Modeling Higher Moments and Risk Premiums for S&P 500 Returns

November 17, 2023

Abstract

Using joint estimation on a large sample of index option prices and the underlying returns, we study how multifactor models capture time-series and cross-sectional patterns in option prices through improved modeling of the dynamics of the first four moments of the return distribution. Including a second and especially a third stochastic volatility factor greatly improves option fit, and the resulting time series of skewness and kurtosis better match non-parametric benchmarks. The third volatility factor is critical in generating larger and more variable skewness and kurtosis risk premiums. Return jumps provide more modest improvements in option fit and a higher equity risk premium, but their impact on higher moment risk premiums is small. All models we investigate struggle to match the unconditional term structure of risk-neutral skewness and kurtosis.

JEL Classification: G12

Keywords: Skewness; Kurtosis; Risk Premiums; Option Valuation; Multi-Factor Model; Particle MCMC.

1 Introduction

A long-standing objective of the option pricing literature has been the development of dynamic models that not only price the cross-section of options with different moneyness and maturities, but also provide consistent prices over time as a function of (latent) state variables. Models that satisfy these criteria are complex, and their estimation is time-intensive. Moreover, Bates (2003) points out that statistical power to distinguish between these models increases when using both returns and options in estimation, because this approach highlights the shortcomings of misspecified models. It is therefore preferable to test and compare these models using long time series of returns and large option panels.

A substantial part of the existing literature has tried to accomplish this difficult task using models with stochastic volatility (SV) and (correlated) jumps in returns and volatility. This literature has evolved towards more complex models with multiple factors. For instance, Ait-Sahalia et al. (2020) and Bardgett et al. (2019) estimate models with two SV factors and jumps in returns and volatility. Andersen et al. (2015b) propose a model with those four factors and an additional tail factor.[1] However, the literature does not yet contain definitive answers on many important questions, such as: How many factors are required, and what is the incremental explanatory power provided by successive factors?; What are the time-series properties and economic content of these factors?; Is it preferable to add an SV or a jump factor to a given one- or two-factor model? Furthermore, while it is well-understood that deviations from normality in the return distribution are critical, and that it is important to allow for differences in the tails of the physical and risk-neutral distributions, the literature has not systematically studied how available models differ in their ability to capture the skewness and kurtosis in the options and return data, and how improved modeling of skewness and kurtosis affects option fit. Also, while the literature has studied variance risk premiums in existing models, we know little about their implications for skewness and kurtosis risk premiums.

The objective of this paper is to address some of these questions by providing a systematic empirical comparison of a large class of option pricing models using S&P 500 options and

[1] For seminal early empirical studies of more parsimonious models, see for example Andersen et al. (2002), Bakshi et al. (1997), Bates (1996), Eraker (2004), and Pan (2002). Duffie et al. (2000) show how to value various derivatives based on these dynamics.

returns, and by documenting the implications of these models for higher moments and higher moment risk premiums. We study these risk premiums based on model estimates from the joint likelihood of options and returns, which is very challenging when using long time series and large cross-sections of options. We rely on advances in estimation methodology, using the particle MCMC framework developed in Andrieu et al. (2010) and the state variable quantile method proposed in Dufays et al. (2023). Our implementation allows for more efficient estimation of complex stochastic volatility models with jumps. We investigate models with up to three SV factors and return jumps. Our most important conclusion is that additional SV factors greatly improve model performance. The third SV factor is critical. It not only provides additional improvements in option fit, but also generates time series patterns in conditional skewness and kurtosis that better match the non-parametric Bakshi et al. (2003) benchmark. A three-factor SV model is also able to generate skewness and kurtosis risk premiums that are much larger and more variable compared to one- and two-factor models. We assess model performance in- and out-of-sample based on statistical as well as economic criteria. In contrast to additional SV factors, jumps in returns provide more modest gains in option fit. They yield higher equity risk premiums, but their impact on the term structure and time series of skewness and kurtosis, and the resulting risk premiums, is small. We conclude that multifactor SV factors can accommodate large negative (and positive) returns, and that they are able to accommodate conditional skewness and kurtosis in the return distribution under the physical and risk-neutral measures. This conclusion is subject to the caveat that we estimate rather simple return jumps with constant jump intensity. It is possible that more flexible jump processes with time-varying intensities, different conditional distributions, or (correlated) return and volatility jumps may perform (much) better. However, consistent with the existing literature, we encounter significant identification problems even with simple return jumps.[2]

Finally, while multifactor SV models, and a three-factor model in particular, provide better option fit and improved modeling of the levels and dynamics of the higher moments, they still exhibit significant shortcomings. Most importantly, these models are unable to match the unconditional term structure of skewness and kurtosis. Presumably this also implies

[2]Also note that our results are application-specific. The existing literature has shown convincingly that jump processes are especially useful at shorter maturities (Bollerslev and Todorov, 2011, 2014; Bollerslev et al., 2015). Jumps may also play a more important role when modeling single-name equity options or high-frequency option prices (Bates, 2019).

that their modeling of higher moment risk premiums falls short, but this is more difficult to ascertain empirically. These remaining shortcomings may prove difficult to address with additional SV factors or existing jump parameterizations, highlighting the need for more realistic dynamics and new modeling approaches.

The paper proceeds as follows. Section 2 presents the SV models, and discusses the data and the estimation procedure. Section 3 presents model estimates and discusses option fit, while Sections 4 and 5 discuss the SV models' implications for the higher moments of the returns distribution and the risk premiums associated with these moments, respectively. Section 6 discusses the performance of models with return jumps. Section 7 concludes. Technical material is collected in an Online Appendix.

2 Models and Data

We present the stochastic volatility models we study and briefly discuss the existing literature. We then summarize the properties of the index option sample and the underlying returns.

2.1 Multi-Factor Stochastic Volatility (SV) Models

The early literature on option pricing in SV models typically considered a single volatility factor, but there is a growing literature that studies models with multiple volatility factors.[3] Additional volatility factors provide additional flexibility to model the longer-horizon return distribution. The literature also combines multiple SV factors with other factors, such as jumps in returns and/or volatility. Because it is not always straightforward to identify the role different factors play in the empirical performance of complex multi-factor models, we first study the implications of a second and third SV factor on the moments of the return distribution and the resulting option fit and risk premiums. We study the impact of jump factors in Section 6. We use a canonical N-factor square-root stochastic volatility model that

[3]See, among others, Andersen et al. (2002), Bakshi et al. (1997), Bates (1996), Chernov and Ghysels (2000), Eraker (2004), and Pan (2002) on the empirical performance of models with a single volatility factor. For models with multiple volatility factors, see for example Ait-Sahalia et al. (2020), Andersen et al. (2015a,b, 2017, 2020), Bardgett et al. (2019), Bates (2000), and Christoffersen et al. (2009).

is an extension of the affine one-factor model in Heston (1993)

$$\frac{dS_t}{S_t} = (r_t - \delta_t + \eta_s \sum_{i=1}^{N} V_{it})dt + \sum_{i=1}^{N} \sqrt{V_{it}} dZ_{it}, \quad (1)$$

$$dV_{it} = \kappa_i(\theta_i - V_{it})dt + \sigma_i \sqrt{V_{it}} dW_{it}, \quad (2)$$

$$\text{Corr}(dW_{it}, dZ_{it}) = \rho_i dt. \quad (3)$$

where S_t is the index level, r_t is the risk-free rate, δ_t is the dividend yield.[4] For each variance factor i, η_{is} measures the contribution to the equity risk premium, κ_i denotes the speed of mean reversion, θ_i the unconditional mean variance, and σ_i determines the variance of variance. dZ_{it} and dW_{it} are Brownian motions with $corr(dZ_{it}, dZ_{jt}) = 0$ and $corr(dW_{it}, dW_{jt}) = 0, i \neq j$. The variance of the stock return is given by $\text{Var}_t(dS_t/S_t) = \sum_{i=1}^{N} V_{it} = \tilde{V}_t$. We refer to the N-factor SV model as SVN.

Consistent with most of the existing literature, we assume that the asset dynamic has the same functional form under the risk neutral measure Q:

$$\frac{dS_t}{S_t} = (r_t - \delta_t)dt + \sum_{i=1}^{N} \sqrt{V_{it}} dZ_{it}^Q, \quad (4)$$

$$dV_{it} = \kappa_i^Q(\theta_i^Q - V_{it})dt + \sigma_i \sqrt{V_{it}} dW_{it}^Q, \quad (5)$$

$$\text{Corr}(dW_{it}, dZ_{it}) = \rho_i dt. \quad (6)$$

where $\kappa_i^Q = \kappa_i - \eta_{iv}$ and $\theta_i^Q = (\kappa_i \theta_i)/\kappa_i^Q$, and the total diffusive variance risk premium is thus equal to $\sum_{i=1}^{N} \eta_{iv} V_{it}$. Denoting $\Theta = (\kappa_1, \theta_1, \sigma_1, \rho_1, \ldots, \kappa_N, \theta_N, \sigma_N, \rho_N, \eta_s, \eta_{1v}, \ldots, \eta_{Nv})$ and $V_t = (V_{1t}, \ldots, V_{Nt})$, the model price of a European call option $C^M(V_t, \Theta)$ with maturity τ and strike price K is given by:

$$C^M(V_t, \Theta) = e^{-r_t \tau} E^Q[\max(S_{t+\tau} - K, 0)]. \quad (7)$$

Because of the affine structure of the models, quasi closed-form solutions for option prices are available. See Appendix A for details. The model price is computed given the current

[4] Most empirical investigations of SV models stay within this affine class to ensure quasi closed-form expressions for European option prices. For studies of non-affine models, see for instance Jones (2003) and Christoffersen et al. (2010).

state V_t and model parameters. We assume that the observed market option price C is equal to the model price plus error:

$$C_{t,h} = C^M_{t,h}(V_t, \Theta) + \varepsilon_{t,h}, \tag{8}$$

where $\varepsilon_{t,h}$ is normally distributed with variance σ_c^2, and $h = 1, 2, .., H_t$, with H_t the total number of options at date $t = 1, \ldots, T$.

2.2 Model Estimation

Our estimation methodology is based on a combination of the particle MCMC framework proposed by Andrieu, Doucet, and Holenstein (2010) and the state variable quantile method introduced in Dufays et al. (2023). The particle MCMC framework leverages an MCMC algorithm to estimate the parameters and latent states. The state variable quantile method facilitates the filtering of latent states within a large option panel data context. We now describe the return and option likelihoods, and the priors which serve as the basis for the MCMC algorithm. Appendix A.3 provides a more detailed discussion.

To estimate the models using returns and options, it is necessary to define their joint likelihood. To balance the relative weight of returns and options, we normalize the option likelihood by taking the H_t root. This normalization also allows us to focus on the composition of the option data, abstracting from any variation in H_t over time.[5] Using a standard Euler discretization, the joint likelihood given the spot variances V_t can be written as:

$$f(R_{t+1}, C_t | V_t, \Theta) = f(R_{t+1} | V_t, \Theta) \left(\prod_{h=1}^{H_t} f(C_{t,h} | C^M_{t,h}(V_t, \Theta)) \right)^{1/H_t}, \tag{9}$$

where $R_{t+1} = \ln(S_{t+1}/S_t)$. The return likelihood can be expressed as:

$$f(R_{t+1} | V_t, \Theta) = \frac{1}{\sqrt{2\pi V_t}} \exp \left\{ -\frac{1}{2} \frac{\left[R_{t+1} - (r_t - \delta_t + (\eta_s - \frac{1}{2}) \sum_{i=1}^N V_{i,t}) \right]^2}{V_t} \right\}. \tag{10}$$

[5] While this rescaling is ad-hoc, the specification of any joint likelihood consisting of option prices and the underlying returns lacks theoretical guidance.

From equation (8), the option likelihood is given by:

$$f(C_{t,h}|C^M_{t,h}(V_t,\Theta)) = \left(\frac{1}{\sqrt{2\pi\sigma_c^2}}\right)\exp\left(-\frac{\sum_{h=1}^{H_t}(C_{t,h}-C^M_{t,h}(V_t,\Theta))^2}{2\sigma_c^2 H_t}\right). \tag{11}$$

We use the following prior distributions: $\ln(\frac{\rho_i}{1-\rho_i}) \sim N(-1.5, 0.5)$ and $\eta_s \sim N(2.5, 0.5)$. The prior on ρ_i is diffuse and encompasses correlation estimates in the existing literature. The prior on η_s is based on simple moment matching, but allows for outliers that might emerge during the MCMC iterations. It is well-known that this parameter is inherently difficult to estimate precisely (**?**).

2.3 Return and Option Data

To estimate the models, we use S&P 500 returns and option prices for the period January 4, 1996 to December 31, 2019, a total of 6037 trading days. The need to use a long sample to identify return dynamics for option valuation is emphasized by, among others, Eraker, Johannes, and Polson (2003), Eraker (2004) and Broadie, Chernov, and Johannes (2007). We obtain index returns and risk-free rates from CRSP, and option prices, zero coupon yields, and dividend yields from OptionMetrics. Panel A of Table 1 provides descriptive statistics for the index returns.

We use both put and call index options and impose the following standard filters: (1) discard options with fewer than 5 days and more than 365 days to maturity; (2) discard options with implied volatility less than 5% and greater than 150%.; (3) discard options with trading volume or open interest less than 5 contracts; (4) discard options with quotes that suggest data errors: options for which the best bid exceeds the best offer, options with a zero bid price, and options with negative put-call parity implied price; (5) discard options with price less than 50 cents.

After imposing these filters, the resulting data set contains 3,052,205 option contracts. The number of traded option series has increased exponentially, especially since 2005, with almost forty times more contracts in 2019 compared to 1996. To avoid the impact of this increase on parameter estimation of the considered models, we create a more balanced panel. First we convert put prices into call prices based on put-call parity. We then use six moneyness bins and five maturity bins to adequately represent the observed daily option surfaces.

Moneyness is defined as strike price divided by index price (K/S). For each moneyness-maturity bin, we include only the most liquid option, defined as the option with the highest trading volume. These are for the most part OTM puts and calls. The data set thus has thirty options per day unless options are not available for certain bins.[6] This procedure yields a data set with 162,761 option contracts for parameter estimation. Panel B of Table 1 provides sample sizes for these moneyness-maturity bins, as well as average option prices and implied volatilities. We use the remaining 2,889,444 option contracts that are not used for model estimation to gauge the 1996-2019 models' out-of sample option fit. In addition, we also use the 2020-2021 OptionMetrics data option prices and apply the same filters to create a balanced panel that we use in an out-of-sample forecasting exercise.

3 Parameter Estimates and Model Fit

All estimates are obtained based on the maximization of the joint likelihood in equation (9) that contains S&P 500 returns and option data. We first discuss the model parameter estimates, and briefly compare them with the existing literature. We also discuss the properties of the (time series of the) stochastic variance. We then analyze the models' pricing performance and discuss in- and out-of-sample option fit.

3.1 Parameter Estimates and Filtered Latent Factors

3.1.1 The One-Factor SV Model

We start with the parameter estimates for the SV1 model in Table 2. This model has been extensively tested using returns, option prices, or returns and options jointly. Our results are mostly consistent with existing estimates. The long run variance θ_1 is close to the sample variance, which is equal to 0.038. The variance process is very persistent ($\kappa_1 = 1.403$). Our estimate of ρ (-0.909) is more negative compared to most other studies, but see Andersen et al. (2015a) for an estimate that is similar to ours. It is well-known that the risk-neutral index return distribution is more negatively skewed than the physical index return distribution (Bakshi et al., 1997; Bates, 2000). A potential explanation for

[6] In the early years of the sample we do not have many observations for short-maturity out-of-the-money calls and/or in-the-money puts.

our more negative estimate of ρ is that, like Andersen et al. (2015a), our option panels are larger compared to most existing work, which puts more emphasis on the risk-neutral return distribution.

The estimates of the risk premium parameters η_s and η_{v1} have the expected positive signs, which implies a positive equity risk premium and a positive $(Q-P)$ variance risk premium. The magnitude of the η_s parameter is plausible; the average of the resulting equity risk premium $\eta_s V_t$ is 0.096, close to the sample average of 0.092. It is more difficult to evaluate the plausibility of the magnitude of the η_{v1} parameter and the resulting variance risk premium, because these quantities are not directly observable from the data. Our estimate of η_{v1} implies that the variance risk premium (VRP) is on average 78% of the physical variance. Most parametric estimates of the SV1 model report larger variance risk premiums, but many of these estimates have much larger standard errors (Eraker, 2004; Hurn et al., 2015; Pan, 2002). Our lower posterior standard deviation may be due to our use of large option panels.

3.1.2 Multifactor SV Models

The BIC criterion in Table 2 provides very strong statistical support for the SV2 model and the second variance factor. Panel B of Figure 1 plots the second filtered variance factor in the SV2 model. Table 2 shows that this second factor is more persistent than the sole variance factor in the SV1 model ($0.434 < 1.403$). The first factor becomes less persistent ($4.705 > 1.403$). This finding is consistent with the existing evidence on two-factor SV models, regardless of whether they are estimated using the underlying index returns or index option data, see for instance Chernov et al. (2003) and Christoffersen et al. (2009). The persistence of the sole factor in the SV1 model strikes a compromise between the different mean reversions of the two factors. Panel A of Figure 1 shows that the path of the total filtered variance for the SV2 model (in green) is highly correlated (0.964) with the path of the filtered variance in the SV1 model (in yellow). However, while the average variance in the two models is very similar, the higher moments (standard deviation, skewness and kurtosis) of the filtered variance in the SV2 model exceed that of the SV1 model.

The BIC criterion in Table 2 also provides very strong statistical support for the SV3 model and the third variance factor. Panel B (C) of Figure 1 plots the second (third) filtered variance factor in the SV3 model. The largest of the SV3 variance factors ($i = 1$) is the least persistent of the three ($\kappa_1 = 0.981$). Table 2 indicates that the two other (smaller)

factors are very persistent, as is also evident from Figure 1. The resulting SV3 variance path is slightly less persistent (0.986) compared to the SV1 (0.992) and SV2 (0.990) models. The option data require persistent state variables, but more flexible models result in slightly lower overall persistence. Despite these differences in persistence, the overall variance path of the SV3 model is very highly correlated with that of the SV1 model (0.964) and the SV2 model (0.997). The SV3 model variance's standard deviation, skewness and kurtosis are much higher than in the SV1 model, and of the same order of magnitude as the SV2 model. As an illustration of these differences, the maximum variance for all three models occurs in November 2008 (the financial crisis), but it is considerably lower for the SV1 model (0.336) compared to the SV2 (0.490) and SV3 (0.538) models.

The higher moments of the return distribution are determined by the estimates of the ρ_i and σ_i parameters. All estimates of the σ_i parameters in the SV2 and SV3 models are positive by construction, and all ρ_i estimates are negative. In the SV2 model, both estimates of the variance risk premium parameter are positive, guaranteeing a positive variance risk premium. However, in the SV3 model η_{v1} and η_{v2} are positive but η_{v3} is negative. We estimate the SV2 and SV3 models with a single equity premium parameter η_s; allowing each factor to have a separate equity premium parameter does not affect the results.

3.2 Option Fit

Table 2 reports the (in-sample) option fit for the entire sample in the row marked by σ_c. More complex models with additional factors greatly improve in-sample option fit. Adding a second volatility factor lowers the average pricing error by \$0.64, from \$3.31 (for the SV1 model) to \$2.67 (for the SV2 model). The RMSE is thus lowered by almost 20 percent (\$0.64/\$3.31). The SV3 model further lowers the average pricing error to \$2.18. The posterior standard deviations are small compared to the posterior means, suggesting that these differences in fit are not only economically but also statistically significant. Table 3 reports on (in-sample) model fit by moneyness and maturity. The most significant improvements in model fit in the SV2 and SV3 models are for longer-maturity options.

Panel A of Figure 2 plots the models' in-sample option fit over time. The more complex models consistently outperform the SV1 model, but the dollar outperformance is relatively larger in the 2008-2009 crisis period. The outperformance also increases at the end of the

sample, which leads up to the Covid-19 recession.

Our estimation setup is designed to avoid overfitting. The parameter estimates are based on returns as well as option prices. Moreover, we filter the state variables and the improved fit does not result from the introduction of additional parameters that are ancillary to the model. Nevertheless, the SV2 and SV3 models' improvements in in-sample performance are perhaps not surprising. We therefore assess model performance in two out-of-sample exercises. First, we use the estimates from the 1996-2019 sample to price options in 2020 and 2021.[7] Table A.1 in the Appendix presents the results. In a second out-of-sample exercise, we consider all available (2,889,444) options in the 1996-2019 OptionMetrics data that are not used in estimation. Table A.2 in the Appendix reports the results. Both out-of-sample exercises confirm the two main conclusions from the in-sample results in Table 3: additional factors significantly improve option fit. Panels B and C of Figure 2 illustrate the differences in fit over time, and confirm that the superior performance of the multifactor models is not due to any particular time period.

4 The Higher Moments of the Return Distribution

Table 1 confirms that index returns are highly non-normal, and time variation in the higher moments of the return distribution critically affects option fit. We therefore investigate (differences in) model-implied higher return moments. We first analyze the time series of model properties that are related to conditional skewness and kurtosis. Next we discuss our estimates of the time series of the level and term structure of conditional model-implied variance, skewness and kurtosis.

4.1 Spot Model Properties

Computing the conditional higher moments is relatively straightforward, but the analytical expressions are lengthy and do not provide much intuition. We therefore start by analyzing some related model properties that are easy to compute and that conveniently capture the intuition for why multifactor models are required to match the data. These

[7] We apply the particle filter using the posterior means of the parameters and the starting variance given by V_t. Each day, we update the filtered variance and we store the mean squared pricing errors given by the mode of the filtered distribution of the spot volatilities.

properties capture various instantaneous or "spot" properties of the models. We discuss leverage correlation, variance-of-variance, and the correlation between returns and variance-of-variance.

