiums than do long-term variance swap contracts.

E. Dissecting the Sample Variation of Different Sources of Volatility Risks

Figure 2 plots the time series of the three state variables extracted from the equity index options data. The financial leverage (X_t) contributes to the return variance through the transformed variance rate $v_t^X = X_t^{-2p}$, which is plotted in the top panel. It can be regarded as a scaled version of the accounting leverage. The time series for the diffusion variance rate v_t^Z and the jump intensity v_t^J are plotted in the middle and bottom panel of Figure 2, respectively.

[Figure 2 about here.]

Inspecting the time-series plots, we can trace various historical events to variations in the three variance rates. For example, the 1997 Asian crises were associated with a moderate spike in the diffusion variance rate (v_t^Z) and a small spike in the jump arrival rate v_t^J , followed by a downward drop in the financial-leverage induced variance rate v_t^X . These variations suggest that the Asian crises increased perceived business risks in the U.S. both in terms of diffusion-type fluctuations and downside jump risks. The heightened business risks are followed by an immediate de-leveraging process that sharply lowered the level of financial leverage and its associated volatility contribution.

The hedge fund crisis in late 1998 was also associated with a moderate spike in diffusion variance rate, but the associated spike in the jump intensity v_t^J was the largest over our sample period, suggesting that market participants were highly concerned with the potential impact of the hedge fund crisis on the U.S. financial system. The 1998 crisis was followed by a de-leveraging process that resulted in another sharp drop in the leverage-induced variance rate v_t^X . Since then, the leverage-induced variance rate increased steadily and reached its maximum during mid 1999. This period of increasing leverage was associated with declining perceived business risks of both types (v_t^Z and v_t^J).

The burst of the Nasdaq bubble in early 2000 started a protracted de-leveraging process, as v_t^X started a slow downward trend. This trend was reverted in mid 2007 as the financial leverage started to build up again right before the most recent financial crisis in 2008. The diffusion variance rate reached the highest level during the 2003 recession, but the market only showed moderate concerns for downside jump risk. The jump risk concern was low during the extended period from mid 2003 and 2007, but it started to pick up in 2008.

The time-series analysis shows that volatility variations at different sample periods come from different risk sources. While the 2003 recession is mostly associated with an increase in diffusion-type economic uncertainty, the 1998 hedge fund crisis is perceived as having the potential of rocking the stability of the financial system. Furthermore, market-wide de-leveraging often follows

when the perceived downside jump risk increases.

The model decomposes the equity index variation into variations in financial leverage and asset value. With equity index options but without actual financial leverage data, the estimation cannot fully identify the equity-to-asset ratio, but can identify a scaled version of it by normalizing the scaling coefficient δ to unity. To validate the interpretations of our estimation results on financial leverage, we investigate whether the state variable X_t extracted from the equity index options actually co-moves with financial leverage estimates compiled from financial reports on the S&P 500 companies.

Bloomberg provides historical computations on the debt-to-asset ratio for the S&P 500 index, based on the total debt and total asset of each company in the index and their associated weights. We label this debt ratio series as DR^A , with the superscript A denoting the accounting source. Since the state variable X_t extracted from the options data is a scaled version of the equity-to-asset ratio, we convert the state variable into a debt ratio and rescale the series to match the mean and standard deviation of the accounting series. We label this time series as DR^E , with the superscript E denoting equity options as the information source.

Figure 3 compares the two time series of debt ratios in Panel A, where the solid line denotes the time series extracted from the equity index options (DR^E) and the dashed line denotes the time series compiled from accounting reports (DR^A). The accounting time series become available at a later date on July 1st, 1998. We plot the two time series over the common sample period. The two time series are obviously not identical, as there are many reasons that they can deviate. On the one hand, the debt ratio extracted from the equity index options can proxy option movements not related to the financial leverage. On the other hand, the accounting measures of debt-to-asset ratio may not truly reflect a company's financial leverage, especially for a company with complex capital structures. Furthermore, since the accounting data on each firm are updated quarterly, the time series compiled from the accounting reports show some discreteness whereas the daily

transactions on equity index options render the variations of the solid line much smoother. Despite these expected deviations, the two time series in Panel A show clear co-movements, especially during notable historical events. For example, both time series show the de-leveraging process around the hedge fund crisis. Both time series also show a build-up of financial leverage before the 2008 financial crisis.

[Figure 3 about here.]

To gauge the co-movements of the two time series statistically, we measure the cross-correlation between changes in the two debt ratios over different horizons, $Corr(\Delta DR_h^E, \Delta DR_h^A)$, with h denoting the horizon. Panel B of Figure 3 plots the correlation estimates across different horizons. The cross-correlation estimates are positive over all horizons. The estimate is 5.4% for one-week changes, and increases to 36.2% for annual changes. Therefore, even without using accounting data, we can use our model to extract financial leverage variations from the equity index options. The extracted financial leverage series co-move positively with the financial leverage compiled from accounting data, and capture the major leverage variations during historical events.

IV. Concluding Remarks

The variation of equity return volatility can come from three distinct economic channels: variations in financial leverage, business risk, and self-exciting market crash risk. This paper develops an equity index dynamics that accommodates all three sources of return volatility variation, and proposes to disentangle the three sources of variation through the variation in equity index options across a wide range of strikes and maturities. Model estimation shows that the volatility feedback effect reveals itself mainly in the variations of short-term options, the self-exciting behavior affects both short-term and long-term option variations, and the financial leverage variation has its largest impact on long-dated options.

The disentangling of the volatility variation reveals economic insights that one would not be able to obtain from the estimation of a standard reduced-form stochastic volatility model. In particular, the model estimation reveals how the market capital structure responds to different types of risks. Contrary to conventional wisdom, financial leverage does not always decline with increased business risk. Instead, companies respond differently to different types of business risk. The financial leverage can increase with increasing business risk if the risk increase is due to small, diffusive market movements. Only when the self-exciting downside jump risk increases do companies become concerned and start the de-leveraging process.

Historically, reduced-form models can price options well due to their flexibility and tractability but these models are not equipped to answer structural questions such as those related to capture structure decisions. Structural models are specifically designed to answer these structural questions, but they lack the flexibility and tractability to price equity options well, and therefore cannot benefit from the rich information content in equity options. This paper proposes a model that bridges the gap between the two types of models. The model specification is motivated by structural models and can be used to address important capital structure questions. At the same time, it retains the flexibility and tractability of reduced-form specifications and can generate equity option pricing performance superior to all existing models estimated in the literature. The structural nature of the model also makes it more parsimonious and better identified than most existing reduced-form specifications, thus facilitating model estimation and the identification of the underlying economic risk sources.

The analysis is on the S&P 500 index, but similar specifications and analyses are equally applicable to the capital structure decisions of individual companies. Thus, our work can become the starting point of a new strand of literature that bridges the gap between derivative pricing and corporate finance. A line for future research is to link the capital structure behavior difference across different companies to the stock options behavior for these companies.

Appendix

A. Option Pricing

Let $c(F_t, K, T)$ denote the time- t forward value of a European call option on the equity index with strike price K and expiry date T , conditional on the time- t values of the index forward F_t and the three state variables (X_t, v_t^Z, v_t^J) . Since the equity index F_t is driven by two orthogonal sources of variations A_t and X_t under the risk-neutral measure, we can perform the valuation through the law of iterated expectations,

$$\begin{aligned} c(F_t, K, T) &\equiv \mathbb{E}[(F_T - K)^+ | (F_t, X_t, v_t^Z, v_t^J)] \\ &= \mathbb{E}[(X_T A_T - K)^+ | (A_t, X_t, v_t^Z, v_t^J)] \\ &= \mathbb{E}[\mathbb{E}[(\mathcal{X} A_T - K)^+ | (X_T = \mathcal{X}, A_t, v_t^Z, v_t^J)] | X_t] \\ &= \mathbb{E}[\mathcal{X} \cdot C(A_t, K/\mathcal{X}, T) | X_t], \end{aligned} \tag{1}$$

where the function $C(A_t, \mathcal{K}, T)$ is defined as

$$C(A_t, \mathcal{K}, T) \equiv \mathbb{E}[(A_T - \mathcal{K})^+], \tag{2}$$

which can be regarded as the forward value of a call option on the asset with strike \mathcal{K} and expiry T . Equation (1) turns the calculation of the call value on the equity index into the computation of the call value on the asset and a numerical integration over all possible equity-to-asset ratios.

A.1. Fourier transform on asset return and FFT valuation of options on asset

To compute the forward call value on asset $C(A_t, \mathcal{K}, T)$, we first derive the generalized Fourier transform of the log asset return $\ln A_T/A_t$,

$$\phi(u) \equiv \mathbb{E}_t \left[e^{iu \ln A_T/A_t} \right], \quad u \in \mathcal{D} \subseteq \mathbb{C}, \quad (3)$$

where \mathcal{D} denotes a subset of the complex plane under which the expectation in equation (3) is well defined. Once we obtain this transform $\phi(u)$, we can compute the option value C via fast Fourier transform (FFT) following the procedure proposed by Carr and Madan (1999).

Our specification for the dynamics on the log asset value can be represented as a time-changed Lévy process with affine activity rates. The generalized Fourier transform is exponential-affine in the state variables (Carr and Wu (2004)),

$$\phi(u) = \exp \left(-a_Z(\tau) - b_Z(\tau)v_t^Z - a_J(\tau) - b_J(\tau)v_t^J \right), \quad \tau = T - t, \quad (4)$$

where the affine coefficients solve the following ordinary differential equations,

$$\begin{aligned} b'_Z(\tau) &= \psi_Z(u) - \kappa_Z^{\mathbb{M}} b_Z(\tau) - \frac{1}{2} \sigma_Z^2 b_Z(\tau)^2, & a'_Z(\tau) &= b_Z(\tau) \kappa_Z \theta_Z, \\ b'_J(\tau) &= \psi_J(u) - (\kappa_J + \sigma_J v_{J^-}) b_J(\tau) + \frac{\sigma_J b_J(\tau) v_{J^-}^{\mathbb{M}}}{1 + \sigma_J b_J(\tau) v_{J^-}^{\mathbb{M}}}, & a'_J(\tau) &= b_J(\tau) \kappa_J \theta_J, \end{aligned} \quad (5)$$

starting at $a_Z(0) = b_Z(0) = a_J(0) = b_J(0) = 0$, and with

$$\begin{aligned} \psi_Z(u) &= \frac{1}{2}(iu + u^2), \\ \psi_J(u) &= \ln(1 + iuv_J) - iu \ln(1 + v_J), \\ \kappa_Z^{\mathbb{M}} &= \kappa_Z - iu\rho\sigma_Z, \\ v_J^{\mathbb{M}} &= v_J/(1 + iuv_J). \end{aligned}$$

The ordinary differential equations governing the coefficients $(a_Z(\tau), b_Z(\tau))$ can be solved analytically,

$$\begin{aligned} b_Z(t) &= \frac{2\psi_Z(u)(1-e^{-\xi\tau})}{2\xi - (\xi - \kappa_Z^M)(1-e^{-\xi\tau})}, \quad \xi = \sqrt{(\kappa_Z^M)^2 + 2\sigma_Z^2\psi_Z(u)}, \\ a_Z(t) &= \frac{\kappa_Z\theta_Z}{\sigma_Z^2} \left[2\ln\left(1 - \frac{\xi - \kappa_Z^M}{2\xi}(1 - e^{-\xi\tau})\right) + (\xi - \kappa_Z^M)\tau \right]. \end{aligned} \quad (6)$$

The ordinary differential equations governing the coefficients $(a_J(\tau), b_J(\tau))$ can be solved numerically using the standard Runge-Kutta 4th order method.

