Time-varying Crash Risk: The Role of Market Liquidity*

Peter Christoffersen Bruno Feunou Yoontae Jeon
Chayawat Ornthanalai

January 2016

Abstract

We study the impact of market liquidity risk on the stock market index by estimating a continuous-time model with time-varying volatility and crash risks. We find that market illiquidity dominates other factors in explaining time-varying market crash risk; it explains 61% of jumps in the S&P 500 index. While we find that the crash probability significantly varies through time, its dynamic depends only weakly on return variance once we include market illiquidity as an economic variable in the model. This finding suggests the relationship between variance and crash probability found in the literature is largely due to their common exposure to market liquidity risk. Our study highlights the importance of market-trading friction in index return dynamics and explains why prior studies find that crash risk increases with market uncertainty level.

JEL Classification: G01, G12

Keywords: Market liquidity; Crash risk; Jump intensity; Options; Filtering

*Preliminary version, please do not cite without permission. We thank the Bank of Canada, SSHRC, and the Global Risk Institute for their financial support. Peter Christoffersen, Yoontae Jeon and Chayawat Ornthanalai are affiliated with the Rotman School of Management, University of Toronto. Bruno Feunou is affiliated with the Bank of Canada. Please address your correspondence to Chayawat Ornthanalai, chay.ornthanalai@rotman.utoronto.ca, 105 St. George Street, Toronto, ON, Canada, M5S 3E6, (416) 946-0669.

1 Introduction

Market liquidity, defined as the ease with which securities can be bought or sold without significant price impact, has become an increasing concern in financial markets. This is evidenced, for example, by the "flash crash" of May 2010, when major US stock indices fell by almost 10%, before recovering quickly. Similarly, market-wide trading halts on August 24th, 2015 generated spikes in asset price volatility across financial markets. These two incidents were quickly identified as symptoms of market illiquidity because they occurred in the absence of major news about fundamentals. Unlike the funding liquidity squeeze witnessed in 2007–2008, current market liquidity risk stems *not* from the banking industry, but perhaps from its absence. In an effort to diminish the chances of reliving the 2008 crisis, regulators and politicians have been working to reduce the role of banks in financial markets, thereby lowering the amount of securities held on bank balance sheets. While this may limit the chances of a subprime crisis repeat, it has the potential to cause investors to increasingly bear the risk borne from financial markets' trading frictions, e.g, market liquidity risk.[1] As a result, the influence of market liquidity on the economy appears to be increasing in importance.

This paper examines the impact of market liquidity risk on the volatility and crash probability of the aggregate stock market — proxied by the S&P 500 index. Our approach is to estimate a continuous-time model with stochastic volatility and dynamic crash probability. The innovation of our method is the introduction of market liquidity risk as an economic factor driving the dynamics of volatility and jump intensity. We measure daily market liquidity risk (or "market illiquidity") using average bid-ask spreads of securities constituting the S&P 500 index, where individual stocks' effective bid-ask spreads are estimated from high-frequency trades.[2] We estimate the model over 2004–2012 using daily S&P 500 index options, realized spot variance, and market illiquidity, and find that 61% of the time-varying crash risk is due to the stock market's exposure to market illiquidity. The influence of market illiquidity dominates other factors including the market's spot variance. During the 2008 crisis, the influence of spot variance dominates and the contribution of market illiquidity falls to about 30%.

Market crashes refer to large, unexpected drops in asset prices. Crashes can occur in the presence of information asymmetry about fundamentals, as well in their absence. In the latter case, market liquidity risk is often the culprit. For instance, Huang and Wang (2008) show in an equilibrium framework that when market participation is costly, potential traders are deterred from being in the market constantly. This causes them to enter the market only when

[1] Chung and Chuwonganant (2014) find that regulatory changes in the US markets have increased the role of public traders in liquidity provision, which has strengthened the relationship between volatility and market liquidity.

[2] This measure is motivated by Ait-Sahalia and Yu (2009), and Goyenko, Holden, and Trzcinka (2009) who find strong empirical supports for using intraday bid-ask as the measure for market illiquidity.

large trading-needs arise, which are often on the selling side.[3] Although there exists empirical evidence suggesting that crashes are often driven by market illiquidity, they are typically anecdotal (e.g. “flash crash”) or indirect. For instance, Lee (2012) and Bradley et al. (2014) find that up to 70% of jumps in equity prices cannot be explained by salient news arrivals, suggesting that trading frictions must play an economically large role in causing stock price jumps.[4]

There also exists an extensive literature on index return models which unanimously agrees that index prices “jump”.[5] In this case, crashes are large and rare jumps in index returns that cannot be captured by the index's volatility level. More recently, several studies have advocated that the probability of observing jumps (or “crashes”) is time-varying.[6] A typical approach is to let the jump arrival rate increase with the level of a stock return variance.[7] Although this modeling framework is parsimonious, it is inconsistent with the notion that crashes are sudden price drops unexplainable by the current volatility level. Therefore, while recent studies in this literature agree that crash risk is time-varying, they are silent on the economic variables driving its dynamic. Our study hopes to contribute by providing economic underpinnings to models with time-varying crash risk, and showing that much of the variations in jump intensity is driven by trading frictions, i.e. illiquidity.

In order to motivate our subsequent modeling framework, we first apply logistic regression analysis linking our market illiquidity measure to a non-parametrically estimated jump probability (e.g., Huang and Tauchen, 2005). We find that market illiquidity significantly increases the ex-post probability of observing jumps in the next day S&P 500 return, and that its effect crowds out the influence of realized variance on the jump probability. We confirm this finding by running predictive OLS regressions on daily realized skewness and find a negative and significant relationship between market illiquidity and next day's realized skewness. Armed with this evidence, we estimate a continuous-time model similar to the stochastic volatility with jump model (SVJ) studied by Pan (2002) and Bates (2006), among others. In this model, the jump arrival rate is affine in return variance. We extend this framework in two aspects. First, we let market illiquidity enter into the dynamic of return variance. Second, we let the time-varying jump intensity dynamic be a function of return variance, market illiquidity, and a latent-state variable. We estimate the model by extracting information embedded in

[3]In an earlier study, Gennotte and Leland (1990) develop a rational expectation model explaining why a large price drop can occur when there is relatively small amount of selling in the market.

[4]See also Jiang, Lo and Verdelhan (2011) who study jumps in the Treasury market.

[5]This literature is too large to cite in full, for some evidence, see Maheu and McCurdy (2004), Andersen, Bollerslev and Lund (2002), Eraker (2004), and Broadie, Chernov and Johannes (2009).

[6]See for examples, Pan (2002), Eraker (2004), Bates (2006, 2012), Christoffersen, Jacobs and Ornthanalai (2012), Ornthanalai (2014), and Andersen, Fusari and Todorov (2015).

[7]Santa-Clara and Yan (2010) is a notable exception for which they model jump intensity as a quadratic function of state variables

index options and intraday trades. We apply the unscented Kalman filter to extract daily latent-state variables in the model. This filtering method allows for sequential learning in the dynamics of latent jump intensity, variance, and illiquidity processes, as well as providing errors in the measurement equation for constructing the model's log-likelihood.

We refer to the most general model that we study as the SJVI. In this model, the jump intensity dynamic, λ_t, is stochastic and affine in the spot variance (V_t), market illiquidity level (L_t), and latent-state variable (Ψ_t). For comparisons, we estimate two other benchmark models with stochastic jump intensity — the SJ and SJV models. The jump intensity dynamic in the SJ model is solely driven by the latent-state variable, while in the SJV model it depends on the spot variance and the latent-state variable. To summarize, the three models that we estimate are characterized by their jump intensity dynamics as follows:

$$\begin{aligned} \text{SJ} \ : \quad & \lambda_t = \Psi_t \\ \text{SJV} \ : \quad & \lambda_t = \Psi_t + \gamma_V V_t \\ \text{SJVI} : \quad & \lambda_t = \Psi_t + \gamma_V V_t + \gamma_L L_t. \end{aligned}$$

In all specifications, we model the spot variance, V_t, as a two-factor square-root process (Heston, 1993) with market illiquidity being one of the factors. Our estimation results show a strong contemporaneous relationship between market illiquidity and spot variance. On average, a one-standard-deviation increase in the level of market illiquidity increases the spot variance by about 12%. This finding lends support to previous studies documenting the positive relationship between return volatility and trading activity (e.g. Lamoureux and Lastrapes, 1990; Chae, 2005).

We find that the nature of jumps that we estimate from the three models reflects stock market crash risk. When a jump occurs, its average size is about -3.7% in return units. Therefore, jumps that we estimate represent large drops in the index price, "crashes," and not market surges. We find that jump intensity dynamics in the SJ model is extremely volatile. The average jump probability of this model is 3.1 jumps per year with an annualized standard deviation of 12.6. The jump component in the SJ model significantly dominates the diffusive variance component in the index return dynamic, particularly during the 2008 crisis where the ex-ante jump probability rises to 60 jumps per year. On the other hand, when we include the spot variance as a variable in the time-varying jump probability, e.g., the SJV and SJVI models, the jump intensity dynamic becomes significantly smoother and less volatile. For instance, the average jump probability levels for the SJV and SJVI models, respectively, are 2.5 and 2.9 per year with annualized standard deviations of 3.3 and 2.2.

Looking at the models' in-sample fit, we find the log-likelihood value from the SJ model is significantly lower (by about 10%) than those from the SJV and SJVI models. Further, the filtered spot variances from the SJ model are significantly downward-biased relative to the realized variance levels calculated from high-frequency data because variations in index returns are dominated by jumps. We believe that the relatively lower log-likelihood value and the volatile nature of jumps observed in the SJ model is due to the difficulty of identifying the latent stochastic jump intensity dynamic in the absence of economic covariates. Estimation of the time-varying jump intensity dynamic is generally difficult, and we refer to Bates (2006) for a brief literature review. This is because jumps, and particularly crashes, are rare events. Therefore, the estimation of time-varying jump dynamics requires that econometricians extract their information from various information-rich sources. Our findings support this view by showing the importance of modeling jump intensity as a function of economic covariates that can be reliably identified from the data, e.g., realized variance or market illiquidity.

We find strong evidence that during our sample period, crash risk in the S&P 500 index mostly reflects investors' fear of market illiquidity. We arrive at this conclusion by examining the contribution of market illiquidity to the jump intensity dynamic in the SJVI model and find a strong contemporaneous positive relationship (t-stat of 13.87). On the other hand, estimation results show the relationship between jump probability and spot variance is positive but statistically weak (t-stat of 1.61). This finding differs from our estimates for the SJV model where market illiquidity is absent in the jump intensity dynamic for which the spot variance level significantly increases with jump intensity (t-stat of 2.17). Collectively, these results show that market illiquidity is the main economic factor driving crash risk and not the level of market's spot variance (e.g., Christoffersen, Jacobs and Ornthanalai, 2012) and support our preliminary evidence found using regression analyses. Further, our results suggest that previous studies find that jump intensity increases with the level of spot variance because of the strong positive relationship between variance and market illiquidity.

In terms of economic magnitude, we find that market illiquidity explains more than half of the S&P 500 index's crash probability during our sample (about 61% on average). On the contrary, the market spot variance's contribution to the jump intensity dynamic is only about 15%, with the remaining 24% coming from the latent jump-intensity-specific factor. However, during the six-month period after the Lehman Brothers' collapse, we find that the market spot variance dominates other factors in explaining the time-varying crash probability, with the contribution as high as 70%. This finding suggests that investors' fear of crash risk during the sub-prime crisis reflects uncertainty about the market's fundamentals, while outside the crisis period, crash risk mostly reflects investors' fear of market illiquidity.

We emphasize that our findings on the relationship between market illiquidity and time-varying volatility and crash risks are not due to market micro-structure noise. This is because the market illiquidity proxy that we use is derived from effective spreads of 500 firms constituting the S&P 500 index and *not* from trades on its ETFs nor its futures contracts. Therefore, the relationship between market illiquidity and index return dynamics that we document is not mechanically generated from market-microstructure noises.

Overall, the findings in this paper emphasize the importance of market liquidity risk in explaining time-varying volatility and crash risks, which is largely missing from prior empirical studies examining index return dynamics. We confirm that our main conclusions hold using various robustness checks. For instance, we show that our estimation results are qualitatively similar before and after the implementation of the "circuit breaker" in 2010. We also re-estimate the models using a different market illiquidity measure besides the effective bid-asks spreads, e.g., Amihud's (2002) measure, and obtain the same conclusions.

The remaining parts of this paper proceed as follows. Section 2 describes the data, sample selection, and reports preliminary evidence. Section 3 describes the model and estimation procedure. Section 4 discusses estimation results and interpretations of our findings. Section 5 demonstrates the robustness of our findings. Finally, Section 6 concludes.

2 Data and Preliminary Evidence

The sample period that we study is from January 1, 2004 through December 31, 2012. We focus on the recent period because the global financial market has gone through a drastic transformation, e.g., new banking regulations, proliferation of algorithmic trading and exchanged-traded funds. Such recent changes has strengthened the relationship between market liquidity and stock market volatility as documented in Chung and Chuwonganant (2014).

The remaining parts of this section describes the construction of main variables that we use and report preliminary evidence found using regression analyses.

2.1 Market Illiquidity

We construct a time-series measure of market liquidity risk at the daily level. Although different illiquidity measures have been proposed in the literature, they do not always capture the same type of market frictions. In this paper, we focus on the trading friction associated with the cost of participating in the stock market. We measure it using effective bid-ask spreads following Goyenko, Holden and Trzcinka (2009) who find strong empirical supports for using intraday bid-ask spreads as the measure of market illiquidity.

We obtain obtain all transactions recorded on securities constituting the S&P 500 index from the TAQ database. Then, for each stock i on day t, we calculate the effective spread of its k^{th} trade as

$$ILQ^i_{t,k} = \frac{2|S^{i,P}_{t,k} - S^{i,M}_{t,k}|}{S^{i,M}_{t,k}}, \qquad (1)$$

where $S_{t,k}^{i,P}$ is the price of the k^{th} trade of stock i on day t, and $S_{t,k}^{i,M}$ is the mid-point of the best prevailing bid and ask at the time of the k^{th} trade. The daily effective spread of stock i on day t is then computed as the dollar-volume weighted average effective spreads over all trades happened during the day

$$ILQ_t^i = \frac{\sum_{k=1}^{K} DolVol_{t,k}^i IL_{t,k}^i}{\sum_{k=1}^{K} DolVol_{t,k}^i}, \tag{2}$$

where $DolVol_{t,k}^i$ is the dollar trading volume of the k^{th} trade. Lastly, we aggregate the effective spreads of firms constituting the S&P 500 index on each day by equally weighting their daily illiquidity measures. This procedure results in a daily market illiquidity measure for the aggregate stock market on day t:

$$ILQ_t = \frac{1}{N} \sum_{i=1}^{N} IL_t^i. \tag{3}$$

We compute the daily market illiquidity measure from January 2, 2004 to December 31, 2012. This results in 2,262 observation-days. The bottom panel of Figure 1 plots the time series of market illiquidity. We see the market illiquidity measure rises significantly during the financial crisis period, but stays relatively stable during other periods, with an occasional few spikes. We also see a sharp spike on May 6, 2010, which is associated with the "flash crash" incident. All numbers reported in 1 are annualized thus 20% of market illiquidity translates to about 0.08% trading cost at the daily level.

2.2 Realized Return Moments

The second set of data we construct is the daily realized variance measure calculated using intraday S&P 500 cash index returns obtained from TickData. Using the latest observation at each minute, we construct a grid of one-minute intraday returns starting from 9:30 am and ending at 4:30 pm. This gives us 390 observations per trading day. After, the daily realized variance is calculated using the *MinRV* estimator of Andersen, Dobrev, and Scaumburg (2012) as follows

$$\text{MinRV}_t^N = \frac{\pi}{\pi - 2} (\frac{N}{N-1}) \sum_{i=1}^{N-1} \min(|r_{i,t}|, |r_{i+1,t}|)^2, \tag{4}$$

where N denotes the number of observations on each day and $r_{i,t}$ denotes the 1-minute log return at i^{th} interval of date t. As the number of interval N goes to infinity, this estimate converges to the diffusive part of the quadratic variation, thus resulting in a jump-robust estimate of the daily variance.