4.1.1 Leverage Correlation

We start with the models' implications for $Corr_t\left(\frac{dS_t}{S_t}, dV_t\right)$, which is related to conditional skewness. Following Christoffersen et al. (2009), we refer to this quantity as leverage correlation. In the SV1 model, the leverage correlation is equal to ρ and thus constant (-0.909). The factor correlations ρ_i are also constant in the N-factor stochastic volatility model in equations (1)-(3). They are all negative and large, highlighting the importance of allowing for negative skewness in the index return distribution. The overall leverage correlation in multifactor models is time-varying:

$$\text{Corr}_t\left(\frac{dS_t}{S_t}, dV_t\right) = \frac{\sum_{i=1}^{N} \sigma_i \rho_i V_{i,t}}{\sqrt{\sum_{i=1}^{N} V_{i,t}} \sqrt{\sum_{i=1}^{N} \sigma_i^2 V_{i,t}}}. \tag{12}$$

Because the left tail of the return distribution critically affects the cross-section of option prices, this additional flexibility is very important. Panel A of Figure 3 plots the leverage correlation for the SV1, SV2, and SV3 models. The leverage correlation in the SV2 model fluctuates between -0.925 and -0.827. The average leverage correlation in the SV3 model is -0.711, much smaller in absolute value compared to the SV1 model (-0.909) and the SV2 model (-0.847). The leverage correlation in the SV3 model is also much more variable over time; it fluctuates between -0.332 and -0.954.[8]

4.1.2 Variance of Variance

Leverage correlations provide intuition about conditional skewness. Similarly, the conditional variance-of-variance provides intuition about the models' ability to capture conditional kurtosis. In the SV1 model, the conditional variance-of-variance is once again quite restric-

[8] Appendix B discusses in more detail the role of the third variance factor in the time variation of the leverage correlation in equation (12)

tive: it is equal to $\sigma^2 V_t$. In multifactor SV models, the variance-of-variance is more flexible:

$$Var_t(\sum_{i=1}^{N} dV_{i,t}) = \sum_{i=1}^{N} \sigma_i^2 V_{i,t}. \tag{13}$$

Panel B of Figure 4 plots the variance-of-variance for the SV1 model, which is by design perfectly correlated with the variance. Panel C of Figure 4 plots the ratio of the variance-of-variance in the SV2 and SV1 models, as well as the SV3/SV1 ratio. On average the variance-of-variance in the SV2 model far exceeds that of the SV1 model (the average ratio is 1.949), and the average SV3 variance-of-variance is even higher (the average SV3/SV1 ratio is 2.276). For all three models, variance-of-variance increases when variance increases and peaks in the financial crisis, but the SV2/SV1 and SV3/SV1 ratios increase more sharply when variance increases. While for the leverage correlation in Figure 3, the SV3 model is the outlier and the SV1 and SV2 models are similar, the variance-of-variance path strongly differs across all three models.

4.1.3 The Correlation between Returns and Variance-of-Variance

Panel A of Figure 4 plots the time series of the correlation between returns and the variance-of-variance, which is given by:

$$Corr_t(\frac{dS_t}{S_t}, \sum_{i=1}^{N} \sigma_i^2 dV_{i,t}) = \frac{\sum_{i=1}^{N} \rho_i \sigma_i^3 V_{i,t}}{\sqrt{(\sum_{i=1}^{N} V_{i,t})(\sum_{i=1}^{N} \sigma_i^6 V_{i,t})}} \tag{14}$$

For the SV1 model, this correlation is once again the constant ρ. For the SV2 model, the correlation in equation (14) is smaller (in absolute value) than the SV1 constant correlation on each day in the sample. For the SV3 model, it is smaller most of the time.

In summary, the additional flexibility in the multifactor models results in very different implications for the time series of leverage correlation, variance-of-variance, and the correlation between returns and variance. This additional flexibility should enhance the models' ability to capture the magnitude and time variation of the higher moments. We now proceed to an analysis of these moments.

4.2 Model-Implied Higher Moments

Due to the affine nature of the characteristic function (CF) in SVN models, we can compute closed-form solutions for the return moments. This is a simple extension of the work of Das and Sundaram (1999) on one-factor models. At each time t, we compute the second, third and fourth centered risk-neutral moment of the log returns for horizon $t+\tau$, defined as $R_{t,t+\tau} =: \ln S_{t+\tau} - \ln S_t$:

$$VAR(R_{t,t+\tau}|V_t) = \sum_{i=1}^{N} \left[A_i^{(2)}(\Theta,\tau) + D_i^{(2)}(\Theta,\tau) V_{it} \right], \tag{15}$$

$$E\Big((R_{t,t+\tau} - E(R_{t,t+\tau}|V_t))^3 | V_t \Big) = \sum_{i=1}^{N} \left[A_i^{(3)}(\Theta,\tau) + D_i^{(3)}(\Theta,\tau) V_{it} \right], \tag{16}$$

$$E\Big((R_{t,t+\tau} - E(R_{t,t+\tau}|V_t))^4 | V_t \Big) = \sum_{i=1}^{N} \left[A_i^{(4)}(\Theta,\tau) + D_i^{(4)}(\Theta,\tau) V_{it} \right] + 3VAR(r_{t,t+\tau}|V_t)^2. \tag{17}$$

We can compute these expressions both under the physical (P) and risk-neutral (Q) measures. While the moments are simple functions of the state variables, the parameters $A_i^{(j)}(\Theta,\tau)$ and $D_i^{(j)}(\Theta,\tau)$ for $i = 1,\ldots,N$ and $j = 2,3,4$ are lengthy and tedious functions of the model parameters. Appendix C provides Matlab code and presents the parameterization of the moments evaluated at the posterior mean of the parameters.

4.3 Time-Variation in Conditional Higher Moments

We start by studying the time-series variation in the level of the conditional risk-neutral higher moments. We report on the 30-day moments; the time-series patterns for other maturities are very similar. The first three columns of Figure 5 plot the time-series of the variance, skewness and kurtosis for the SV1, SV2, and SV3 models. To assess how the model-implied moments match the data, the fourth column reports nonparametric measures of risk-neutral variance, skewness and kurtosis, computed following Bakshi and Madan (2000) and Bakshi et al. (2003), henceforth BKM. Appendix D discusses the computation of the BKM moments.

All three model-implied variance series are similar to the BKM variance time series, but

the SV3 model more closely captures the very high 30-day BKM variance in the financial crisis.[9] Model-implied skewness is negative on every day in the sample for all models, while model-implied kurtosis is positive. Moreover, skewness becomes more negative when kurtosis becomes more positive. Model-implied skewness and kurtosis are thus complements; they jointly capture higher tail risk in some periods. These increases in kurtosis and (the absolute value of) skewness mainly occur in two periods: the first time in the May 2005 to May 2007 period, and the second time in the March 2017 to January 2018 period.

These two periods with higher (absolute value of) skewness and kurtosis are low-volatility periods, and we do not see increases in tail risk in the financial crisis, when the variance sharply increases. These stylized facts are not model-driven. While the BKM figures in the fourth row of Figure 5 are much more noisy, they show a similar pattern. These findings are perhaps a bit surprising, and to the best of our knowledge they have not been extensively discussed in the existing literature. Figure A.3 in the Appendix plots the non-standardized moments instead and shows that these findings can be traced back to the time series patterns in the variance.[10] During high variance periods, non-standardized skewness and kurtosis are higher than usual, but standardized skewness and kurtosis are lower.

Do additional volatility factors improve the modeling of the time-series variation in higher moments? Panel C of Table 4 reports the correlations between the model-implied and BKM moments. The variance correlations are very high, even for the SV1 model. For skewness, the correlation is 35% for the SV1 model and 33% for the SV2 model, but it increases to 63% for the SV3 model. These improvements are due to the SV3 model's ability to accommodate fluctuations in tail risk. Skewness in the SV3 model fluctuates between -0.340 and -2.682 and kurtosis between 3.374 and 16.326; these ranges are much narrower for the SV2 and especially the SV1 models.[11] While the third SV factor also improves the modeling of kurtosis, the improvements are more significant for skewness.

We conclude that based on the correlation with the nonparametric BKM models, the SV3

[9]Note that the model-implied and BKM variance time series are also very similar to that of the spot variance in Panel A of Figure 1.

[10]The non-standardized third and fourth moments are highly correlated with the spot variance; the standardized moments much less so. For example, for the SV1 model, the correlation between the spot variance and the non-standardized third moment is -1 by definition. Even in the SV3 model, this correlation is very high (-0.956).

[11]Note that this finding confirms the intuition from the plots of the leverage correlation in Panel A of Figure 3 and the correlation between the return and the variance of variance in Panel A of Figure 4.

model clearly outperforms the SV1 and SV2 models. The importance of the third volatility factor is also evidenced by the correlations between the higher moments. The correlation between the BKM variance and skewness (kurtosis) is 0.177 (-0.160). For the SV3 model, the correlation between the spot variance and the risk-neutral skewness (kurtosis) are 0.228 and -0.270 respectively, similar to the BKM correlations. For the SV2 model, these correlations are 0.714 and -0.609 respectively; for the SV1 model, they are 0.774 and -0.662. Additional variance factors can address the inherent limitations of the SV1 model and capture the correlations between the higher moments. However, consistent with the evidence in Panel C of Table 4, the second volatility factor does not lead to substantial improvements, and a third factor is required.

4.4 The Term Structure of Model-Implied Higher Moments

Panel C of Table 4 also reports the correlations between the model-implied and BKM 180-day moments, and the conclusions are similar. In fact, similar conclusions conclusions obtain for horizons up to one year. We conclude that the SV3 model significantly outperforms the SV1 and SV2 models in modeling the time-variation in higher moments. Can it also match the unconditional term structure? Panel A of Table 4 reports the time-series averages of the higher moments for the 30- and 180-day horizon. [12] Average skewness (kurtosis) is negative (positive) at all maturities and becomes more negative (positive) as a function of maturity (up to one year) in all three models. At maturities less than one year, the kurtosis of the SV3 model exceeds that of the SV2 model, which in turn exceeds that of the SV1 model. The same remark applies to the absolute value of the skewness, except for very short maturities, where SV2 skewness is most negative. Variance (annualized) also increases with maturity and the number of volatility factors.

All models fail to capture important aspects of the (nonparametric) BKM term structure. The average BKM variance term structure is very slightly upward sloping; all three models underpredict the 30-day variance and imply a steeper slope. Similarly, the models cannot capture the BKM term structures of (the absolute value of) skewness and kurtosis. If anything, these term structures are downward sloping for the BKM moments, but the

[12] Figure A.4 in the Appendix plots the average model-implied term structure of the risk-neutral moments up to one year, and Figure A.6 in the Appendix plots the time series of the difference between the 180-day and 30-day risk-neutral moments.

models predict an upward sloping term structure. Because the skewness and kurtosis in the SV3 model exceed the SV1 and SV2 levels at all maturities, the SV3 model fits the 30-day moments better. It better captures the time-variation at the 180-day horizon, but on average it overpredicts tail risk. This may be due to model misspecification. Another potential explanation is the composition of the option sample; it is possible that the SV3 model can match these moments when more long-term options are included, or when we separately estimate the model for long-maturity options only.

4.5 The Third Variance Factor and Higher Moments

The third variance factors is especially important for explaining the time series of the higher moments. Figure A.5 in the Appendix plots the impact of a one standard deviation change to the posterior mean of the third factor on the risk-neutral moments. The impact is very high in those periods when tail risk is high, and the effect is economically large compared to the average skewness and kurtosis, as well as the conditional skewness and kurtosis during those periods.

5 Moment Risk Premiums

We now discuss how the models differ with respect to their implications for the moment risk premiums. We first discuss the time series and term structure of the equity, variance, skewness and kurtosis risk premiums separately. Then we analyze the co-movement between these risk premiums and the relation between the model factors and risk premiums.

We report annualized numbers for the first two moments. Rather than reporting on Q minus P for all moments or vice versa, we consistently construct risk premiums that are predicted by theory to be positive; Thus, for the first and third moment, we report $P - Q$; for the second and fourth moment, we report $Q - P$. This notational convention facilitates the discussion of risk premiums and specifically the relation between risk premiums.

5.1 The Equity and Variance Risk Premiums

For the first and second moments, it is straightforward to compute the time series of the so-called spot risk premiums. They are simple functions of the state variables and the

parameters and have been extensively studied in the existing literature. For instance, in the SV1 model, the spot equity premium at t is equal to $\eta_s V_t$ and the spot variance risk premium $(Q-P)$ is equal to $\eta_v V_t$. Given the positive estimates of η_s and η_v in Table 2, the SV1 model yields a positive spot equity premium and a negative variance risk premium at each time t. After we discuss the spot and variance risk premiums, we presents results for the integrated risk premiums, which are horizon-dependent.

5.1.1 The Spot Equity Risk Premium

For the SV models in Table 2, the estimate of the equity risk premium parameter η_s slightly changes with the number of factors, from 2.548 in the SV1 model to 2.706 in the SV2 model and 2.606 in the SV3 model. This results in a economically small differences in the diffusive spot equity premiums for these three models; the differences are also small in view of the posterior standard deviations of the estimated parameters.

5.1.2 The Spot Variance Risk Premium

Panel B of Figure 3 plots the filtered time path of the spot VRP for the SV1 model. The spot VRP is perfectly correlated with the filtered variance, which is very restrictive. Panel C of Figure 3 plots the ratio of the SV2 spot VRP versus that of the SV1 in blue and the SV3-SV1 ratio in red. The SV3 average spot VRP is 0.047, compared to 0.040 and 0.029 respectively for the SV2 and the SV1 models. These differences are economically significant. The SV2 VRP and especially the SV3 VRP are also more variable than the SV1 VRP, and contain more outliers. The SV2 and SV3 models generate a (much) higher spot VRP than the SV1 model in high-volatility periods and the SV3 model allows for a (much) smaller spot VRP in low volatility periods. Panel C of Figure 3 shows that the SV3 VRP strongly deviates from the two others, and clarifies that these differences mainly occur prior to the financial crisis. At the start of the sample period, the SV3 VRP briefly becomes negative. This can be traced back to the negative estimate of η_{v1} in Table 2.[13] Despite these differences, the time

[13]The SV3 model can be re-estimated subject to a positive sign restriction on this parameter, but it is unclear if this is desirable. The VRP for the SV3 model is very small or negative from the start of the sample in 1996 to May 1997 and between November 2003 and May 2007, which are both low-volatility periods that should exhibit smaller risk premiums.

series of the model VRPs are very highly correlated.[14] Moreover, the correlation between the spot VRPs and the Aruoba-Scotti-Diebold (ADS) business index (Aruoba et al., 2009) is approximately -50% for all three models.

Are the higher VRPs consistent with the data and the existing literature? It is well-known that it is difficult to precisely estimate the variance risk premium, and the literature does not contain a lot of evidence on the impact of model specification on spot variance risk premiums in multifactor SV models. The most closely related finding in the literature is Ait-Sahalia et al. (2020), who report a (diffusive) spot VRP of 8% for their benchmark model, 1.7% for the SV1 model, and 4% for other nested models. Their estimated variance jump risk premium is small. Our VRP estimates are close to the one they report for nested models, but our estimates are much more similar across models. Bardgett et al. (2019) study related multifactor models, but they do not discuss the magnitude of the spot VRP. The risk premiums in the multifactor models of Andersen et al. (2015a,b) and Gruber et al. (2021) are harder to compare to ours, because they estimate the models under the risk-neutral measure, and subsequently retrieve risk premiums. Bollerslev et al. (2009) report a nonparametric estimate of 2.2% for a sample with a realized variance of 1.8%.

5.1.3 Integrated Equity and Variance Risk Premiums

The integrated equity and variance risk premiums are very highly correlated with the corresponding spot risk premiums, but they are horizon-dependent and can therefore be used to study the unconditional and conditional term structure of risk premiums. They can be defined in different ways. One approach is to use the difference between the physical and risk-neutral variance of the log returns in equation (15). This is the approach we use for the integrated skewness and kurtosis risk premiums below. However, for consistency with existing papers such as Ait-Sahalia et al. (2020) and Bardgett et al. (2019), for the first two moments we instead use expressions based on the quadratic variation of the log-price

[14] The correlation between the SV2 and SV1 VRP paths is 0.954. The correlation between the SV3 and SV1 VRP paths is 0.941 and between SV2 and SV3 it is 0.976.

process:

$$\begin{aligned}
\text{IERP}(t,t+\tau) &= \frac{1}{\tau}\left(E^P\big(\frac{S_{t+\tau}}{S_t}|V_t\big] - E^Q\big(\frac{S_{t+\tau}}{S_t}|V_t\big)\right), \\
&= \eta_s\left(\sum_{i=1}^{N}\theta_i + \tilde{\rho}_i(V_{it}-\theta_i)\right), \qquad (18)\\
\text{IVRP}(t,t+\tau) &= \frac{1}{\tau}\left(E^Q\big(\text{QV}_{t,t+\tau}|V_t\big) - E^P\big(\text{QV}_{t,t+\tau}|V_t\big)\right), \\
&= \sum_{i=1}^{N}\Big((1-\tilde{\rho}_i^Q)\theta_i^Q - \big(1-\tilde{\rho}_i\big)\theta_i + \big(\tilde{\rho}_i^Q - \tilde{\rho}_i\big)V_{it}\Big), \qquad (19)
\end{aligned}$$

where $\tilde{\rho}_i = \frac{(1-\exp(-\kappa_i\tau))}{\kappa_i\tau}$, $\tilde{\rho}_i^Q = \frac{(1-\exp(-\kappa_i^Q\tau))}{\kappa_i^Q\tau}$ for $i = 1,\ldots,N$ and where $\text{QV}_{t,t+\tau}$ denotes the quadratic variation of the log-price process.

Panel B of Table 4 shows that the average IERP for the SV2 and SV3 models slightly exceeds that of the SV1 model. Consistent with the findings of Ait-Sahalia et al. (2020), the variability of the IERP is larger at shorter horizons. For example, the standard deviation of the 30-day IERP in the SV3 model is 0.11, compared to 0.04 for the 180-day IERP.

Figure 6 plots the time series of the 30-day IERP. The dynamics of the IERP are largely determined by the patterns in the diffusive spot variance and consequently they are strongly countercyclical, because the IERP increases when volatility goes up. Figure 7 plots the difference between the 180- and 30-day IERPs. On most days in the sample, the term structure is upward sloping.[15] Panel B of Table 4 also shows that on average the term structure of the IERP is slightly upward sloping, regardless of the number of SV factors. However, in crisis periods the term structure switches signs, and it becomes sharply downward sloping. This captures the intuition that in crisis periods, investors require large equity risk premiums for short-horizon risky investments. The multifactor models, which better fit the option data, accentuate this stylized fact. Somewhat surprisingly, the existing literature on parametric option pricing models does not contain a wealth of evidence on these patterns in the level and term structure slope of the IERP. Our results are consistent with the findings in Ait-Sahalia et al. (2020), who rely on a more complex multifactor model with jumps. We find that this pattern is present in all models, including the simple SV1.

[15]For the SV1 model, the term structure is upward sloping on 64.3% of the sample days. For the SV2 model this is 69.79%, andfor the SV3 model it is 72.63%.

Panel B of Table 4 reports the average level of the IVRP. Consistent with theory and with the results in Ait-Sahalia et al. (2020), the average variance risk premiums, computed as the Q variance minus the P variance, are positive for all maturities in all models. Figure 6 confirms that the SV3 model can accommodate a higher IVRP in crisis periods and a lower IVRP in calm periods. We also confirm the finding in Ait-Sahalia et al. (2020) and Bardgett et al. (2019) that the (annualized) average VRP is increasing (in absolute value) as a function of maturity. The second row in Figure 7 plots the difference between the 180- and 30-day IVRPs. Unlike for the IERPs in the first row, the term structure slope of the IVRP is countercyclical. For the maturities in Figure 7, it remains positive on each day for all models. While this finding that the slope of the IVRP term structure is positive obtains most of the time, it is not ubiquitous.[16] Finally, for a given model and maturity, the IERP and IVRP levels are highly positively correlated. However, the IERP and IVRP term structures are highly negatively correlated.

5.2 The Skewness and Kurtosis Risk Premiums

Panel B of Table 4 reports the skewness and kurtosis risk premiums for the 30-day and 180-day horizons.[17] Except for the SV1 model at short maturities, skewness risk premiums are on average positive. A positive skewness risk premium is consistent with theory: It means that the Q-skewness is more negative than the P-skewness, which reflects the aversion towards left tail risk and the resulting fatter tails of the risk-neutral distribution. The kurtosis risk premium is defined as Q-kurtosis minus P-kurtosis, and we therefore once again expect it to be positive. This is indeed what we find for the SV2 and SV3 models, but the SV1 model once again generates an implausible (negative) risk premium for short and intermediary maturities.

Figure 6 plots the time series of the 30-day risk premiums. The time series at other horizons are highly correlated with the 30-day risk premiums. It is difficult to generate reasonable skewness and kurtosis risk premiums in the SV1 model. Not only is the sign of the average risk premium implausible at short maturities, the risk premiums are very small

[16] For instance, the difference between the 365- and 180-day IVRP is briefly negative during the financial crisis for the SV3 model (but not for the SV1 and SV2 models).

[17] Figure A.7 in the Appendix plots the average term structure of these risk premiums for horizons up to one year.

and do not change much over time. The SV2 model generates somewhat larger skewness risk premiums, while the SV3 model generates much larger risk premiums that fluctuate more over time.[18] In the SV3 model, the skewness and kurtosis risk premiums are elevated in the 2004-2007 period and during the last three years of the sample.

Figure 7 plots the time series of the term structure slope for the skewness and kurtosis risk premiums. For a given model, the two time series are clearly highly correlated. The term structure of these risk premiums is much more persistent than the levels of the risk premium in Figure 6. There are several differences with the IERP and the IVRP term structures. First, the difference between the SV3 model on the one hand and the SV1 and SV2 models on the other hand is much more pronounced for the skewness and kurtosis risk premiums. Second, the term structure for the skewness and kurtosis risk premiums is upward sloping most of the time for all models. It rarely inverts, and not necessarily when the IERP term structure inverts.

Figure 8 further illustrates this by plotting the risk-neutral term structure on three sample days, one with low volatility, one with median volatility, and one with high volatility. While for all three models the variance term structure strongly depends on the volatility regime, the skewness and kurtosis term structure is relatively similar.