With the generalized Fourier transform $\phi(u)$ on the asset return, we first re-scale the forward call value on the asset $c(k) = C(A_t, \mathcal{K}, T)/A_t$ with $k = \ln \mathcal{K}/A_t$ to represent the forward call value in percentages of the forward asset value as a function of moneyness defined as the log strike over forward. Then, we derive the Fourier transform on the re-scaled forward call $c(k)$ in terms of the Fourier transform on the asset return,

$$\chi(u) \equiv \int_{-\infty}^{\infty} e^{iuk} c(k) dk = \frac{\phi(u-i)}{(iu)(iu+1)}, \quad (7)$$

which is well-defined when u contains an imaginary component $u = u_r - i\alpha$, with u_r being real and α being a real positive number. With the transform in (7), the call value can be computed via the following inversion,

$$c(k) = \frac{e^{-\alpha k}}{\pi} \int_0^{\infty} e^{-iu_r k} \chi(u_r - i\alpha) du_r. \quad (8)$$

We perform the inversion numerically by discretizing the integral using the trapezoid rule:

$$c(k) \approx \frac{e^{-\alpha k}}{\pi} \sum_{m=0}^N \delta_m e^{-iu_m k} \chi(u_m - i\alpha) \Delta u, \quad (9)$$

where $\delta_m = \frac{1}{2}$ when $m = 0$ and 1 otherwise. We cast the operation in (9) in the form of discrete fast Fourier transform (FFT), which is an efficient algorithm for computing discrete Fourier coefficients. The discrete Fourier transform is a mapping of $\mathbf{f} = (f_0, \dots, f_{N-1})^\top$ on the vector of Fourier

coefficients $\mathbf{d} = (d_0, \dots, d_{N-1})^\top$, such that

$$d_j = \frac{1}{N} \sum_{m=0}^{N-1} f_m e^{-jm\frac{2\pi}{N}i}, \quad j = 0, 1, \dots, N-1. \quad (10)$$

FFT allows the efficient calculation of \mathbf{d} if N is an even number, say $N = 2^n, n \in \mathbb{N}$. The algorithm reduces the number of multiplications in the required N summations from an order of 2^{2n} to that of $n2^{n-1}$, a very considerable reduction.

To map the operation in equation (9) to the FFT form in (10), we set the summation grid by $\eta = \Delta u$ and $u_m = \eta m$, and we set the relative strike grid by $k_j = -b + \lambda j$ with $\lambda = 2\pi/(\eta N)$ and $b = \lambda N/2$. Then, the call value becomes

$$c(k_j) \approx \frac{1}{N} \sum_{m=0}^{N-1} f_m e^{jm\frac{2\pi}{N}i}, \quad f_m = \delta_m \frac{N}{\pi} e^{-\alpha k_j + iu_m b} \chi(u_m - i\alpha)\eta, \quad (11)$$

with $j = 0, 1, \dots, N-1$. The inversion has the FFT form and can hence be computed efficiently across the whole spectrum of strikes k_j .

A.2. Numerical integration with Gauss-Hermite quadrature

Once we have computed the forward call value on asset across the whole spectrum of strikes using the FFT method, we approximate the integration in equation (1) by a weighted sum of a finite number (M) of forward asset call values at equity-to-asset ratio values X_j ,

$$c(F_t, K, T) = \int_0^\infty f(X|X_t) \mathcal{X}C(A_t, K/X, T) dX \approx \sum_{j=1}^M \mathcal{W}_j X_j C(A_t, K/X_j, T), \quad (12)$$

where $f(X|X_t)$ denotes the transition density of X from X_t at time t to X at time T . The points X_j and their corresponding weights in the approximation are chosen according to the Gauss-Hermite quadrature rule.

The constant elasticity of variance process in equation (2) is related to a standard Bessel process of order $\nu = 1/(2p)$ through the change of variable, $z_t = X_t^p/(\delta p)$. From the well-known expression for the transition density of the Bessel process (see Borodin and Salminen (1996) and Revuz and Yor (1999) for details on Bessel processes), we can derive the probability transition density as,

$$f(X|X_t) = \frac{X^{2p-\frac{3}{2}}X_t^{\frac{1}{2}}}{\delta^2 p(T-t)} \exp\left(-\frac{X_t^{2p} + X^{2p}}{2\delta^2 p^2(T-t)}\right) I_\nu\left(\frac{X_t^p X^p}{\delta^2 p^2(T-t)}\right), \quad (13)$$

where $I_\nu(x)$ is the modified Bessel function of the first kind of order ν .

The Gauss-Hermite quadrature rule is designed to approximate the integral $\int_{-\infty}^{\infty} h(x) e^{-x^2} dx$, where $h(x)$ is an arbitrary smooth function. After some re-scaling, the integral can be regarded as an expectation of $h(x)$ where x is a normally distributed random variable with zero mean and variance of one half. See Davis and Rabinowitz (1984) for details.

To apply the quadrature rule, we need to map the quadrature nodes and weights $\{x_i, w_j\}_{j=1}^M$ to our choice of X_j and the weights \mathcal{W}_j . Given the constant elasticity of variance dynamics, one reasonable choice is,

$$X(x) = X_t e^{\sqrt{2V_X}x - \frac{1}{2}V_X}, \quad V_X = X_t^{-2p}(T-t). \quad (14)$$

The choice is motivated by a log-normal approximation of the density of X by assuming that the instantaneous return variance $\delta^2 X_t^{-2p}$ is fixed. Then, given the Gauss-Hermite quadrature $\{w_j, x_j\}_{j=1}^M$, we choose the X_j points as

$$X_j = X_t e^{\sqrt{2V_X}x_j - \frac{1}{2}V_X}, \quad (15)$$

and the summation weights as

$$\mathcal{W}_j = \frac{f(X_j|X_t)X'(x_j)}{e^{-x_j^2}} w_j = \frac{f(X_j|X_t)X_j\sqrt{2V_X}}{e^{-x_j^2}} w_j. \quad (16)$$

B. Statistical Behaviors of the SPX Implied Volatility Surface

This appendix summarizes the statistical behaviors of the SPX implied volatility surface based on the over-the-counter data set and explains how our model can capture these behaviors. The data set contains 40 implied volatility series from a combination of eight time-to-maturities from one month to five years and five relative strikes per maturity from 80 to 120 percent of the spot level.⁵ The time series are sampled weekly from January 8, 1997 to March 5, 2008, for a total of 583 weeks with 23,320 implied volatility observations.

B.1. The mean implied volatility surface

Figure 4 plots the mean implied volatility surface in Panel A as a function of the relative strikes and time to maturities. At each maturity, the mean implied volatility declines with increasing strike, forming the well-known negative implied volatility skew (e.g., Bakshi, Cao, and Chen (1997), Carr and Wu (2003), and Foresi and Wu (2005)). Across different maturities, since an 80% relative strike is far more out-of-the-money at one-month maturity than at five-year maturity, the implied volatility skew looks much steeper at short than at long maturities.

[Figure 4 about here.]

To adjust for the maturity and implied volatility level difference, we define a standardized moneyness measure, $d \equiv \ln(K/S)/(IV\sqrt{\tau})$, where IV denotes the implied volatility at the relative strike (K/S) and the time to maturity $\tau = T - t$ (in years). With the standardized moneyness, we construct an implied volatility skew measure that is comparable across dates and maturities,

$$SK_{t,T} = \frac{IV_{t,T}(80\%) - IV_{t,T}(120\%)}{|d_{t,T}(80\%) - d_{t,T}(120\%)|}. \quad (17)$$

⁵Instead of quoting dollar invoice prices, the OTC market quotes in terms of the Black-Scholes implied volatility. The invoice dollar price is calculated based on the Black-Scholes formula using the implied volatility quote, a reference spot price level, and the agreed-upon financing details, e.g., interest rates, borrowing cost, and dividend yields.

which measures the implied volatility difference at 80% and 120% strikes, divided by the absolute difference in their respective standardized moneyness. Since the 80% strike implied volatility is universally higher than the corresponding 120% strike implied volatility across all dates and maturities, the skew measure generates positive estimates. The higher the estimate, the steeper the negative slope of the implied volatility plot against the standardized moneyness measure. Panel B of Figure 4 plots the sample averages of the skew measure at different maturities in the solid line and their 10th and 90th percentiles in the two dashed lines. The skew measure is on average lower at short maturities due to the smile pattern (curvature). As the option maturity increases, the smile becomes a pure skew and the slope increases in absolute magnitude.

B.2. Principal component analysis

To understand the driving forces underlying the implied volatility surface variation, we perform principal component analysis (PCA) on the 40 implied volatility series. We estimate the covariance matrix of the weekly changes on the 40 series, and compute the eigenvalues and eigenvectors of the covariance matrix. The normalized eigenvalues can be interpreted as the percentage variation explained by each principal component, and the mimicking portfolio for each component can be formed by using the eigenvector corresponding to the eigenvalue as the portfolio weight. The bar chart in Panel A of Figure 5 represents the normalized eigenvalues for the first ten principal components. The first three principal components explain 85.09%, 8.24%, and 3.30% of the weekly variation in the implied volatility surface, respectively, for a total of 96.62%. The PCA evidence suggests that our three-factor stochastic volatility structure has the potential to explain the majority of the variation in the implied volatility surface.

[Figure 5 about here.]

The surface in Panel B of Figure 5 represents the eigenvector corresponding to the largest eigenvalue, which captures the loading of the first principal component (PC) on the 40 implied volatility series. The loadings are positive and relatively uniform across all 40 series. Thus, this first principal component captures the overall level variation of the implied volatilities. The loadings decline with increasing maturity, reflecting the declining variation of the implied volatility series at longer maturities.

Panel C of Figure 5 plots the loadings of the second principal component, which are relatively flat across the strike dimension at each maturity, but decline monotonically as the time to maturity increases. The loadings are positive at short maturities, but become negative when the maturities are six months and longer, thus creating a term structure slope factor.

Panel D of Figure 5 shows that the loadings of the third principal component exhibit strong variation along the strike dimension. At each maturity, the loadings are positive at low strikes but negative at high strikes. The strike dimension variation is particularly strong at short maturities. The smaller variation at longer maturities is partly due to the scaling because the same relative strike range implies a smaller range of standard deviations at longer maturities.