We calculate the daily realized variance measure, RV, constructed as the sum of squared 1-minute log returns: $RV_t^N = \sum_{i=1}^N r_{i,t}^2$. This method measures the total quadratic variation in returns that are to the diffusive component, measured by $MinRV$, and the jump component. We use this RV measure to perform ex-post daily jump detection in the next section. We construct the daily realized skewness, $RSkew$, and realized kurtosis, $RKurt$, measure following the method in Amaya, Christoffersen, Jacobs, and Vasques (2015). These realized higher moments are calculated using 1-minute log returns data as follows:

$$RSkew_t^N = \frac{\sqrt{N}\sum_{i=1}^N r_{i,t}^3}{(RV_t^N)^{3/2}} \tag{5}$$

$$RKurt_t^N = \frac{N\sum_{i=1}^N r_{i,t}^4}{(RV_t^N)^2}, \tag{6}$$

where N is the number of time intervals in a trading day. As N goes to infinity, the above two measures converge to the cubic and quadratic variations of jump component in the daily return, i.e., the diffusive component is excluded in their measurement.

We emphasize the market illiquidity proxy that we use is computed from effective spreads of 500 firms constituting S&P 500 index while all realized return moments are constructed from the 1-minute log returns of the S&P 500 cash index. Therefore, these two sets of measure are derived from transactions of securities traded under different names in the stock exchanges. This makes our subsequent analyses free from the concern that market illiquidity and realized return moments are endogenously related due to common market micro-structure noises.

2.3 Predicting Ex-post Jumps

Ex-post jump detection using intraday returns has been studied extensively in the recent literatures.[8] Following the conventional approaches, we detect jumps in daily index returns using the test statistics constructed from the difference between RV_t and $MinRV_t$. We label each day as a "jump day" if the test statistic, which asymptotically converges to a standard normal distribution, falls beyond the 99.9% critical value. Out of 2,262 sample days, we find that 538 days (23.75%) are classified as jump days. We note that this method is an ex-post detection of a jump because it informs us that a jump has occurred after we have already observed return on that day.

Using results from the daily jump detections, we run a predictive logit regression model examining to which variables successfully predict the occurrence of jumps the next day. We examine three variables of interest and their combinations, namely the market illiquidity

[8] See Huang and Tauchen (2005) for concise summary.

measure ILQ_t, the diffusive quadratic variation measure $MinRV_t$, and the realized skewness measure $RSkew_t$. The most general logit regression includes all three variables as specified below:

$$Pr(J_{t+1}) \quad = \quad \text{logit}(\beta_0 + \beta_1 MinRV_t + \beta_2 ILQ_t + \beta_3 RSkew_t), \tag{7}$$

where J_{t+1} is an indicator variable equal to one if a jump is detected on day $t+1$, and zero otherwise. All other variables are as defined in previous sections.

Table 1 summarizes the logit regression results. We report the estimated coefficients and their t-stats in parentheses. When each variable is regressed individually, it appears statistically significant in predicting the occurrence of jumps the next day. Columns (1)–(3) in Table 1 show that both $MinRV_t$ and ILQ_t are positive and statistically significant at the 99% level, while $RSkew_t$ is significant at 95% level. Also, both $MinRV_t$ and ILQ_t have positive coefficient loading where $RSkew_t$ has a negative coefficient. The positive coefficients on $MinRV_t$ and ILQ_t confirm the intuition that jumps are more likely going to occurred following on a day of more volatile and illiquid market conditions. On the other hand, the negative coefficient on $RSkew_t$ is consistent with the well-documented evidence that jump arrivals cluster in time. A jump in index return is generally on the negative side and therefore is associated with a negative skewness in the index return distribution. Overall, the negative coefficient on $RSkew_t$ in column (3) is an evidence of jumps clustering.

Columns (4)–(7) report results found on various combinations of the independent variables of interest. Strikingly, when more than one variables are included in the logit regression, we find the market illiquidity measure ILQ_t dominates in the predictive power while other variables lose their predictive ability. Importantly, we find the coefficient estimate on $MinRV_t$ turns negative in columns (4) and (7), although not statistically significant, whenever the market illiquidity proxy is included. Lastly, we find that $RSkew_t$ plays no role in predicting ex-post jump probability when all three independent variables are included.

Overall, simple logistic regression analyses indicate the importance of market illiquidity in explaining the time-varying jump probability of index returns. Further, it shows that omission of the market illiquidity measure can lead to a different conclusion on the role of spot variance, $MinRV_t$, in predicting jumps probability.

2.4 Predicting Realized Higher Moments

As an additional analysis, we estimate predictive OLS regressions on the realized higher moments. As discussed previously, the realized skewness and realized kurtosis measures are proxies for the fat-tailed characteristics of the index return distributions. However, they provide an informative way to identify a crash from a stock market surge.

Because crashes are large sudden drops in asset prices, the more negative the realized skewness signals the higher probability that a crash has occurred. On the other hand, we expect the realized kurtosis measure to increase with the probability that either a crash or a surge in the stock market has occurred. Therefore, if market illiquidity is a strong predictor of the stock market crashes, we expect that ILQ_t would negatively predict the realized skewness measure. On the other hand, if market illiquidity does not predict the stock market surges, we expect that ILQ_t would not positively predict the realized kurtosis measure. Based on this intuition, we estimate the following two regression models:

$$\begin{aligned} RSkew_{t+1} &= \alpha + \beta_1 MinRV_t + \beta_2 ILQ_t + \beta_3 RSkew_t + \epsilon^1_{t+1} \quad (8) \\ RKurt_{t+1} &= \alpha + \beta_1 MinRV_t + \beta_2 ILQ_t + \beta_3 RSkew_t + \epsilon^2_{t+1}. \quad (9) \end{aligned}$$

Table 2 reports the results. Consistent with previous findings, Panel A shows that $MinRV_t$ positively predicts the skewness while ILQ_t negatively predicts the realized skewness. A negative coefficient on ILQ_t indicates that market illiquidity positively predicts a more downward jump probability in index returns. Panel B shows the results from a predictive regression model on realized kurtosis. We find that $MinRV_t$ strongly predicts the realized kurtosis while ILQ_t now loads negatively. This finding supports our conjecture that market illiquidity does not predict the stock market surges. We further note that the reported t-stats in Table 2 are quite small because daily realized returns moments are known to be noisy variables.

Overall, evidence from Sections 2.3 and 2.4 suggest that market illiquidity is an important economic predictor forthe stock market crashes. On the other hand, the diffusive variance component does not strongly predict negative jumps in index returns.

Motivated by the above non-parametric evidence, we develop a continuous-time model that captures the importance of market illiquidity in explaining the time-varying volatility and crash risks. We discuss this in the next section.

3 Model and Estimation

3.1 The SJVI Model

We begin by specifying the processes governing the log stock price, spot variance, spot illiquidity, and latent component of jump intensity dynamic under the risk-neutral measure ($\mathbb{Q}$). We use the notations S_t and V_t to denote stock price and spot variance at time t. We let L_t represent the spot market illiquidity which measures the liquidity risk of the stock market at time t, with a higher value indicating a more illiquid market. We include a stochastic process Ψ_t that is designed to capture the latent time-varying jump intensity in index returns. Thus, the model consists of the four factors that fully describe the return dynamics under $\mathbb{Q}$:

$$\begin{aligned} d\log(S_t) &= (r - \frac{1}{2}V_t - \xi\lambda_t)dt + \sqrt{V_t}(\sqrt{1-\rho^2}dW_t^1 + \rho dW_t^2) + q_t dN_t && (10) \\ dV_t &= \kappa_V(\theta_V - V_t)dt + \gamma dL_t + \xi_V\sqrt{V_t}dW_t^2 && (11) \\ dL_t &= \kappa_L(\theta_L - L_t)dt + \xi_L\sqrt{L_t}dW_t^3 && (12) \\ d\Psi_t &= \kappa_\Psi(\theta_\Psi - \Psi_t)dt + \xi_\Psi\sqrt{\Psi_t}dW_t^4, && (13) \end{aligned}$$

where r denotes the risk-free rate and all Brownian motions dW_t^i, for $i = 1$ to 4, are independent to each other.

We assume the market illiquidity process, L_t, and the latent jump intensity process, Ψ_t, in equations (12) and (13) follow the standard square-root process with long-run mean levels of θ_L and θ_Ψ, respectively. The variance dynamic in equation (11) is almost identical to the Heston's (1993) square-root process with an exception of an additional term γdL_t. Equation (11) shows the evolution of spot variance depends on its own mean-reverting drift, the diffusive component, and the market illiquidity process L_t. The long-run mean of the spot variance, as we will later show, is equal to θ_V, and the mean-reversion speed to the long-run variance is denoted by κ_V.

The log stock price dynamic described in equation (10) follows a standard jump-diffusion process where $q_t dN_t$ denotes the discontinuous jump component. Following the extant literature on index return models, we assume that jumps follow a compound Poisson process with intensity λ_t and each individual jump is i.i.d. normal with the jump mean size θ and the jump size standard deviation δ. In order to ensure the discounted log stock price is martingale, we include the jump compensation term $\xi = e^{(\theta+\frac{\delta^2}{2})} - 1$ in equation (10). Lastly, to complete the model, we specify the dynamic of the time-varying jump intensity λ_t as follows:

$$\text{SJVI model:} \qquad \lambda_t = \Psi_t + \gamma_V V_t + \gamma_L L_t. \tag{14}$$

The above specification for jump intensity is motivated by numerical tractability and for ease of interpreting results. Equation (14) shows the time-varying jump arrival rate is determined jointly by the levels of spot variance V_t, spot market illiquidity L_t, and state variable Ψ_t. The state variable Ψ_t is designed to capture the portion of jump intensity dynamic not explained by the covariates V_t and L_t.

Our jump intensity specification is more general than those examined in prior studies which estimated a continuous-time model with time-varying jump intensity that is affine in the level of spot variance (e.g. Andersen, Benzoni and Lund, 2002; Pan, 2002; Eraker, 2004; Bates, 2006). The model stays in the class of affine jump diffusion models. We therefore have the closed-form solution to the characteristic function of log stock price using the results in

Duffie, Pan and Singleton (2000).

The variance dynamic that we consider in equation (11) also falls into the class of two-factor stochastic volatility models, which have been shown to effectively explain the term structure of index option prices.[9] Our model differs from the existing two-factor volatility literature in that we allow the expected future variance to depend on the levels of spot variance, V_t, and spot market illiquidity, L_t, as shown in the equation below:

$$E_t[V_T] = \theta_V + (V_t - \theta_V)e^{-\kappa_V(T-t)} + [(L_t - \theta_L)\frac{\gamma\kappa_L}{\kappa_V - \kappa_L}](e^{-\kappa_V(T-t)} - e^{-\kappa_L(T-t)}). \tag{15}$$

The above equation shows the current level of market illiquidity positively affects the shape of the expected term structure of variance. However, its impact dissipates as the horizon increases. This is seen from the third term on the right-hand side of equation (15), which coverges to 0 as time T goes to infinity. Without the market illiquidity term γL_t, the spot variance process reduces to the Heston's (1993) model, and the expected future variance is given by the first two terms on the right-hand side of equation (15).

For the remaining parts of this paper, we refer to the general model that we introduced as the stochastic jump with variance and illiquidity (SJVI) model.

3.2 Benchmark Models

We consider two nested specifications of the SJVI model. In the first specification, we shut off the influence of the illiquidity channel in the time-varying jump intensity dynamic, i.e. by setting $\gamma_L = 0$ in equation (14). As the result, the probability of observing jumps depends on the level of spot variance and the latent state component as follows

$$\text{SJV model:} \qquad \lambda_t = \Psi_t + \gamma_V V_t. \tag{16}$$

We refer to the model with jump intensity specification described in equation (16) as the stochastic jump intensity with variance (SJV) model. This functional form of jump intensity specification is nests the affine jump intensity dynamic, $\lambda_t = \gamma_0 + \gamma_V V_t$, that is commonly adopted in the time-varying jump studies (e.g. Pan, 2002; Bates, 2006; Ornthanalai, 2014). Equation (16) shows that when we let Ψ_t be a constant, the jump intensity dynamic becomes affine in the spot variance.

The second nested specification we study shuts off the impact of both market illiquidity as and the spot variance from influencing jump probability. That is, we set γ_V and γ_L equal

[9] See for examples, Christoffersen et al. (2008), Christoffersen, Heston and Jacobs (2009), Egloff et al. (2010), Bates (2012), and Andersen, Fusari and Todorov (2015).

to zero in equation (14). This yields

$$\text{SJV model:} \qquad \lambda_t = \Psi_t, \tag{17}$$

which is simply the latent stochastic process Ψ_t. We refer to the model with the jump intensity specification in equation (17) as the stochastic jump intensity model (SJ).

Besides the jump intensity specification, we keep all other aspects of the three models that we study identical. This approach allows us to focus solely on the role of market illiquidity and spot variance in determining the time-varying jump risk.

3.3 Filtering

As in all continuous-time stochastic volatility models, the model that we study features unobserved state variables to be filtered. All the three models we study contain three latent state variables: $V_t, L_t.$ and Ψ_t. This section describes the filtering method that we use.

We extract the latent state variables using the square-root Unscented Kalman Filter (UKF) of Van der Merwe and Wan (2001). We apply the UKF method because the observed data, including option prices, that we fit the models to are highly non-linear in the state variables. The UKF method has been shown to perform well for solving non-linear filtering and has been applied widely in the finance literature.[10] We refer to Christoffersen et al. (2012) for technical details and comparison between different filtering methods.

As the state variables in the filtering equations evolve under the physical probability ($\mathbb{P}$) measure, we need to define their under the physical measure. We do not impose any risk premiums on the L_t and Ψ_t processes for simplicity and also because the literature has not yet provided a clear guidance on how to model their risk premiums. As a result, there is no change to these two processes from the $\mathbb{Q}$ to $\mathbb{P}$ measures. We apply the commonly used functional form of the variance price of risk to the spot variance process, which given by $\nu_V\sqrt{V_t}$ as in Heston (1993). This price of risk specification shifts the Brownian shock in equation (11) by $dW_t^{2,\mathbb{P}} = dW_t^2 - \nu_V\sqrt{V_t}dt$, where superscript $\mathbb{P}$ denotes that it is evaluated under the physical probability measure. Applying this transformation, the resulting variance process under the $\mathbb{P}$-measure can be written as

$$dV_t = \kappa_V^{\mathbb{P}}(\theta_V^{\mathbb{P}} - V_t)dt + \gamma dL_t + \xi_V\sqrt{V_t}dW_t^{2,\mathbb{P}}, \tag{18}$$

where we have the following parameter mappings $\kappa_V^{\mathbb{P}} = \kappa_V - \nu_V\xi_V$ and $\theta_V^{\mathbb{P}} = \theta_V\kappa_V/\kappa_V^{\mathbb{P}}$.

We discretize the $\mathbb{P}$-measure state dynamics using the conventional Euler scheme at the

[10]For recent papers using UKF as the filtering method, see Bakshi, Carr, and Wu (2008) and Filipović, Gourier, and Mancini (2015)

daily interval. The evolution of the full state-space system in the discretized form can be written as follows

$$V_{t+1} = V_t + \kappa_V^{\mathbb{P}}(\theta_V^{\mathbb{P}} - V_t)\Delta t + \gamma\kappa_L(\theta_L - L_t)\Delta t + \xi_V\sqrt{V_t}\epsilon_{t+1}^1 + \gamma\xi_L\sqrt{L_t}\epsilon_{t+1}^2 \tag{19}$$

$$L_{t+1} = L_t + \kappa_L(\theta_L - L_t)\Delta t + \xi_L\sqrt{L_t}\epsilon_{t+1}^2 \tag{20}$$

$$\Psi_{t+1} = \Psi_t + \kappa_\Psi(\theta_\Psi - \Psi_t)\Delta t + \xi_\Psi\sqrt{\Psi_t}\epsilon_{t+1}^3, \tag{21}$$

where all error terms ϵ_{t+}^i, for $i = 1$ to 3, are i.i.d. standard normal. In the above state-space system, we set the time step $\Delta t = 1/252$ to reflect the daily discretization interval. In order to keep our notations to minimum, we apply the superscript $\mathbb{P}$ only to parameters under the physical measure that differ in values from their corresponding risk-neutral parameters.