5.3 The Co-Movement of Moment Risk Premiums

We now discuss how the risk premiums co-move, with emphasis on how the models differ in this dimension. We again focus on the 30- and 180-day horizons. By definition the integrated ERP and VRP are perfectly correlated in the SV1 model, as well as the spot ERP and VRP. Despite the additional flexibility provided by the SV2 and SV3 models, this correlation remains very high at both the 30- and 180-day horizons. The correlations between the first two moment risk premiums and skewness (kurtosis) are also rather similar. These correlations are always negative, except for the correlation between the ERP and VRP with kurtosis in the SV1 model at the 30-day horizon. This again highlights the restrictive nature of this model.

The correlation between the skewness and kurtosis risk premiums strongly differs across models at short horizons. For the SV3 model, this correlation is 0.956 at the 30-day horizon

[18] Figure A.3 in the Appendix illustrates the critical role of the non-standardized skewness and kurtosis in generating time-variation and outliers in the risk premiums.

and 0.963 at the 180-day horizon. For the SV2 model, we get a similar correlation at the 180-day horizon (0.985), but for the 30-day horizon the correlation is much lower (0.751). For the SV1 model, the correlation at 180 days is similar to that of the other two models, but it is negative at the 30-day horizon.

It is tempting to conclude that the additional flexibility afforded by the SV3 model results in improved modeling of this correlation. However, it is very difficult to confirm whether this matches the patterns in the data. The time series of the BKM risk-neutral skewness and kurtosis are noisy and may not be estimated precisely, especially for longer maturities. The estimation of the conditional counterparts under the physical measure is even more difficult. While the improved option fit suggests that the co-movement of the risk premiums in the SV3 model captures the data better, we therefore conclude that we require improved nonparametric modeling of higher moment risk premiums to definitively assess model performance in this dimension.

6 Jump Factors in Multifactor SV Models

We investigate how adding jumps to SV models improves option fit. We enrich the physical and risk-neutral specification of the return process with a jump process, as follows:

$$\frac{dS_t}{S_t} = (r_t - \delta_t + \eta_s \sum_{i=1}^{N} V_{it} - \lambda \bar{\mu}_s^Q)dt + \sum_{i=1}^{N} \sqrt{V_{it}} dZ_{it} + (e^{J_t^s} - 1)dN_t, \tag{20}$$

$$\frac{dS_t}{S_t} = (r_t - \delta_t - \lambda \bar{\mu}_s^Q)dt + \sum_{i=1}^{N} \sqrt{V_{it}} dZ_{it}^Q + (e^{J_t^{sQ}} - 1)dN_t^Q, \tag{21}$$

where N_t is a Poisson process with constant jump intensity λ and J_t^s is the jump size parameter related to returns. We assume $J_t^s \sim N(\mu_s, \sigma_s^2)$. The term $\lambda \bar{\mu}_s^Q$ is the compensation of the jump component, with $\bar{\mu}_s^Q = e^{(\mu_s^Q + \sigma_s^2/2)} - 1$. The jump risk premium is equal to $\eta_{J_s} = \mu_s - \mu_s^Q$. This SVJR specification has four additional parameters $(\lambda, \mu_s, \sigma_s^2, \mu_s^Q)$ compared to its SV model counterpart.[19]

[19] We use the following prior distributions for the jump parameters: $\ln(\frac{\lambda - 0.5}{252 - 0.5}) \sim N(-9, 0.5)$, $\mu_s \sim N(-0.03, 0.01^2)$ and $\ln(\sigma_s) \sim N(\ln(0.02) - \frac{0.25^2}{2}, 0.25^2)$. The narrow prior for λ guides the sampler towards fewer jumps, consistent with the literature's findings. Note that the priors on μ_s and σ_s are broad enough to leave the posterior outcomes largely unaffected.

6.1 Parameter Estimates and Model Fit

Table 5 reports on the SVJR models. The structure of the jump significantly differs across models. Our estimates for the SV1JR model are largely consistent with the existing literature. The jump occurs on average once every two years and the average risk-neutral jump size is negative and similar in size (-6.2%) to several existing studies. However, while the jump intensities in the SV2JR and SV3JR models are similar to the SV1JR model, the risk-neutral jump size estimates are positive. This is inconsistent with the usual motivation given for these jumps that they capture sudden large drops in the stock price. A potential explanation is that the second and third variance factors capture these downward movements, and are also able to capture the additional skewness and kurtosis that the jump factor helps capture in the SV1JR model. Finally, Table 5 also indicates that when we add a jump in returns to the stochastic volatility models, the κ and θ estimates are somewhat impacted, but the resulting diffusive variance factors are largely similar to those from the corresponding SV models.

Adding a return jump in the SV1JR model lowers the pricing error from \$3.31 (in the SV1 model) to \$3.16 (15 dollarcents), but the improvement in fit from adding a return jump to the SV2JR is only 4 dollarcents, and for the SV3JR model is is negative. While this may seem surprising, note that the estimation maximizes the log-likelihood, and the BIC of the SV3JR model far exceeds that of the SV3 model. An analysis of the (in-sample) model fit by moneyness and maturity confirms the insights of Das and Sundaram (1999) that the improvements from return jumps are due to short-maturity options.

How do these findings compare to the existing literature? There are relatively few papers that study multifactor SV models with return jumps. Bates (2000) provides an early study of a model with two SV factors and a Poisson-normal return jump. More recently, Ait-Sahalia et al. (2020) and Bardgett et al. (2019) study more complex models with two SV factors and jumps in returns and volatility. Andersen et al. (2015a,b) add an additional tail factor to various models with multiple SV factors and jumps in returns and volatility. Gruber et al. (2021) estimate a different three-factor specification. The literature generally confirms that when fitting large option panels, additional SV factors lead to large improvements in option fit, while the overall improvements in fit from modeling jumps are more modest. Jumps help improve model fit during turbulent periods (Andersen et al., 2015b) and result in improved

modeling of risk premiums, which we study below.

6.2 Higher Moments and Risk Premiums

The total spot equity premium in the SV1JR model is exactly the same as in the SV1 model. The jump risk premium accounts for 18% (0.017/(0.079 + 0.017)) of the overall spot equity risk premium in the SV1JR model. This estimate is similar to other estimates of the SV1JR model in the literature (Pan, 2002; Broadie et al., 2007; Aït-Sahalia et al., 2021) However, our findings are once again dramatically different for the SV2JR and SV3JR models because the sign of the risk premium estimates is implausible. This confirms our earlier findings that additional SV components capture the stylized facts captured by the return jump in a SV1JR model.

The higher moments and risk premiums for the SVJR models can be computed using the same approach used for the SV models. They are given by

$$VAR(R_{t,t+\tau}|\boldsymbol{V}_t) = \sum_{i=1}^{N}\left[A_i^{(2)}(\Theta,\tau)+D_i^{(2)}(\Theta,\tau)V_{it}\right]+\lambda P^{(2)}(\Theta,\tau), \quad (22)$$

$$E\Big((R_{t,t+\tau}-E(R_{t,t+\tau}|\boldsymbol{V}_t))^3|\boldsymbol{V}_t\Big) = \sum_{i=1}^{N}\left[A_i^{(3)}(\Theta,\tau)+D_i^{(3)}(\Theta,\tau)V_{it}\right]+\lambda P^{(3)}(\Theta,\tau), \quad (23)$$

$$\begin{aligned} E\Big((R_{t,t+\tau}-E(R_{t,t+\tau}|\boldsymbol{V}_t))^4|\boldsymbol{V}_t\Big) &= \sum_{i=1}^{N}\left[A_i^{(4)}(\Theta,\tau)+D_i^{(4)}(\Theta,\tau)V_{it}\right]+\lambda P^{(4)}(\Theta,\tau) \\ &\quad +3VAR(R_{t,t+\tau}|\boldsymbol{V}_t)^2, \end{aligned} \quad (24)$$

where

$$\begin{aligned} P^{(2)}(\Theta,\tau) &= \tau\left((\mu^Q)^2+\sigma_s^2\right), \\ P^{(3)}(\Theta,\tau) &= \tau\,\mu^Q\big((\mu_Q)^2+3\sigma_s^2\big), \\ P^{(4)}(\Theta,\tau) &= \tau\left((\mu^Q)^4+6\,(\mu^Q)^2\,\sigma_s^2+3\,(\sigma_s)^4\right). \end{aligned}$$

Recall that the estimate of μ_s^Q in the SV1JR model is negative, consistent with existing estimates in the literature. This negative estimate means that returns jumps induce additional negative skewness in the risk-neutral distribution. However, the estimates of μ_s^Q in

the SV2JR and SV3JR models are positive. Moreover, the estimates of σ_s in the SV2JR and SV3JR models are much smaller compared to the SV1JR model, which suggests that the return jumps will affect kurtosis less in these models. Figure 9 confirms that the return jump affects the higher moments in the SV1JR model, but not in the SV2JR and SV3JR models. A similar conclusion holds for the higher moment risk premiums.

An important caveat is that our comparison between SV and jump components is limited to very simple Poisson-normal return jumps with constant jump intensity. This leaves open the possibility that more general and flexible jump factors can outperform additional SV factors. The literature contains several such models, including jumps with time-varying intensity, infinite-activity Levy jump processes, time-changed Levy processes, and correlated return and volatility jumps.[20] However, we note that jump components significantly add to the computational burden, because more paths are required when using simulation-based models and/or because solutions to differential equations may not be available in closed form. We encountered significant identification problems even with simple return jumps, and these identification problems seem to worsen when we implement more complex jump processes. We believe that the implementation of more complex jump factors that can outperform additional SV factors, or that can add significant explanatory power to a three-factor SV model, constitutes an interesting outstanding research challenge when estimation is based on the joint likelihood of returns and options.[21]

7 Conclusion

The dynamics of the higher moments of the return distribution are critically important for modeling time-series and cross-sectional patterns in option prices. We exploit novel estimation techniques to estimate multi-factor models and explore if these models capture the first four return moments and how (co-movements between) these moments affect index option fit. Adding a third volatility factor greatly improves option fit by allowing for much

[20] See for example Aït-Sahalia et al. (2015), Bates (2006, 2012), Bakshi et al. (2008), Bakshi and Wu (2010), Carr and Wu (2003, 2004), Fulop et al. (2015), Yu et al. (2011), Ornthanalai (2014), Pan (2002), and Santa-Clara and Yan (2010).

[21] The relative importance of such jump processes will of course depend on the underlying securities as well as the data frequency. They are more likely to be important for single-name equity options and for high-frequency-data. See Bates (2019) for an application to high-frequency data.

richer skewness and kurtosis patterns that better match non-parametric benchmarks. Both jumps and additional volatility factors allow for more complex patterns in risk premiums, but for modeling large cross-sections of index options, an additional volatility factor adds much more explanatory power than a jump factor, and it allows for richer and more realistic patterns in higher moments and the associated risk premiums. All models we investigate have difficulty to match the unconditional term structure of skewness and kurtosis, highlighting the need for more realistic dynamics and new modeling approaches.

We leave several important questions for future research. Our empirical analysis is limited to jump factors with a relatively simple structure. While it is computationally intensive to estimate more complex jump processes, the literature contains several interesting candidates, including jumps with time-varying intensity, infinite-activity Levy jump processes, time-changed Levy processes, and correlated return and volatility jumps. These models may be able to address the remaining shortcomings of multifactor SV models. A second important avenue for future research is to use our estimation approach to compare multifactor SV models with the models in Gruber et al. (2021) and Andersen et al. (2015a, 2017). While these studies convincingly show that multifactor SV models can be improved upon, this conclusion is primarily based on fitting option prices under the risk-neutral measure.

References

Yacine Aït-Sahalia, Julio Cacho-Diaz, and Roger Laeven. Modeling Financial Contagion Using Mutually Exciting Jump Processes. *Journal of Financial Economics*, 117:585–606, 2015.

Yacine Ait-Sahalia, Mustafa Karaman, and Loriano Mancini. The Term Structure of Equity and Variance Risk Premia. *Journal of Econometrics*, 219(2):204–230, 2020.

Yacine Aït-Sahalia, Chenxu Li, and Chen Xu Li. Closed-form Implied Volatility Surfaces for Stochastic Volatility Models with Jumps. *Journal of Econometrics*, 222(1, Part B): 364–392, 2021.

Torben G. Andersen, Luca Benzoni, and Jesper Lund. An Empirical Investigation of Continuous-Time Models for Equity Returns. *Journal of Finance*, 57:1239–1284, 2002.

Torben G. Andersen, Nicola Fusari, and Viktor Todorov. Parametric Inference and Dynamic State Recovery From Option Panels. *Econometrica*, 83(3):1081–1145, 2015a. URL http://doi.wiley.com/10.3982/ECTA10719.

Torben G. Andersen, Nicola Fusari, and Viktor Todorov. The Risk Premia Embedded in Index Options. *Journal of Financial Economics*, 117(3):558–584, 2015b. URL http://dx.doi.org/10.1016/j.jfineco.2015.06.005.

Torben G. Andersen, Nicola Fusari, and Viktor Todorov. Short-Term Market Risks Implied by Weekly Options. *Journal of Finance*, 72(3):1335–1386, 2017.

Torben G. Andersen, Nicola Fusari, and Viktor Todorov. The Pricing of Tail Risk and the Equity Premium: Evidence From International Option Markets. *Journal of Business & Economic Statistics*, 38(3):662–678, 2020.

Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342, 2010.

S. B. Aruoba, F. X. Diebold, and C. Scotti. Real-Time Measurement of Business Conditions. *Journal of Business and Economic Statistics*, 27(4):417–427, 2009.

Gurdip Bakshi and Dilip Madan. Spanning and Derivative-security Valuation. *Journal of Financial Economics*, 55(2):205–238, 2000.

Gurdip Bakshi and Liuren Wu. The Behavior of Risk and Market Prices of Risk Over the Nasdaq Bubble Period. *Management Science*, 56(12):2251–2264, 2010.

Gurdip Bakshi, Charles Cao, and Zhiwu Chen. Empirical Performance of Alternative Option Pricing Models. *Journal of Finance*, 52(5):2003–2049, 1997.

Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options. *Review of Financial Studies*, 16 (1):101–143, 2003. ISSN 0893-9454.

Gurdip Bakshi, Peter Carr, and Liuren Wu. Stochastic Risk Premiums, Stochastic Skewness in Currency Options, and Stochastic Discount Factors in International Economies. *Journal of Financial Economics*, 87(1):132–156, 2008.

Chris Bardgett, Elise Gourier, and Markus Leippold. Inferring Volatility Dynamics and Risk Premia from the S&P 500 and VIX Markets. *Journal of Financial Economics*, 131 (3):593–618, 2019.

David S. Bates. Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options. *Review of Financial Studies*, 9(1):69–107, 1996. URL http://www.nber.org/papers/w4596.

David S. Bates. Post-'87 Crash Fears in S&P 500 Futures Option Market. *Journal of Econometrics*, 94(1-2):181–238, 2000.

David S. Bates. Empirical Option Pricing: A Retrospection. *Journal of Econometrics*, 116 (1-2):387–404, 2003.

David S. Bates. Maximum Likelihood Estimation of Latent Affine Processes. *Review of Financial Studies*, 19(3):909–965, 2006.

David S. Bates. U.S. Stock Market Crash Risk, 1926-2010. *Journal of Financial Economics*, 105(2):229–259, 2012. URL http://dx.doi.org/10.1016/j.jfineco.2012.03.004.

David S. Bates. How Crashes Develop: Intradaily Volatility and Crash Evolution. *Journal of Finance*, 74(1):193–238, 2019.

Tim Bollerslev and Viktor Todorov. Tails, Fears, and Risk Premia. *The Journal of Finance*, 66(6):2165–2211, 2011.

Tim Bollerslev and Viktor Todorov. Time-varying Jump Tails. *Journal of Econometrics*, 183(2):168–180, 2014.

Tim Bollerslev, George Tauchen, and Hao Zhou. Expected Stock Returns and Variance Risk Premia. *Review of Financial Studies*, 22(11):4463–4492, 2009.

Tim Bollerslev, Viktor Todorov, and Lai Xu. Tail Risk Premia and Return Predictability. *Journal of Financial Economics*, 118(1):113–134, 2015.

Mark Broadie, Mikhail Chernov, and Michael Johannes. Model Specification and Risk Premia: Evidence From Futures Options. *Journal of Finance*, 62(3):1453–1490, 2007.

Peter Carr and Liuren Wu. What Type of Process Underlies Options? A Simple Robust Test. *The Journal of Finance*, 58(6):2581–2610, 2003.

Peter Carr and Liuren Wu. Time-changed Lvy Processes and Option Pricing. *Journal of Financial Economics*, 71(1):113–141, 2004.

Mikhail Chernov and Eric Ghysels. A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation. *Journal of Financial Economics*, 56(3):407–458, 2000.

Mikhail Chernov, A. Ronald Gallant, Eric Ghysels, and George Tauchen. Alternative Models for Stock Price Dynamics. *Journal of Econometrics*, 116(1-2):225–257, 2003.

Peter Christoffersen, Steven Heston, and Kris Jacobs. The Shape and Term Structure of the Index Option Smirk: Why Multifactor Stochastic Volatility Models Work So Well. *Management Science*, 55(12):1914–1932, 2009.

Peter Christoffersen, Kris Jacobs, and Karim Mimouni. Volatility Dynamics for the S&P500: Evidence from Realized Volatility, Daily Returns, and Option Prices. *Review of Financial Studies*, 23(8):3141–3189, 2010.

Sanjiv Ranjan Das and Rangarajan K. Sundaram. Of Smiles and Smirks: A Term Structure Perspective. *Journal of Financial and Quantitative Analysis*, 34(2):211–239, 1999. URL http://www.jstor.org/stable/2676279?origin=crossref.

Arnaud Dufays, Yuguo Liu, Kris Jacobs, and Jeroen Rombouts. Fast Filtering with Large Option Panels: Implications for Asset Pricing. *Journal of Financial and Quantitative Analysis*, xx(xx):xx–xx, 2023.

Darrell Duffie, Jun Pan, and Kenneth Singleton. Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica*, 68(6):1343–1376, 2000. URL http://doi.wiley.com/10.1111/1468-0262.00164.

Bjørn Eraker. Do Stock Prices and Volatility Jump? Evidence From Spot and Option Prices. *Journal of Finance*, 59(3):1367–1403, 2004.

Bjørn Eraker, Michael Johannes, and Nicholas Polson. The Impact of Jumps in Equity Index Volatility and Returns. *Journal of Finance*, 58(3):1269–1300, 2003.

Andras Fulop, Junye Li, and Jun Yu. Self-Exciting Jumps, Learning, and Asset Pricing Implications. *The Review of Financial Studies*, 28(3):876–912, 2015.

Peter H. Gruber, Claudio Tebaldi, and Fabio Trojani. The Price of the Smile and Variance Risk Premia. *Management Science*, 67(7):4056–4074, 2021.

Steven L. Heston. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies*, 6(2):327–343, 1993. URL https://academic.oup.com/rfs/article-lookup/doi/10.1093/rfs/6.2.327.

A. S. Hurn, K. A. Lindsay, and A. J. McClelland. Estimating the Parameters of Stochastic Volatility Models Using Option Price Data. *Journal of Business and Economic Statistics*, 33(4):579–594, 2015.

Christopher S. Jones. The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets. *Journal of Econometrics*, 116(1-2):181–224, 2003.

Chayawat Ornthanalai. Levy Jump Risk: Evidence from Options and Returns. *Journal of Financial Economics*, 112(1):69–90, 2014.

Jun Pan. The Jump-risk Premia Implicit in Options: Evidence from an Integrated Time-series Study. *Journal of Financial Economics*, 63(1):3–50, 2002.

Pedro Santa-Clara and Shu Yan. Crashes, Volatility, and the Equity Premium: Lessons from S&P 500 Options. *The Review of Economics and Statistics*, 92(2):435–451, 2010.

Cindy L. Yu, Haitao Li, and Martin T. Wells. MCMC Estimation of Levy Jump models using Stock and Option Prices. *Mathematical Finance*, 21(3):383–422, 2011.