To gauge the contribution of long-dated options, we have also performed the PCA with the 20 OTC implied volatility series with maturities from one month to one year, while excluding the series with maturities greater than one year. In this case, the loading patterns on the second and third principal components become less clear but become a mix of strike and maturity dimension variations. Thus, it is crucial to include long-dated options to identify the term structure effect, which is crucial in separating the different sources of stochastic volatility in our model specification.

Compared to the OTC implied volatility quotes, the exchange-listed index options trade over a wider strike range but a narrower maturity range. We have also performed the PCA on the listed options. Since the listed options are at fixed strike prices and expiries, we need to perform

interpolation to obtain implied volatility time series at fixed moneyness and time to maturities. From the interpolated implied volatility data on listed options, we obtain a similar three-factor structure, but the strike-dimension variation contributes more to the total variance than does the term structure variation.

B.3. Dynamic interactions between index returns and implied volatilities

To understand the interactions between the stock index and the option implied volatility levels and skews, we use the average of the eight at-the-money implied volatility series to proxy the volatility level and the average of the eight implied volatility skew series defined in equation (17) to proxy the volatility skew. The weekly index returns and the weekly changes in the volatility level and skew show strongly negative correlations, estimated at -0.8114 and -0.707 , respectively. When the stock index drops, the implied volatility level tends to go up and the negative implied volatility skew tends to become steeper. Accordingly, the risk-neutral return distribution becomes both wider and more negatively skewed.

Figure 6 plots the weekly changes in the volatility level and the skew in Panel A and Panel B, respectively, against weekly returns on the equity index. The circles denote data points and the solid lines denote local linear regression fits with a Gaussian kernel and a default bandwidth choice. The plots confirm the correlation estimates in showing strongly negative co-movements between the index return and changes in the volatility level and the volatility skew.

[Figure 6 about here.]

The plot in panel B also shows asymmetry in the volatility skew response to the index return. The response is stronger to large negative returns than to large positive returns. The slope of the fitted line is steeper when the index return is highly negative, but becomes virtually flat when the index return is highly positive. When measuring the correlation between the index return and the

skew changes conditioning on different return levels, we obtain an estimate of -0.6612 conditioning on the index return being lower than -5% , but a virtually zero estimate at 0.03 conditioning on the index return being greater than 5% .

B.4. The model's structural capability in capturing the implied volatility behaviors

Backus, Foresi, and Wu (1997) show that the shape of the implied volatility plot against the standardized moneyness measure reflects the deviation of the underlying return risk-neutral distribution from the normality assumption under the Black-Scholes model. A negative implied volatility skew reveals negative skewness in the return distribution under the risk-neutral measure. The fact that the slope becomes steeper at longer option maturities indicates that the risk-neutral index return distribution becomes more negatively skewed at longer conditioning horizons.

Many mechanisms can generate negative skewness in the return distribution at a fixed horizon, but generating increasing skewness as the horizon increases up to five years can be difficult. By virtue of the central limit theorem, if the return distribution is stationary and has finite variance, the distribution should converge to normality with time aggregation. Most pure jump models or stochastic volatility models with mean-reverting variance dynamics satisfy the conditions for the central limit theorem and thus imply flatter implied volatility skews at long option maturities.⁶

Our model incorporates four distinct mechanisms that can generate a negative implied volatility skew: (i) the leverage effect via the power dependence on leverage ($p > 0$); (ii) the negative jump in the asset value; (iii) the volatility feedback effect via the negative instantaneous correlation ($\rho < 0$) between the Brownian motion in the asset return and the Brownian motion in its variance rate; and (iv) the self-exciting behavior via the synchronized jump in the asset return and its intensity. Among the four sources, jumps in asset return induce return skewness at short-maturities. The volatility feedback effect and the self-exciting mechanism induce negative skewness at inter-

⁶One exception is the finite-moment log stable pure jump model of Carr and Wu (2003), where the central limit theorem no longer applies as the return variance is infinite under the risk-neutral measure.

mediate maturities. The persistence of the skewness depends on the risk-neutral mean-reversion speed of the corresponding variance rate dynamics. With stationary variance rate dynamics, all three sources of skewness will eventually die out at long maturities. However, the leverage effect induces a negative skewness that does not die out at long maturities. The $3/2$ process derived from the financial leverage variation in (13) is mean-repelling under the risk-neutral measure and can thus circumvent the convergence behavior dictated by the central limit theorem. Therefore, our model has the structural capability of generating the observed term structure behavior for the implied volatility skew.

The principal component analysis shows that one needs a three-factor volatility structure to capture the level, the term structure, and the skew variation of the implied volatility surface. Our model is the only known model estimated in the literature that can accommodate all three dimensions of variation and their dynamic interactions with the index return.

The diffusion variance rate v_t^Z , the jump variance rate v_t^J , and the financial leverage variation X_t (or its transformed variance rate v_t^X) generate three distinct sources of variation in the implied volatility surface. Shocks to the three variance rates (v_t^Z, v_t^J, v_t^X) can have different impacts across the implied volatility term structure when their risk-neutral mean-reversion speeds are different. The different responses generate variations in the term structure of the implied volatility that dominate the second principal component.

Furthermore, downside jumps in the asset value generate instantaneous negative skewness in the equity index return. The relative variation between the two variance rates (v_t^Z, v_t^J) alters the relative composition of continuous versus discontinuous components in the asset return and thus generates variations in the return skewness, particularly at short to moderate horizons. By contrast, the return skewness generated from the financial leverage affects both short and long maturities. These variations in skewness can explain the variation of the third principal component.

Finally, all three sources of variations (v_t^X, v_t^Z, v_t^J) have negative co-movements with the index return and thus contribute to the observed negative correlation between index returns and implied volatility level changes. The self-exciting behavior embedded in the v_t^J dynamics generates the observed asymmetry in the relation between the index return and the volatility skew.

C. Quasi-Maximum Likelihood Estimation with Unscented Kalman Filter

Let $V_t \equiv [X_t, v_t^Z, v_t^J]^\top$ denote the state vector at time t . We specify the state propagation equation based on an Euler approximation of the statistical dynamics of the three state variables:

$$V_t = f(V_{t-1}; \Theta) + \sqrt{Q_{t-1}} \varepsilon_t, \quad (18)$$

where ε_t denotes the standardized forecasting error vector. The forecasting function $f(V_{t-1}; \Theta)$ and the forecasting error covariance matrix are given by,

$$f(V_{t-1}; \Theta) = \begin{bmatrix} X_t + X_t^{1-p}(a - \kappa_L^\top V_{t-1})\Delta t \\ \kappa_Z \theta_Z \Delta t + (1 - \kappa_Z^{\mathbb{P}} \Delta t) v_{t-1}^Z \\ \kappa_J \theta_J \Delta t + (1 - \kappa_J^{\mathbb{P}} \Delta t) v_{t-1}^J \end{bmatrix}, \quad Q_{t-1} = \left\langle \begin{matrix} X_{t-1}^{2-2p} \\ \sigma_Z^2 v_{t-1}^Z \\ \sigma_J^2 (v_{t-1}^{\mathbb{P}})^2 v_{t-1}^J \end{matrix} \right\rangle \Delta t \quad (19)$$

with $\Delta t = 7/365$ denoting the weekly frequency of the data, $\langle \cdot \rangle$ denoting a diagonal matrix, and Θ denoting the parameter set.

The measurement equations are specified on the implied volatility quotes, with additive, normally-distributed measurement errors:

$$y_t = h(V_t; \Theta) + \sqrt{R} e_t, \quad (20)$$

where y_t denotes the time- t forward value of the out-of-the-money options computed from the implied volatility quotes, scaled by the Black-Scholes vega of the option, $h(V_t; \Theta)$ denotes the

corresponding model value as a function of the state vector V_t and the parameter set Θ . Given the assumed risk-neutral dynamics, the option value $h(V_t; \Theta)$ can be computed efficiently via a combination of quadrature integration and fast Fourier inversion as described in Appendix A. We assume that the pricing errors on the scaled option prices are i.i.d. normal with zero mean and constant variance. Hence, we can write the pricing error covariance matrix as, $R = \sigma_e^2 I_{40}$, with σ_e being a scalar and I_{40} denoting an identity matrix of dimension 40, corresponding to the 40 implied volatility series across the five relative strikes at each maturity and eight time to maturities.

When the state propagation and the measurement equation are Gaussian linear, the Kalman (1960) filter provides efficient forecasts and updates on the mean and covariance of the state vector and observations. Our state-propagation equations and measurement equations do not satisfy the Gaussian and linear conditions. We use an extended version of the Kalman filter, the unscented Kalman filter, to handle the deviations.

Let $\bar{V}_t, \bar{y}_t, \bar{\Sigma}_{xy,t}$ denote the time- $(t-1)$ ex ante forecasts of time- t values of the state vector, the measurement series, and the covariance between series x and y ; let $\hat{V}_t, \hat{y}_t, \hat{\Sigma}_{xy,t}$ denote the corresponding ex post update on the state vector, the measurement, and the covariances. The unscented Kalman filter uses a set of deterministically chosen (sigma) points to approximate the state distribution. At each time t , if we use k to denote the number of states (three in our model) and use $\eta > 0$ to denote a control parameter, we first generate a set of $2k+1$ sigma vectors χ_{t-1} from the time $(t-1)$ updated mean \hat{V}_{t-1} and covariance $\hat{\Sigma}_{VV,t-1}$ of the state vector according to the following equations,

$$\begin{aligned} \chi_{t-1,0} &= \hat{V}_{t-1}, \\ \chi_{t-1,i} &= \hat{V}_{t-1} \pm \sqrt{(k+\eta)(\hat{\Sigma}_{VV,t-1})_j}, \quad j=1, \dots, k; \quad i=1, \dots, 2k, \end{aligned} \tag{21}$$

with the corresponding weights w_i given by,

$$w_0 = \eta/(k + \eta), \quad w_i = 1/[2(k + \eta)], \quad i = 1, \dots, 2k. \quad (22)$$

These sigma vectors form a discrete distribution with w_i being the corresponding probabilities.