We next describe the functional relationships linking the latent state variables to the observed data used in the estimation. The first observable is the illiquidity measure denoted by ILQ_t, which we introduced earlier in Section 2. We recall that it is is calculated as the daily aggregate effective spreads of S&P 500 constituents, i.e., weighted sum of relative bid-ask spreads of all trades occurred during the day. The other observables that we filter the state variables from are daily at-the-money (ATM) and out-of-the-money (OTM) S&P 500 index options. These three sets of observables are used in the measurement equations in the UKF procedure. We write the system of measure equations as follows

$$\log(\mathrm{ILQ}_{t+1}) = \log(E_t[\int_t^{t+1} L_s ds]) + u_{t+1}^1 \tag{22}$$

$$ATM_{t+1}^O = ATM_{t+1}^M(V_{t+1}, L_{t+1}, \Psi_{t+1}) + u_{t+1}^2 \tag{23}$$

$$OTM_{t+1}^O = OTM_{t+1}^M(V_{t+1}, L_{t+1}, \Psi_{t+1}) + u_{t+1}^3, \tag{24}$$

where measurement errors u_{t+1}^i, for $i = 1$ to 3, are independent normal random variables with constant variances. The above filtering equations are applied to all trading days from January 2, 2004 to December 31, 2012 which results in 2,262 days of observations.

The latent spot illiquidity process in the state-space dynamic describes the instantaneous level of illiquidity at each moment and not at the aggregated daily level. In order to filter L_t from the daily observed measure of market illiquidity, we integrate the spot illiquidity process over the day as shown in equation (22). The spot illiquidity measure follows a square-root process. Integrating L_t from over day $t + 1$ shows that the equation for filtering the spot illiquidity, L_t, can be written as

$$\log(\mathrm{ILQ}_{t+1}) = \log\left[\theta_L\Delta t + (L_t - \theta_L)(-\frac{1}{\kappa_L}(e^{-\kappa_L\Delta t} - 1)\right] + u_{t+1}^1. \tag{25}$$

We use the log of effective spread instead in the measurement equation because the empirical distribution of effective spreads is close to being a log-normal.

Following Pan (2002), we collect two time-series of closing mid-price of options quotes which we label ATM and OTM. ATM refers to the price of the call option that has the moneyness, defined as the ratio of forward-to-strike price, being closest to 1.00. Similarly, OTM refers to the price of the put option that has the moneyness closest to 0.95. Both options are chosen to have time to maturity as close as possible to 30 calendar days. Figure 2 plots daily Black-Scholes option-implied volatilities calculated from the ATM and OTM contracts that we use in our study. As argued by Pan (2002), it is important to use OTM options in the measurement equation as it provides the richest information about investors' expectation of the crash probability of the stock market.

We follow Trolle and Schwartz (2009) and use Black-Scholes vega-weighted price as the functional form in the measurement equations for options fitting; see equations (23)–(24). This method scales the value of options across time making their prices more comparable, which in turn, facilitates the assumption of the normally distributed errors in the measure equations. Therefore, ATM^O_{t+1} and OTM^O_{t+1} in equations (23)–(24) represent the scaled ATM and OTM option prices observed at the end of day t. Similarly, the variables ATM^M_{t+1} and OTM^M_{t+1} denote the model-implied option price scaled by the market Black-Schole vega. The model-implied option price is a function of the three state variables V_{t+1}, L_{t+1}, and Ψ_{t+1}, as well as the model parameters, which are simultaneously estimated. All measurement errors are assumed to be uncorrelated.

The models that we study in this paper fall under the affine jump-diffusion framework. Therefore, the conditional characteristic function of log stock price is available in an exponential affine form. Following Duffie, Pan, and Singleton (2000), we derive the log affine functional form of the characteristic function in Appendix B. The coefficients in the characteristic function are not all available in terms of elementary functions, thus, we solve for them numerically in the Ricatti system of equations. The fact that we must solve the coefficients in the characteristic function numerically together with the large sample period of 2,262 days pose some computational challenge for evaluating option prices using the commonly used numerical integration technique. In addition, the UKF method that we apply the filter state variables require that option prices are calculated each day sequentially. To facilitate the computational challenge, we use the Fast Fourier Transform (FFT) approach first developed by Carr and Madan (1999).

Lastly, at this stage, we do not need to specify the risk premiums associated with the first Brownian motion, dW^1_t, and the compound Poisson jumps, qdN_t because they only alter the drift term of returns dynamics that is not part of the estimation. We will return to discuss the specification of the equity and jump risk premium in the later section where they are estimated using a time-series of daily index returns.

3.4 Estimation

We estimate the models by maximizing the log-likelihood function resulting from the UKF step. We assume the measurement errors are conditionally normal, therefore, the time t conditional log-likelihood takes the following form:

$$l_t(\Theta) = -\frac{3}{2}\log(2\pi) - \frac{1}{2}\log(\det|\Omega_t|) - \frac{1}{2}(Y_t - \bar{Y}_t)^T(\Omega_t)^{-1}(Y_t - \bar{Y}_t), \tag{26}$$

where $\bar{Y}_t$ and Ω_t denote the ex-ante forecasts of the mean and covariance conditional on time $t-1$ information of observables Y_t. We let Θ denotes the set of all parameters to be estimated. All vectors are 3 dimensional and matrices are 3-by-3 symmetric matrix.

In addition to the log-likelihood resulting from the measurement error equations, we follow Andersen, Fusari, and Todorov (2015) and add a penalization term that compares the filtered spot variance component, V_t, to the model-free estimates of spot variance calculated from the high-frequency data. Incorporating this penalizing term, the conditional log-likelihood function that we estimate at time t is

$$L_t(\Theta) = l_t(\Theta) + \omega \log(\ (\sqrt{V_t^n} - \sqrt{V_t})^2\ ), \tag{27}$$

where $l_t(\Theta)$ is given in equation (26), V_t^n is the realized spot variance computed using 1-minute grid returns from S&P 500 index and V_t is the filtered spot variance from the UKF procedure. We describe the construction of the realized spot variance measure in more details in Appendix C.

The tuning parameter ω in equation (27) is set equal to 0.05 following Andersen, Fusari, and Todorov (2015). This parameter determines to weight of the penalization term from the fitting the realized spot variance. As a robustness check, we verify that our main results remain virtually unchanged when picking different values of ω. The model parameters are then estimated by maximizing the sum of conditional log-likelihoods over the sample period from January 2, 2004 to December 31, 2012.

4 Results

4.1 MLE Estimates

Table 4 reports parameter estimates for the three models. The first, second and third columns report results for the SJ, SJV and SJVI models, respectively. We report log-likelihood values of the three models in the bottom row.

We find that parameters governing the square-root dynamic of spot variance are well estimated. Their parameter estimates are fairly consistent across the models. The correlation estimates of the two Brownian shocks in return and spot variance, ρ, are about -35% confirming the asymmetric return-variance relationship found in the extant literature. We find that the spot market illiquidity level, L_t, significantly impacts the level of spot variance, V_t. This is seen from the estimates of γ which measure the contemporaneous relationship between market illiquidity and spot variance. We find that across the three models, the estimates γ are about 0.12. This suggests that a one-standard deviation increase in the spot market illiquidity, L_t, would increase the spot variance level by about 12% after controlling for the persistence dynamic of the variance process.

The strong relationship we find between market illiquidity and return variance lends support to previous studies examining the relationship between return volatility and market liquidity. In particular, motivated by the mixture of distribution hypothesis (MDH), which assumes that volatility and volume simultaneously depends on a latent information process, past research effort has been devoted to studying the relationship between stock return volatility and trading volume (e.g., Clark 1973; Epps and Epps, 1976; Tauchen and Pitts, 1983). Nevertheless, the findings in this literature have been mixed and the understanding of relationships between information flows and trading activity has been an on-going active research area. For instance, Lamoureux and Lastrapes (1990) estimates a GARCH volatility model and find that trading volume is the main driver of stock return volatility and that past stock return innovations became insignificant once trading volume is included in the model.[11] While we find that market illiquidity significantly drives the dynamic of spot variance, its effect does not eliminate the strong persistence in the variance dynamic. Further, the recent literature agrees that trading volume is an inadequate measure of market liquidity.[12] Given the recent availability of intraday trading data, we can more precisely measure market liquidity risk by calculating the cost of participating in the stock market (i.e., transaction cost). Our results estimated using a continuous-time model documenting a strong relationship between market illiquidity and return variance therefore contribute to this stream of literature.

Estimates of the jump-size mean, θ, and the jump-size standard deviation, δ, in Table 4 are very similar across the three models. This finding shows that the nature of jump size identified by the three models are similar in magnitude. The average jump mean size is about 0.37. This suggests that the jump we identify corresponds to a drop of 3.7% in daily S&P 500 index returns, indicating a crash in the stock market index.

[11] In contrary, several studies find evidence conflicting with the MDH specification. These studies include Heirnstra and Jones (1994), Lamoureux and Lastrapes (1994), Richardson and Smith (1994), Anderson (1996).

[12] See for examples, Lee, Mucklow, and Ready (1993), Jones (2002), Fleming (2003), Fujimoto (2004).

We next examine parameter estimates governing the time-varying jump intensity. First, we look at the dynamic of the latent jump-intensity specific factor, Ψ_t. Table 4 shows the magnitude of parameters driving the Ψ_t dynamic in the SJ model differs significantly from those in the other two models. For instance, the long-run mean θ_Ψ, the mean-reversion speed κ_ψ, and the volatility ξ_Ψ of the jump-intensity specific factor are significantly larger for the SJ model. These findings are expected because in the SJ model, jump intensity dynamic solely depends on the latent state variable Ψ_t. Further, these results confirm that the dynamic of jump intensity is time-varying and follows mean-reverting process.

Table 4 shows that when we add covariates to the jump intensity dynamic, e.g., the SJV and SJVI models, the log-likelihood values of the model fit increases substantially. The improvement is large with an increase of about 10% relative to the SJ model. We therefore find a strong support for modeling jump intensity as a function of economic covariates. We further discuss the sources of improvement in the model fit in a later section. Looking at the SJV model, we find the impact of spot variance on jump intensity, γ_V, is positive and statistically significant at the five percent level (t-stat is 2.17). This finding is consistent with Pan (2002), Bates (2006), and Andersen, Fusari and Todorov (2015).

For the SJVI model, we find that when we add market illiquidity to the jump intensity specification, the estimate of γ_V substantially decreases in magnitude and its statistically significance drops to 0.11 in term of p-value (t-stat is 1.61). On the other hand, the impact of spot market illiquidity loads very strong (t-stat is 13.87). This finding shows that the inclusion of market illiquidity as an economic covariate significantly weakens the relationship between jump intensity and spot variance. This finding is consistent with our conclusions from Table 1 which we obtained using regression analyses on non-parametrically estimated jumps.

4.2 Time-varying Volatility and Crash Risks

This section examines the time-series dynamics of market spot volatility and jump intensity that we estimated. Table 5 report descriptive statistics of the daily jump intensity, λ_t, spot variance, V_t, and spot illiquidity, L_t, dynamics that we obtained using the UKF from 2004–2012. We find that the time-series statistics of the spot illiquidity are almost identical across the three models. This suggests that its dynamic is well identified when we extract their information from the daily market illiquidity measure ILQ_t calculated using intraday bid-ask spreads.

We find some distinct differences across the spot variance and jump intensity dynamics in Table 5. To facilitate the visualization, we plot their annualized time-series dynamics. Figures 3 and 4 plot the daily annualized spot volatility and jump intensity, respectively, for the three models. We find that the spot volatilities filtered from the SJ model are smaller in magnitude relative those from the SJV and SJVI models. On the other hand, the jump intensity dynamic of the SJ model is very volatile relative to the other two models. For instance, looking at the time-series statistics of λ_t in Table 5, we find the average expected number of jumps implied by this model is 3.13 per year, but with a median of 0.67 and a standard deviation of 12.62. This shows that the distribution of jump intensities filtered from the SJ model is highly skewed and dispersed. This finding is confirmed when looking at the top panel in Figure 4 which shows the expected number of jumps in the SJ model tremendously increase during the 2008–2009 crisis period. The volatile nature of jump risk estimated from the SJ model explains why its filtered spot volatility dynamic are relatively smaller in magnitude (see Figure 3) than in other two models — because variations in index price dynamics are predominantly captured by jumps.

We find the jump intensity dynamic estimated from the SJV and SJVI models are relatively smooth with the annualized jump-size standard deviation of 3.35 and 2.73, respectively. In these two models, the levels of jump intensity are relatively stable before mid-2007, but increasingly rises after and peaks in the fall of 2008. We believe the relatively stable jump intensity dynamics observed in the SJV and SJVI models are due to the improved identifications resulting from the use of covariates in the jump intensity specification. This argument is supported by looking at the models' log-likelihood performance, which is substantially worse under the SJ model where there is no covariate in the jump intensity specification. We therefore find support for the modeling approach of letting jump intensity be a function of economic covariates that can be identified using observable data.

We next examine the economic contribution of the spot variance and market illiquidity to the jump intensity dynamics. Figure 5 plots the decomposition of daily jump intensity dynamics. Here, we decompose daily jump intensities filtered from the SJV model (top panel) and from the SJVI model (bottom panel) into their respective components.

For the SJV model, the top panel of Figure 5 shows that the market's spot variance is the main component driving jump intensity dynamic. This corresponds to the $\gamma_V V_t$ in the jump intensity dynamic. In order the see how much each component contributes through time, we plot their daily percentage contributions in Figure 6. We find that on average, more than half the jump intensity level is explained by the market's spot variance. The time-series average of its contribution is about 61.55%. We find the jump-intensity-specific factor Ψ_t explains a substantially large portion of time-varying jump intensity. Its average contribution is about 38%, which is two-third in importance relative to the spot variance. This finding shows that a non-trivially large portion of jump intensity cannot be explained by the dynamic of market's spot variance.

The bottom panel of Figure 5 shows the decomposition of daily jump intensities estimated from the SJVI model. Here, we find the jump intensity dynamic is heavily dominated by its co-movement with the spot market illiquidity. We plot daily percentage contributions of each jump intensity component in Figure 7. The results shown are largely consistent with the findings in the bottom panel of Figure 5. We find that, on average, the market illiquidity factor explains about 61% of jump probability in the SJVI model. In contrast to our findings for the SVJ model in Figure 6, we find the market's spot variance explains, on average, only 15%, with the remaining 24% contribution coming from the jump intensity factor, Ψ_t. Therefore, the market's spot variance is the least important factor in explaining the jump intensity dynamic for the SJVI model. This shows that the explanatory power of return variance in time-varying jump risk mostly comes from its relationship with market illiquidity. Once we control for market illiquidity as an economic variable driving time-varying jump risk, the relative contribution of spot variance diminishes.

The above findings offers important insights to the existing literature on index return models which has increasingly documented the importance of time-varying crash risk (e.g. Bates, 2006, 2012; Christoffersen et al., 2012). The common practice in this literature is to let jump intensity be an affine function of spot variance. This modeling approach is appealing because it is parsimonious. It identifies time-varying jump intensity as a constant multiple of the spot variance, thereby eliminating the need to introduce an additional state variable to the model. We find that our estimation results for the SVJ model provide some support for this modeling approach. However, we emphasize that the key economic variable that matters most from our results for modeling jump intensity dynamics is not the market spot variance, but the market illiquidity factor. Lastly, our findings suggest the reason previous studies find a positive relationship between the stock market's time-varying crash risk and spot variance is because of their common exposure to market liquidity risk.

4.3 Option Fits

We also compare the three models based on their in-sample option fits. We define in-sample option pricing error as the sum of squared errors (SSE) in fitting the observed Black-Scholes vega-weighted option prices obtained from the UKF procedure as shown below

$$\text{SSE(ATM)} = \sum_{t=1}^{T} (ATM^{O}_{t+1} - A\bar{T}M^{M}_{t+1})^2, \tag{28}$$

where $A\bar{T}M^{M}_{t+1}$ denotes the ex-ante forecast of vega-weighted ATM option price at time $t+1$. Option pricing error for OTM options are computed in a similar way.