Figure 1: Filtered Variance Paths. SV Models

Panel (A): Total Variance for the SV1, SV2 and SV3 Models

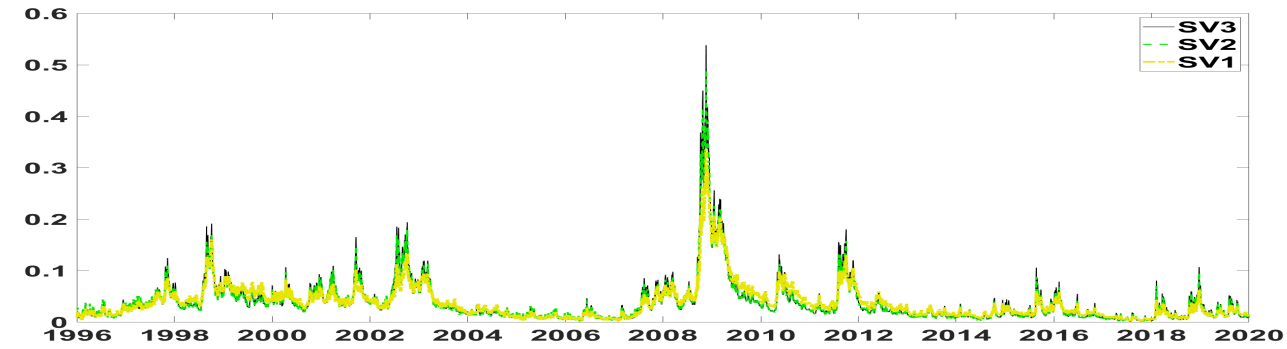

Panel (B): Second Variance Component for the SV2 and SV3 Models

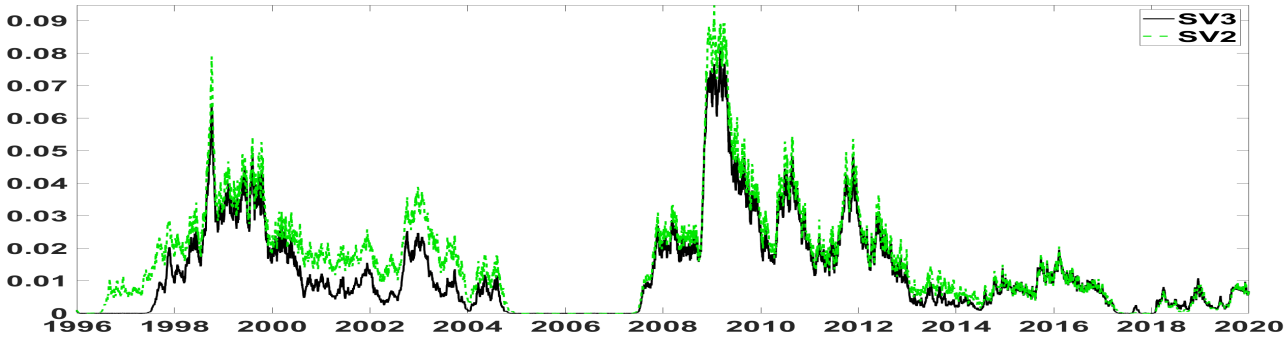

Panel (C): Third Variance Component for the SV3 Model

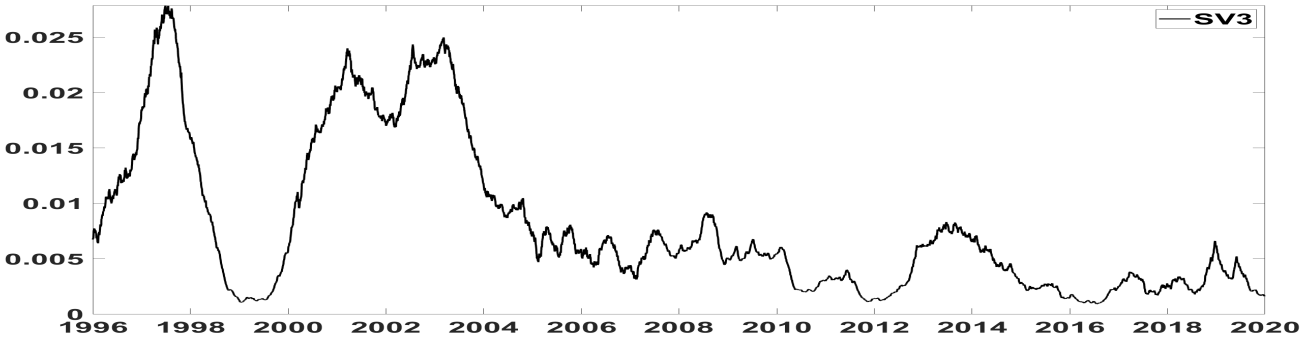

Notes: The top panel plots the path of the total filtered variance in the SV1, SV2 and SV3 models. The middle panel plots the second variance component in the SV2 model, and the bottom panel plots the third variance factor in the SV3 model. The sample period is from January 4, 1996 to December 31, 2019.

Figure 2: In-Sample and Out-of-Sample Option Fit

Panel (A): In-Sample Option Fit SV1, SV3 and SV3JR

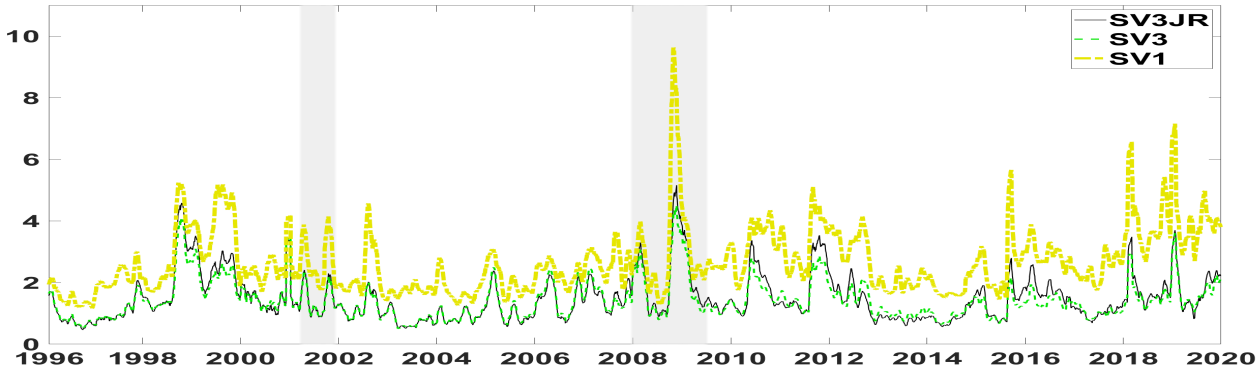

Panel (B): Out-of-Sample Option Fit SV1, SV3 and SV3JR

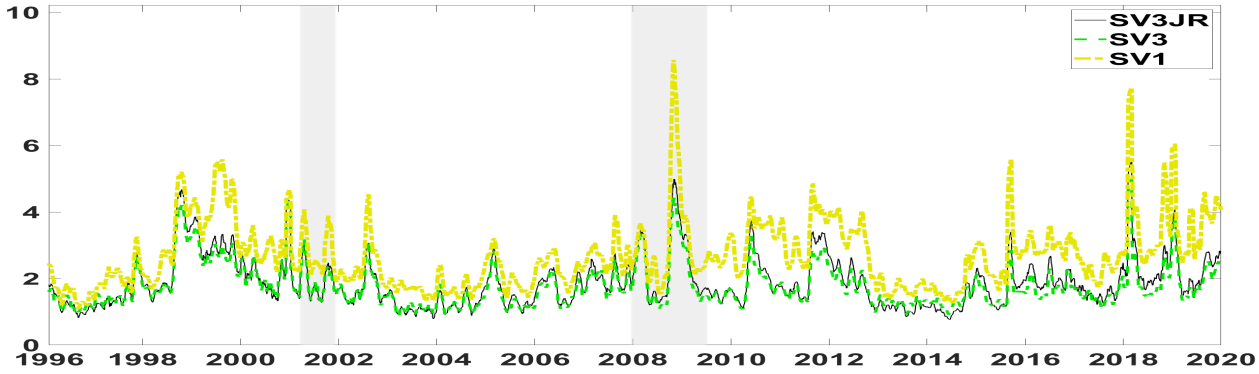

Panel (C): Out-of-Sample Option Fit SV2, SV3 and SV3JR

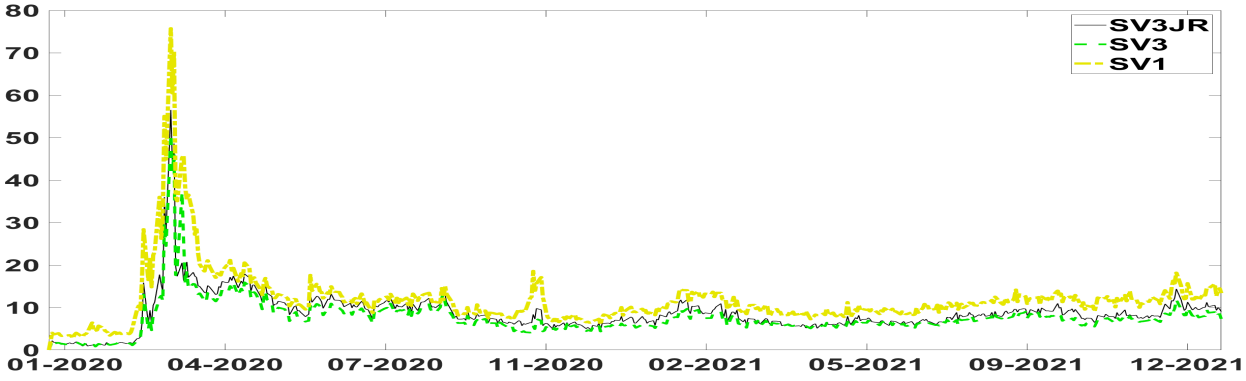

Notes: We plot the option fit root mean squared error over time. The series have been smoothed using a simple moving average over one month.

Figure 3: Leverage Correlations and Variance Risk Premiums in SV Models.

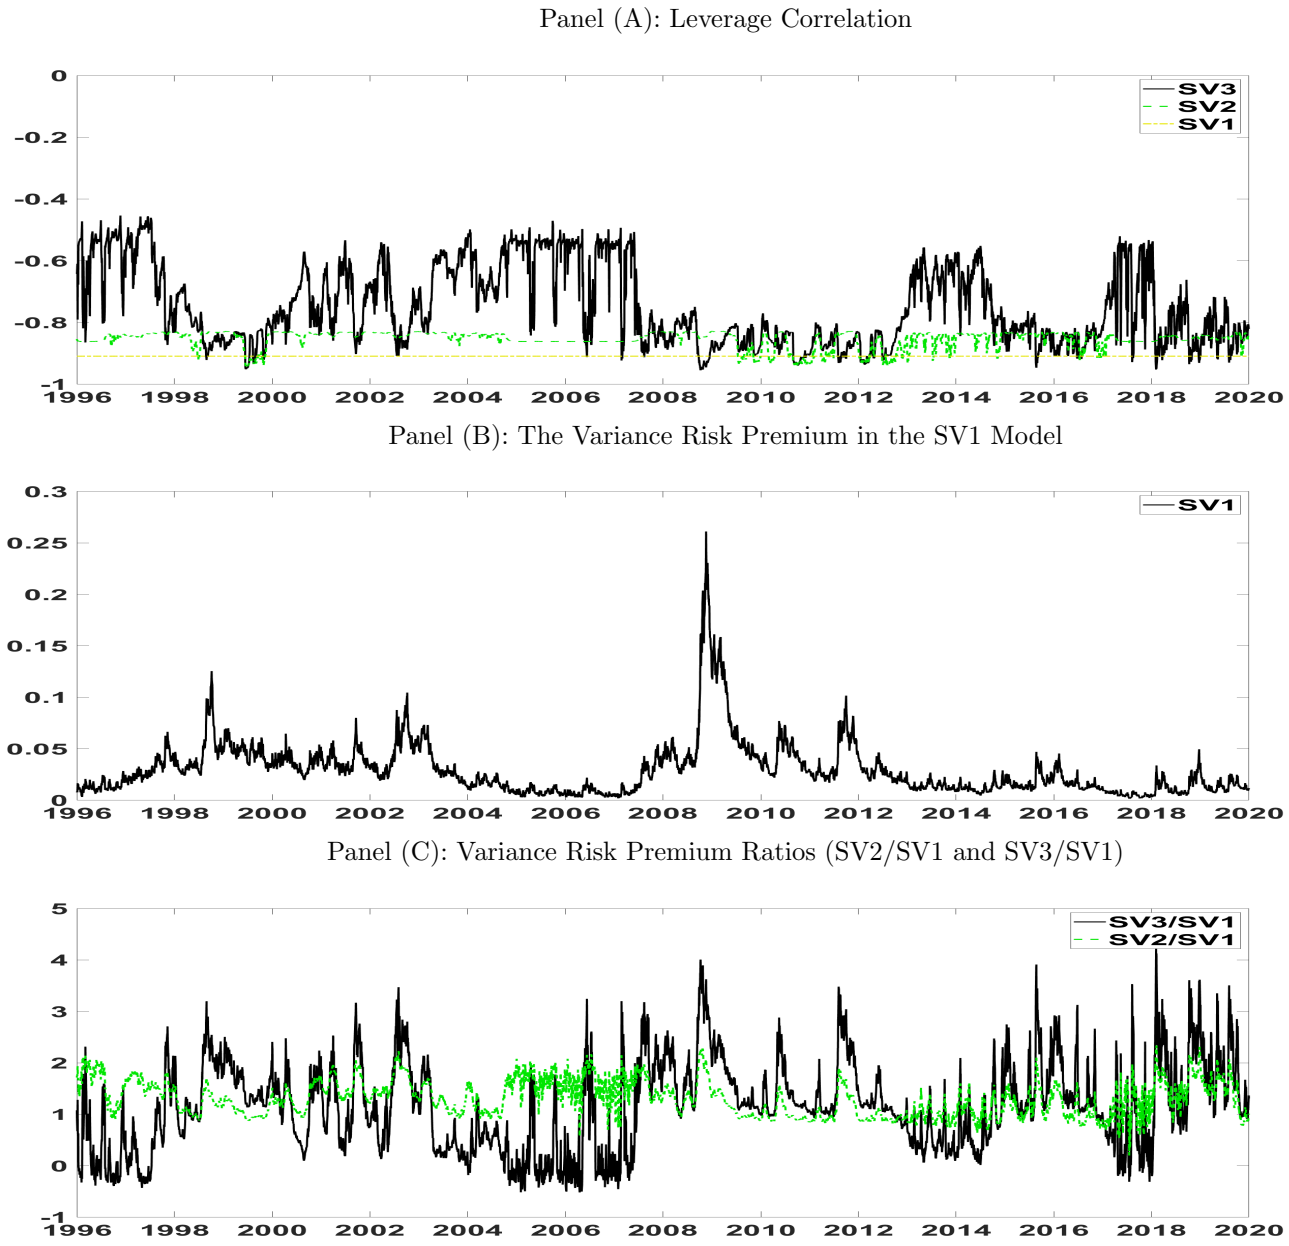

Notes: The top panel plots the leverage correlations in the SV1, SV2 and SV3 models. The middle panel plots the variance risk premium in the SV1 model. The bottom panel plots the ratios of the variance risk premium (SV2/SV1 and SV3/SV1). The sample period is from January 4, 1996 to December 31, 2019.

Figure 4: Properties of the Variance of Variance in SV Models.

Panel (A): Correlations between returns and variance of variance

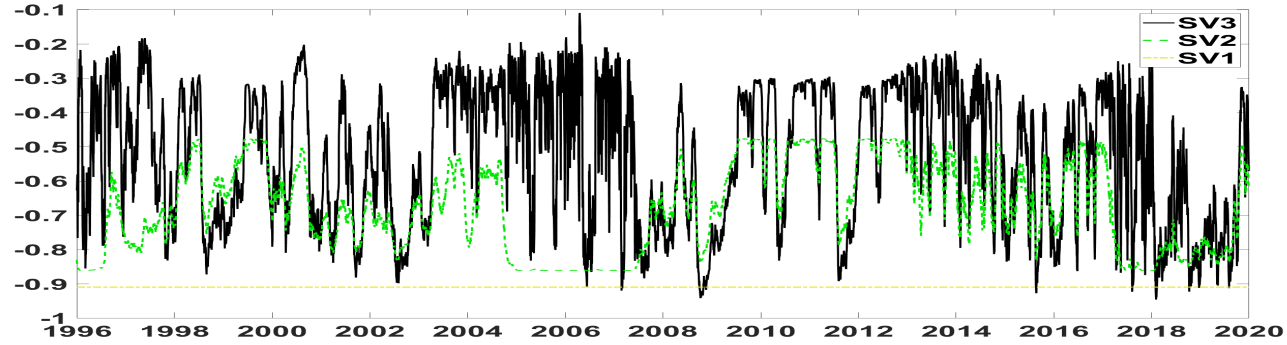

Panel (B): Conditional variance of variance - SV1 model

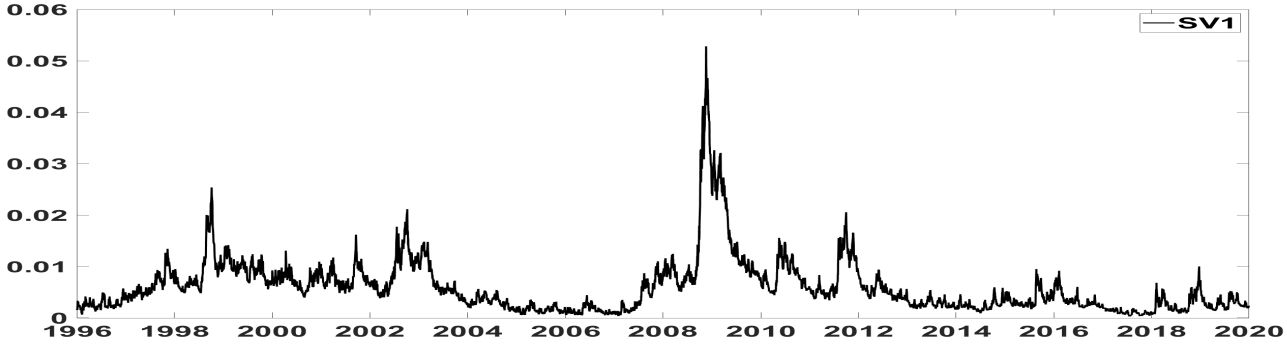

Panel (C): Conditional variance of variance (SV2/SV1) and (SV3/SV1)

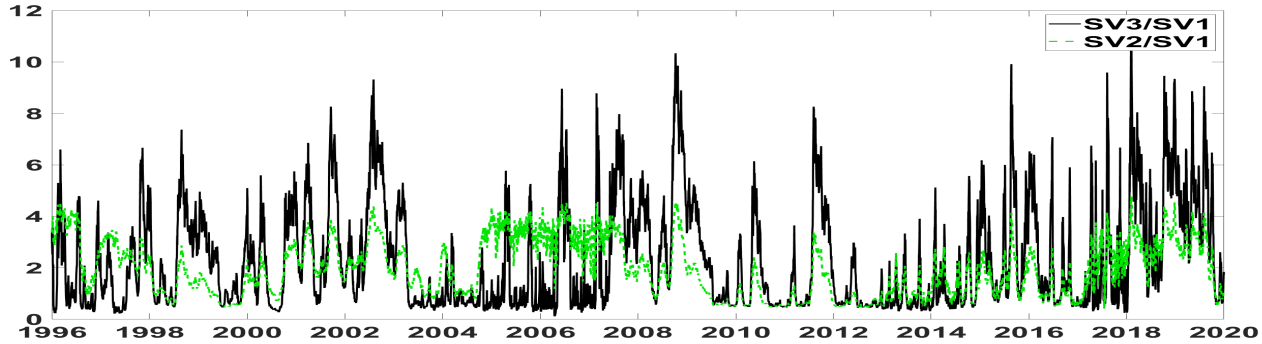

Notes: The top panel plots the correlations between returns and variance of variance in the SV1, SV2 and SV3 models. The middle panel plots the conditional variance of variance in the SV1 model. The bottom panel plots the ratios of the conditional variance of variance (SV2/SV1 and SV3/SV1). The sample period is from January 4, 1996 to December 31, 2019.

Figure 5: Risk-Neutral Variance, Skewness and Kurtosis. 30-Day Maturity

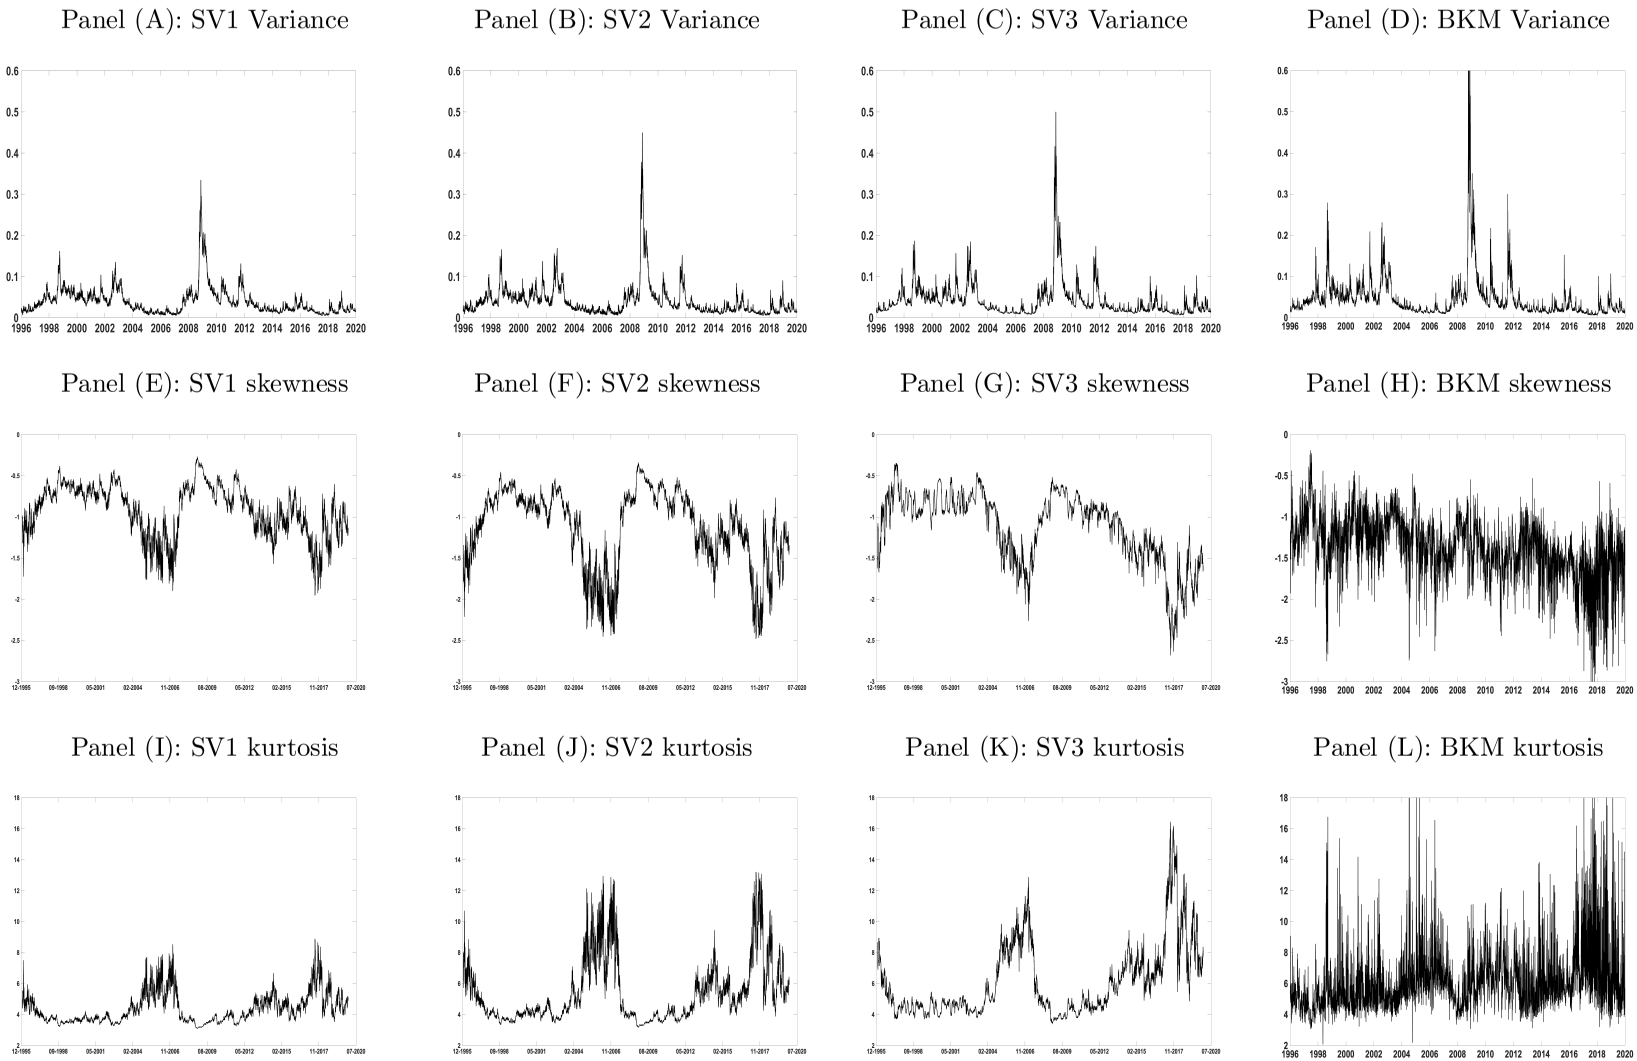

Notes: The first three columns report on the SV1, SV2, and SV3 models. The fourth column reports on the risk-neutral moments computed according to Bakshi, Kapadia, and Madan (2003) (BKM). We report on the 30-day maturity. The top row plots the variance, the middle row plots the skewness, and the bottom row plots the kurtosis. The sample period is from January 4, 1996 to December 31, 2019