We propagate these sigma points through the propagation equation (18) to compute the forecasted mean and covariance of the state vector at time t ,

$$\begin{aligned} \bar{\chi}_{t,i} &= f(\chi_{t-1,i}; \Theta), \quad \bar{V}_t = \sum_{i=0}^{2k} w_i \bar{\chi}_{t,i}, \\ \bar{\Sigma}_{VV,t} &= \sum_{i=0}^{2k} w_i (\bar{\chi}_{t,i} - \bar{V}_t)(\bar{\chi}_{t,i} - \bar{V}_t)^\top + Q_{t-1}. \end{aligned} \quad (23)$$

We then re-generate the sigma points $\tilde{\chi}_t$ based on the forecasted mean \bar{V}_t and covariance $\bar{\Sigma}_{VV,t}$, and compute the forecasted mean and covariances of the measurements,

$$\begin{aligned} \bar{\xi}_{t,i} &= h(\tilde{\chi}_{t,i}; \Theta), \quad \bar{y}_t = \sum_{i=0}^{2k} w_i \bar{\xi}_{t,i}, \\ \bar{\Sigma}_{yy,t} &= \sum_{i=0}^{2k} w_i (\bar{\xi}_{t,i} - \bar{y}_t) (\bar{\xi}_{t,i} - \bar{y}_t)^\top + R, \\ \bar{\Sigma}_{Vy,t} &= \sum_{i=0}^{2k} w_i (\tilde{\chi}_{t,i} - \bar{V}_t) (\bar{\xi}_{t,i} - \bar{y}_t)^\top. \end{aligned} \quad (24)$$

With these moment conditions, we perform the filtering step the same as in the the Kalman filter,

$$\hat{V}_t = \bar{V}_t + \mathcal{K}_t (y_t - \bar{y}_t), \quad \hat{\Sigma}_{VV,t} = \bar{\Sigma}_{VV,t} - \mathcal{K}_t \bar{\Sigma}_{yy,t} \mathcal{K}_t^\top, \quad (25)$$

where the Kalman gain is $\mathcal{K}_t = \bar{\Sigma}_{Vy,t} (\bar{\Sigma}_{yy,t})^{-1}$. We refer the reader to Wan and van der Merwe (2001) for general treatments of the unscented Kalman filter.

Given the forecasted option prices \bar{y}_t and their conditional covariance matrix $\bar{\Sigma}_{yy,t}$ obtained from the unscented Kalman filtering, we compute the quasi-log likelihood value for each week's

observation on the option prices assuming normally distributed forecasting errors,

$$l_t(\Theta) = -\frac{1}{2} \log |\bar{\Sigma}_{yy,t}| - \frac{1}{2} \left((y_t - \bar{y}_t)^\top (\bar{\Sigma}_{yy,t})^{-1} (y_t - \bar{y}_t) \right). \quad (26)$$

We estimate the model parameters by numerically maximizing the sum of the conditional log likelihood value on each date,

$$\Theta \equiv \arg \max_{\Theta} \mathcal{L}(\Theta, \{y_t\}_{t=1}^N), \quad \text{with} \quad \mathcal{L}(\Theta, \{y_t\}_{t=1}^N) = \sum_{t=1}^N l_t(\Theta), \quad (27)$$

where N denotes the number of weeks in the sample.

D. Option Pricing Performance Comparison with A Reduced-Form Benchmark

In theory, a reduced-form model designed from the affine framework of Duffie, Pan, and Singleton (2000) can accommodate any number of factors. In practice, without the discipline of an economic structure, such models often experience identification issues. As a result, most empirical studies in the literature limit the specifications to one stochastic volatility factor, with only a few studies exploring the estimation of two volatility factors. By contrast, by imposing economic structures, our model not only can answer structural economic questions that cannot be addressed by a reduced-form specification, but it is also very parsimonious and highly identifiable, even with three volatility factors. The three-factor structure is supported by the evidence from the principal component analysis. By design, it shall perform better than any one- or two-factor specifications in matching the observed option behavior. To gauge the statistical significance of the superior performance, we compare the pricing performance of our model with a two-factor reduced-form benchmark.

We create the two-factor reduced-form benchmark by setting $X_t = 1$ and hence $F_t = A_t$ in

our specification. This way, we can regard the benchmark as a restricted version of our model that ignores the leverage effect. By comparison, Bates (2000) specifies two jump-diffusion return components, each driven by a separate stochastic volatility process. The model can experience identification issues as the two return components can rotate. Huang and Wu (2004) enhance the identification by limiting one return component to a pure diffusion and the other component to be a pure jump process. Our benchmark specification is similar to Huang and Wu (2004), but with an additional layer of interaction between the index return and return volatility through the self-exciting behavior of downside return jumps. Thus, the reduced-form benchmark that we construct represents an enhanced version of the state of the art in reduced-form equity index option pricing.

Table II reports the parameter estimates for the reduced-form benchmark model. Without the financial leverage factor, the estimated risk-neutral dynamics for both the diffusion variance and the jump intensity become more persistent to help capturing the behavior of long-term options. In particular, the risk-neutral mean-reversion speed for the jump intensity κ_J becomes virtually zero. Compared to the full model, the estimation of the reduced-form benchmark experiences some stability issues as the model struggles between fitting the short-term versus the long-term options.⁷

Table III reports the summary statistics of the pricing errors from the two estimated models. Panel A reports the sample averages of the pricing errors, defined as the difference between the implied volatility quotes and the corresponding model values. The mean pricing errors from our model are mostly small except at the one-month maturity and do not show any obvious patterns. The mean pricing errors from the reduced-form benchmark are larger overall and show an obvious term structure pattern: The mean pricing errors are mostly positive at short and long maturities but are negative at intermediate maturities. The pattern suggests that the benchmark model generates more curvature on the term structure than observed from the data.

⁷The stability issue becomes even more severe when one estimates a one-factor stochastic volatility model such as Heston (1993) to options across a wide range of maturities. In particular, the mean-reversion speed of the variance process becomes larger when the estimation is only on short-term options, but the estimate can approach zero or even negative when long-term options are used in the model estimation.

[Table III about here.]

Panel B reports the mean absolute pricing error in implied volatility points. The estimates from our model are mostly smaller than those from the reduced-form benchmark. The average mean absolute pricing error from the 40 implied volatility series is 0.63 for our model and 1.05 for the reduced-form benchmark. Thus, our model reduces the mean absolute pricing error of the benchmark by 40%. The bid-ask spreads for OTC implied volatility quotes are about half a volatility point for the liquid series, and wider for less liquid series. The mean absolute pricing error from our model is just about as wide as the average bid-ask spread. The mean absolute pricing error from the benchmark, however, is over one volatility point.

Panel C reports the mean absolute percentage error, where the pricing errors are represented in percentages of the observed option value. The mean absolute pricing error averages at 5.68% for the reduced-form benchmark and 3.43% for our model, against a 40% reduction. Thus, our model generates vastly better option pricing performance than does the benchmark.

The last row of the table reports the maximized log likelihood values from the two models. The reduced-form benchmark can be regarded as a constrained version of our model, with five fewer parameters and one fewer state variable. One can construct a likelihood ratio static between the two models as twice the difference between the log likelihood values, which has a Chi-squared distribution with 588 degrees of freedom. The p -value from this likelihood ratio test is virtually zero. The reduced-form benchmark is strongly rejected.

To gauge the stability of the model and its out-of-sample performance, we divide the sample into two subsamples. The first subsample contains 291 weeks from January 8, 1997 to July 31, 2002, and the second subsample contains 292 weeks from August 7, 2002 to March 5, 2008. We re-estimate the models using the first subsample and use the subsample parameter estimates to price options over the whole sample period while allowing the states to propagate weekly. Table V re-

ports the mean absolute pricing error from this exercise. The three panels report the mean absolute pricing error over the whole sample period, the in-sample period, and the out-of-sample period, respectively. For each model, the performances over the different sample periods are similar, with the out-of-sample performance slightly worse than the in-sample performance. Comparing the pricing performance of the two models, the mean absolute pricing error for our model is about half of the two-factor benchmark, for both in sample and out of sample. Likelihood ratio tests generate virtually zero p -values for both the in-sample period and the out-of-sample period.

[Table V about here.]

Taken together, compared to the two-factor reduced-form benchmark, our full model not only enjoys better economic interpretation, but also generates vastly superior performance, both in sample and out of sample. Statistically, the performance difference is significant beyond any reasonable confidence level. Economically, our model can reduce the mean absolute pricing error of the benchmark by 40% or more.

E. Estimation with Exchange-Traded Options

For the main results of the paper, we have chosen to use OTC options quotes to estimate the model because they span a wider range of option maturities and can thus enhance the identification of the more persistent risk factors such as the self-exciting jump intensity and the financial leverage variation. For robustness, we investigate how the results vary if we estimate the model using options quotes from the listed market, the Chicago Board of Options Exchange (CBOE). We obtain the listed market data from OptionMetrics.

As in the estimation with the OTC data, we sample the listed options weekly over the same sample period. On the exchange, both a call and a put option trade at each strike and maturity.

The two options contain the same information about the underlying equity index dynamics as their differences merely reflect the current forward level. Therefore, at each strike and maturity, we pick the out-of-the-money (OTM) option (call when the strike is greater than the forward and put otherwise) for model estimation because OTM options tend to be more actively traded than their in-the-money counterparts. We further filter the data by requiring that (i) the time-to-maturities of the chosen options are no less than 21 days and (ii) the log strike deviation from the log forward is within two standard deviations of its mean.⁸ With such filtering, 113,940 observations are retained for model estimation. The number of observations is about five fold of the OTC data set because the strikes are much more densely populated in the listed market.

Out of the 113,940 options, 71,519 are OTM put options and 42,421 are OTM call options. Figure 7 plots the histogram of the option maturities in Panel A and the histogram of the relative strikes in percentages of the forward price in Panel B. Panel A shows that the exchange-traded options span a narrower range of maturities than the OTC data. The longest time-to-maturity in the listed data set is less than three years, and the time-to-maturities of the options are concentrated at short maturities. Panel B shows that the strikes on the exchange-traded options are not only more densely populated, but also cover a wider range, even after our truncations of deep out-of-the-money options.

[Figure 7 about here.]

We employ the same estimation procedure on the listed options data set. Since the number of option contracts on the exchange varies over time, the dimension of the measurement equation varies. Our estimation procedure can readily accommodate the varying dimension of the measurement equation. Table VI report the parameter estimates and standard deviations (in parentheses).

⁸Specifically, for each option contract, we define the standardized variable $z = (\ln K/F_t + \frac{1}{2}IV^2\tau)/(IV\sqrt{\tau})$, with τ denoting the time-to-maturity and IV denoting the corresponding implied volatility of the contract. Then, we exclude options with z greater than two in absolute magnitude.

The parameter estimates are similar in magnitude to those obtained from the OTC data in Table I. The standard errors for most parameters are even smaller given the much larger number of observations. The similarities between the two sets of parameter estimates highlight the robustness of the model specification and the estimation procedure. The discussions based on the two sets of estimates would be qualitatively similar.

The three state variables extracted from the listed options also follow similar historical patterns as those extracted from the OTC data, but with slightly more idiosyncratic movements. When we measure the cross-correlations on the two sets of state variables, the correlation estimates on the levels are at 0.8705, 0.9526, 0.9744, respectively, for the three state variables (X_t, v_t^Z, v_t^J) . The correlation estimates on the corresponding weekly changes are, 0.5755, 0.9313, 0.7915, respectively. The correlation between the two sets of estimates is the highest for v_t^Z and the lowest for X_t .

This ranking for the correlation estimates is expected. The impacts of the highly mean-reverting diffusion variance rate v_t^Z are mainly on short-term options, where the two sets of data overlap. Therefore, the extracted movements from the two markets are similar. On the other hand, since the X_t -transformed variance rate v_t^X is mean-repelling under the risk-neutral measure, its variation mainly shows up on long-dated options, on which we have more observations from the OTC data set than from the listed options data set. As a result, the series extracted from the OTC data can be more accurate than that from the listed options data. When we measure the correlation between changes in the accounting debt ratio (DR^A) and changes in the corresponding measure (DR^E) extracted from the listed options, the estimates are smaller over one-week changes at 1.44%. The estimates even become negative at intermediate horizons, suggesting that the leverage variations extracted from the listed options data are not as accurate because leverage variations mostly affect long-term options. Only when we measure the correlation over annual changes do we obtain a similar estimate at 36.93%.