Table 6 reports the in-sample option pricing errors of three models. As expected, the SJ model performs poorly in fitting both ATM and OTM options, having the largest pricing errors of all three. Both SJV and SJVI models produce superior option price fits with comparable magnitude. Pricing errors of ATM options are very similar in magnitude between SJV and SJVI models and the difference is quite small. However, the improvement in fitting OTM options is much larger using the SJVI model suggesting that its jump intensity specification is more well-suited for capturing the jump intensity dynamic embedded in the index options.

4.4 Risk Premiums

So far, we have focused on the risk-neutral estimates. The UKF estimation procedure that we use does not rely on daily returns data. Therefore, we did not have to assume a specific risk premium specification for the jump and the diffusive components in the return dynamic. This section describes the procedure for extracting the jump and diffusive risk premiums from our models and discuss the findings.

Using the risk-neutral parameter estimates in Table 4 and daily filtered states variables $\{\hat{V}_t, \hat{L}_t, \hat{\Psi}_t\}$ estimated previously, we infer the risk premium parameters. This is done by estimating the model on daily S&P 500 index returns from 2004–2012, and keeping the parameters that are not affected by the change of probability measures fixed. This approach of identifying risk premiums was also employed in Andersen, Fusari, and Todorov (2015).

We assume the conventional form of the pricing kernel that preserves the affine structure of the model under the physical measure. The prices of risk associated with the four Brownian motions are given by

$$\begin{aligned} dW_t^{1,\mathbb{P}} &= dW_t^1 - \nu_1 \sqrt{V_t} dt &(29)\\ dW_t^{2,\mathbb{P}} &= dW_t^2 - \nu_V \sqrt{V_t} dt &(30)\\ dW_t^{3,\mathbb{P}}, dW_t^{4,\mathbb{P}} &= dW_t^3, dW_t^4. &(31) \end{aligned}$$

The parameter ν_1 in equation (29) corresponds to the price of risk parameter for the first Brownian innovation in the return process. We recall that ν_V is the price of risk parameter for the volatility innovation which we estimated from options and realized spot variance as part of the UKF steps. Its estimate is report in Table 4. We recall that we do not impose any risk premium assumptions on the third and fourth Brownian motions corresponding to the liquidity and latent jump intensity innovation, respectively.

We follow Pan (2002) and assume the difference between jump distributions under the physical and risk-neutral measures derives from the jump-size risk premium, ν_θ, defined as the difference between jump-size means, $\theta^P - \theta$. The dynamic of log-stock price under the physical probability measure can be written as

$$
\begin{aligned}
d\log(S_t) &= (r - \frac{1}{2}V_t - \xi^P\lambda_t + (\sqrt{1-\rho^2}\nu_1 + \rho\nu_v)V_t)dt + \\
&\quad \sqrt{V_t}(\sqrt{1-\rho^2}dW_t^{1,\mathbb{P}} + \rho dW_t^{2,\mathbb{P}}) + q_t dN_t^{\mathbb{P}} \qquad (32)
\end{aligned}
$$

where $\xi^P = \exp(\theta^P + \frac{1}{2}\delta^2)$ is the jump compensator under the physical measure. Comparing the $\mathbb{P}$-measure return dynamic in equation (32) to the $\mathbb{Q}$-measure return dynamic in equation (10) shows that the equity risk premium, π_t, can be written as

$$
\begin{aligned}
\pi_t &= (\xi^P - \xi)\lambda_t + (\sqrt{1-\rho^2}\nu_1 + \rho\nu_V)V_t \qquad (33) \\
&= (\xi^P - \xi)\lambda_t + \nu_S V_t, \qquad (34)
\end{aligned}
$$

where we define $\nu_S = \sqrt{1-\rho^2}\nu_1 + \rho\nu_v$ in equation (34).

Using the filtered state variables, $\{\hat{V}_t, \hat{L}_t, \hat{\Psi}_t\}$, we apply daily discretization to the return process and estimate the risk premium parameters ν_θ and ν_S using MLE while fixing all other parameters. The estimate for ν_1 are then inferred from ν_S. Section C in the Appendix shows the discretization of the continuous-time model, and presents the log-likelihood function for fitting the return process.

Table 7 reports estimation results of the risk premium parameters. We find that the jump risk premium parameter ν_θ is well identified in all models. The estimates for ν_θ are reported with statistical significant of one percent or greater. On the other hand, estimates of the diffusive risk premium parameter ν_S are marginally significant. These findings are consistent with Pan (2002) who find that jump risk premium is easily identified from index options data, while risk premiums associated with the diffusive and variance risks are more difficult to precisely estimate. Table 7 also reports estimates for the price of risk coefficient ν_1 associated with the first Brownian motion. They are inferred from their corresponding estimates of ν_S in Table 7, and ν_V in Table 4. Because ν_1 is indirectly inferred, we do not report its t-statistic. This parameter can be usefully thought as the price risk for exposure to the diffusive component in index return.

Using the estimates reported in Tables 7 and 4, we quantify the economic magnitude of each risk premium component in terms of annualized excess returns. Equation (34) shows that the equity risk premium can be decompose into two main components. The first component represents the compensation for bearing the stock market's crash risk, $(\xi^P - \xi)\lambda_t$. The second component represents the compensation for bearing the stock market's diffusive return and variance risks, $\nu_S V_t$. For brevity, we refer to $\nu_S V_t$ as the diffusive risk in the equity risk premium.

We first look at the compensation for bearing the stock market's crash risk. For each model, we calculate the long-run jump risk premium level $(\xi^P - \xi)\bar{\lambda}_t$, where $\bar{\lambda}_t$ is the annualized time-series mean of the jump intensity dynamic reported in Table 5. We find the compensation for bearing the market's crash risk for the SJ, SJV and SJVI models are 8.04%, 5.42%, 4.91% in annualized excess returns, respectively. The jump risk premium implied by the SJ model is relatively high. This is likely because jump intensities from the SJ model are more volatile and larger in magnitude than in the other models. We find that jump risk premium estimates implied by the SJV and SJVI models are mostly consistent with prior studies that estimated a time-varying jump risk model the on S&P 500 index. Pan (2002) estimates the jump risk premium using index options over the 1989–1996 period and find that it is about 3.5% in annualized excess return. In a more recent study, Ornthanalai (2014) estimates the jump risk premium implied by the compound Poisson jump process over the 1996–2012 period and finds that its magnitude is 4.52% per year.

We next look at the compensation for bearing the stock market's diffusive risk. This is calculated as $\nu_S \bar{V}_t$, where $\bar{V}_t$ is the time-series mean of the annualized variance reported in Table 5. We find the compensation for bearing the diffusive risk for the SJ, SJV and SJVI models are 8.11%, 7.31%, 8.60% in annualized excess returns, respectively. The magnitudes of the diffusive risk premiums are fairly stable across the three models. They are relatively higher than the magnitude of 4.7% reported in Ornthanalai (2014). However, we recall that estimates of ν_S are marginally significant as shown in Table 7, suggesting that the diffusive risk premiums of the three models have reasonably large standard errors.

The realized equity premium calculated using daily index returns data over the 2004–2012 period is 8.7% per year. The total equity premiums that we find for the SJ, SJV and SJVI models are 16.15%, 12.73% and 13.60%, in annualized returns, respectively. Our estimates of the total equity premium are larger than the value calculated using daily returns data. This finding is expected as the total equity risk premiums that we find are estimated from option prices. They represent investors' ex-ante demand for bearing risks after taking into account which their risk aversion. We conclude that the magnitudes of equity risk premium implied by our estimates are economically plausible.

5 Robustness

5.1 Circuit Breakers

Following the Flash Crash incident on May 6, 2010, the SEC has installed a "circuit breakers" on 404 NYSE-listed S&P500 stocks on June 16, 2010 to halt the trading for 5 minutes if any stock experiences more than 10 percent movement, either up or down, in a 5-minute period. This new trading rule potentially affects our aggregate illiquidity measure constructed from individual firm's effective spreads, and hence alters the impact of market illiquidity on jump probability. We test whether this change in the market-trading rules alter our findings on the impact of market illiquidity on jump intensity dynamic.

We take June 16, 2010 as the date of exogenous shift in the market structure. Specifically, we allow the parameter γ_L to be different prior to and after the implementation date of the circuit breaker as below:

$$\lambda_t = \Psi_t + \gamma_V V_t + 1_{\text{Before}} \cdot \gamma_L^b L_t + 1_{\text{After}} \cdot \gamma_L^a L_t, \tag{35}$$

where 1_{Before} is an indicator function equal to one for all dates t before June 16, 2010, and zero otherwise. Similarly, 1_{After} is an indicator function equal to one for all dates on and after June 16, 2010, and zero otherwise.

We re-estimate the SJVI model using the above augmented jump intensity specification using the identical procedure on the same dataset. Our objective is the test whether γ_L^b significantly differs from γ_L^a. If the circuit breaker has any material impact on the relationship between market illiquidity and the market's jump probability, then we would expect to see a difference between these two coefficients. Below equation summarizes the estimation results from estimating this augmented SJVI model. For brevity, we only report estimates for γ_L^a and γ_L^b in the jump intensity specification. The t-stat for each parameter is reported in parentheses underneath its estimate.

$$\lambda_t = \Psi_t + \gamma_V V_t + 1_{\text{Before}} \cdot \underset{(7.44)^{***}}{9.72}\, L_t + 1_{\text{After}} \cdot \underset{(3.31)^{***}}{8.91}\, L_t.$$

We find the estimate of γ_L^a is slightly lower, being 8.91, relative to the estimate of 9.72 for γ_l^b. Thus, introduction of circuit breakers has slightly reduced the impact of market illiquidity on jump intensity by eliminating possible sudden movements but does materially impact the importance of the illiquidity channel. Both coefficients are statistically significant at 99% level with a lower t-stat for γ_L^a because of the much smaller sample size (2010–2012) in period after the circuit breaker implementation.

5.2 Alternate Iliquidity measure

We have so far defined market illiquidity as an aggregate effective spreads of S&P 500 constituents. This measure captures the aggregate transaction cost of participating in the stock market and has been shown in Ait-Sahalia and Yu (2009), and Goyenko, Holden, and Trzcinka (2009) to be a good proxy for market illiquidity. This section tests whether our results are robust to other ways of defining market illiquidity.

We construct a daily market illiquidity measure in the spirit of Amihud (2002). On each day t, we compute the Amihud illiquidity measure for each firm i constituting the S&P 500 index as a fraction of absolute return, $|r_i|$, over dollar trading volume, $DVol_{i,t}$, per that day.

$$ALIQ_{i,t} = \sum_{i=1}^{N} \frac{|r_{i,t}|}{DVol_{i,t}}. \tag{36}$$

The daily Amihud market illiquidity measure for the stock market is then calculated as an equally-weighted average of individual firms' Amihud illiquidity measure.

Figure 8 plots the market Amihud illiquidity measure in comparison to the effective spread measure. Both measures are normalized to have the same mean over the sample period. This normalization method does not impact our results because the absolute level does not matter for our specification. We immediately see that the Amihud illiquidity measure is much noisier than the effective spread measure. Therefore, we expect the structural estimation exercise using the Amihud illiquidity measure to be very sensitive to the filtering of spot illiquidity level L_t, as well as its measurement error variance. Consequently, a direct comparison between the two measures via statistical inference is not straightforward.

We reestimate the SJVI model using the market Amihud illiquidity measure and report the results in Table 8. Interestingly, the jump intensity contribution from the spot variance V_t almost disappears while the market illiquidity component dominates. The estimated coefficient γ_V is small in magnitude is not statistically significant. Meanwhile, the estimated coefficient γ_L is now much larger and remains statistically significant. Accordingly, mean jump size parameter θ is estimated to be roughly 1% lower than the original case indicating that Amihud measure results in more frequent jumps with smaller magnitude. In-sample option pricing errors reported in the bottom panel of Table 8 are also lower using the Amihud illiquidity measure with a notable improvement in fitting OTM option prices. We believe that these findings are influenced by the noisiness of the raw Amihud Illiquidity measure, thus, the filtering estimation favors the frequent small-sized jumps in the return process.

Overall, our main conclusions remain the same, or perhaps, even stronger with the alternative definition of market illiquidity measure constructed in the spirit of Amihud (2002). We conclude our main results are robust to a differing definition of market illiquidity.

6 Conclusion

We study the role of market illiquidity in explaining the time-varying market volatility and crash risk in the S&P 500 index. We estimate a continuous-time model with stochastic volatility and crash probability. We introduced market liquidity as an observable variable to the model by allowing it to affect the dynamics of spot variance and jump risk intensity. We follow the recent empirical literature in market illiquidity risk (e.g. Ait-Sahalia and Yu, 2009; Goyenko, Holden and Trzcinka, 2009) and measure the daily stock market illiquidity level using volume-weighted intraday bid-ask spreads of all securities constituting the S&P500 index. We estimate the model over 2004–2012 using daily S&P 500 index options, realized spot variance and market illiquidity measure, and find that 61% of time-varying crash risk is due to the stock market's exposure to market illiquidity. The influence of market illiquidity dominates other factors that we examined including the market's spot variance. This is with an exception of the 2008 crisis period when the influence of spot variance dominates and the contribution of market illiquidity falls to about 30%. Overall, our paper highlights the importance of market-trading frictions in index return models and suggests that the time-varying crash risk mostly reflects investors' fear of market illiquidity.

7 Appendix

A High Frequency Measures

Following Andersen, Fusario, and Todorov (2015) and Mancini (2009), we construct the consistent estimator of spot variance at the end of each trading day using the 1-minute grid of S&P 500 futures returns as follows.

$$\hat{V}_t^{(n,m_n)} = \frac{n}{m_n} \sum_{i=n-m_n+1}^{n} (r_{i,t})^2 \mathbb{I}(|r_{i,t}| \leq \alpha n^{-\omega})$$

We use 1-minute-grid returns over 6.5 hours in a trading day, thus resulting in $n = 390$ observations. The value of m_n is set to be 75% of n for each day. Other tuning parameters are set as follows: $\alpha = 4\sqrt{BPV_t}$ and $\omega = 0.49$ where BPV denotes the bi-power variation of day t computed using full 1-minute grid of returns.

B Affine Coefficients in the Characteristic Function

Since our model is casted in affine form, the conditional characteristic function is exponential affine in the state variables following Duffie, Pan, and Singleton (2000).

$$E_t[\exp(i\phi \log(S_T))] = \exp(\alpha(\tau) + \beta_0(\tau)\log(S_t) + \beta_1(\tau)V_t + \beta_2(\tau)L_t + \beta_3(\tau)\Psi_t)$$

We use the notation $\tau = T - t$ for simplicity. The coefficients satisfy following system of Ricatti ODE with the boundary conditions $\beta_0(0) = i\phi$ and $\alpha(0) = \beta_1(0) = \beta_2(0) = \beta_3(0) = 0$

$$\begin{aligned}
\frac{d\beta_0}{d\tau} &= 0 \\
\frac{d\alpha}{d\tau} &= ir\phi + (\kappa_V\theta_V + \gamma\kappa_L\theta_L)\beta_1 + \kappa_L\theta_L\beta_2 + \kappa_\Psi\theta_\Psi\beta_3 \\
\frac{d\beta_1}{d\tau} &= \frac{1}{2}\xi_V^2\beta_1^2 + (\xi_V\rho i\phi - \kappa_V)\beta_1 + (\frac{1}{2}(i\phi)^2 - (\frac{1}{2} + \gamma_v\xi)i\phi + \gamma_v\theta_u) \\
\frac{d\beta_2}{d\tau} &= \frac{1}{2}\xi_L^2\beta_2^2 + (\gamma\xi_L^2\beta_1 - \kappa_L)\beta_2 + (\frac{1}{2}\gamma^2\xi_L^2\beta_1^2 - \gamma\kappa_L\beta_1 - \gamma_l\xi i\phi + \gamma_l\theta_u) \\
\frac{d\beta_3}{d\tau} &= \frac{1}{2}\xi_\Psi^2\beta_3^2 - \kappa_\Psi\beta_3 + \theta_u - \xi i\phi
\end{aligned}$$

where $\theta_u = (e^{\theta i\phi + \frac{1}{2}\delta^2(i\phi)^2} - 1)$. Equations for β_0, β_1, and β_3 can be solved analytically in terms of elementary functions while α and β_2 need to be solved numerically. We employ 4^{th} order Runge-Kutta method with the step size of $\Delta t = 1/252$.