Figure 6: 30-Day Moment Risk Premiums

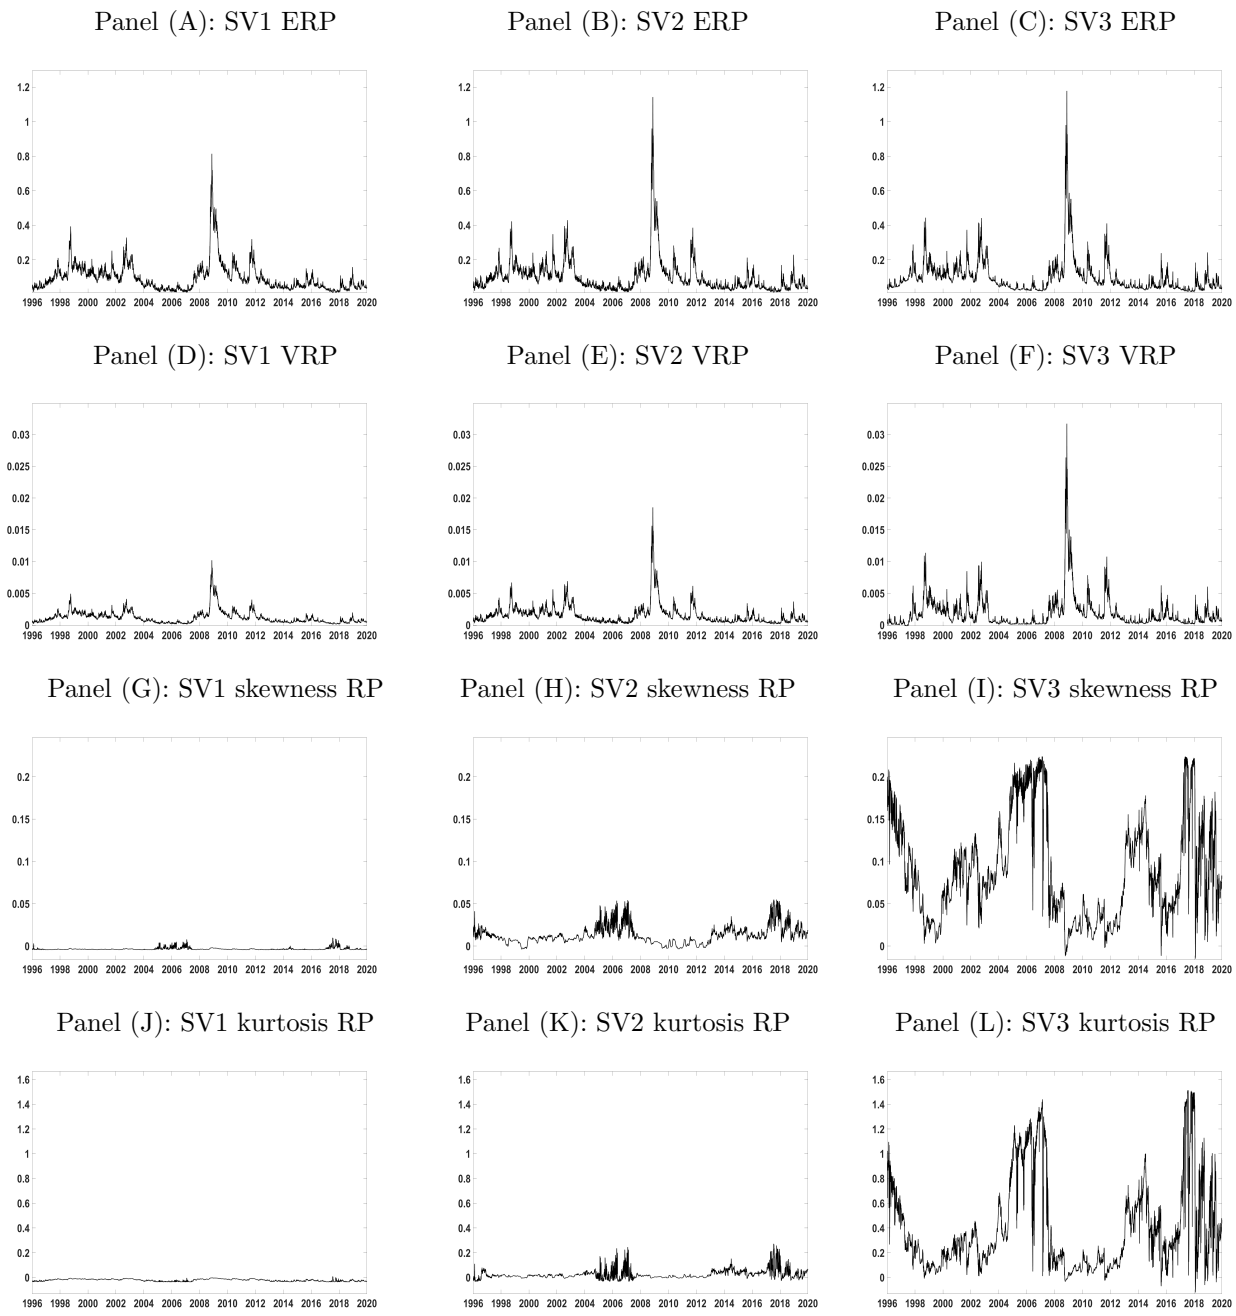

Notes: The first column reports on the SV1 model, the second column on the SV2 model, and the third column on the SV3 model. The top row plots the 30-day equity risk premium (ERP), the second row plots the variance risk premium, the third row plots the skewness risk premium, and the bottom row plots the kurtosis risk premium. The sample period is from January 4, 1996 and December 31, 2019.

Figure 7: The Term Structure of Moment Risk Premiums.

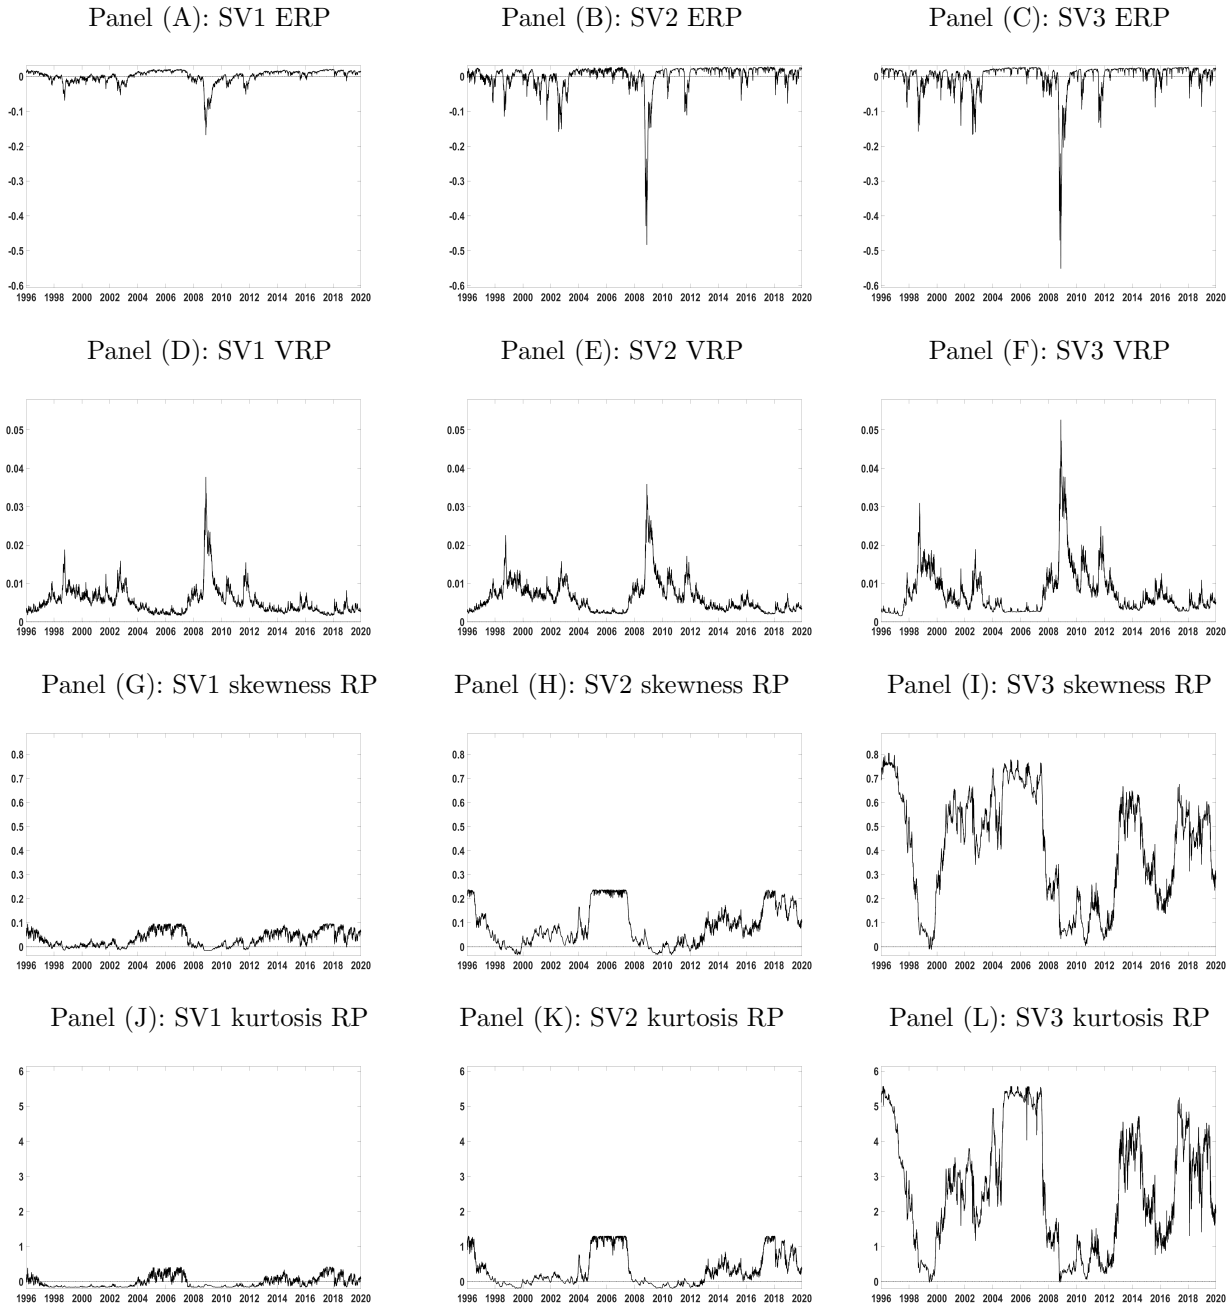

Notes: We plot the difference between the 180-day and 30-day risk premiums. The sample period is from January 4, 1996 to December 31, 2019.

Figure 8: The Risk-Neutral Term Structure on Three Selected Days

Panel (A): Variance - LowVol

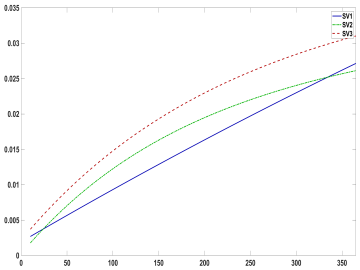

Panel (B): Variance - MeanVol

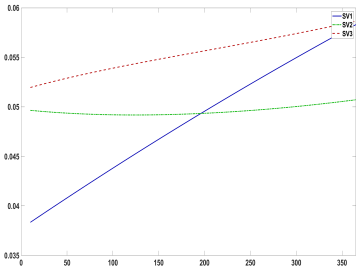

Panel (C): Variance - HighVol

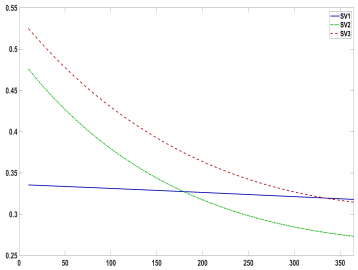

Panel (D): Skewness - LowVol

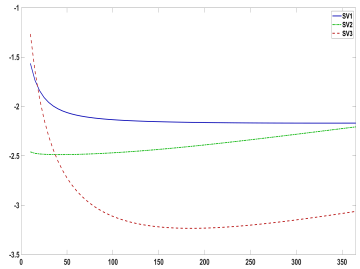

Panel (E): Skewness - MeanVol

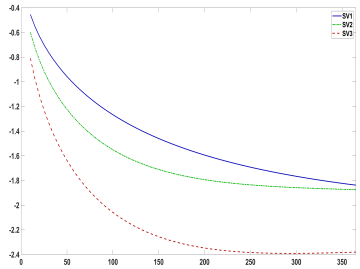

Panel (F): Skewness - HighVol

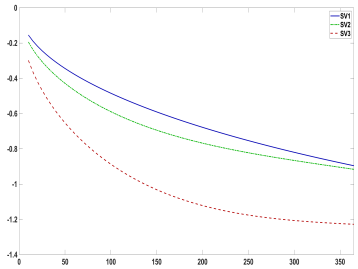

Panel (G): Kurtosis - LowVol

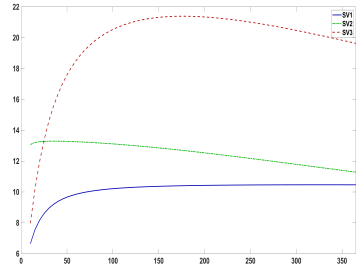

Panel (H): Kurtosis - MeanVol

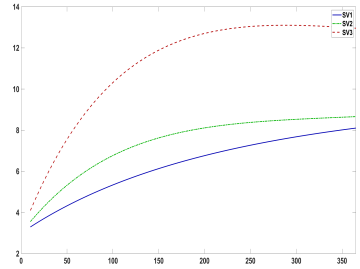

Panel (I): Kurtosis - HighVol

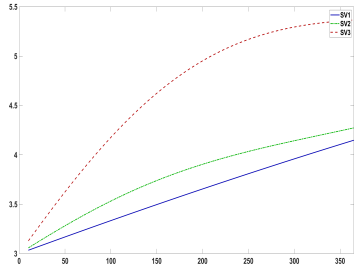

Notes: We plot the term structure (up to one year) of risk-neutral variance, skewness and kurtosis on three selected days in the sample that belong to different volatility regimes: July 19, 2017, a low volatility (LowVol) day in the left column; August 01, 2007, an average volatility (MeanVol) day in the middle column; and November 20, 2008, a high volatility (HighVol) day in the right column. The sample period is from January 4, 1996 and December 31, 2019.

Figure 9: Ratio of 30-day Risk-Neutral Moments.

Panel (A): SV1JR/SV1 Variance

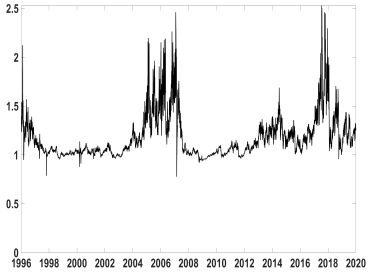

Panel (B): SV2JR/SV2 Variance

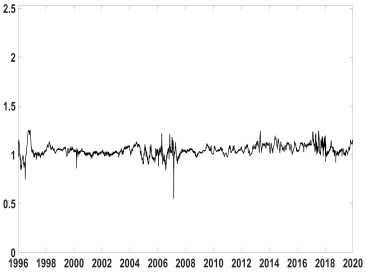

Panel (C): SV3JR/SV3 Variance

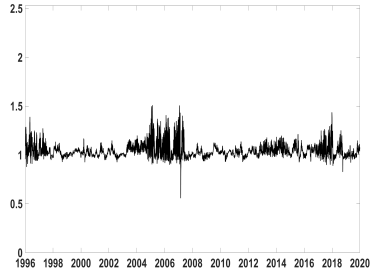

Panel (D): SV1JR/SV1 Skewness

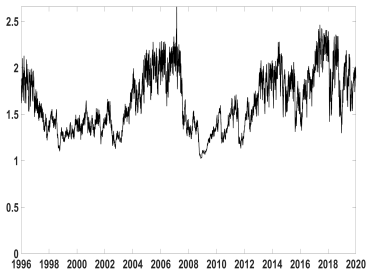

Panel (E): SV2JR/SV2 Skewness

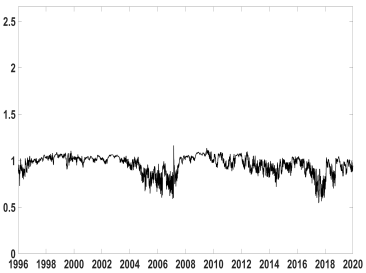

Panel (F): SV3JR/SV3 Skewness

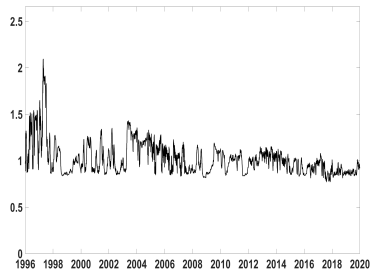

Panel (G): SV1JR/SV1 Kurtosis

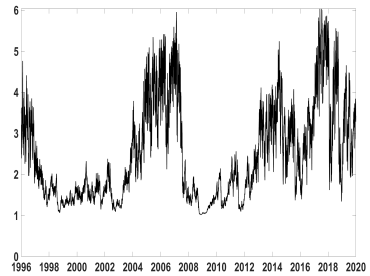

Panel (H): SV2JR/SV2 Kurtosis

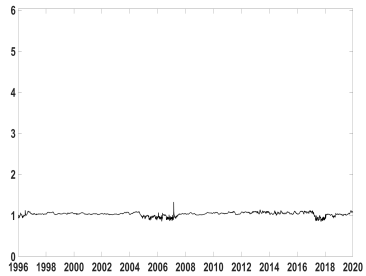

Panel (I): SV3JR/SV3 Kurtosis

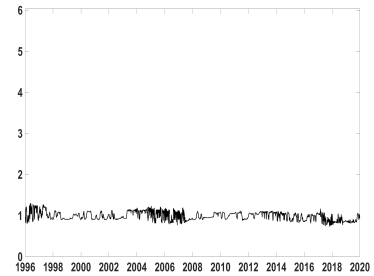

We plot the ratio between the SVJR and SV 30-day risk-neutral variance, skewness and kurtosis. The sample period is from January 4, 1996 to December 31, 2019.

Table 1: Return and Option Data

Panel A: Return Data

	Mean	StdDev	Skewness	Kurtosis	Max	Min
Index Returns	0.0919	0.1951	-0.1833	13.1768	0.1158	-0.1198

Panel B: Option Data

	Maturity (days)					
Moneyness (K/S)	5-30	30-60	60-90	90-180	180-365	All
Number of Option Contracts						
0.85-0.90	4626	6026	5556	5956	5647	27811
0.90-0.95	5389	6094	5820	6158	5966	29427
0.95-1.00	5483	6103	5926	6226	6220	29958
1.00-1.05	5482	6107	5935	6269	6321	30114
1.05-1.10	3836	5952	5701	6044	5924	27457
1.10-1.15	1228	2818	3376	5197	5375	17994
All	26044	33100	32314	35850	35453	162761
Average Call Prices						
0.85-0.90	201.19	196.08	207.73	213.75	239.36	211.62
0.90-0.95	116.91	125.23	137.39	150.04	177.24	141.36
0.95-1.00	46.12	59.12	69.73	87.83	118.08	76.18
1.00-1.05	11.63	22.78	34.38	49.6	82.32	40.14
1.05-1.10	2.29	5.29	10.34	20.57	43.96	16.49
1.10-1.15	1.98	2.96	5.13	9.33	24.28	8.74
All	63.35	68.58	77.45	88.52	114.21	82.42
Average Implied Volatilities						
0.85-0.90	0.33	0.27	0.26	0.25	0.24	0.27
0.90-0.95	0.26	0.24	0.23	0.23	0.22	0.23
0.95-1.00	0.20	0.20	0.20	0.20	0.20	0.20
1.00-1.05	0.17	0.17	0.17	0.18	0.19	0.18
1.05-1.10	0.18	0.15	0.15	0.16	0.17	0.16
1.10-1.15	0.29	0.19	0.17	0.16	0.16	0.19
All	0.24	0.20	0.20	0.19	0.20	0.21

Notes: Panel A reports descriptive statistics for the sample of daily index returns between January 4, 1996 and December 31, 2019. The mean and standard deviation are annualized. Panel B reports the number of contracts, average call price, and average implied volatility in the option data set used for estimation of the models where we choose the most liquid (highest trading volume) option within each moneyness-maturity range. Moneyness is defined as K/S. Due to the fact that OTM options are more heavily traded, this data set mainly consists of OTM call and OTM put options.