REFERENCES

- Adrian, T., and H. S. Shin, 2010, “Liquidity and Leverage,” *Journal of Financial Intermediation*, 19(3), 418–437.
- Aït-Sahalia, Y., J. Cacho-Diaz, and R. J. Laeven, 2010, “Modeling Financial Contagion Using Mutually Exciting Jump Processes,” Nber working paper no. 15850, Princeton University.
- Azizpour, S., and K. Giesecke, 2008, “Self-Exciting Corporate Defaults: Contagion or Frailty,” working paper, Stanford University.
- Backus, D., S. Foresi, and L. Wu, 1997, “Accounting for Biases in Black-Scholes,” working paper, New York University.
- Bakshi, G., C. Cao, and Z. Chen, 1997, “Empirical Performance of Alternative Option Pricing Models,” *Journal of Finance*, 52(5), 2003–2049.
- Bakshi, G., N. Ju, and H. Ou-Yang, 2006, “Estimation of Continuous-time Models with an Application to Equity Volatility,” *Journal of Financial Economics*, 82(1), 227–249.
- Bakshi, G., and N. Kapadia, 2003a, “Delta-Hedged Gains and the Negative Market Volatility Risk Premium,” *Review of Financial Studies*, 16(2), 527–566.
- Bakshi, G., and N. Kapadia, 2003b, “Volatility Risk Premium Embedded in Individual Equity Options: Some New Insights,” *Journal of Derivatives*, 11(1), 45–54.
- Bakshi, G., and D. Madan, 2006, “A Theory of Volatility Spread,” *Management Science*, 52(12), 1945–1956.

- Bates, D. S., 2000, "Post-'87 Crash Fears in the S&P 500 Futures Option Market," *Journal of Econometrics*, 94(1), 181–238.
- Beckers, S., 1980, "The Constant Elasticity of Variance Model and Its Implications for Option Pricing," *Journal of Finance*, 35(3), 661–673.
- Black, F., 1976, "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3(1-2), 167–179.
- Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81(3), 637–654.
- Bollerslev, T., M. Gibson, and H. Zhou, 2011, "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities," *Journal of Econometrics*, 160(1), 235–245.
- Bollerslev, T., G. Tauchen, and H. Zhou, 2009, "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22(11), 4463–4492.
- Borodin, A. N., and P. Salminen, 1996, *Handbook of Brownian Motion*. Birkhauser, Boston, MA.
- Broadie, M., M. Chernov, and M. Johannes, 2007, "Model Specification and Risk Premia: Evidence from Futures Options," *Journal of Finance*, 62(3), 1453–1490.
- Carr, P., H. Geman, and D. Madan, 2001, "Pricing and Hedging in Incomplete Markets," *Journal of Financial Economics*, 62(1), 131–167.
- Carr, P., and D. Madan, 1999, "Option Valuation Using the Fast Fourier Transform," *Journal of Computational Finance*, 2(4), 61–73.
- Carr, P., and J. Sun, 2007, "A New Approach for Option Pricing under Stochastic Volatility," *Review of Derivatives Research*, 10(2), 87–150.

- Carr, P., and L. Wu, 2003, “Finite Moment Log Stable Process and Option Pricing,” *Journal of Finance*, 58(2), 753–777.
- Carr, P., and L. Wu, 2004, “Time-Changed Lévy Processes and Option Pricing,” *Journal of Financial Economics*, 71(1), 113–141.
- Carr, P., and L. Wu, 2007, “Stochastic Skew in Currency Options,” *Journal of Financial Economics*, 86(1), 213–247.
- Carr, P., and L. Wu, 2009, “Variance Risk Premiums,” *Review of Financial Studies*, 22(3), 1311–1341.
- Chacko, G., and L. Viceira, 2003, “Spectral GMM Estimation of Continuous-Time Processes,” *Journal of Econometrics*, 116(1), 259–292.
- Choi, J., and M. Richardson, 2008, “The Volatility of the Firm’s Assets,” working paper, New York University.
- Christie, A. A., 1982, “The Stochastic Behavior of Common Stock Variances: Value, Leverage, and Interest Rate Effects,” *Journal of Financial Economics*, 10(4), 407–432.
- Cox, J. C., 1996, “The Constant Elasticity of Variance Option Pricing Model,” *Journal of Portfolio Management*, 23(1), 15–17.
- Cremers, M., J. Driessen, and P. J. Maenhout, 2008, “Explaining the Level of Credit Spreads: Option-Implied Jump Risk Premia in a Firm Value Model,” *Review of Financial Studies*, 21(5), 2209–2242.
- Davis, P. J., and P. Rabinowitz, 1984, *Methods of Numerical Integration*. Academic Press, New York.

- Ding, X., K. Giesecke, and P. I. Tomecek, 2009, "Time-Changed Birth Processes and Multi-Name Credit Derivatives," *Operations Research*, 57(4), 990–1005.
- Drechsler, I., and A. Yaron, 2010, "Whats Vol Got to Do With It," *Review of Financial Studies*, forthcoming.
- Duffie, D., J. Pan, and K. Singleton, 2000, "Transform Analysis and Asset Pricing for Affine Jump Diffusions," *Econometrica*, 68(6), 1343–1376.
- Dupire, B., 1994, "Pricing with a Smile," *Risk*, 7(1), 18–20.
- Egloff, D., M. Leippold, and L. Wu, 2010, "The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments," *Journal of Financial and Quantitative Analysis*, 45(5), 1279–1310.
- Emanuel, D. C., and J. D. MacBeth, 1982, "Further Results on the Constant Elasticity of Variance Call Option Pricing Model," *Journal of Financial and Quantitative Analysis*, 17(4), 533–554.
- Eraker, B., 2004, "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices," *Journal of Finance*, 59(3), 1367–1404.
- Eraker, B., M. Johannes, and N. Polson, 2003, "The Impact of Jumps in Equity Index Volatility and Returns," *Journal of Finance*, 58(3), 1269–1300.
- Foresi, S., and L. Wu, 2005, "Crash-O-Phobia: A Domestic Fear or A Worldwide Concern?," *Journal of Derivatives*, 13(2), 8–21.
- Heston, S., 1997, "A Simple New Formula for Options with Stochastic Volatility," working paper, University of Maryland.
- Heston, S. L., 1993, "Closed-Form Solution for Options with Stochastic Volatility, with Application to Bond and Currency Options," *Review of Financial Studies*, 6(2), 327–343.

- Huang, J., and L. Wu, 2004, "Specification Analysis of Option Pricing Models Based on Time-Changed Lévy Processes," *Journal of Finance*, 59(3), 1405–1440.
- Hull, J., and A. White, 1987, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42(2), 281–300.
- Hurd, T., and C. Li, 2008, "In Search of Hybrid Models for Credit Risk: from Leland-Toft to Carr-Linetsky," working paper, McMaster University.
- Ishida, I., and R. F. Engle, 2002, "Modeling Variance of Variance: The Square Root, the Affine, and the CEV GARCH models," working paper, New York University.
- Jackwerth, J. C., and M. Rubinstein, 1996, "Recovering Stochastic Processes from Option Prices," working paper, University of California–Berkeley.
- Javaheri, A., 2005, *Inside Volatility Arbitrage: The Secrets of Skewness*. John Wiley & Sons, London.
- Jones, C. S., 2003, "The Dynamics of Stochastic Volatility: Evidence From Underlying and Options Markets," *Journal of Econometrics*, 116(1), 181–224.
- Kalman, R. E., 1960, "A New Approach to Linear Filtering and Prediction Problems," *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D), 35–45.
- Kou, S. G., 2002, "A Jump-Diffusion Model for Option Pricing," *Management Science*, 48(8), 1086–1101.
- Leland, H. E., 1994, "Risky Debt, Bond Covenants and Optimal Capital Structure," *Journal of Finance*, 49(4), 1213–1252.
- Leland, H. E., and K. B. Toft, 1996, "Optimal Capital Structure, Endogenous Bankruptcy and the Term Structure of Credit Spreads," *Journal of Finance*, 51(3), 987–1019.

- Lewis, A. L., 2000, *Option Valuation under Stochastic Volatility*. Finance Press, Newport Beach, California, USA.
- Madan, D., and E. Seneta, 1990, “The Variance Gamma (V.G.) Model for Share Market Returns,” *Journal of Business*, 63(4), 511–524.
- Madan, D. B., P. P. Carr, and E. C. Chang, 1998, “The Variance Gamma Process and Option Pricing,” *European Finance Review*, 2(1), 79–105.
- Medvedev, and O. Scaillet, 2003, “A Simple Calibration Procedure of Stochastic Volatility Models with Jumps by Short Term Asymptotics,” working paper, HEC, University of Geneva.
- Merton, R. C., 1974, “On the Pricing of Corporate Debt: The Risk Structure of Interest Rates,” *Journal of Finance*, 29(1), 449–470.
- Merton, R. C., 1976, “Option Pricing When Underlying Stock Returns Are Discontinuous,” *Journal of Financial Economics*, 3(1), 125–144.
- Pan, J., 2002, “The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study,” *Journal of Financial Economics*, 63(1), 3–50.
- Revuz, D., and M. Yor, 1999, *Continuous Martingales and Brownian Motion*. Springer, Berlin, Germany, 3rd edn.
- Schroder, M., 1989, “Computing the Constant Elasticity of Variance Option Pricing Formula,” *Journal of Finance*, 44(1), 211–219.
- Wan, E. A., and R. van der Merwe, 2001, “The Unscented Kalman Filter,” in *Kalman Filtering and Neural Networks*, ed. by S. Haykin. Wiley & Sons Publishing, New York, chap. 7, pp. 221–280.
- Zhou, H., 2010, “Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty,” working paper, Federal Reserve Board.

Table I
Maximum Likelihood Estimates of Model Parameters

Entries report the maximum likelihood estimates of the model parameters and their standard errors (in parentheses). The parameters are grouped into three panels, each describing the dynamics of one source of volatility risk.

Panel A. Leverage Effect					
	ρ	κ_{XX}	κ_{XZ}	κ_{XJ}	a_X
Coef.	2.9952	0.0001	17.5341	-0.0774	0.0003
Std. Err.	(0.0086)	(0.0000)	(0.4735)	(0.0000)	(0.0003)
Panel B. Volatility Feedback					
	ρ	κ_Z	θ_Z	σ_Z	γ'
Coef.	-0.8388	3.0112	0.0244	0.5987	-17.4512
Std. Err.	(0.0038)	(0.0151)	(0.0001)	(0.0033)	(0.4603)
Panel C. Self-Exciting Market Disruptions					
	ν_J	κ_J	θ_J	σ_J	γ^J
Coef.	0.1926	0.0009	113.8402	5.6356	0.4468
Std. Err.	(0.0004)	(0.0000)	(0.0773)	(0.0321)	(0.0033)

Table II
Maximum Likelihood Estimates of Model Parameters for the Reduced-form Benchmark

Entries report the maximum likelihood estimates of the model parameters and their standard errors (in parentheses) for the two-factor reduced-form benchmark model. The parameters are grouped into two panels, each describing the dynamics of one source of volatility risk.