C Discretization of daily returns and estimation

We apply daily discretization to the physical return process in (32). This yields

$$r_{t+1} \simeq (r + (\nu_S - \frac{1}{2})\hat{V}_t - \xi^P \hat{\lambda}_t)\Delta t + \sqrt{\hat{V}_t}\sqrt{\Delta t}\epsilon_t + \sum_{i=1}^{N_t} y_{i,t}, \tag{32}$$

where $\nu_S = \sqrt{1-\rho^2}\nu_1 + \rho\nu_v$, and ϵ_t is the standard normal innovation. The jump component is represented a compound Poisson process $\sum_{i=1}^{N_t} y_{i,t}$, where N_t is the number of jump arrival with intensity $lambda_t$ on day t, and $y_{i,t}$ is i.i.d. normal with mean θ^P and variance δ^2. Conditional on the number of jumps $N_t = j$, we can write the likelihood as conditionally normal, thus, the daily return likelihood can be analytically computed.

References

[1] Aït-Sahalia, Y. and J. Yu, 2009, High Frequency Market Microstructure Noise Estimates and Liquidity Measures, The Annals of Applied Statistics 3, 422–457.

[2] Amaya, D., P. Christoffersen, K. Jacobs, and A. Vasquez, 2015, Does Realized Skewness Predict the Cross-Section of Equity Returns? Journal of Financial Economics 118, 135-167.

[3] Amihud, Y., 2002, Illiquidity and Stock Returns: Cross–Section and Time–Series Effects, Journal of Financial Markets 5, 31–56.

[4] Andersen, T., 1996, Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility, Journal of Finance 51, 160–204.

[5] Andersen, T., L. Benzoni, and J. Lund, 2002, An Empirical Investigation of Continuous-Time Equity Return Models, Journal of Finance 57, 1239–1284.

[6] Andersen, T., D. Dobrev, and E. Schaumburg, 2012, Jump-Robust Volatility Estimation Using Nearest Neighbor Truncation, Journal of Econometrics 169, 75-93.

[7] Andersen, T., N. Fusari, and V. Todorov, 2015a, Parametric Inference and Dynamic State Recovery from Option Panels, Econometrica 83, 1081–1145.

[8] Andersen, T., N. Fusari, and V. Todorov, 2015b, The Risk Premia Embedded in Index Options, Journal of Financial Economics 117, 558–584.

[9] Bakshi, G., P. Carr, and L. Wu, 2008, Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies, Journal of Financial Economics 87, 132–156.

[10] Bates, D., 2006, Maximum Likelihood Estimation of Latent Affine Processes, Review of Financial Studies 19, 909–965.

[11] Bates, D., 2012, U.S. Stock Market Crash Risk, 1926–2006. Journal of Financial Economics 105, 229–259.

[12] Bradley, D., J. Clarke, S. Lee, and C. Ornthanalai, 2014, Are Analysts' Recommendations Informative: Evidence from the Impact of Timestamp Delays, Journal of Finance 69, 645–673.

[13] Broadie, M., M. Chernov, and M. Johannes, 2007, Model Specification and Risk Premiums: The Evidence from the Futures Options, Journal of Finance 62, 1453–1490.

[14] Chae, J., 2005, Trading Volume, Information Asymmetry, and Timing Information, Journal of Finance 60, 413–442.

[15] Christoffersen, P., C. Dorion, K. Jacobs, and L. Karoui, 2014, Nonlinear Filtering in Affine Term Structure Models, Management Science, 60, 2248–2268.

[16] Christoffersen, P., R. Goyenko, K. Jacobs, and M. Karoui, 2014, Illiquidity Premia in the Equity Options Market, Working Paper.

[17] Christoffersen, P., S. Heston, and K. Jacobs, 2009, The Shape and Term Structure of the Index Option Smirk: Why Multifactor Stochastic Volatility Models work so Well, Management Science 55, 1914–1932.

[18] Christoffersen, P., K. Jacobs, C. Ornthanalai, Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options, Journal of Financial Economics 106, 447-472.

[19] Chung, K.H., and C. Chuwonganant, 2014, Uncertainty, market structure, and liquidity, Journal of Financial Economics 113, 476–499.

[20] Clark, P., 1973, A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices, Econometrica 31, 135–155.

[21] Duffie, D., J. Pan and K. Singleton, 2000, Transform Analysis and Asset Pricing for Affine Jump-Diffusions, Econometrica 68, 1343–1376.

[22] Egloff, D., M. Leippold, and L. Wu, The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments, Journal of Financial and Quantitative Analysis 45, 1279–1310.

[23] Epps, T and Epps, M., 1976, The Stochastic Dependence of Security Price Changes and Transaction Volumes: Implications for the Mixture-of-Distribution Hypothesis, Econometrica 44, 305–321.

[24] Eraker, B., 2004, Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices, Journal of Finance 59, 1367–1403.

[25] Filipović, D., E. Gourier, and L. Mancini, 2015, Quadratic variance swap models, Journal of Financial Economics, forthcoming.

[26] Fleming, M., 2003, Measuring Treasury Market Liquidity, FRBNY Economic Policy Review, 83–108.

[27] Fujimoto, A., 2004, Macroeconomic Sources of Systematic Liquidity, Working Paper.

[28] Gennotte G., and H. Leland, 1990, Market Liquidity, Hedging, and Crashes, American Economic Review 80, 999–1021.

[29] Goyenko, R., C. Holden and C. Trzcinka, 2009, Do Liquidity Measures Measure Iiquidity?, Journal of Financial Economics 92, 153–181.

[30] He, X. and R. Velu, 2014, Volume and Volatility in a Common-Factor Mixture of Distribution Model, Journal of Financial and Quantitative Analysis 49, 33–49.

[31] Heimstra, C. and Jones, J., 1994, Testing for Linear and Nonlinear Granger Causality in the Stock Price-Volume Relation, Journal of Finance 49, 1639–1664.

[32] Hendershott, T., Jones, C., and Menkveld, A., 2011, Does Algorithmic Trading Improve Liquidity?, Journal of Finance 66, 1–33.

[33] Heston, S., 1993, A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, Review of Financial Studies 6, 327–343.

[34] Huang, X. and G. Tauchen, 2005, The Relative Contribution of Jumps to Total Price Variance, Journal of Financial Econometrics 3, 456–499.

[35] Huang, J. and J. Wang, 2009, Liquidity and Market Crashes, Review of Financial Studies 22, 2607–2643.

[36] Jiang, G., I. Lo and A. Verdelhan, 2011, Information Shocks, Liquidity Shocks, Jumps, and Price Discovery: Evidence from the U.S. Treasury Market, Journal of Financial and Quantitative Analysis 46, 527–551.

[37] Jones, C., 2002, A Century of Stock Market Liquidity and Trading Costs, Working Paper, Columbia University.

[38] Lamoureux, C. and Lastrapes, W., 1990, Heteroskedasticity in Stock Return Data: Volume versus GARCH Effects, Journal of Finance 45, 221–229.

[39] Lamoureux, C. and Lastrapes, W., 1994, Endogenous Trading Volume and Momentum in Stock-Return Volatility, Journal of Business and Economic Statistics 12, 253–260.

[40] Lee, S., 2012, Jumps and Information Flow in Financial Markets, Review of Financial Studies 25, 439–479.

[41] Lee, C., Mucklow, B., and Ready, M., 1993, Spreads, Depths, and the Impact of Earnings Information: An Intraday Analysis, Review of Financial Studies 6, 345–374.

[42] Maheu, J., McCurdy, T., 2004, News arrival, jump dynamics and volatility components for individual stock returns, Journal of Finance 59, 755–793.

[43] Mancini, C., 2009, Non-parametric threshold estimation for models with stochastic diffusion coefficients and jumps, Scandinavian J. Statistics 36, 270–296.

[44] Ornthanalai, C., 2014, Levy jump risk: Evidence from Options and Returns, Journal of Financial Economics 112, 69–90.

[45] Richardson, M. and Smith, T., 1994, A Direct Test of the Mixture of Distribution Hypothesis: Measuring the Daily Flow of Information, Journal of Financial and Quantitative Analysis 29, 101–116.

[46] Pan, J., 2002. The jump risk premia implicit in options: Evidence from an integrated time series study, Journal of Financial Economics 63, 3–50.

[47] Richardson, M. and Smith, T., 1994, A Direct Test of the Mixture of Distribution Hypothesis: Measuring the Daily Flow of Information, Journal of Financial and Quantitative Analysis 29, 101-116.

[48] Santa-Clara, P., Yan, S., 2010. Crashes, Volatility and the Equity Premium: Lessons from S&P 500 Options. Review of Economics and Statistics 92, 435–451.

[49] Tauchen, G. and Pitts, M., 1983, The Price Variability-Volume Relationship on Speculative Markets, Econometrica 51, 485–505.

[50] Tookes, H., 2008, Information, Trading, and Product Market Interactions: Cross-Sectional Implications of Informed Trading, Journal of Finance 63, 379–413.

[51] Trolle, A., and Schwartz, E., 2009, Unspanned Stochastic Volatility and the Pricing of Commodity Derivatives, Review of Financial Studies 22, 4423–4461.

[52] Wan, E., R. van der Merwe, 2000, Dual Estimation and the Unscented Transformation, Neural Information Processing Systems 12, 666–672.

Table 1: Logit Regression Results on Non-parametrically detected Jumps

	Probability of observing a jump the next day						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
MinRV	2.17			−1.43		2.10	−1.44
	(1.91)*			(−1.24)		(1.85)*	(−1.25)
ILQ		3.66		4.82	3.60		4.77
		(3.54)***		(3.32)***	(3.48)***		(3.29)***
RSkew			−0.02		−0.02	−0.02	−0.02
			(−1.07)		(−0.77)	(−0.91)	(−0.78)

The t-stats are in parenthesis.

Notes: We report estimated coefficients and t-stats from the predictive logit regression on non-parametrically detected jumps from daily S&P 500 index returns. The sample period is from January 2, 2004 to December 31, 2012. The dependent variable is an indicator function that is equal to one on day t if jump is detected, and zero otherwise. The independent variables include lagged realized variance estimator, *MinRV*, from Andersen, Dobrev, and Scaumburg (2012); market illiquidity proxy, *ILQ*, measured by daily averaged effective spreads across firms in the S&P 500 constituents; and realized skewness measure, *RSkew*, calculated following Amaya, Christoffersen, Jacobs, and Vasquez (2015). All indepdent variables are expressed in annualized terms by multiplying their daily values by 252. Ex-post daily jumps are detected at the 99.9% confidence level. Year and day-of-the week fixed effects are included. We report robust t-statistic in bracket below each parameter estimate. ***, **, * indicates statistical significance at the one, five, and ten percent confidence levels.

Table 2: OLS regression results on daily realized skewness and kurtosis

Panel A. Realized skewness: $RSkew_{t+1}$

	(1)	(2)	(3)
$MinRV_t$	-0.15		1.29
	(-0.16)		(1.12)
ILQ_t		-1.17	-2.34
		(-1.46)	$(-1.91)*$
$RSkew_t$	-0.05	-0.05	-0.05
	$(-1.85)^{*}$	$(-1.94)^{**}$	$(-1.97)^{**}$
$Return_t$	-7.31	-7.36	-7.10
	$(-2.05)^{**}$	$(-2.06)^{**}$	$(-1.97)^{**}$
Intercept	0.35	0.49	0.72
	$(3.30)^{***}$	$(0.55)^{**}$	$(3.27)^{***}$

Panel B. Realized kurtosis: $RKurt_{t+1}$

	(1)	(2)	(3)
$MinRV_t$	0.73		8.13
	(0.19)		$(1.84)^{*}$
ILQ_t		-4.55	-11.81
		(-1.21)	$(-2.10)^{**}$
$RKurt_t$	-0.04	-0.04	-0.04
	(-1.80)	$(-1.82)^{*}$	$(-1.88)^{*}$
$Return_t$	11.45	9.93	11.10
	(0.80)	(0.71)	(0.797)
Intercept	11.70	12.50	13.52
	$(20.70)^{***}$	$(14.48)^{***}$	$(12.85)^{***}$

Notes: We report predictive OLS estimates on daily realized skewness (Panel A) and realized kurtosis (Panel B) measures of S&P 500 index returns. The sample period is from January 2, 2004 to December 31, 2012. Daily measures of realized skewness, $RSkew_{t+1}$, and realized kurtosis, $RKurt_{t+1}$, are constructed from high-frequency data following the method in Amaya et al. (2015). The independent variables include lagged realized variance estimator, *MinRV*, from Andersen, Dobrev, and Scaumburg (2012); market illiquidity proxy; *ILQ*, measured by daily averaged effective spreads across firms in the S&P 500 constituents; *Return*, log S&P 500 return. Lagged dependent variable is also included in each regression. Day-of-the week fixed effects are included. We report robust t-statistic clustered in bracket below each parameter estimate. ***, **, * indicate statistical significance at the one, five, and ten percent confidence levels.

Table 3: ARMA regresion model on change in realized skewness and kurtosis

Panel A. Realized skewness: $\Delta RSkew_{t+1}$

	(1)	(2)	(3)
$\Delta MinRV_t$	0.30		2.04
	(0.37)		$(1.90)^{*}$
ΔILQ_t		-1.91	-4.29
		$(-1.68)^{*}$	$(-2.53)^{**}$
$\Delta RSkew_t$	-0.08	-0.08	-0.08
	$(-3.83)^{***}$	$(-3.89)^{***}$	$(-3.89)^{***}$
$Return_t$	-0.15	-0.57	-0.90
	(-0.35)	(-1.33)	$(-1.91)^{*}$
AICC	1.566	1.564	1.563

Panel B. Realized kurtosis: $\Delta RKurt_{t+1}$

	(1)	(2)	(3)
$\Delta MinRV_t$	1.48		5.87
	(0.37)		(1.06)
ΔILQ_t		-6.29	-14.75
		(-0.89)	(-1.38)
$\Delta RKurt_t$	-0.10	-0.10	-0.10
	$(-4.84)^{***}$	$(-4.82)^{***}$	$(-4.78)^{***}$
$Return_t$	-0.44	-3.22	-4.44
	(-0.14)	(-0.85)	(-1.05)
AICC	4.767	4.794	4.795

Notes: We report predictive regression results on daily changes in realized skewness (Panel A), and changes in realized kurtosis (Panel B) measures of S&P 500 index returns. Each specification is estimated using MLE assuming a ARMA(1,1) structure in the residuals. The sample period is from January 2, 2004 to December 31, 2012. Daily changes in realized skewness, $\Delta RSkew_{t+1}$, and daily changes in realized kurtosis, $\Delta RKurt_{t+1}$, are constructed from high-frequency data following the method in Amaya et al. (2015). The independent variables include lagged change in realized variance estimator, $\Delta MinRV$, from Andersen, Dobrev, and Scaumburg (2012); change in market illiquidity proxy; ΔILQ, measured by daily averaged effective spreads across firms in the S&P 500 constituents; *Return*, log S&P 500 return. Lagged dependent variable is also included in each regression. We report Heteroskedasticity-adjusted t-statistic below each parameter estimate. ***, **, * indicate statistical significance at the one, five, and ten percent confidence levels.

Table 4: Maximum Likelihood Estimates: 2004–2012

	(1) SJ Model $\lambda_t = \Psi_t$	(2) SJV Model $\lambda_t = \Psi_t + \gamma_V V_t$	(3) SJVI Model $\lambda_t = \Psi_t + \gamma_V V_t + \gamma_L L_t$
Parameter	Estimate	Estimate	Estimate
κ_V	3.5792	3.5652	3.5490
	(9.97)	(8.39)	(4.73)
θ_V	0.0311	0.0309	0.0312
	(11.94)	(11.00)	(5.66)
ξ_V	0.3431	0.3425	0.3456
	(59.84)	(22.97)	(42.49)
κ_L	2.3767	2.3445	2.3532
	(3.69)	(1.27)	(4.61)
θ_L	0.1829	0.1822	0.1709
	(7.90)	(5.16)	(6.69)
ξ_L	0.1543	0.1510	0.1583
	(38.50)	(6.10)	(36.80)
κ_Ψ	1.6727	0.6606	0.6616
	(3.71)	(0.83)	(2.25)
θ_Ψ	2.1142	0.1022	0.1014
	(3.50)	(3.34)	(1.66)
ξ_Ψ	0.4038	0.2036	0.2037
	(4.23)	(7.22)	(1.79)
ρ	-0.3505	-0.3511	-0.3527
	(17.88)	(2.90)	(5.92)
θ	-0.0368	-0.0367	-0.0373
	(12.03)	(10.41)	(26.28)
δ	0.0368	0.0318	0.0314
	(19.59)	(13.48)	(29.84)
γ	0.1151	0.1167	0.1180
	(7.06)	(3.52)	(7.95)
ν_V	1.5961	1.5592	1.5536
	(1.34)	(0.42)	(0.71)
γ_V		52.7228	18.3801
		(2.17)	(1.61)
γ_L			9.2590
			(13.87)
Log-Likelihood:	7,085.19	7,641.24	7,707.29

The t-stats are in parenthesis.