Table 2: Parameter Estimates for SV Multifactor Models

	SV1	SV2	SV3
κ_1	1.403	4.705	5.292
	(0.269)	(0.687)	(0.896)
θ_1	0.041	0.025	0.021
	(0.009)	(0.004)	(0.004)
σ_1	0.397	0.690	0.993
	(0.004)	(0.010)	(0.022)
ρ_1	-0.909	-0.861	-0.977
	(0.005)	(0.008)	(0.005)
κ_2		0.434	0.603
		(0.251)	(0.307)
θ_2		0.006	0.002
		(0.080)	(0.056)
σ_2		0.298	0.305
		(0.009)	(0.007)
ρ_2		-0.957	-0.976
		(0.008)	(0.006)
κ_3			0.671
			(0.147)
θ_3			0.008
			(0.002)
σ_3			0.056
			(0.004)
ρ_3			-0.569
			(0.043)
η_s	2.548	2.706	2.606
	(0.573)	(0.539)	(0.577)
η_{v1}	0.777	1.123	1.824
	(0.270)	(0.676)	(0.897)
η_{v2}		0.997	1.435
		(0.251)	(0.309)
η_{v3}			-0.367
			(0.150)
σ_c	3.312	2.666	2.176
	(0.032)	(0.028)	(0.024)
Diffusive ERP	0.096	0.101	0.099
Loglikelihood	3020.53	3864.38	4195.62
BIC	-5980.12	-7624.28	-8243.24

Notes: We report parameter estimates based on the joint likelihood from returns and options for the SV multifactor models. Parameters are annualized and reported under the physical measure. In parentheses, we report the posterior standard deviation for each parameter. Diffusive ERP represents the average equity risk premium.

Table 3: In-Sample Option Fit

	Maturity (days)					
Moneyness (K/S)	5-30	30-60	60-90	90-180	180-365	All
SV1						
0.85-0.90	1.759	2.833	3.369	3.641	4.260	3.172
0.90-0.95	2.234	2.952	3.078	3.204	4.222	3.138
0.95-1.00	3.103	3.292	3.038	2.786	4.380	3.320
1.00-1.05	3.258	3.671	3.421	3.156	4.724	3.646
1.05-1.10	1.949	2.442	2.837	3.322	4.563	3.022
1.10-1.15	2.044	2.171	2.265	2.727	4.427	2.727
All	2.391	2.894	3.001	3.139	4.429	3.171
SV2/SV1						
0.85-0.90	0.845	0.787	0.818	0.882	0.884	0.843
0.90-0.95	0.764	0.735	0.803	0.884	0.832	0.804
0.95-1.00	0.715	0.695	0.774	0.898	0.755	0.767
1.00-1.05	0.742	0.743	0.804	0.872	0.710	0.774
1.05-1.10	0.857	0.799	0.847	0.885	0.747	0.827
1.10-1.15	1.006	0.914	0.905	0.926	0.784	0.907
All	0.822	0.779	0.825	0.891	0.785	0.820
SV3/SV1						
0.85-0.90	0.727	0.607	0.621	0.655	0.643	0.651
0.90-0.95	0.662	0.624	0.675	0.716	0.658	0.667
0.95-1.00	0.689	0.660	0.705	0.810	0.663	0.705
1.00-1.05	0.710	0.658	0.654	0.705	0.637	0.673
1.05-1.10	0.813	0.660	0.604	0.597	0.606	0.656
1.10-1.15	0.975	0.775	0.660	0.619	0.572	0.720
All	0.762	0.664	0.653	0.684	0.630	0.679
SV1JR/SV1						
0.85-0.90	0.878	0.894	0.878	0.927	1.051	0.926
0.90-0.95	0.917	0.933	0.905	0.915	0.987	0.932
0.95-1.00	0.969	0.986	1.000	0.996	0.963	0.983
1.00-1.05	0.909	0.930	0.963	0.997	0.974	0.954
1.05-1.10	0.933	0.922	0.910	0.930	0.970	0.933
1.10-1.15	0.954	0.942	0.939	0.921	0.921	0.935
All	0.927	0.934	0.933	0.948	0.978	0.944
SV2JR/SV1						
0.85-0.90	0.819	0.729	0.757	0.864	0.901	0.814
0.90-0.95	0.729	0.676	0.755	0.888	0.858	0.781
0.95-1.00	0.732	0.706	0.771	0.900	0.765	0.775
1.00-1.05	0.793	0.770	0.799	0.847	0.696	0.781
1.05-1.10	0.886	0.810	0.798	0.815	0.714	0.805
1.10-1.15	0.996	0.909	0.877	0.843	0.722	0.869
All	0.826	0.767	0.793	0.859	0.776	0.804
SV3JR/SV1						
0.85-0.90	0.745	0.616	0.631	0.728	0.732	0.690
0.90-0.95	0.653	0.594	0.666	0.778	0.728	0.684
0.95-1.00	0.711	0.686	0.731	0.847	0.697	0.734
1.00-1.05	0.769	0.714	0.706	0.744	0.660	0.718
1.05-1.10	0.849	0.714	0.642	0.632	0.626	0.692
1.10-1.15	0.979	0.818	0.702	0.610	0.576	0.737
All	0.784	0.690	0.680	0.723	0.670	0.709

Notes: The top panel reports the root mean squared option pricing error for the SV1 model for each moneyness-maturity bucket. For the other models, we report the ratio between the respective model RMSE and the SV1 RMSE.

Table 4: Model-Implied and Non-Parametric Risk-Neutral Moments and Risk Premia

Model / moment	Variance	Skewness	Kurtosis
Panel A: Moment Average			
	Horizon 30 days		
BKM	0.039	-1.375	6.180
SV1	0.040	-0.916	4.367
SV2	0.040	-1.133	5.385
SV3	0.041	-1.119	6.121
	Horizon 180 days		
BKM	0.049	-1.323	4.814
SV1	0.048	-1.642	7.206
SV2	0.049	-1.744	8.079
SV3	0.053	-2.074	11.462

	Mean	Variance	Skewness	Kurtosis
Panel B: Risk Premiums Average				
	Horizon 30 days			
SV1	0.096	0.001	-0.004	-0.021
SV2	0.102	0.002	0.013	0.024
SV3	0.100	0.002	0.090	0.400
	Horizon 180 days			
SV1	0.098	0.006	0.030	-0.040
SV2	0.102	0.008	0.100	0.357
SV3	0.101	0.009	0.505	3.026

	Variance	Skewness	Kurtosis
Panel C: Correlation with BKM			
	Horizon 30 days		
SV1	92.28	35.00	36.84
SV2	97.12	33.33	35.61
SV3	97.43	62.74	42.23
	Horizon 180 days		
SV1	97.99	35.76	24.84
SV2	98.05	34.83	24.91
SV3	98.11	51.26	29.30

Notes: The top panel reports the average model-implied risk-neutral moments, as well as the nonparameteric BKM moment estimate. The middle panel reports the model-implied moment risk premiums, computed as the difference between the risk-neutral and physical moments for the even moments, and the reverse for the uneven moments. The bottom panel reports the correlation coefficients (in percentages) between the model implied risk-neutral moments and the nonparametric BKM risk-neutral moments (Bakshi, Kapadia, and Madan, 2003).

Table 5: Parameter Estimates for SV Multifactor Models with Jumps in Returns

	SV1JR	SV2JR			SV3JR	
κ	0.952	3.171	0.276	4.146	0.295	0.347
	(0.262)	(0.582)	(0.195)	(0.693)	(0.215)	(0.160)
θ	0.034	0.032	0.002	0.027	0.010	0.007
	(0.015)	(0.006)	(0.038)	(0.005)	(0.593)	(0.023)
σ	0.361	0.682	0.239	0.799	0.269	0.052
	(0.006)	(0.012)	(0.009)	(0.015)	(0.010)	(0.003)
ρ	-0.913	-0.923	-0.965	-0.963	-0.966	-0.407
	(0.006)	(0.005)	(0.008)	(0.005)	(0.007)	(0.145)
η_v	0.667	0.159	1.190	0.639	1.350	-0.229
	(0.261)	(0.571)	(0.197)	(0.690)	(0.216)	(0.199)
η_s	2.479	3.017			2.995	
	(0.602)	(0.550)			(0.547)	
λ	0.508	0.534			0.534	
	(0.004)	(0.022)			(0.022)	
μ_s	-0.027	-0.009			-0.012	
	(0.010)	(0.006)			(0.007)	
σ_s	0.118	0.022			0.026	
	(0.002)	(0.003)			(0.003)	
μ_s^Q	-0.062	0.041			0.033	
	(0.003)	(0.002)			(0.004)	
σ_c	3.156	2.629			2.305	
	(0.031)	(0.028)			(0.025)	
Diffusive ERP	0.079	0.115			0.116	
Jump ERP	0.017	-0.027			-0.025	
Loglikelihood	3330.48	4003.58			4457.57	
BIC	-6565.19	-7867.87			-8732.31	

Notes: We report parameter estimates based on the joint likelihood from returns and options for the SV multifactor models with jumps in returns. Parameters are annualized and reported under the physical measure. In parentheses, we report the posterior standard deviation for each parameter. Diffusive (Jump) ERP represents the average risk premium due to the diffusive (jump) components.

Appendix

A Estimation

A.1 Characteristic function

The characteristic function of $X_t = \ln S_t$ is given by $\phi_{\ln S_t}(\omega) = E_t(\exp(j\omega X_{t+\tau}))$ in which j is the imaginary number. In particular, the closed form function is equal to

$$\begin{aligned}
E_t(\exp(j\omega X_{t+\tau})) &= \exp\big(j\omega X_t + \sum_{i=1}^{N} A_i(\tau) + \sum_{i=1}^{N} D_i(\tau)V_{it}\big), \qquad (25)\\
A_i(\tau) &= \kappa_i\theta_i\big(r_{i,-}\tau - \frac{2}{\sigma_i^2}\ln\big[\frac{1-g_i\exp(-d_i\tau)}{1-g_i}\big]\big),\\
D_i(\tau) &= r_{i,-}\frac{1-\exp(-d_i\tau)}{1-g_i\exp(-d_i\tau)},\\
r_{i,-} &= \frac{1}{\sigma_i^2}(-\rho_i\sigma_i(j\omega)+\kappa_i - d_i),\\
r_{i,+} &= \frac{1}{\sigma_i^2}(-\rho_i\sigma_i(j\omega)+\kappa_i + d_i),\\
g_i &= \frac{r_{i,-}}{r_{i,+}},\\
d_i &= \sqrt{(-\rho_i\sigma_i(j\omega)+\kappa_i)^2 - j\omega(j\omega-1)\sigma_i^2}.
\end{aligned}$$

A.2 Option price evaluation

To compute option prices, we rely on a formula proposed in Carr and Madan (1998) who start from the call price:

$$C_0(K,T) = \exp(-rT)\int_{\ln K}^{\infty}(\exp(x)-\exp(\ln K))f(x)dx, \quad \text{where } x = \ln S_T,$$

where $f(x)$ is the risk-neutral density of $\ln S_T$. Applying a Fourier transform and a dampening parameter α given by $c_T(k) = \exp(\alpha k)C_T(k)$ (since the call price is not square-integrable),

where $k = \ln K$ and $C_T(k) = C_0(K,T)$, this leads to

$$\psi_T(\omega) = \int_{-\infty}^{\infty} \exp(j\omega k) c_T(k) dk.$$

The inverse Fourier transform gives back the call price:

$$c_T(k) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(-j\omega k) \psi_T(\omega) d\omega, \tag{26}$$

$$C_T(k) = \frac{\exp(-\alpha k)}{2\pi} \int_{-\infty}^{\infty} \exp(-j\omega k) \psi_T(\omega) d\omega, \tag{27}$$

$$= \frac{\exp(-\alpha k)}{\pi} \int_{0}^{\infty} \exp(-j\omega k) \psi_T(\omega) d\omega, \tag{28}$$

where the last equality holds because $C_T(k)$ is real and so the Fourier transform is odd in its imaginary part and even in its real part. Carr and Madan (1998) derive an analytical expression of $\psi_T(\omega)$ such that the call price in Equation (26) exhibits one integral only:

$$\psi_T(\omega) = \frac{\exp(-rT)\phi_{\ln S_t}(\omega - (\alpha+1)j)}{\alpha^2 + \alpha - \omega^2 + j(2\alpha+1)\omega}. \tag{29}$$

Substituting the previous equation into the call price leads to

$$C_T(k) = \frac{\exp(-\alpha k)}{\pi} \int_{0}^{\infty} \exp(-j\omega k) \frac{\exp(-rT)\phi_{\ln S_t}(\omega - (\alpha+1)j)}{\alpha^2 + \alpha - \omega^2 + j(2\alpha+1)\omega} d\omega. \tag{30}$$

We apply Simpson's rule to evaluate the call price given by Equation (30).

A.2.1 How to make option evaluation faster ?

We discuss two simple numerical ways we employ to make the option evaluation faster.

1. **Evaluating the option formula for multiple strikes at the same maturity.** In Equation (30), all but one term within the integral do not depend on the strike. So we

can numerically compute the term:

$$\frac{\exp(-rT)\phi_{\ln S_t}(\omega-(\alpha+1)j)}{\alpha^2+\alpha-\omega^2+j(2\alpha+1)\omega},$$

and then integrate for all strikes simultaneously.

2. **Evaluating the option formula for multiple spot volatilities.** Note that the spot volatilities appear linearly within the exponential function of the characteristic function since it is given by

$$\phi_{\ln S_t}(\omega) \;=\; \exp\big(j\omega X_t + A(\tau) + \sum_{i=1}^{N} D_i(\tau)V_{it}\big).$$

To compute the option prices of these Z different spot volatilities, we proceed as follows

(a) We start by computing the characteristic function for $\Big(V_{11} \quad \dots \quad V_{N1}\Big)$, denoted by $\phi_{\ln S_t}(\omega, V_{11}, \dots, V_{N1})$.

(b) We rewrite the characteristic function as follows:

$$\phi_{\ln S_t}(\omega, V_{1z}, \dots, V_{Nz}) \;=\; \phi_{\ln S_t}(\omega, V_{11}, \dots, V_{N1}) \exp\big(\sum_{i=1}^{N} D_i \Delta V_{iz}\big), \text{ for } z = 2, \dots, Z,$$

where $\Delta V_{iz} = V_{iz} - V_{i1}$.

A.3 MCMC sampler

Using the standard Euler discretization, the N-factor stochastic volatility model with jumps in mean is specified under the physical measure as

$$\begin{aligned} y_{t+1} &= \mu_t + \sum_{i=1}^{N} \sqrt{V_{i,t}} z_{i,t+1} + J^s_{t+1} B_{t+1}, \\ V_{i,t+1} - V_{i,t} &= \kappa_i(\theta_i - V_{i,t}) + \sigma_i \sqrt{V_{i,t}} w_{i,t+1}, \\ &= \underbrace{\phi_i}_{\kappa_i \theta_i} - \kappa_i V_{i,t} + \sigma_i \sqrt{V_{i,t}} w_{i,t+1}, \\ \text{Corr}(w_{i,t}, z_{i,t}) &= \rho_i, \end{aligned}$$

in which $\mu_t = r_t - \delta_t - \frac{1}{2}(\sum_{i=1}^{N} V_{i,t}) - \lambda \bar{\mu}_{Js} + \eta_s(\sum_{i=1}^{N} V_{i,t})$, $J^s_{t+1} \sim N(\mu_{Js}, \sigma^2_{Js})$, B_{t+1} is a Bernoulli distribution with parameter λ, $z_{i,t+1}$ and $w_{i,t+1}$ are each standard Normally distributed. We sample the model parameters and the latent variables within an MCMC algorithm consisting of six steps. To simplify the notation, we first define the following variables:

$$\begin{aligned} \boldsymbol{\rho}^{V'}_t &= \begin{pmatrix} \rho_1 \sigma_1 V_{1,t} & \rho_2 \sigma_2 V_{2,t} & \dots & \rho_N \sigma_N V_{1,t} \end{pmatrix}, \\ \boldsymbol{\Sigma}^V_t &= \begin{pmatrix} \sigma_1^2 V_{1,t} & 0 & \dots & 0 \\ 0 & \sigma_2^2 V_{2,t} & \dots & 0 \\ & & \dots & \\ 0 & 0 & \dots & \sigma_N^2 V_{N,t} \end{pmatrix}, \\ \tilde{\boldsymbol{\omega}}_{t+1} &= \begin{pmatrix} V_{1,t+1} - V_{1,t} - \kappa_1(\theta_1 - V_{1,t}) \\ V_{2,t+1} - V_{2,t} - \kappa_2(\theta_2 - V_{2,t}) \\ \dots \\ V_{N,t+1} - V_{N,t} - \kappa_N(\theta_N - V_{N,t}) \end{pmatrix}. \end{aligned}$$

MCMC sampler:

1. Sampling $\{\phi_i, \sigma_i, \rho_i, \kappa_i^Q, \lambda, \mu_{Js}, \mu_{Js}^Q, \sigma_{Js}\}$ for $i = 1, \dots, N$. We alternate between the DREAM algorithm, see Vrugt et al. (2009), and the adaptive random walk (RW)

Metropolis method of Atchadé and Rosenthal (2005) for drawing those parameters.

(a) Every 2 out of 3 MCMC iterations, we use the DREAM algorithm. Multiple MCMC algorithms are run in parallel and we use the current state of each chain for making a proposal. The new parameter values are accepted according to a standard Metropolis ratio (since the DREAM proposal distribution is symmetric). To improve the acceptance rate, we apply the Multiple-try DREAM approach of Laloy and Vrugt (2012) based on the current states of the MCMC algorithms rather than the historical draws. In our implementation, the number of trials amounts to 4 at each iteration and we run 10 MCMC algorithms in parallel.

(b) Every 1 out of 3 MCMC iterations, we sample the model parameters one at a time using the RW method of Atchadé and Rosenthal (2005). Let us denote by $\zeta^{(h)}$, the realization of a model parameter at the h-th MCMC iteration. First, the algorithm samples a proposal from a Normal distribution $\tilde{\zeta} \sim N(\zeta^{(h)}, \sigma^2_{h-1,\zeta})$. Second, the draw $\tilde{\zeta}$ is accepted or rejected following a Metropolis acceptance ratio. Eventually, the variance $\sigma^2_{h,\zeta}$ is updated as follows: $\sigma^2_{h,\zeta} = \sigma^2_{h-1,\zeta} + \frac{\text{acc. rate}_\zeta - 0.44}{h^{0.6}}$, in which acc. rate stands for the current acceptance rate of the parameter ζ. To avoid negative values, we set $\sigma^2_{h,\zeta} = 0.0001$ when $\sigma^2_{h,\zeta} < 0$.