Panel A. Diffusion Variance Dynamics					
	ρ	κ_Z	θ_Z	σ_Z	γ'
Coef.	-0.7055	1.5275	0.0307	0.5412	-4.6827
Std. Err.	(0.0022)	(0.0085)	(0.0002)	(0.0023)	(0.0486)
Panel B. Self-Exciting Jump Intensity Dynamics					
	ν_J	κ_J	$\kappa_J\theta_J$	σ_J	γ'
Coef.	0.0710	0.0000	1.1325	23.1720	0.0338
Std. Err.	(0.0001)	(0.0000)	(0.0126)	(0.1447)	(0.0004)

Table III
Comparing the Model Pricing Performance to a Reduced-form Benchmark

Entries report the summary statistics of the pricing errors from our model (left side) and the reduced-form benchmark (right side). The pricing errors are defined as the difference between market observations and model values. The last row of the table reports the maximized log likelihood values for the two models.

$\frac{K}{S}/m$	Full Model						Two-factor Reduced-form Benchmark									
	1	3	6	12	24	60	1	3	6	12	24	60				
A. Mean Pricing Error in Implied Volatility Points																
80	-0.61	0.51	0.19	0.04	-0.09	-0.22	-0.27	-0.21	0.44	0.69	0.14	-0.03	0.12	0.18	0.22	0.29
90	0.41	0.26	0.20	0.14	0.05	-0.01	-0.01	0.07	-0.52	-0.21	-0.15	-0.05	0.06	0.10	0.15	0.23
100	-0.20	0.07	0.15	0.10	0.03	0.03	0.09	0.22	-0.36	-0.19	-0.06	-0.05	-0.06	-0.04	0.01	0.11
110	-0.53	-0.01	0.10	0.02	-0.06	-0.01	0.11	0.27	2.64	0.24	0.10	-0.03	-0.22	-0.21	-0.15	-0.04
120	-1.27	0.34	0.35	0.14	-0.11	-0.07	0.06	0.25	3.35	0.12	-0.19	-0.19	-0.36	-0.40	-0.34	-0.23
Average								0.01								0.13
B. Mean Absolute Error in Implied Volatility Points																
80	1.63	0.78	0.78	0.67	0.49	0.48	0.67	0.88	2.76	1.49	1.15	0.99	0.79	0.70	0.79	1.02
90	0.88	0.52	0.52	0.50	0.33	0.20	0.37	0.63	1.78	0.89	0.70	0.62	0.51	0.47	0.63	0.88
100	0.76	0.29	0.38	0.44	0.32	0.20	0.31	0.57	1.10	0.63	0.75	0.68	0.40	0.33	0.51	0.76
110	1.01	0.54	0.58	0.57	0.47	0.35	0.36	0.58	3.42	1.04	1.18	1.07	0.59	0.40	0.42	0.65
120	2.34	0.94	0.89	0.81	0.61	0.48	0.44	0.60	5.12	1.79	1.27	1.21	0.84	0.61	0.49	0.57
Average								0.63								1.05
C. Mean Absolute Percentage Error																
80	4.97	2.53	2.79	2.64	2.03	1.98	2.82	3.79	8.15	4.76	4.09	3.80	3.14	2.82	3.24	4.16
90	3.43	2.07	2.07	2.20	1.51	0.91	1.71	2.87	6.80	3.65	2.87	2.62	2.24	2.10	2.78	3.79
100	4.32	1.66	1.94	2.19	1.67	1.00	1.48	2.68	6.13	3.72	4.19	3.54	2.04	1.73	2.46	3.51
110	8.40	3.64	3.55	3.32	2.68	1.92	1.80	2.77	26.52	6.74	7.45	6.15	3.11	2.26	2.32	3.28
120	22.21	6.71	5.98	5.15	3.73	2.84	2.40	2.96	37.31	13.38	8.79	7.23	4.55	3.47	2.92	3.22
Average																5.68
\mathcal{L}												46,604			38,583	

Table IV
Subsample Parameter Estimates

Entries report the maximum likelihood estimates of the model parameters and their standard errors (in parentheses) based on the first half of the sample from January 1997 to July 2002.

Panel A. Full Model					
(i) Leverage Effect					
	ρ	κ_{XX}	κ_{XZ}	κ_{XJ}	a_X
Coef.	4.3998	0.0000	15.4266	-0.0864	0.0003
Std. Err.	(0.0286)	(0.0000)	(0.0001)	(0.0040)	(0.0000)
(ii) Volatility Feedback					
	ρ	κ_Z	θ_Z	σ_Z	γ^v
Coef.	-0.7139	3.3060	0.0342	0.6754	-26.2424
Std. Err.	(0.0017)	(0.0001)	(0.0001)	(0.0001)	(0.0044)
(iii) Self-Exciting Market Disruptions					
	ν_J	κ_J	θ_J	σ_J	γ^J
Coef.	0.2419	0.0011	55.9159	3.6447	0.3127
Std. Err.	(0.0004)	(0.0000)	(0.3087)	(0.0001)	(0.0031)

Panel B. Two-factor Reduced-form Benchmark					
(i) Diffusion Variance Dynamics					
	ρ	κ_Z	θ_Z	σ_Z	γ^v
Coef.	-0.6606	2.0120	0.0381	0.6939	-5.9088
Std. Err.	(0.0067)	(0.0338)	(0.0005)	(0.0096)	(1.5486)
(ii) Self-Exciting Jump Intensity Dynamics					
	ν_J	κ_J	$\kappa_J \theta_J$	σ_J	γ^J
Coef.	0.0709	0.0000	1.3026	30.3586	0.7404
Std. Err.	(0.0003)	(0.0018)	(0.0358)	(0.4476)	(0.0256)

Table V
Comparing In-sample and Out-of-sample Pricing Performance

Entries report the mean absolute pricing error in implied volatility points from our model (left side) and the two-factor reduced-form benchmark (right side). The parameters of the two models are estimated using the first half of the data from January 8, 1997 to July 31, 2002. The three panels report the pricing performance over the whole sample from January 8, 1997 to March 5, 2008, the in-sample period (first half), and the out-of-sample period (second half), respectively.

$\frac{\kappa}{\sigma}$ /m	Full Model						Two-factor Reduced-form Benchmark									
	1	3	6	12	24	36	48	60	1	3	6	12	24	36	48	60
A. Full-sample Mean Absolute Pricing Error in Implied Volatility Points, January 1997 – March 2008																
80	1.94	0.82	0.84	0.74	0.59	0.54	0.66	0.93	3.61	1.26	1.07	1.00	0.93	0.92	1.03	1.19
90	1.00	0.50	0.57	0.59	0.46	0.38	0.51	0.76	1.53	0.68	0.57	0.60	0.67	0.76	0.95	1.14
100	0.92	0.43	0.43	0.53	0.47	0.41	0.50	0.72	1.20	0.76	0.74	0.53	0.58	0.67	0.91	1.11
110	0.91	0.54	0.56	0.58	0.59	0.53	0.55	0.71	2.06	0.86	0.95	0.78	0.78	0.75	0.89	1.09
120	2.42	0.95	0.79	0.71	0.71	0.66	0.64	0.74	4.89	1.65	1.16	1.09	1.02	0.98	0.99	1.10
Average																1.13
B. In-sample Mean Absolute Pricing Error in Implied Volatility Points, January 1997 – July 2002																
80	1.50	1.00	0.98	0.66	0.49	0.41	0.47	0.67	2.54	1.67	1.47	1.09	0.84	0.85	1.09	1.39
90	0.93	0.66	0.67	0.54	0.31	0.16	0.31	0.57	1.75	0.88	0.80	0.63	0.47	0.55	0.85	1.17
100	0.70	0.29	0.53	0.55	0.29	0.16	0.31	0.57	1.00	0.45	0.63	0.62	0.31	0.30	0.62	0.95
110	0.96	0.51	0.63	0.66	0.37	0.21	0.29	0.54	2.85	1.14	1.15	1.06	0.59	0.28	0.40	0.73
120	1.44	0.72	0.69	0.75	0.49	0.28	0.27	0.48	4.02	1.42	1.09	1.30	0.92	0.54	0.39	0.56
Average																1.03
C. Out-of-sample Mean Absolute Pricing Error in Implied Volatility Points, August 2002 – March 2008																
80	2.39	0.63	0.70	0.82	0.69	0.67	0.85	1.19	4.67	0.84	0.68	0.90	1.02	1.00	0.97	0.99
90	1.06	0.34	0.46	0.64	0.62	0.60	0.70	0.96	1.32	0.48	0.34	0.57	0.87	0.97	1.05	1.12
100	1.15	0.57	0.32	0.52	0.66	0.65	0.69	0.87	1.39	1.07	0.85	0.44	0.85	1.04	1.20	1.28
110	0.87	0.57	0.50	0.51	0.81	0.86	0.82	0.88	1.28	0.58	0.76	0.51	0.96	1.22	1.38	1.45
120	3.41	1.18	0.89	0.68	0.93	1.04	1.01	0.99	5.75	1.87	1.24	0.89	1.11	1.41	1.60	1.63
Average																1.24

Table VI
Model Parameters Estimated Using Exchange-Traded Options

Entries report the maximum likelihood estimates of the model parameters and their standard errors (in parentheses). The parameters are classified into three panels, each describing the dynamics on one source of volatility risk. The parameters are estimated using the S&P 500 index options quotes from the Chicago Board of Options Exchange.