Notes: We report MLE estimates of the three time-varying jump models: (1) SJ, (2) SJV, and (3) SJVI. The sample period is from January 2, 2004 to December 31, 2012. Each model is estimated using daily out-of-the-money (OTM) and at-the-money (ATM) S&P 500 index options, averaged effective spreads of S&P 500 constituents, and spot variance estimated from 1-minute high-frequency S&P 500 futures data. We maximize the log likelihood function in equation (27) where the state variables are estimated using the Unscented Kalman Filter (UKF). We report t-statistic calculated following the BHHH method using the outer product of the gradient in bracket below each parameter estimate.

Table 5: Descriptive statistics of filtered jump intensities and spot variances

	(1) SJ Model $\lambda_t = \Psi_t$	(2) SJV Model $\lambda_t = \Psi_t + \gamma_v V_t$	(3) SJVI Model $\lambda_t = \Psi_t + \gamma_v V_t + \gamma_l L_t$
Panel A. Jump intensity λ_t			
Mean	3.1296	2.5273	2.9172
Median	0.6732	1.6013	2.3124
Std. Dev.	12.6212	3.3500	2.1723
25 percentile	0.3261	0.9441	1.8025
75 percentile	1.2662	2.7791	3.2946
Panel B. Spot variance V_t			
Mean	0.0211	0.0290	0.0316
Median	0.0155	0.0156	0.0146
Std. Dev.	0.0178	0.0429	0.0561
25 percentile	0.0098	0.0095	0.0086
75 percentile	0.0267	0.0297	0.0273
Panel C. Spot illiquidity L_t			
Mean	0.1681	0.1681	0.1681
Median	0.1539	0.1539	0.1538
Std. Dev.	0.0558	0.0557	0.0558
25 percentile	0.1374	0.1375	0.1374
75 percentile	0.1741	0.1742	0.1742

Notes: We report the descriptive statistics of filtered jump intensities, λ_t, spot variances, V_t, and spot illiquidity, L_t for three models: SJ, SJV, and SJVI. The variables are reported in annualized term by multiplying their daily values with 252. We obtain the filtered state variables from Unscented Kalman Filter (UKF) step in the MLE estimation. Models' parameter estimates are reported in Table 4.

Table 6: Option Pricing Errors of Different Models

	(1) SJ Model $\lambda_t = \Psi_t$	(2) SJV Model $\lambda_t = \Psi_t + \gamma_v V_t$	(3) SJVI Model $\lambda_t = \Psi_t + \gamma_v V_t + \gamma_l L_t$
Option moneyness	Estimate	Estimate	Estimate
OTM	288.30	227.34	208.87
ATM	216.63	153.39	150.53

Notes: We report the in-sample option pricing errors of three models estimated from the MLE. Numbers reported are sum of squared errors (SSE) from the measurement equations in the UKF step. The sum of squared errors for ATM options is calculated as shown below

$$\text{SSE(ATM)} = \sum_{t=1}^{T}(ATM_{t+1}^{O} - A\bar{T}M_{t+1}^{M})^2$$

where $A\bar{T}M_{t+1}^{M}$ denotes the ex-ante forecast of vega-weighted ATM option price at time $t+1$, and ATM_{t+1}^{O} denotes the vega-weighted ATM option price observed in the data. The sum of squared errors for OTM options is computed in a similar way.

Table 7: Risk Premium Parameters Estimated from daily Returns: 2004–2012

	(1) SJ Model $\lambda_t = \Psi_t$	(2) SJV Model $\lambda_t = \Psi_t + \gamma_v V_t$	(3) SJVI Model $\lambda_t = \Psi_t + \gamma_v V_t + \gamma_l L_t$
Parameter	Estimate	Estimate	Estimate
$\nu_s = \sqrt{1-\rho^2}\nu_1 + \rho\nu_v$	3.842	2.519	2.750
	(1.89)	(1.39)	(1.54)
ν_θ	0.0263	0.022	0.017
	(11.76)	(4.23)	(2.24)
ν_1	3.505	2.105	2.353
Log-Likelihood:	7,093.69	7,197.63	7,204.14

The t-stats are in parenthesis.

Notes: We report MLE estimates of the risk premium parameters for the three time-varying jump models: (1) SJ, (2) SJV, and (3) SJVI. Each model is fitted to daily S&P 500 return daily returns data from January 2, 2004 to December 31, 2012. We obtain daily state values V_t, L_t, and Ψ_t, as well as $\mathbb{Q}$-measure parameters from the first-stage estimation results reported in Table 4. The parameter ν_θ is the difference between jump-size means under the physical and risk-neutral measures, i.e., $\theta^P - \theta$. The parameter ν_1 corresponds the price of risk coefficient associated with the Brownian innovation in the return process; see equation (29). To facilitate econometric identification, we estimate $\nu_s = \sqrt{1-\rho^2}\nu_1 + \rho\nu_V$ from daily returns MLE and later infer ν_1 from its estimate, and using the value of ν_V reported in Table 4.

Table 8: Parameter Estimates using Liquidity, Options, and SpotVariance. 2004–2012

	(1) SJVI : Amihud-ILQ	(2) SJVI : ES-ILQ
Parameter	Estimate	Estimate
κ_V	4.0642	3.5490
	(1.52)	(4.73)
θ_V	0.0303	0.0312
	(2.09)	(5.66)
ξ_V	0.3273	0.3456
	(10.03)	(42.49)
κ_L	2.5795	2.3532
	(2.28)	(4.61)
θ_L	0.1057	0.1709
	(2.42)	(6.69)
ξ_L	0.2636	0.1583
	(9.16)	(36.80)
κ_Ψ	1.5589	0.6616
	(1.23)	(2.25)
θ_Ψ	1.1071	0.1014
	(1.21)	(1.66)
ξ_Ψ	0.1177	0.2037
	(1.57)	(1.79)
ρ	-0.3679	-0.3527
	(1.12)	(5.92)
θ	-0.0277	-0.0373
	(5.60)	(26.28)
δ	0.0278	0.0314
	(6.93)	(29.84)
γ	0.0993	0.1180
	(1.77)	(7.95)
ν_v	1.5752	1.5536
	(0.22)	(0.71)
γ_V	3.5039	18.3801
	(0.05)	(1.61)
γ_L	235.5151	9.2590
	(9.40)	(13.87)
Option moneyness	Option pricing error	
OTM	186.00	208.87
ATM	143.84	150.53

The t-stats are in parenthesis.

Notes: We report the parameter estimates, t-stats, and option pricing errors of SJVI models estimated using daily OTM and ATM options, spot variance estimated from 1-min grid of high-frequency returns data, and two different measures of illiquidity, for the period starting January 2, 2004 and ending December 31, 2012. The first column reports the estimation result using Amuihud illiquidity measure while the second column reports the result using effective spread as a measure of illiquidity. Amuhud illiquidity measure is computed as the equally weighted average of all individual firm's Amihud illiquidity measure that constitutes S&P500 index each day. Individual firm's daily Amihud illiquidity measure is defined by its absolute daily return divided by daily dollar volume, $|r_{i,t}|/\text{DVol}_{i,t}$. All models are estimated by maximizing log likelihood from Unscented Kalman Filter (UKF). T-stats are computed following the BHHH method using the outer product of the gradient.

Figure 1: S&P 500 returns, Spot Variance, and Market Illiquidity

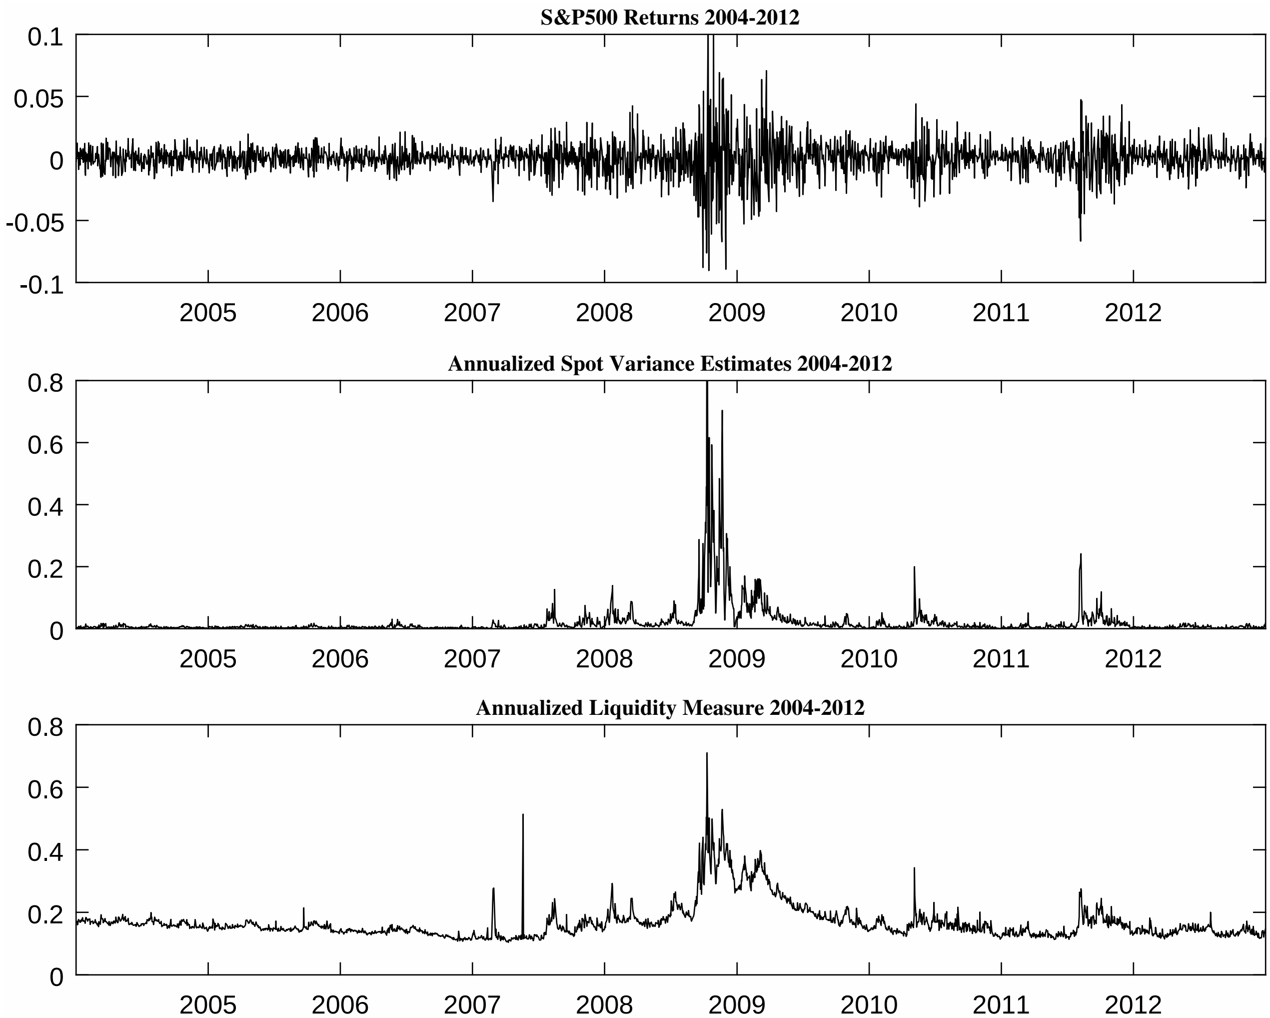

Notes: In the top panel, we plot the daily returns on the S&P 500 index from January 2, 2004 to December 31, 2012. In the middle panel, we plot the annualized spot variance estimates computed using the 1-minute grid of returns following the approach of Andersen, Fusari, and Todorov (2015). In the bottom panel, we plot the illiquidity measure defined by equally weighted average of annualized effective spread from all firms constituting S&P 500 index each day.

Figure 2: Implied volatilities of OTM and ATM options

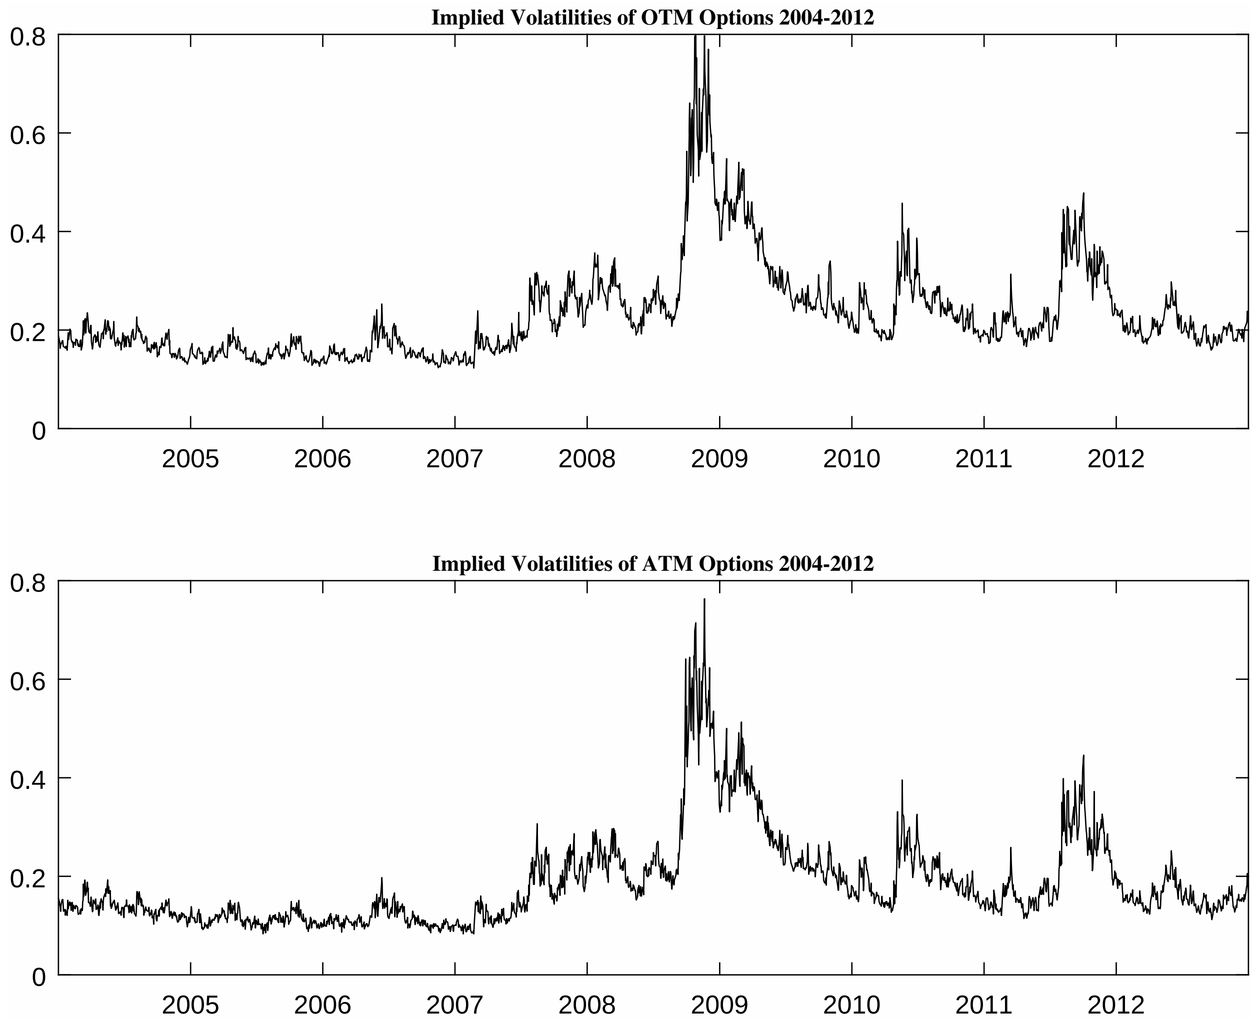

Notes: In the top panel, we plot the daily implied volatilities of out of the money (OTM) options written on S&P500 index from January 2, 2004 to December 31, 2012. In the bottom panel, we plot the implied volatilities of at the money (ATM) options for the same underlying. Both options are chosen to have the time to maturity to be closest to 30 calendar days. OTM options are chosen to have forward price to strike ratio to be closest to 0.95 while ATM options have the same ratio being closest to 1.