2. Sampling $\boldsymbol{\kappa} = \begin{pmatrix} \kappa_1 \\ \kappa_2 \\ \cdots \\ \kappa_N \end{pmatrix}$. Assuming the prior distribution $\boldsymbol{\kappa} \sim N(\underline{\mu}_\kappa, \underline{\Sigma}_\kappa)\mathbb{1}_{\boldsymbol{\kappa}>0}$, the conditional posterior distribution is given by

$$\begin{aligned} \boldsymbol{\kappa}|\Theta_{\setminus\boldsymbol{\kappa}}, \boldsymbol{V}_{1:T}, \boldsymbol{B}_{1:T} &\sim N(\bar{\mu}_\kappa, \bar{\Sigma}_\kappa)\mathbb{1}_{\boldsymbol{\kappa}>0}, \\ \bar{\Sigma}_\kappa &= (\sum_{t=1}^{T-1} \boldsymbol{X}_t \Omega^{-1}_{\tilde{V},t} \boldsymbol{X}_t + \underline{\Sigma}^{-1}_\kappa)^{-1} \\ \bar{\mu}_\kappa &= \bar{\Sigma}_\kappa (\sum_{t=1}^{T-1} \boldsymbol{X}'_t \Omega^{-1}_{\tilde{V},t} \boldsymbol{\Lambda}_{t+1} + \underline{\Sigma}^{-1}_\kappa \underline{\mu}_\kappa). \end{aligned}$$

in which $N(\bar{\mu}_\kappa, \bar{\Sigma}_\kappa)\mathbb{1}_{\boldsymbol{\kappa}>0}$ denotes a multivariate truncated Normal distribution and

$$\begin{aligned}
\boldsymbol{X}_t &= \begin{pmatrix} -V_{1,t} & 0 & \dots & 0 \\ 0 & -V_{2,t} & \dots & 0 \\ & & \dots & 0 \\ 0 & 0 & \dots & -V_{N,t} \end{pmatrix}, \\
\boldsymbol{\Lambda}_{t+1} &= \begin{pmatrix} V_{1,t+1} - V_{1,t} - \phi_1 \\ \dots \\ V_{N,t+1} - V_{N,t} - \phi_N \end{pmatrix} - \boldsymbol{\rho}_t^V \frac{(y_{t+1} - \mu_{t+1} - \mu_{Js}\mathbb{1}_{B_{t+1}=1})}{\sum_{i=1}^N V_{i,t} + \sigma_{Js}^2 \mathbb{1}_{B_{t+1}=1}}, \\
\Omega_{\tilde{V},t} &= \boldsymbol{\Sigma}_t^V - \frac{\boldsymbol{\rho}_t^V \boldsymbol{\rho}_t^{V'}}{\sum_{i=1}^N V_{i,t} + \sigma_{Js}^2 \mathbb{1}_{B_{t+1}=1}},
\end{aligned}$$

3. Sampling σ_c^2. Assuming an improper prior, we sample σ_c^2 from its conjugate inverse-gamma distribution. If we denote the weighting scheme in the option log-likelihood function by λ_t (so $\lambda_t = H_t$ when we estimate under both options and returns, and $\lambda_t = 1$ when estimation is done with options only), then the conditional posterior distribution of σ_c^2 leads to

$$\begin{aligned}
\sigma_c^2 | \Theta_{\setminus \sigma_c^2}, \boldsymbol{V}_{1:T}, \boldsymbol{B}_{1:T} &\propto \prod_{t=1}^T (2\pi\sigma_c^2)^{-\frac{H_t}{2\lambda_t}} \exp\big(-\frac{1}{\sigma_c^2} (\frac{\sum_{i=1}^{H_t} (C_{t,i} - \hat{C}_{t,i})^2}{2\lambda_t}) \big), \\
&= (2\pi\sigma_c^2)^{-\frac{1}{2}\sum_{t=1}^T \frac{H_t}{\lambda_t}} \exp\big(-\frac{1}{\sigma_c^2} (\sum_{t=1}^T \frac{\sum_{i=1}^{H_t} (C_{t,i} - \hat{C}_{t,i})^2}{2\lambda_t}) \big), \\
&\sim IG(\sum_{t=1}^T \frac{H_t}{2\lambda_t} - 1, \sum_{t=1}^T \frac{\sum_{i=1}^{N_t} (C_{t,i} - \hat{C}_{t,i})^2}{2\lambda_t}).
\end{aligned}$$

4. Sampling η_s assuming a prior distribution denoted by $\eta_s \sim N(\underline{\mu}_{\eta_s}, \underline{\sigma}_{\eta_s}^2)$. The conditional

posterior distribution is given by

$$\begin{aligned}
\eta_s|\Theta_{\setminus\eta_s}, \boldsymbol{V}_{1:T}, \boldsymbol{B}_{1:T} &= N(\bar{\mu}_{\eta_s}, \bar{\sigma}^2_{\eta_s}), \\
\bar{\sigma}^2_{\eta_s} &= \Big(\frac{1}{\underline{\sigma}^2_{\eta_s}} + \sum_{t=1}^{T} \frac{\bar{V}^2_{t-1}}{(\sum_{i=1}^{N}(1-\rho_i^2)V_{i,t-1} + \sigma^2_{Js}\mathbb{1}_{B_t=1}))}\Big)^{-1}, \\
\bar{\mu}_{\eta_s} &= \bar{\sigma}^2_{\eta_s}\Big(\frac{\underline{\mu}_{\eta_s}}{\underline{\sigma}^2_{\eta_s}} + \sum_{t=1}^{T} \frac{(y_t - \tilde{\mu}_t - \mu_{Js}\mathbb{1}_{B_t=1} - \boldsymbol{\rho}^{V\prime}_{t-1}(\boldsymbol{\Sigma}^V_{t-1})^{-1}\tilde{\boldsymbol{\omega}}_t)}{(\sum_{i=1}^{N}(1-\rho_i^2)V_{i,t-1} + \sigma^2_{Js}\mathbb{1}_{B_t=1})}\bar{V}_{t-1}\Big),
\end{aligned}$$

in which $\tilde{\mu}_t = \mu_t - \eta_s(\sum_{i=1}^{N} V_{i,t}) = r_t - \delta_t - \frac{1}{2}(\sum_{i=1}^{N} V_{i,t}) - \lambda\bar{\mu}_{Js}$ and $\bar{V}_{t-1} = (\sum_{i=1}^{N} V_{i,t-1})$.

5. Sampling jump times $\boldsymbol{B}_{1:T}$. To sample the jump parameters, we operate similarly as Eraker (2004). The probability of a jump at time $t+1$ is given by

$$f(B_{t+1} = 1|\Theta, \boldsymbol{V}_{1:T}, B_{\setminus t+1}, y_{1:T}) = \frac{f(y_{t+1}|\boldsymbol{V}_{t+1}, \Theta, B_{t+1} = 1)f(B_{t+1} = 1|\lambda)}{\sum_{i=0}^{1} f(y_{t+1}|\boldsymbol{V}_{t+1}, \Theta, B_{t+1} = i)f(B_{t+1} = i|\lambda)},$$

in which $f(B_{t+1} = 1|\lambda) = \lambda$ and $f(y_{t+1}|\boldsymbol{V}_{t+1}, \Theta, B_{t+1} = 1)$ is a normal density function specified as,

$$y_{t+1}|\boldsymbol{V}_{t+1}, \Theta, B_{t+1} = 1 \sim N\big(\mu_{t+1} + \boldsymbol{\rho}^{V\prime}_t(\boldsymbol{\Sigma}^V_t)^{-1}\tilde{\boldsymbol{\omega}}_{t+1} + \mu_{J_s}, \sum_{i=1}^{N} V_{i,t} - \boldsymbol{\rho}^{V\prime}_t(\boldsymbol{\Sigma}^V_t)^{-1}\boldsymbol{\rho}^V_t + \sigma^2_{J_s}\big).$$

When there are no jumps ($B_{t+1} = 0$), we have the same distribution with $\mu_{Js} = \sigma_{Js} = 0$.

6. Sampling the spot variances, $\boldsymbol{V}_{1:T}$. We apply a conditional sequential Monte Carlo (cSMC) algorithm (see, Andrieu et al., 2010) with backward sampling to sample the spot variances. For each factor, a cSMC algorithm is run to sample one factor at a time using the state variable quantile method presented in Appendix **??**. Drawing one factor has also the advantage to simplify the implementation of the full truncation method, (see, Lord et al., 2010) in order to deal with potentially negative spot variances.

A.4 State Variable Quantile method

While the algorithm of Appendix A.3 is in principle general enough to perform full inference on all considered models, it is not sufficient to handle the large option panel data setting we face in this paper. The computational time required for step 6 is simply too high because the particle filter implies the computation of option prices for each of the J particles, i.e. V_t^j $j = 1, \ldots, J$, at each time t and for each option $h \in [1, H_t]$. We solve this issue with our State Variable Quantile (SVQ) method.

The SVQ method reduces this computational burden to Q computations where Q does not depend on J. To do so, it approximates each option price $C_{t,h}^M(V_t^j, \Theta)$ as a function of the particles using a polynomial regression calibrated on Q option prices. For each option price $C_{t,h}$, our approach proceeds as follows.

1. Take Q quantiles $\{V_t^1, ..., V_t^Q\}$ evenly spaced over [0,100] from the predictive filtered spot variance distribution, i.e. $V_t|Y_{1:t-1}$.

2. Evaluate the option prices $\{C_{t,h}^M(V_t^1, \Theta), ..., C_{t,h}^M(V_t^Q, \Theta)\}$.

3. Minimize the sum of squared pricing errors using a polynomial function of order p:

$$(\hat{\beta}_1, \ldots, \hat{\beta}_{p+1}) = \underset{(\beta_1,\ldots,\beta_{p+1})}{\arg\min} \sum_{q=1}^{Q} \left(C_{t,h}^M(V_t^q, \Theta) - \beta_1 - \beta_2 V_t^q - \ldots - \beta_{p+1}(V_t^q)^p \right)^2. \quad (31)$$

4. For each particle $j \in [1, J]$, set the option price to:

$$C_{t,h}^M(V_t^j, \Theta) = \hat{\beta}_1 + \hat{\beta}_2 V_t^j + \ldots + \hat{\beta}_{p+1}(V_t^j)^p. \quad (32)$$

The implementation of the SVQ method requires a choice regarding the number Q of quantiles and the order p of the polynomial regression. Based on extensive experimentation with these tuning parameters, we set $p = 3$ and $Q = 12$ in our empirical implementation.

B Partial Leverage Correlations

The critical role of the third volatility factor in describing the leverage correlation can also be seen by inspecting the contribution of the individual variance factors in equation (12). To this end, we define the partial leverage correlations as:

$$\text{Corr}_t\left(\frac{dS_t}{S_t}, dV_{j,t}\right) = \frac{\rho_j\sqrt{V_{j,t}}}{\sqrt{\sum_{i=1}^N V_{i,t}}}. \tag{33}$$

Figure A.1 in the Appendix plots these partial leverage correlations for the SV2 and SV3 models. It is striking that in the SV2 model, the two leverage correlations are strongly negatively correlated (-0.920). In the SV3 model, the correlation between the first and the second partial leverage correlations is -0.229, the correlation between the first and third is 0.459, and it is -0.675 between the second and third leverage correlations. These patterns again illustrate the additional flexibility provided by the third SV factor.

C Computing Model-Implied Moments

C.1 Matlab code for computing the cumulants

The centered moments given in (15) are derived from the cumulant of the log characteristic function, which amounts to

$$\log E_t(\exp(j\omega X_{t+\tau})) \quad = \quad j\omega X_t + \sum_{i=1}^N \big(A_i(\tau) + D_i(\tau)V_{it}\big), \tag{34}$$

where the terms are defined in Appendix A.1. The log characteristic function is a linear function of the variance factors and the terms $A_i(\tau)$, $D_i(\tau)$ have the same expression across factors. Therefore, the cumulants are also linear functions of the variance factors. In this section, we present the Matlab code for computing the coefficients $A_i^{(j)}(\Theta,\tau)$ and $D_i^{(j)}(\Theta,\tau)$ of Equation (15) (respectively denoted by A_j and D_j in the code below) for $j = 2,\ldots,4$.

```
%%% Variable initialization
syms kappaQ thetaQ sigma rho tau q V_t r_t omega real
%%% Characteristic function of the SV model
omega_j = j*omega;
d = sqrt((-rho*sigma*(omega_j) + kappaQ)^2 - (omega_j)*(omega_j-1)*sigma^2);
r_neg = (-rho*sigma*(omega_j) + kappaQ - d)/sigma^2;
r_pos = (-rho*sigma*(omega_j) + kappaQ + d)/sigma^2;
g = r_neg/r_pos;
A = kappaQ*thetaQ*(r_neg*tau - (2/sigma^2)*log((1-g*exp(-d*tau))/(1-g)));
D = r_neg*(1-exp(-d*tau))/(1-g*exp(-d*tau));

%%% Computing the coefficients:
A_2 = simplify(subs(diff(diff(A,omega),omega),omega,0)*(j)^(-2));
D_2 = simplify(subs(diff(diff(D,omega),omega),omega,0)*(j)^(-2));
A_3 = simplify(subs(diff(diff(diff(A,omega),omega),omega),omega,0)*(j)^(-3));
D_3 = simplify(subs(diff(diff(diff(D,omega),omega),omega),omega,0)*(j)^(-3));
A_4 = simplify(subs(diff(diff(diff(diff(A,omega),omega),omega),omega),omega,0)*(j)^(-4));
D_4 = simplify(subs(diff(diff(diff(diff(D,omega),omega),omega),omega),omega,0)*(j)^(-4));
```

C.1.1 Moment Estimates

To provide some intuition, consider the moments evaluated at $\Theta = \bar{\Theta}$, i.e. the posterior mean of the parameters, and consider the maturity $\tau = 30/365$. This gives the following risk-neutral moments:

1. Risk-neutral variance of the three SV models:

$$\begin{aligned} SV1 &= 0.00019 + 0.0813V_t, \\ SV2 &= 0.00035 + 0.0729V_{1t} + 0.0851V_{2t}, \\ SV3 &= 0.00037 + 0.0743V_{1t} + 0.0861V_{2t} + 0.0789V_{3t}. \end{aligned}$$

2. Risk-neutral third centered moment of the three SV models:

$$\begin{aligned} SV1 &= -0.00001 - 0.0036V_t, \\ SV2 &= -0.00002 - 0.0051V_{1t} - 0.0030V_{2t}, \\ SV3 &= -0.00003 - 0.0086V_{1t} - 0.0032V_{2t} - 0.0003V_{3t}. \end{aligned}$$

3. Risk-neutral fourth centered moment of the three SV models:

$$\begin{aligned} \text{SV1: } \text{kt}_Q &= 3.64 \times 10^{-7} + 0.0003V_t + 0.0198V_t^2, \\ \text{SV2: } \text{kt}_Q &= 1.67 \times 10^{-6} + \begin{pmatrix} 0.00069 & 0.00033 \end{pmatrix} \boldsymbol{V}_t + 3 \begin{pmatrix} 0.0729 & 0.0851 \end{pmatrix} \boldsymbol{V}_t \boldsymbol{V}_t' \begin{pmatrix} 0.0729 \\ 0.0851 \end{pmatrix}, \\ \text{SV3: } \text{kt}_Q &= 3.55 \times 10^{-6} + \begin{pmatrix} 0.0015 \\ 0.0004 \\ 0.0002 \end{pmatrix}' \boldsymbol{V}_t + 3 \begin{pmatrix} 0.0743 & 0.0861 & 0.0789 \end{pmatrix} \boldsymbol{V}_t \boldsymbol{V}_t' \begin{pmatrix} 0.0743 \\ 0.0861 \\ 0.0789 \end{pmatrix}, \end{aligned}$$

Based on our observations, the spot variances of the various SV models exhibit a high degree of similarity. Assuming that they are identical across SV models allows us to gain insights into the impact of each factor on the risk-neutral moment. Specifically, when comparing SV3 to SV1 and assuming identical total variance, we can identify the differences in the risk-neutral variance at a 30-day horizon between SV3 and SV1 as follows:

$$\begin{aligned} SV3 - SV1 &= 0.00037 + 0.0743V_{1t} + 0.0861V_{2t} + 0.0789V_{3t} - 0.00019 - 0.0813V_t, \\ &= 0.00018 - 0.0070V_{1t} + 0.0048V_{2t} - 0.0024V_{3t}. \end{aligned}$$

We conclude that the second factor has an amplifying effect on the risk-neutral variance compared to the SV1 variance. On the other hand, the remaining two factors have a diminishing effect relative to the SV1 variance.

D Computing Nonparametric Risk-Neutral Moments

Building on Bakshi and Madan (2000), Bakshi, Kapadia, and Madan (2003) put forward a nonparametric procedure for computing risk-neutral moments from options. They show that under the martingale pricing measure, the discounted risk-neutral expectations of squared, cubic, and quartic stock returns $r_{t,t+\tau} \equiv (\ln S_{t+\tau} - \ln S_t)$ at time t with the forward-looking horizon τ are given by V, W, X:

$$\begin{aligned}
V \equiv E_t^*[r_{t,t+\tau}^2] &= \int_S^\infty \frac{2\left(1 - \ln\frac{K}{S}\right)}{K^2} C dK + \int_0^S \frac{2\left(1 - \ln\frac{K}{S}\right)}{K^2} P dK \\
W \equiv E_t^*[r_{t,t+\tau}^3] &= \int_S^\infty \frac{6\ln\frac{K}{S} - 3\left(\ln\frac{K}{S}\right)^2}{K^2} C dK + \int_0^S \frac{6\ln\frac{K}{S} - 3\left(\ln\frac{K}{S}\right)^2}{K^2} P dK \\
X \equiv E_t^*[r_{t,t+\tau}^4] &= \int_S^\infty \frac{12\left(\ln\frac{K}{S}\right)^2 - 4\left(\ln\frac{K}{S}\right)^3}{K^2} C dK + \int_0^S \frac{12\left(\ln\frac{K}{S}\right)^2 - 4\left(\ln\frac{K}{S}\right)^3}{K^2} P dK
\end{aligned} \tag{35}$$

where S is the underlying spot price and K is the strike price. The OTM call and put prices are denoted by C and P, respectively, with maturity τ and strike K.

Option-implied moments (volatility, skewness, and kurtosis) can be inferred from these risk-neutral expectations. The τ-period risk-neutral moments are given by:

$$\begin{aligned}
\text{Vol}_{t,t+\tau} &= \sqrt{e^{r\tau}V - \mu^2} \\
\text{Skew}_{t,t+\tau} &= (e^{r\tau}W - 3\mu e^{r\tau}V + 2\mu^3)/\text{Vol}(t,\tau)^3 \\
\text{Kurt}_{t,t+\tau} &= (e^{r\tau}X - 4\mu e^{r\tau}W + 6e^{r\tau}\mu^2 V - 3\mu^4)/\text{Vol}(t,\tau)^4
\end{aligned} \tag{36}$$

where r is the risk-free rate (linearly interpolated from the OptionMetrics zero-curve data) and $\mu = e^{r\tau} - 1 - \frac{e^{r\tau}}{2}V - \frac{e^{r\tau}}{6}W - \frac{e^{r\tau}}{24}X$ is the expected return over period τ.

Since there is no continuity of option prices, we calculate the integrals using cubic splines. Furthermore, we only compute the moments for days that have at least two OTM call prices and two OTM put prices available. For each maturity on a given day, we interpolate implied volatilities using a cubic spline across moneyness levels (K/S) to obtain a continuum of implied volatilities. For moneyness levels below (above) the available moneyness in the data, we use the implied volatility of the lowest (highest) available strike price. Doing so, we

obtain a grid of 1000 implied volatilities for moneyness levels between 0.01 % and 300 %. These implied volatilities are converted into call and put prices as follows: moneyness levels smaller than 100 % are used to generate put prices and moneyness levels larger than 100 % are used to generate call prices. All integrals are computed using trapezoidal numerical integration. We use standard linear interpolation to calculate the moments at a given future horizon (30-day, 180-day in this paper). To compute the risk-neutral moments in this paper, we use data from OptionMetrics on S&P 500 index call and put options between 1996 and 2019 and apply the following filters so that we discard options (1) with fewer than 5 days and more than 365 days to maturity; (2) with implied volatility less than 5% and greater than 150%.; (3) with volume or open interest less than 5 contracts; (4) with quotes that suggest data errors: options for which the best bid exceeds the best offer, options with a zero bid price, and options with negative put-call parity implied price; (5) with price less than 50 cents.

E Additional Estimation Details

Figure A.1: Partial Leverage Correlations. SV Models

Panel (A): SV2: returns and variance factor 1

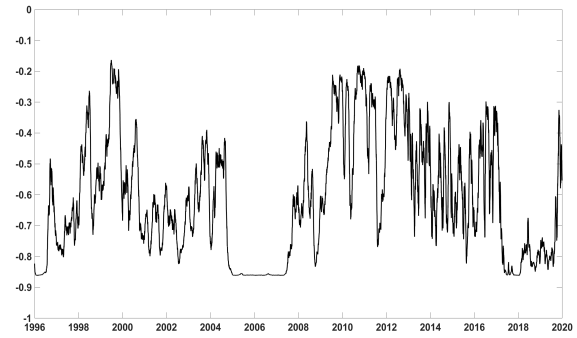

Panel (B): SV2: returns and variance factor 2

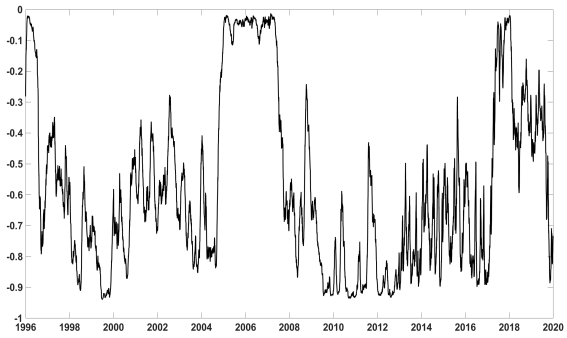

Panel (C): SV3: returns and variance factor 1

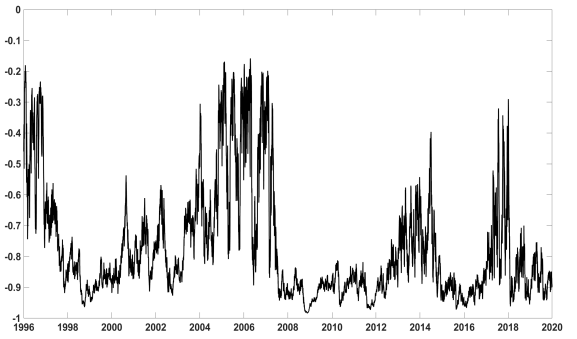

Panel (D): SV3: returns and variance factor 2

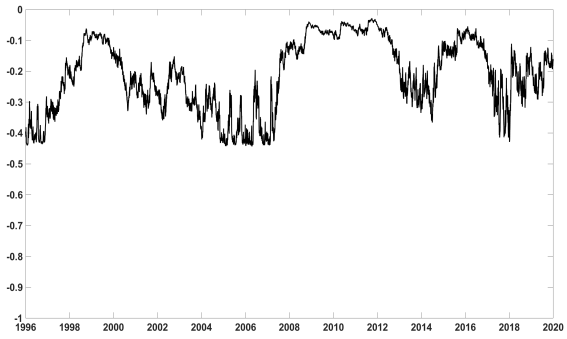

Panel (E): SV3: returns and variance factor 3

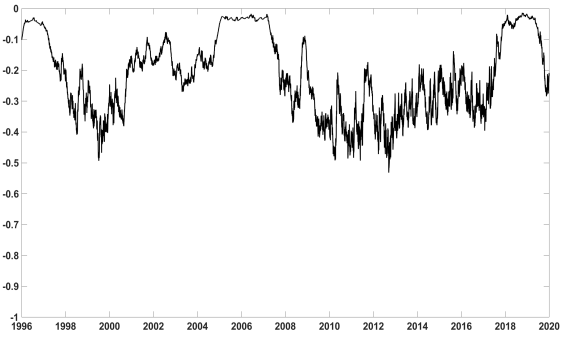

Notes: We plot the filtered partial leverage correlation path estimated from the different factor in the SV2 and SV3 models .