Panel A. Leverage Effect					
	p	κ_{XX}	κ_{XZ}	κ_{XJ}	a_X
Coef.	2.6336	0.0002	18.7207	-0.0832	0.0083
Std. Err.	(0.0013)	(0.0000)	(0.4735)	(0.0000)	(0.0173)
Panel B. Volatility Feedback					
	ρ	κ_Z	θ_Z	σ_Z	γ^V
Coef.	-0.8009	3.4349	0.0227	0.5952	-17.7505
Std. Err.	(0.0006)	(0.0020)	(0.0000)	(0.0006)	(0.1340)
Panel C. Self-Exciting Market Disruptions					
	ν_J	κ_J	θ_J	σ_J	γ^J
Coef.	0.2039	0.0014	89.3706	6.4851	0.4270
Std. Err.	(0.0002)	(0.0001)	(3.4729)	(0.0238)	(0.0356)

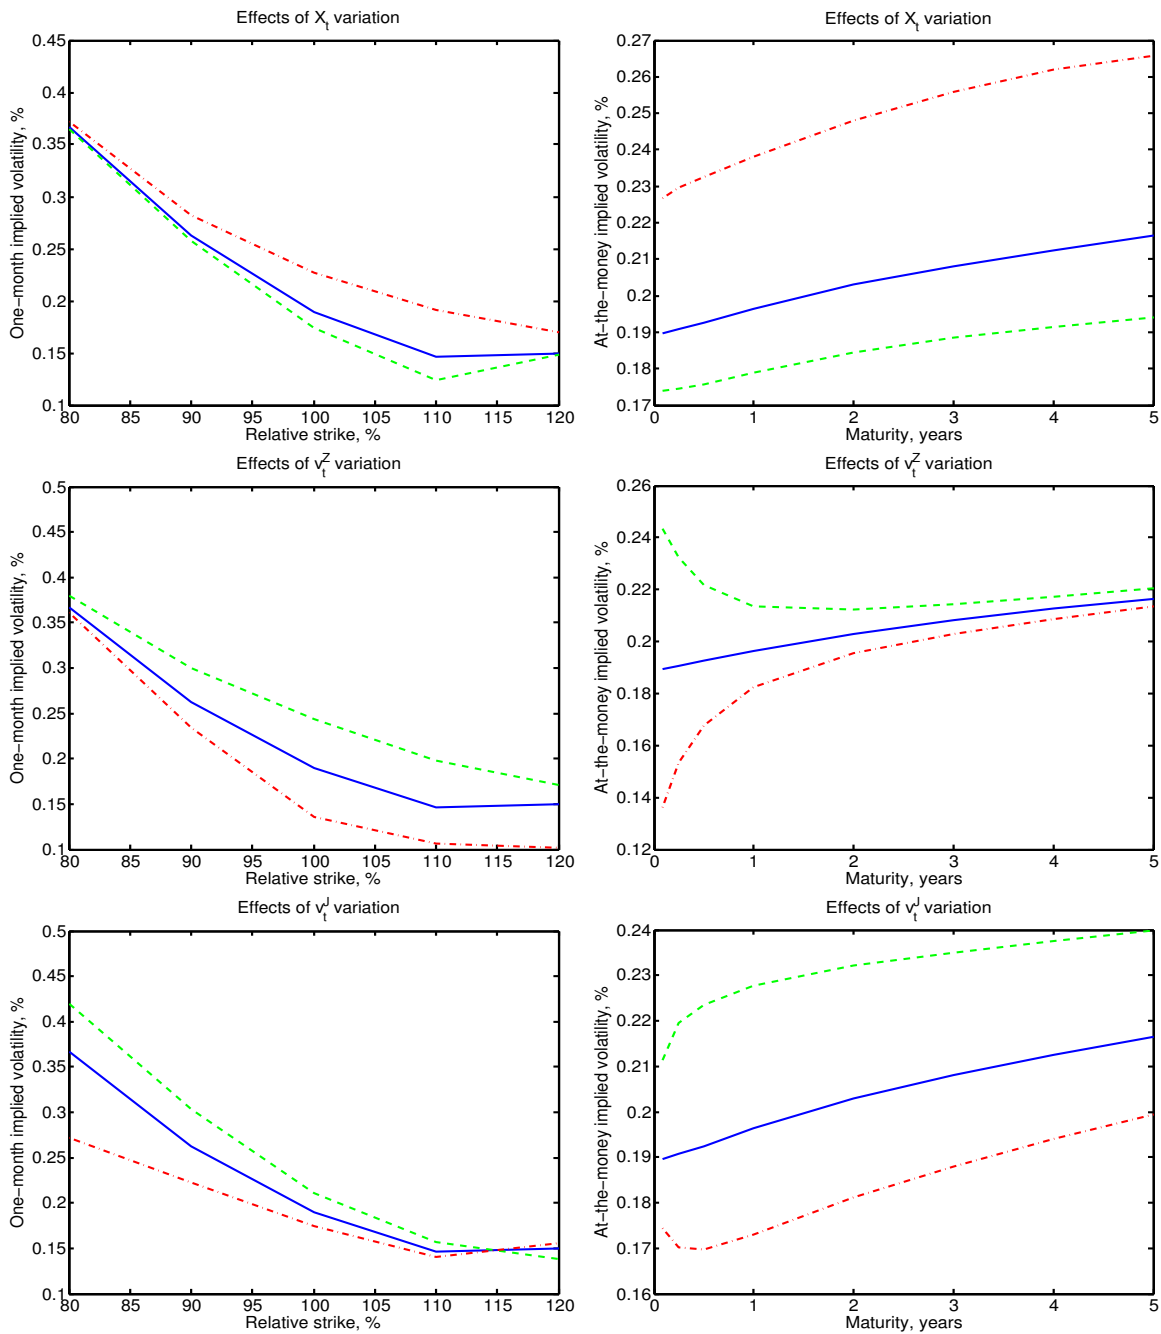


Figure 1. Shocks and implied volatility responses.

The solid lines in each panel represent the implied volatility generated from the estimated model when evaluated at the sample average of the state variables. The dashed lines are obtained by setting one state variable to its 90th percentile while holding the other two to their average values. The dashed-dotted lines are obtained by setting one state variable to its 10th percentile while holding the other two to their average. The three rows represent implied volatility responses to the three state variables. The left columns plot the responses of the one-month implied volatility skew, whereas the right columns plot the responses of the at-the-money implied volatility term structure.

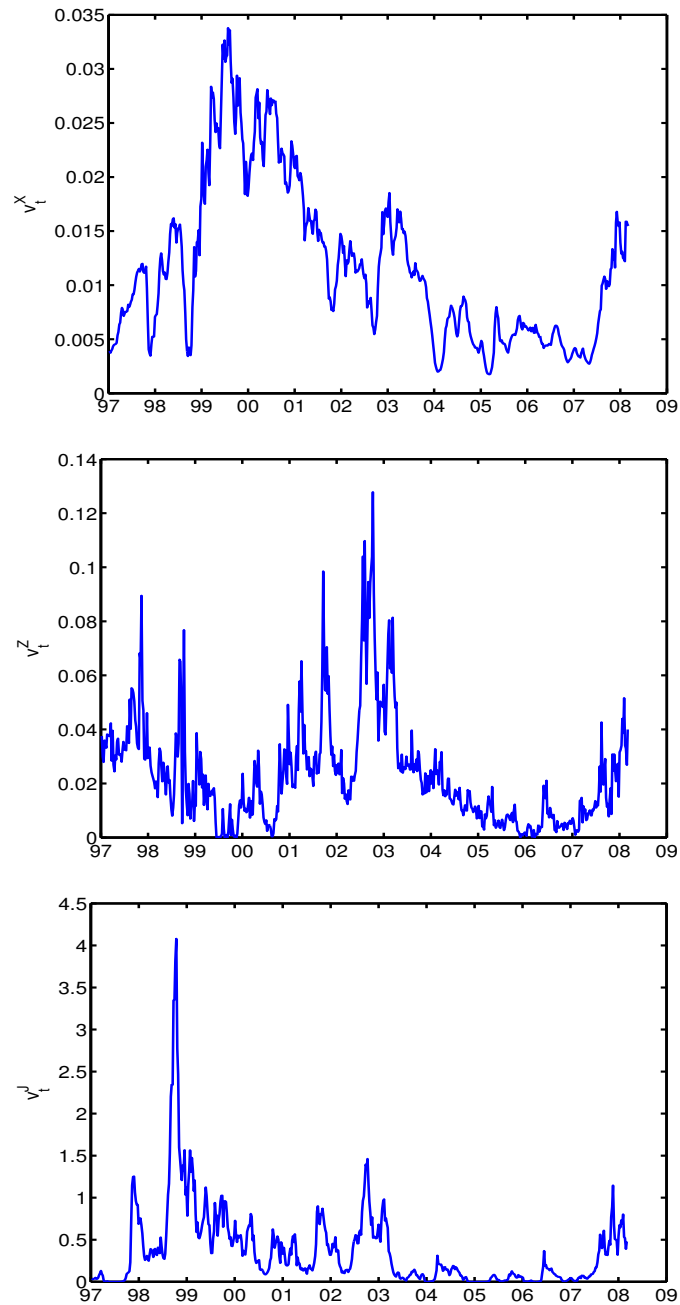


Figure 2. The time series of the three sources of return variance.

The time series of the three variance rates (v_t^X, v_t^Z, v_t^J) are extracted from the observed index option implied volatility quotes using the unscented Kalman filter under the estimated model parameters. Each panel is for one variance rate.

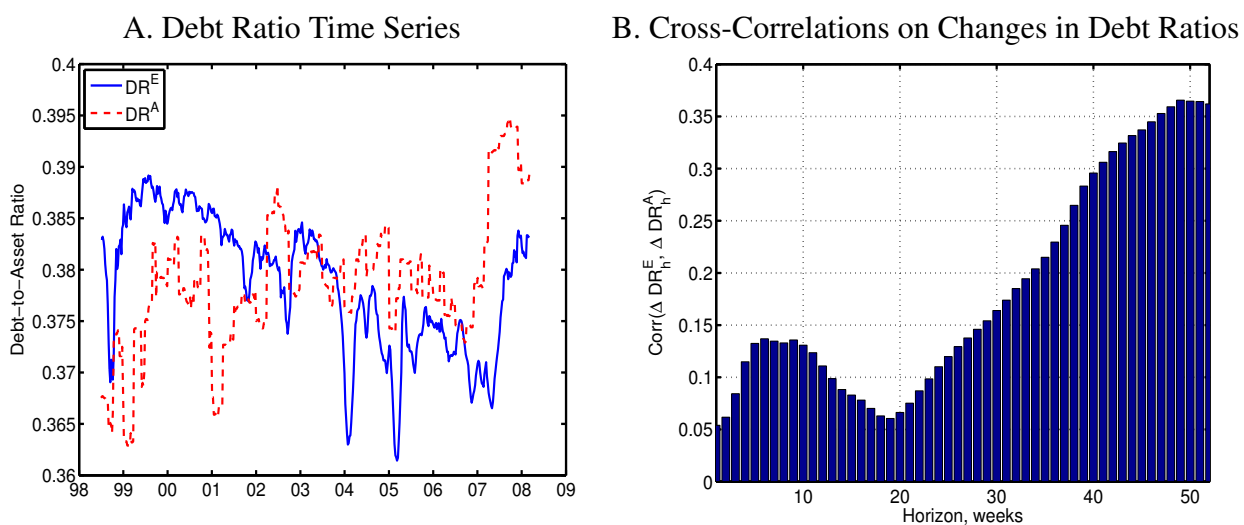


Figure 3. Debt-to-asset ratios extracted from equity index options and accounting reports. Panel A plots the time series of two debt-to-asset ratios, with the solid line denoting the series extracted from equity index options (DR^E) based on our model specification and parameter estimates and the dashed line denoting the corresponding estimates compiled from accounting reports (DR^A). The lines are scaled to have the same sample mean and standard deviation. The bars in panel B represent the cross-correlation estimates between changes in the two sets of debt ratios over different horizons.