Figure 3: Filtered Spot Volatility

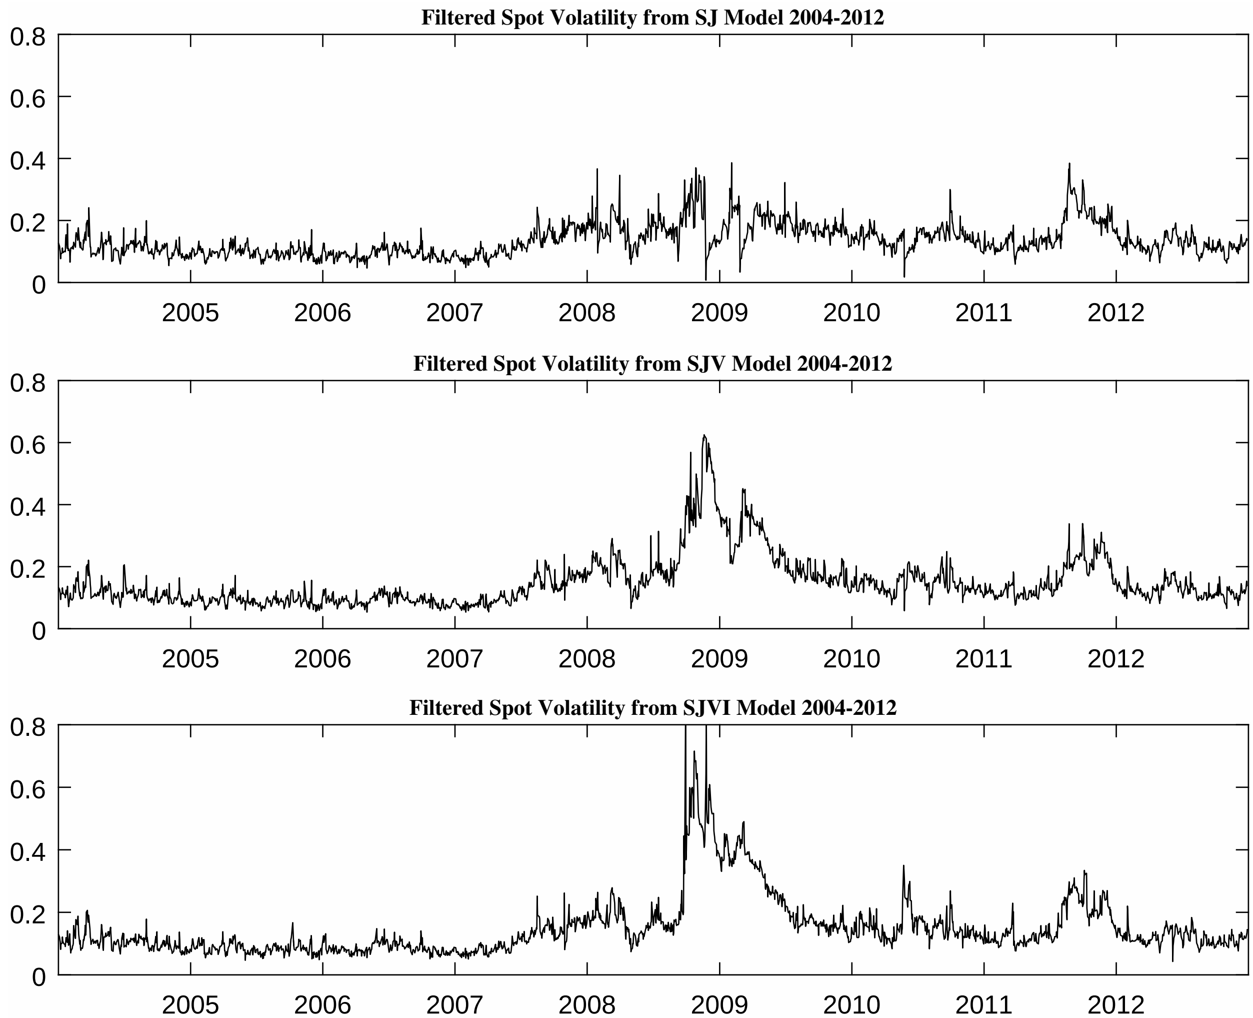

Notes: We plot the daily annualized spot volatility $\sqrt{V_t}$ filtered from three models we consider from January 2, 2004 to December 31, 2012. The top panel corresponds to SJ model that has jump intensity purely driven by latent stochastic jump intensity process, the middle panel corresponds to SJV model that has jump intensity being driven by latent stochastic jump intensity and variance, and the bottom panel corresponds to SJVI model that has jump intensity being driven by latent stochastic jump intensity, variance, and illiquidity.

Figure 4: Filtered Jump Intensity λ_t

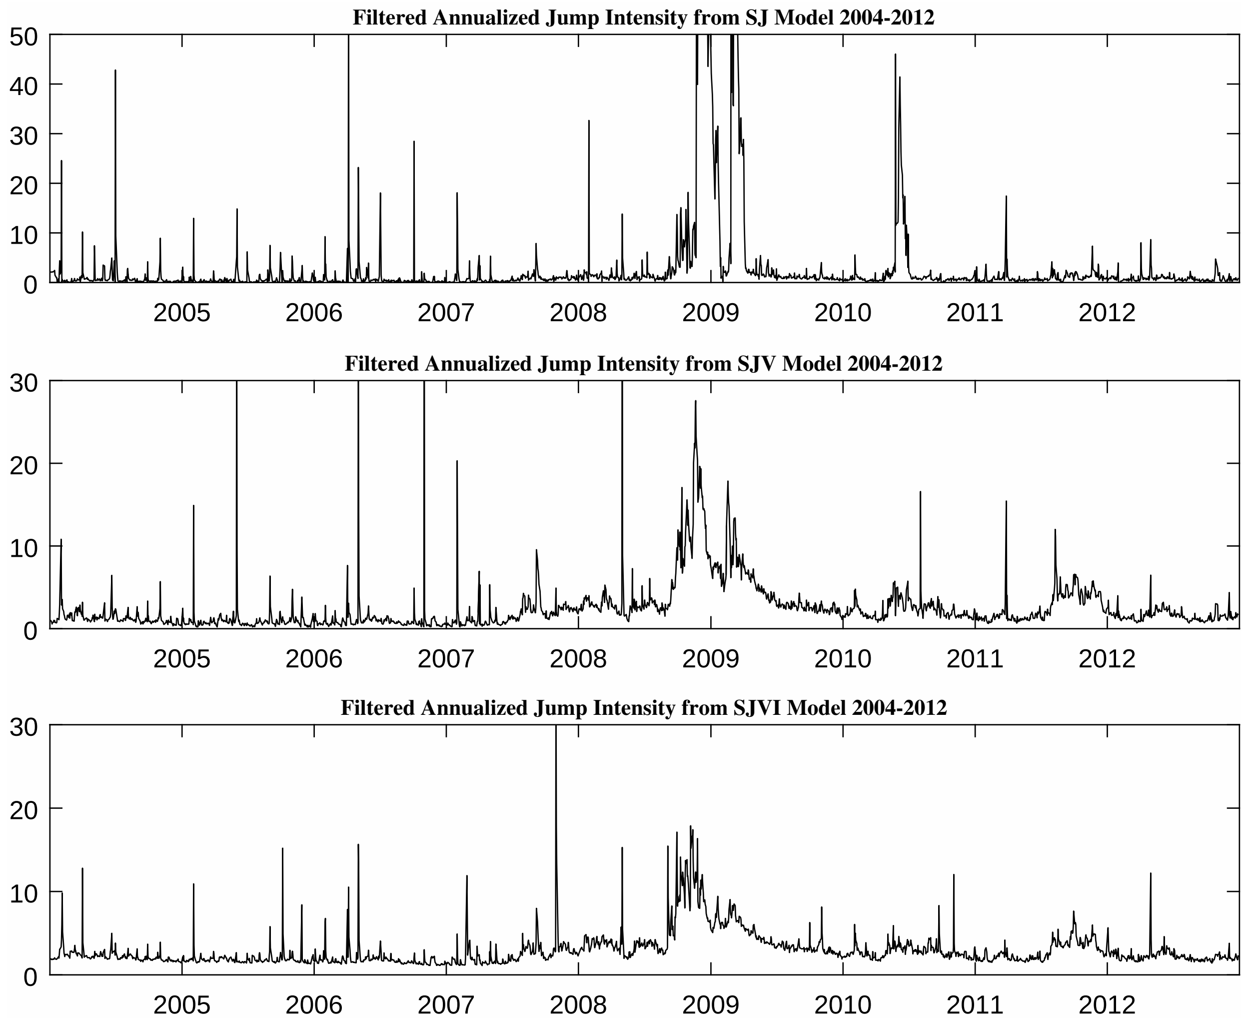

Notes: We plot daily annualized jump intensities λ_t filtered for the three models that we study from January 2, 2004 to December 31, 2012. The jump intensity specifications in the three models can be summarized as follows:

$$\begin{aligned} \text{SJ} \quad &: \quad \lambda_t = \Psi_t \\ \text{SJV} \quad &: \quad \lambda_t = \Psi_t + \gamma_V V_t \\ \text{SJVI} \quad &: \quad \lambda_t = \Psi_t + \gamma_V V_t + \gamma_L L_t. \end{aligned}$$

The top panel corresponds to the SJ model that has jump intensities solely driven by a latent jump-intensity term, the middle panel corresponds to SJV model that has jump intensity being driven by latent stochastic jump intensity and variance, and the bottom panel corresponds to SJVI model that has jump intensity being driven by latent stochastic jump intensity, variance, and illiquidity.

Figure 5: Decomposition of Jump Intensity λ_t: SJV vs. SJVI

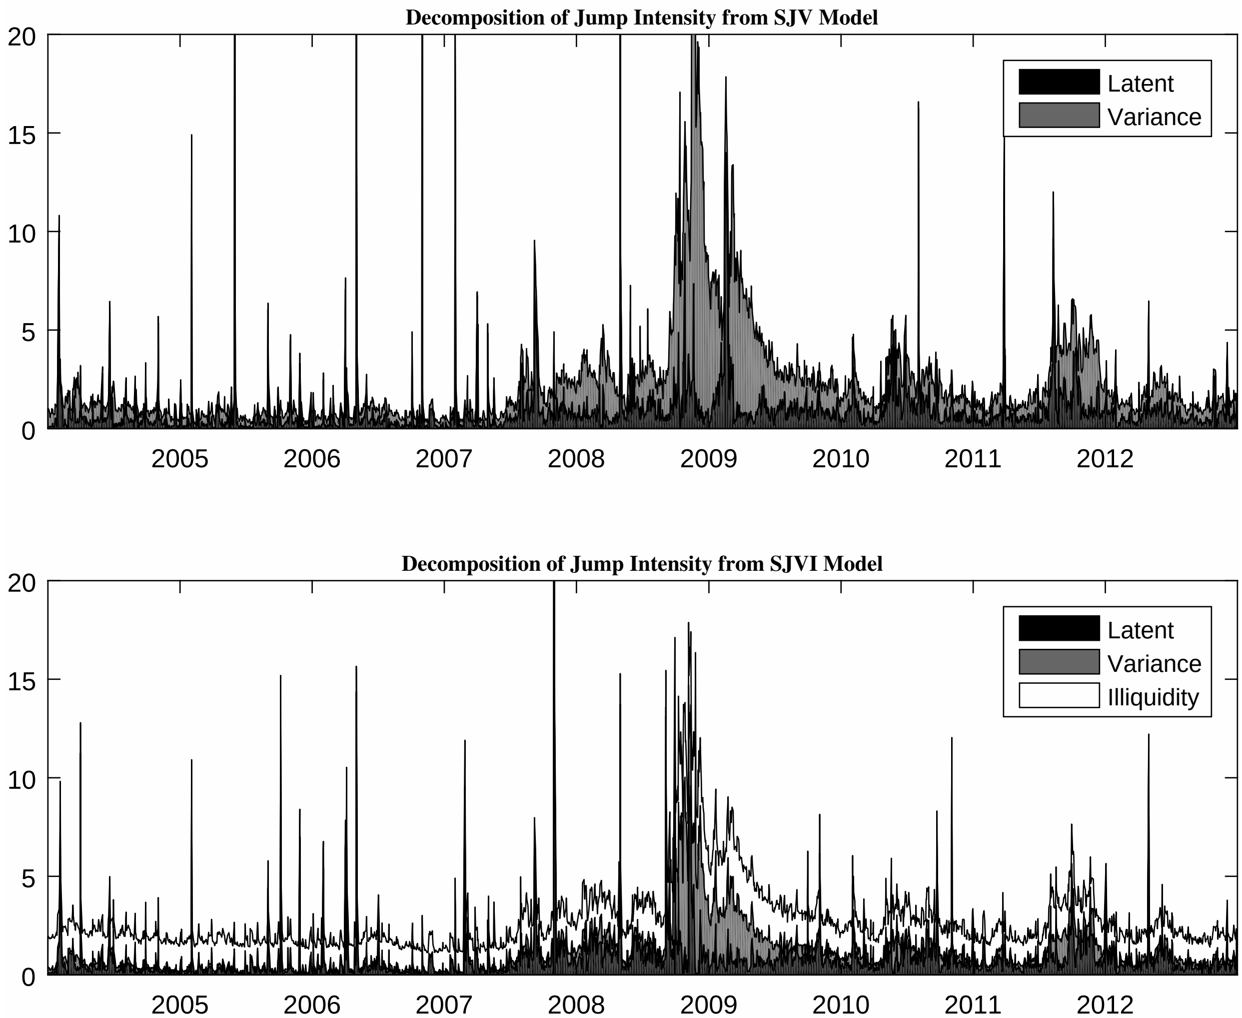

Notes: We plot the decomposition of daily annualized jump intensities λ_t filtered from the SJV model (top panel) and the SJVI model (bottom panel) from January 2, 2004 to December 31, 2012. The top panel decomposes daily jump intensity dynamics of the SJV model into the portion coming from the latent stochastic jump-intensity term, Ψ_t, and the portion that is due to the daily spot variance, $\gamma_V V_t$. In the bottom panel, we decompose daily jump intensity dynamics of the SJVI model into the portion coming from the latent stochastic jump-intensity specific term Ψ_t, the portion that is due to the daily spot variance, $\gamma_V V_t$, and the portion that is due to daily spot market illiquidity, $\gamma_L L_t$.