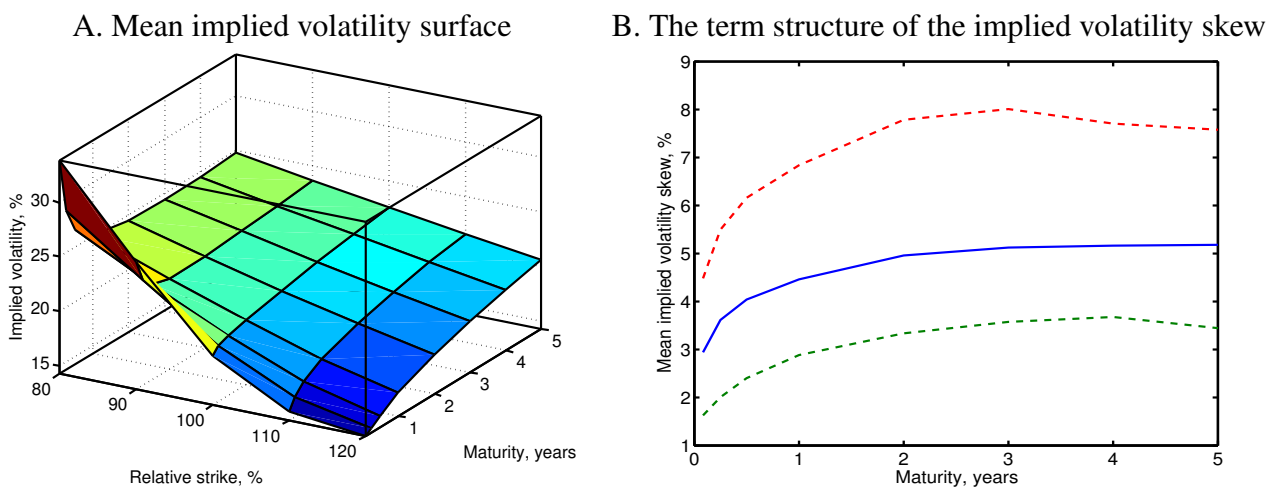


Figure 4. The mean implied volatility surface and implied volatility skew.

Panel A plots the sample averages of the implied volatility surface as a function of the relative strike and time to maturity. Panel B plots the sample averages (solid line) and the 10th and 90th percentiles (dashed lines) of the implied volatility skew estimates across different maturities. The skew is defined as the difference between 80% and 120% strike implied volatilities divided by the absolute difference in the corresponding standardized moneyness measure $d \equiv \ln(K/S)/(IV\sqrt{\tau})$, where IV is the implied volatility at the relative strike K/S and time to maturity τ .

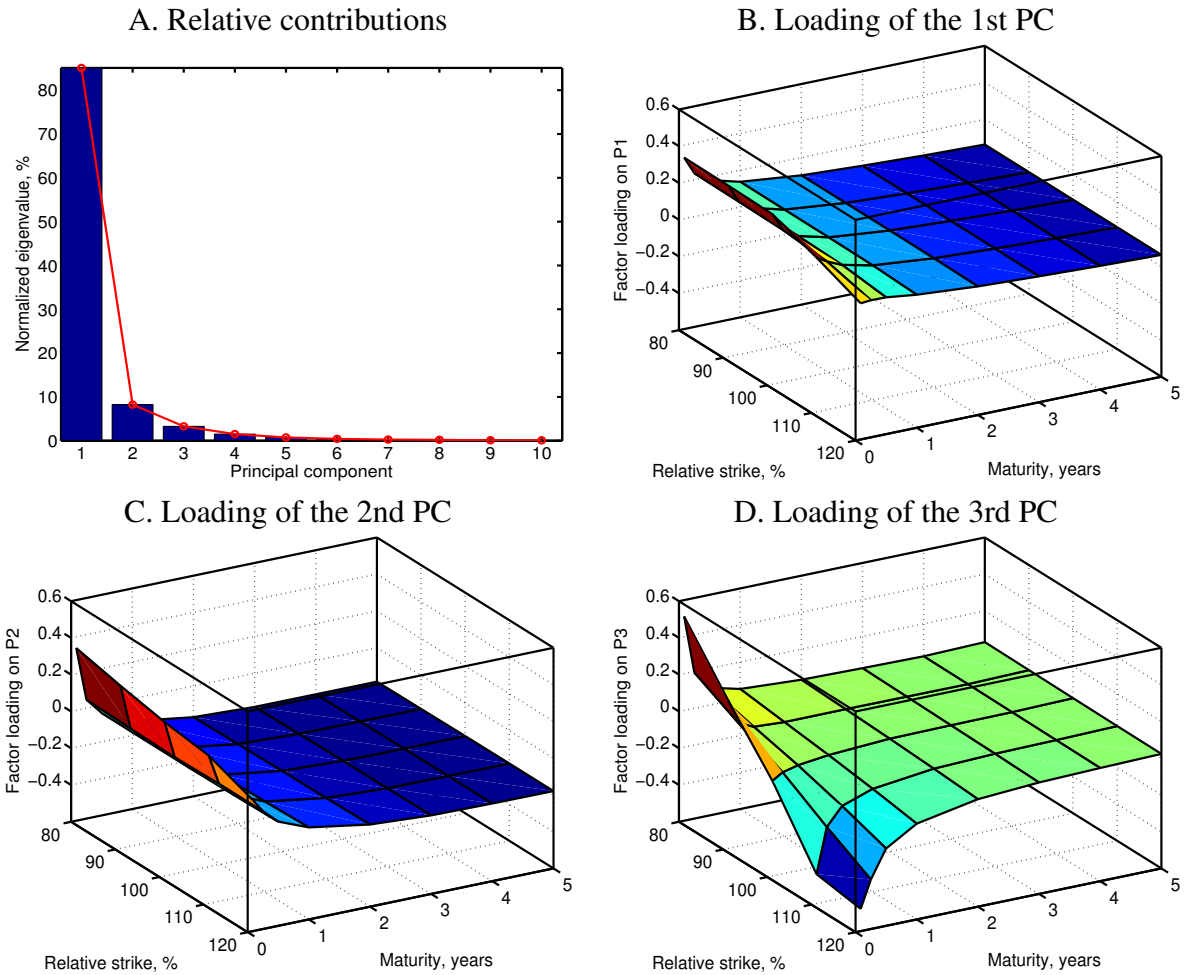


Figure 5. Principal component analysis on the implied volatility surface.

The bar chart in Panel A plots the first ten normalized eigenvalues of the covariance matrix of weekly changes on the 40 implied volatility series. They can be interpreted as the percentage variation explained by each principal component. Panels B to C plot the eigenvectors of the first three principal components, respectively. They can be interpreted as the loading coefficients of the three principal components on the 40 implied volatility series.

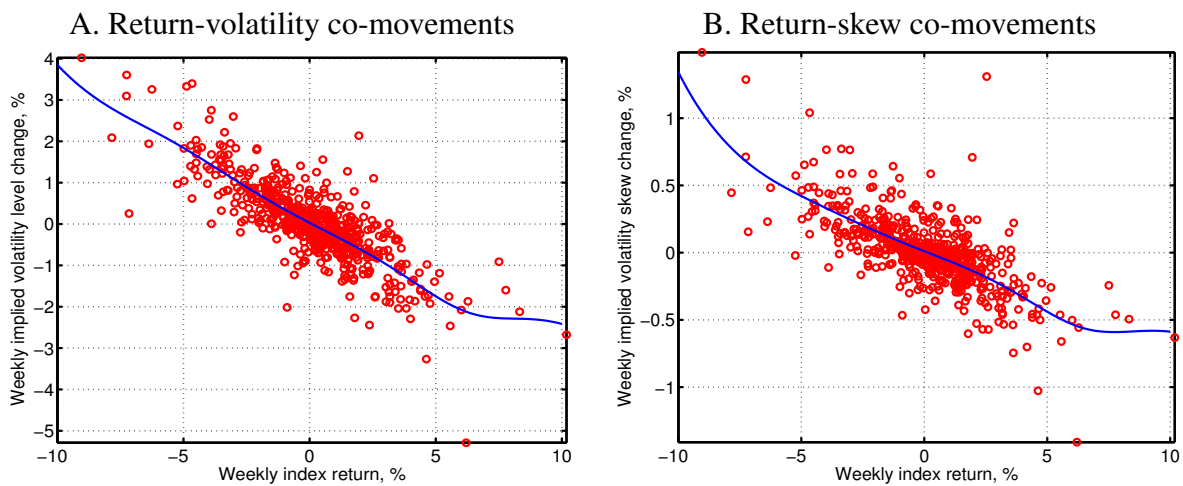


Figure 6. Asymmetric co-movements between index returns and the implied volatilities.

The two panels plot weekly changes in the average at-the-money implied volatility level (panel A) and the average implied volatility skew (panel B) against weekly index returns. Circles denote data points and solid lines represent local linear regression fits with a Gaussian kernel and a default bandwidth choice.

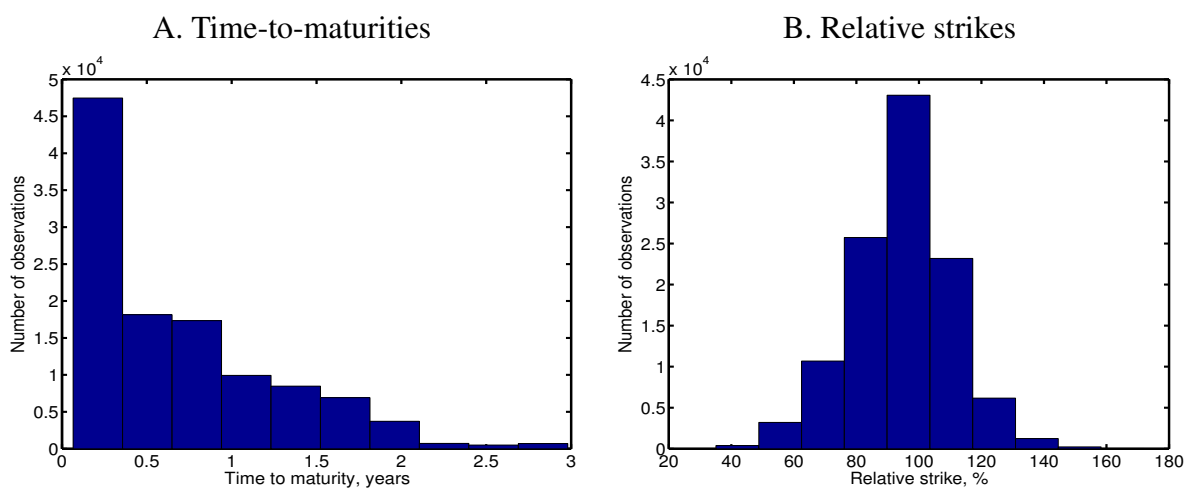


Figure 7. Histograms on the time-to-maturities and relative strikes on listed options.
 The two panels plot the histograms of the time to maturities (panel A) and relative strikes (in percentage of forward, panel B) of the listed options sample.