Figure 6: Relative Contribution (%) to Jump Intensity λ_t: SJV Model

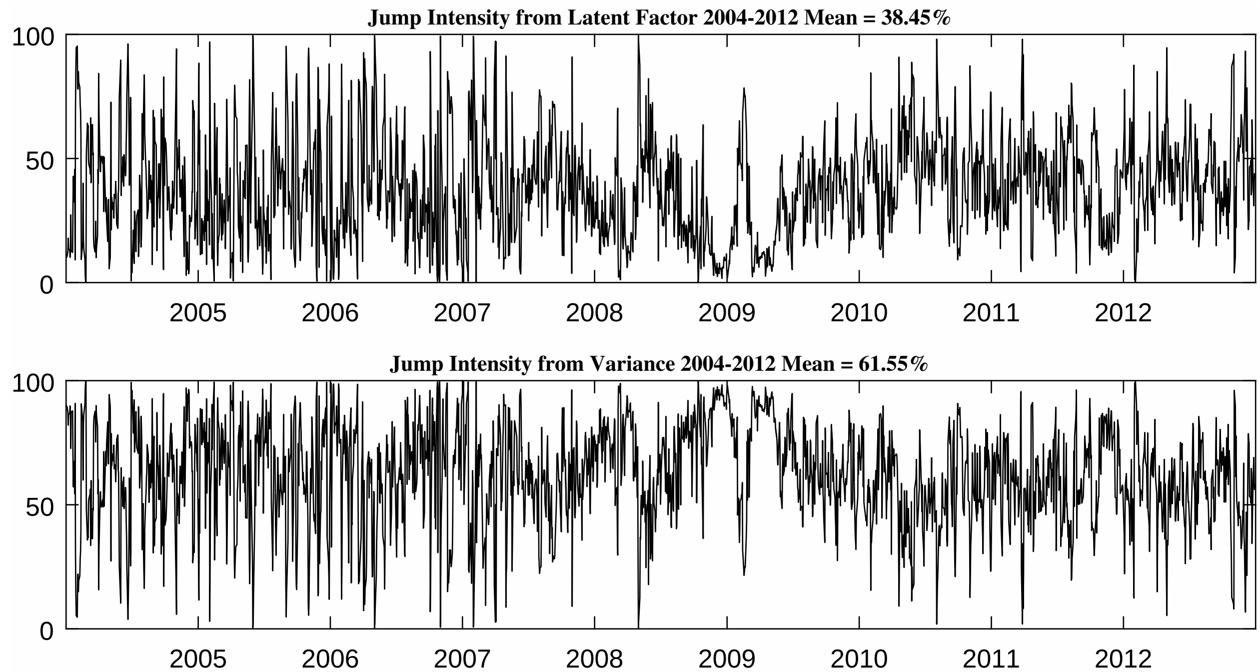

Notes: We plot the breakdown of daily annualized jump intensity, $\lambda_t = \Psi_t + \gamma_V V_t$, filtered from the SJV model from January 2, 2004 to December 31, 2012. The top panel plots the percentage contribution coming from the latent stochastic jump intensity term, Ψ_t/λ_t, while the bottom panel plots the contribution coming from the variance term, $\gamma_V V_t/\lambda_t$.

Figure 7: Relative Contribution (%) to Jump Intensity λ_t: SJVI Model

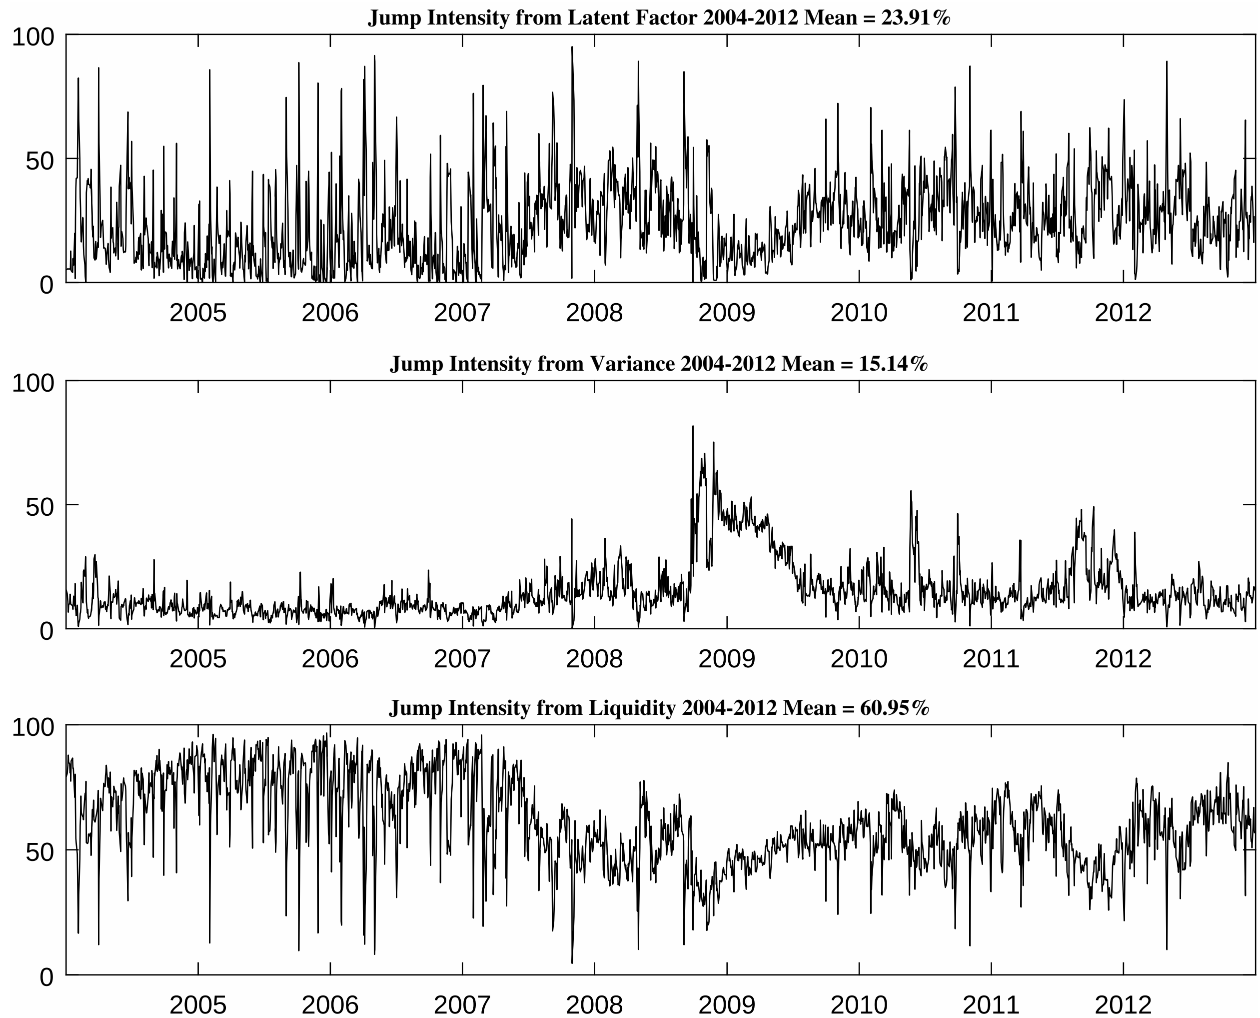

Notes: We plot the breakdown of daily annualized jump intensity, $\lambda_t = \Psi_t + \gamma_V V_t + \gamma_L L_t$, filtered from the SJVI model from January 2, 2004 to December 31, 2012. The top panel plots the percentage contribution coming from the latent stochastic jump intensity term, Ψ_t/λ_t. The middle panel plots the contribution coming from the variance term, $\gamma_V V_t/\lambda_t$. The bottom panel plots the contribution coming from the illiquidity term, $\gamma_l L_t$.

Figure 8: Time Series of Liquidity Measures

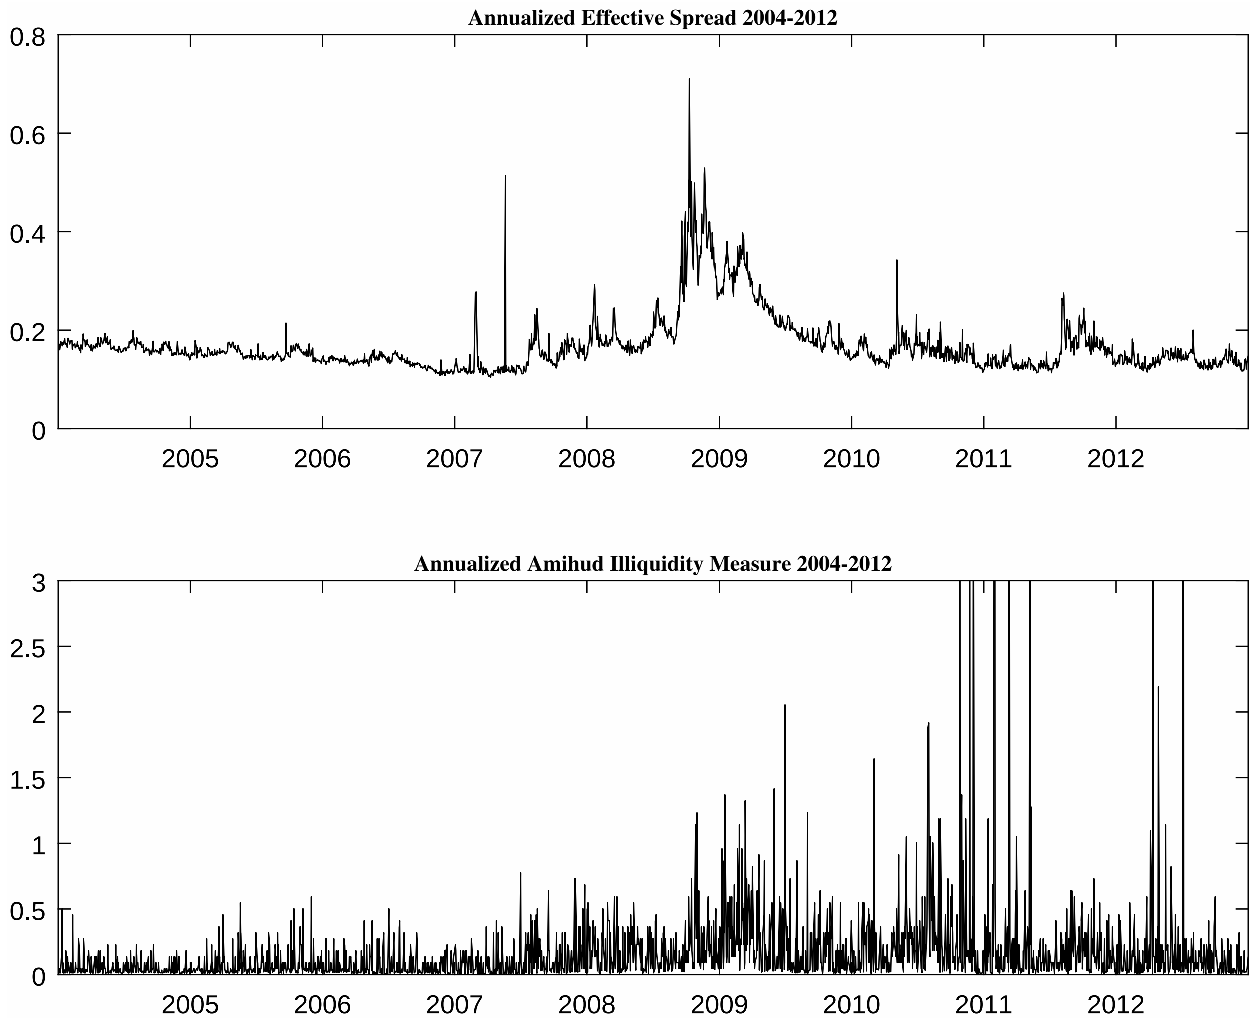

Notes: We plot the alternative measure of daily market illiquidity from January 2, 2004 to December 31, 2012. The top panel plots the annualized effective spread measure and the bottom panel plots the Amihud's (2002) illiquidity measure. We compute the daily Amuhud illiquidity measure for the aggregate stock market as the equally-weighted average Amihud illiquidity measure of all securities constituting the S&P 500 index on each day. An individual firm's daily illiquidity measure is calculated as the absolute daily return divided by the daily dollar volume, $|r_{i,t}|/\mathrm{DVol}_{i,t}$. We normalize the Amuhud illiquidity measure to have same in-sample mean as the illquidity measure that we calculated using intraday effective spreads.

Figure 9: Model-implied Variance Risk Premium

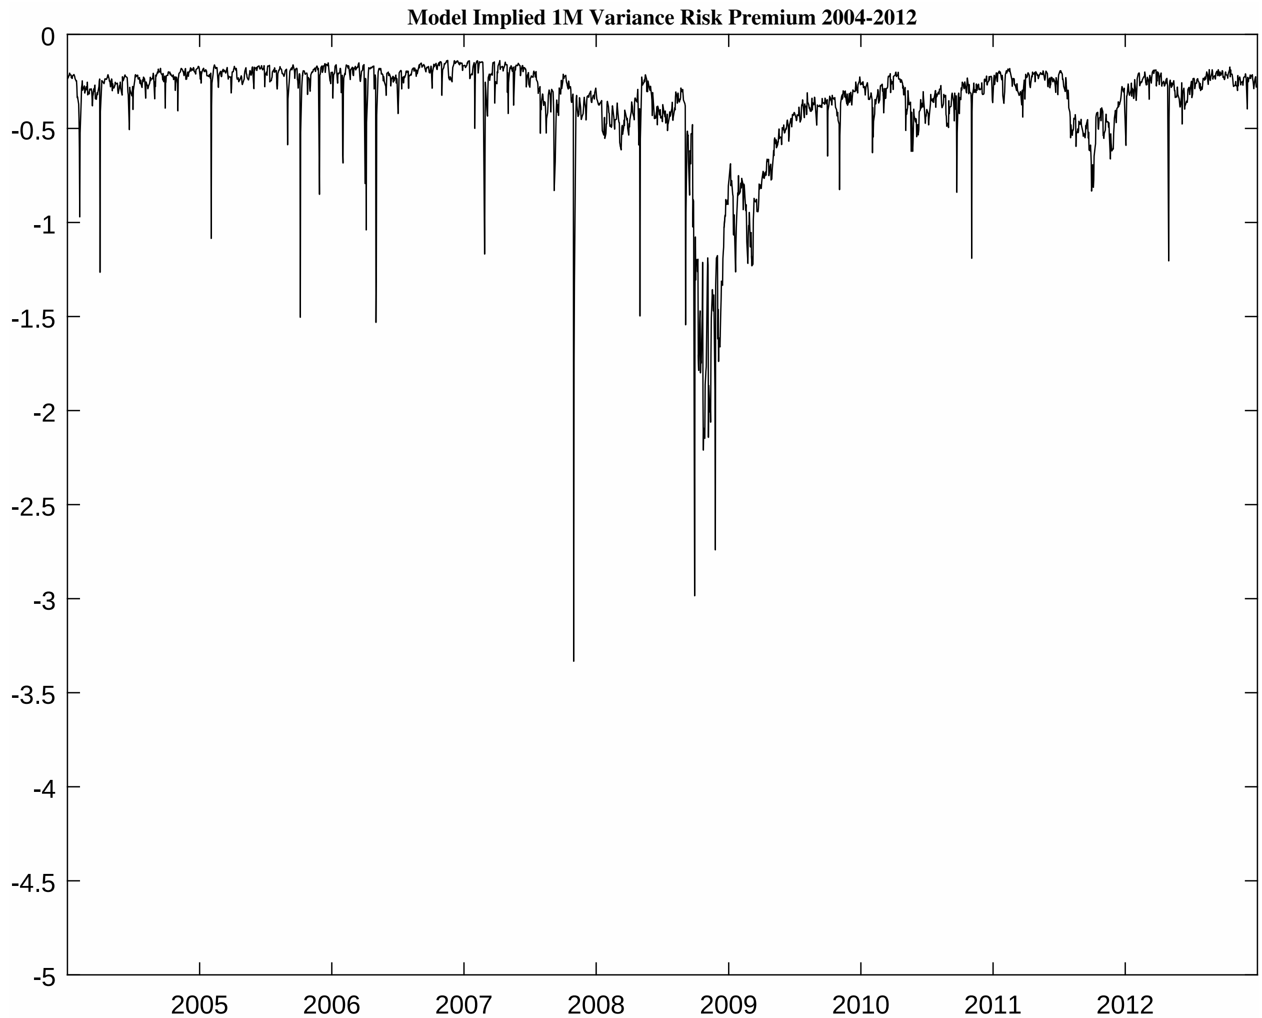

Notes: We plot the model-implied one-month variance risk premium (1M VRP) from January 2, 2004 to December 31, 2012. 1M VRP is computed as the model implied difference between expected total quadratic variation under $\mathbb{P}$ and $\mathbb{Q}$ measures (1M VRP $= (E_t^{\mathbb{P}}[QV_{t,t+1M}] - E_t^{\mathbb{Q}}[QV_{t,t+1M}]) \times 12 \times 100$). All numbers are annualized then multiplied by 100.