Figure A.2: Riskneutral Volatility

Panel (A): Volatility SV1-SV1JR: 30-days

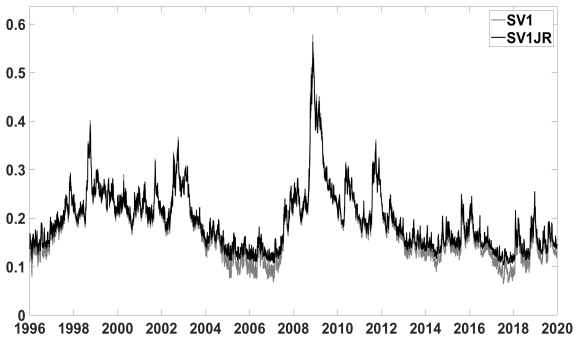

Panel (B): Volatility SV1-SV1JR: 180-days

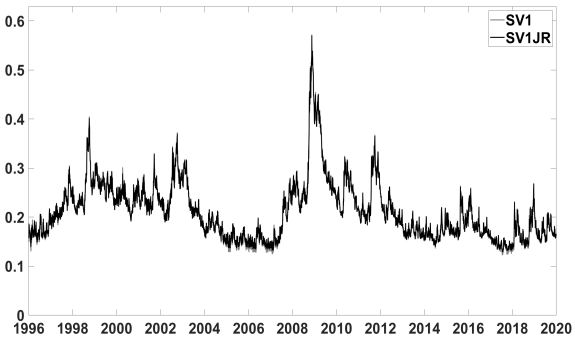

Panel (C): Volatility SV2-SV2JR: 30-days

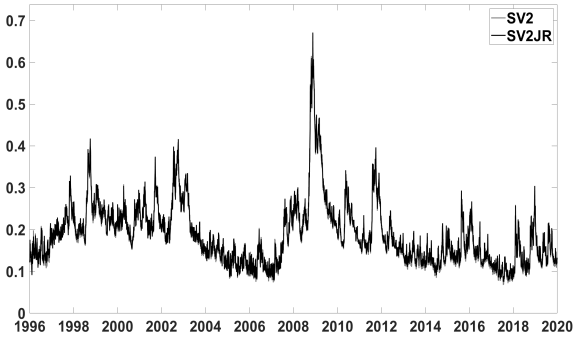

Panel (D): Volatility SV2-SV2JR: 180-days

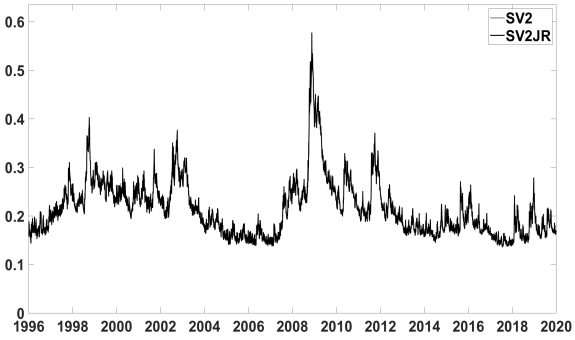

Panel (E): Volatility SV3-SV3JR: 30-days

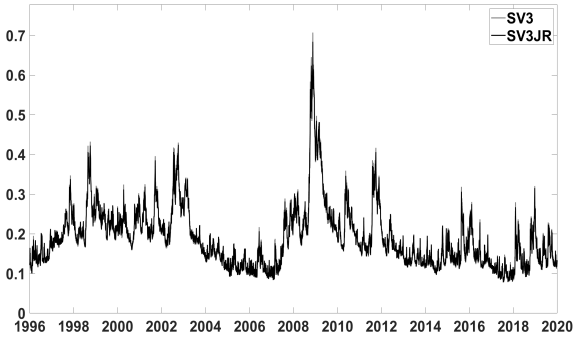

Panel (F): Volatility SV3-SV3JR: 180-days

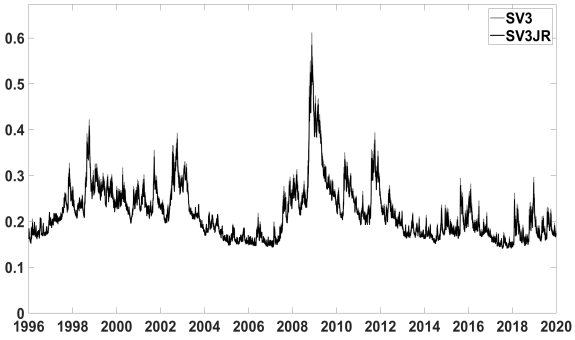

Notes: We plot the 30 and 180 days riskneutral volatilty.

Figure A.3: Time Series of Non-Standardized Moments and Risk Premia (30 day Maturity)

Panel (A): SV1 moment 3

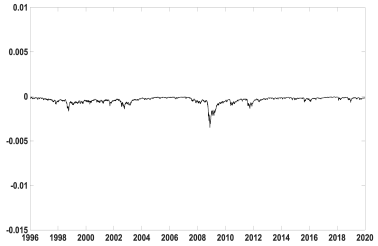

Panel (B): SV2 moment 3

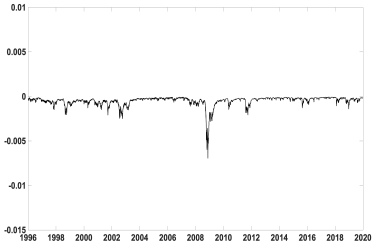

Panel (C): SV2 moment 3

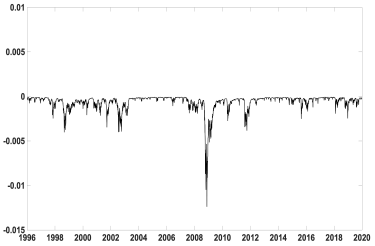

Panel (D): SV1 moment 3 RP

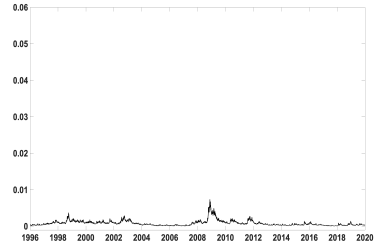

Panel (E): SV2 moment 3 RP

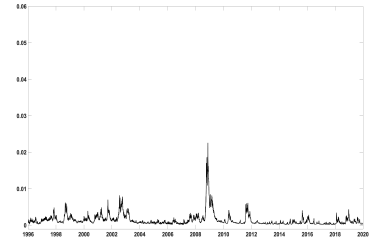

Panel (F): SV3 moment 3 RP

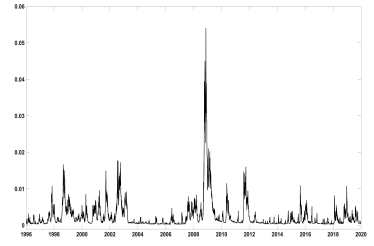

Panel (G): SV1 moment 4

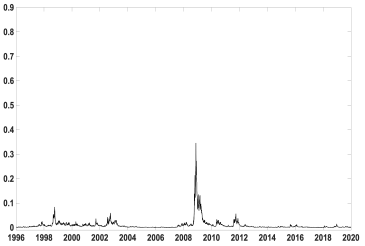

Panel (H): SV2 moment 4

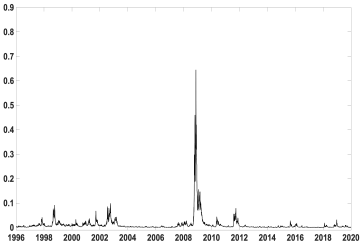

Panel (I): SV3 moment 4

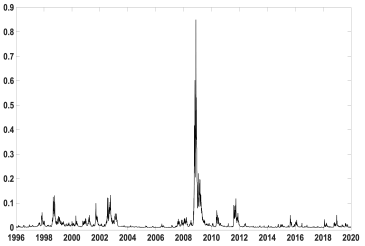

Panel (J): SV1 moment 4 RP

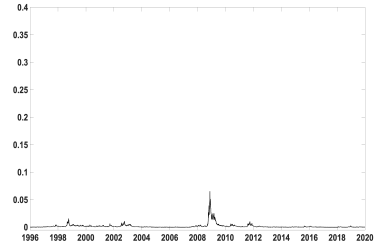

Panel (K): SV2 moment 4 RP

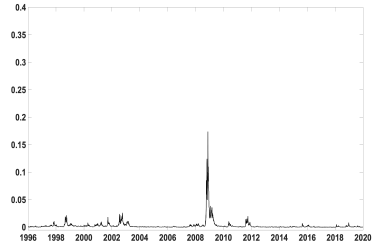

Panel (L): SV3 moment 4 RP

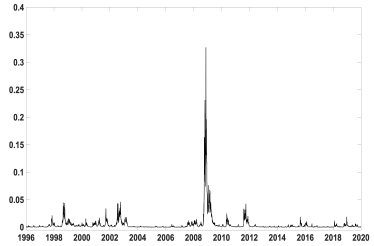

Notes: We plot the 30-day riskneutral third moment paths estimated from respectively the SV1, SV2 and SV3 models in the top row, the third moment risk premium paths in the second row. The third row plots the riskneutral fourth moment paths, and the last row plots the fourth moment risk premium paths. The sample period runs between January 4, 1996 and December 31, 2019.

Figure A.4: Term Structure of Risk-Neutral Moments. SV Models

Panel (A): Variance

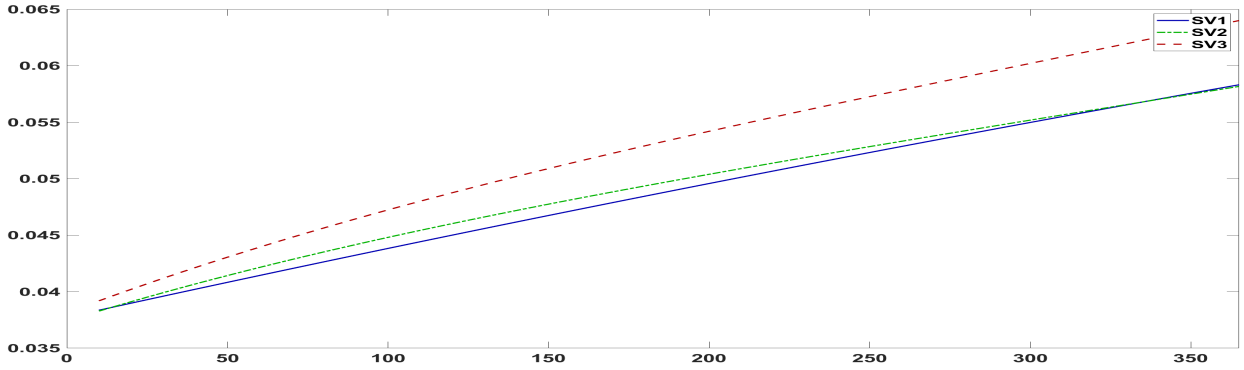

Panel (B): Skewness

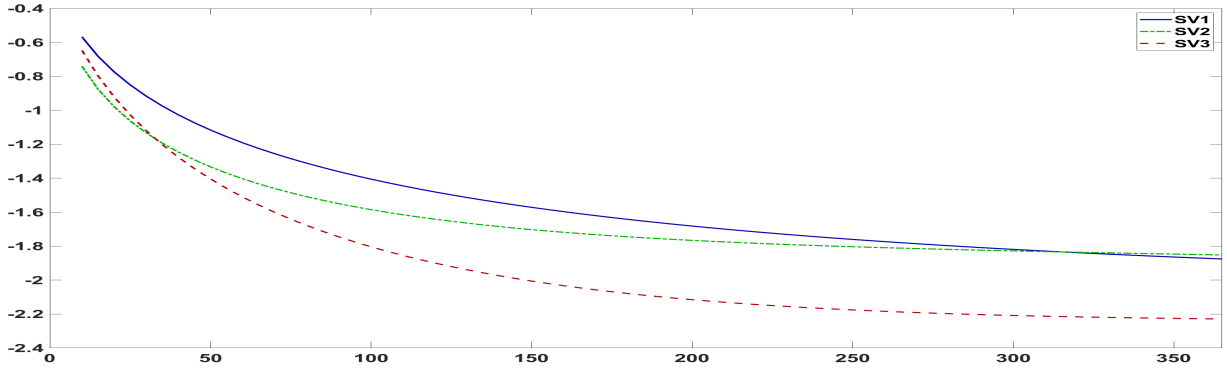

Panel (C): Kurtosis

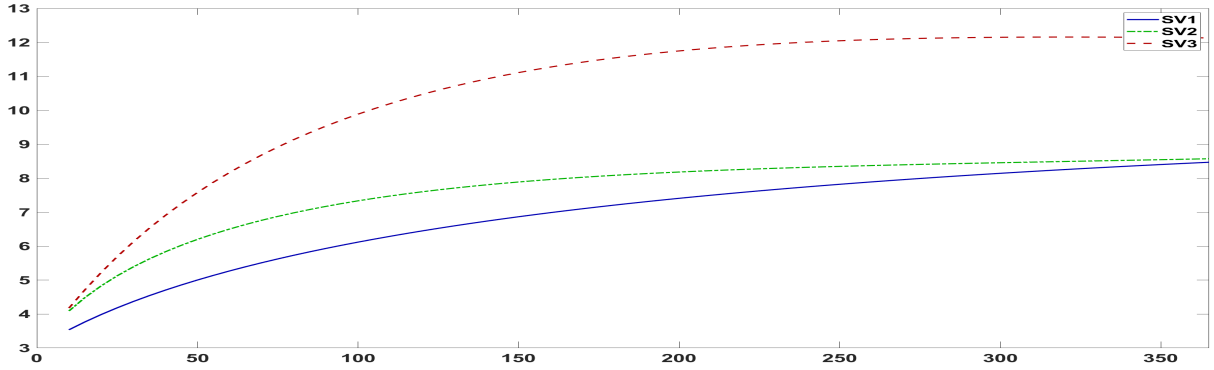

Notes: We plot the unconditional term structure (from one day up one year) of the riskneutral variance, skewness and kurtosis of the SV models.

Figure A.5: Impact of the Third Factor in SV3 Model on Risk-Neutral Moments.

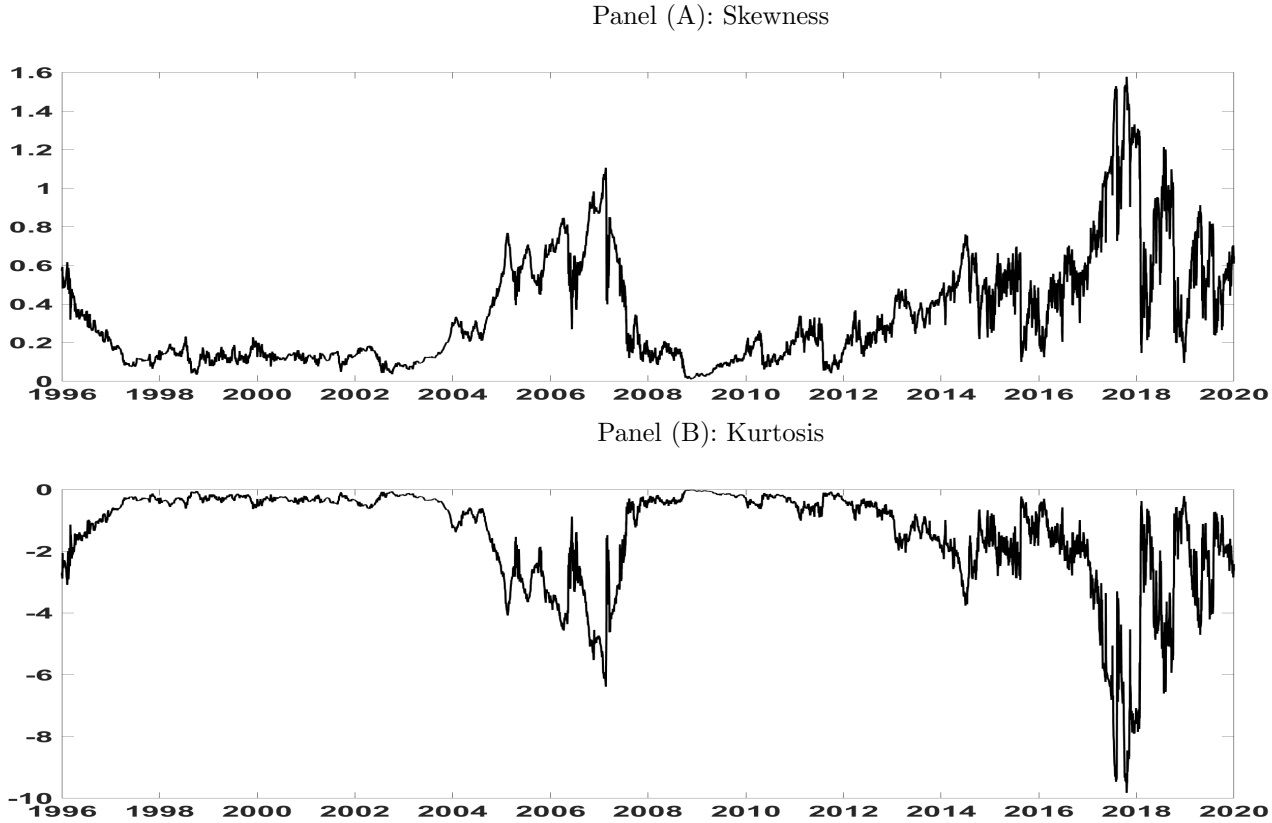

Notes: We consider the difference of the risk-neutral moments computed with the posterior mean of the third factor increased by one standard deviation with respect to the those moments computed with the posterior mean of the third factor. The sample period is from January 4, 1996 to December 31, 2019.

Figure A.6: The Risk-Neutral Term Structure.

Panel (A): SV1 - Variance

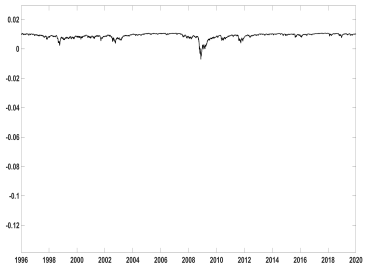

Panel (B): SV2 - Variance

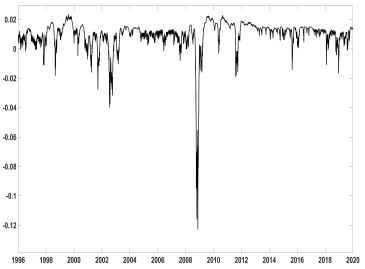

Panel (C): SV3 - Variance

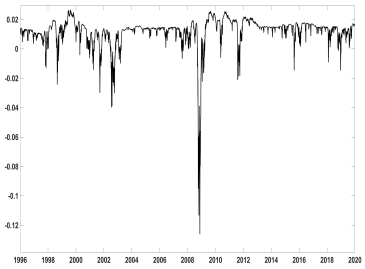

Panel (D): SV1 - Skewness

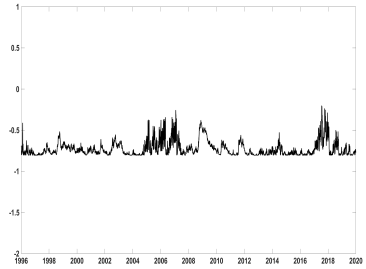

Panel (E): SV2 - Skewness

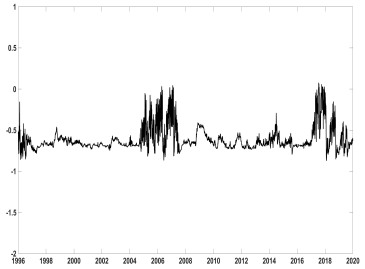

Panel (F): SV3 - Skewness

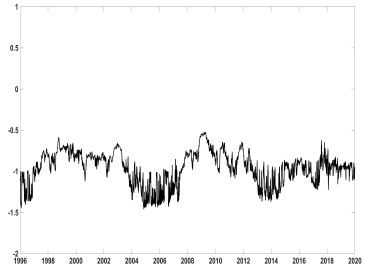

Panel (G): SV1 - Kurtosis

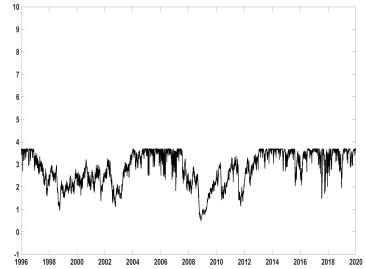

Panel (H): SV2 - Kurtosis

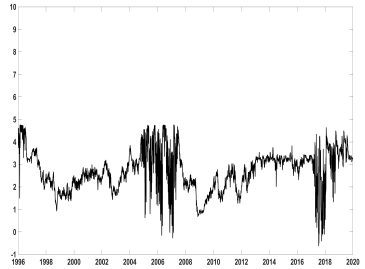

Panel (I): SV3 - Kurtosis

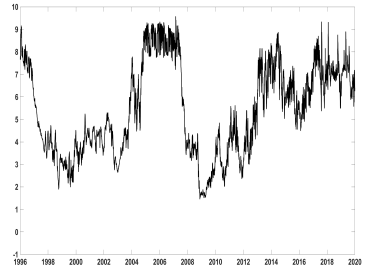

We plot the difference between the 180-day and 30-day risk-neutral variance, skewness and kurtosis. The sample period is from January 4, 1996 to December 31, 2019.

Figure A.7: Term Structure of Risk Premiums. SV Models

Panel (A): Variance RP

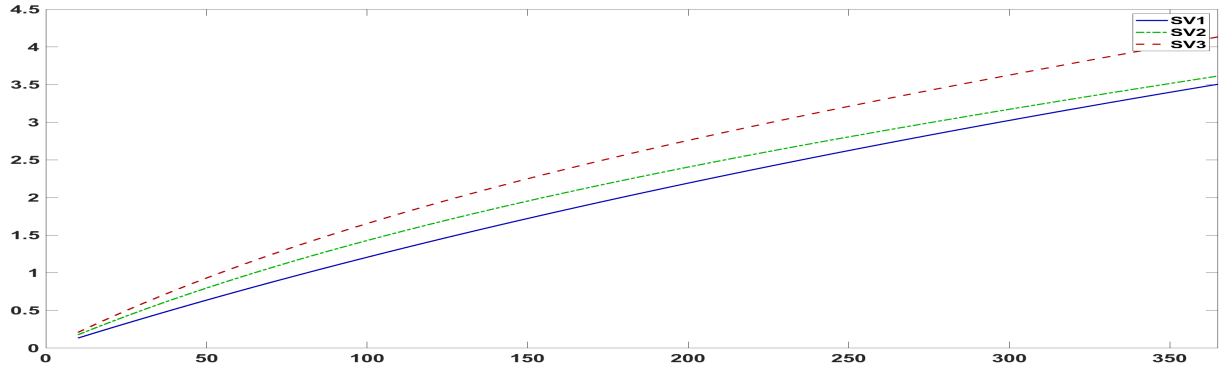

Panel (B): Skewness RP

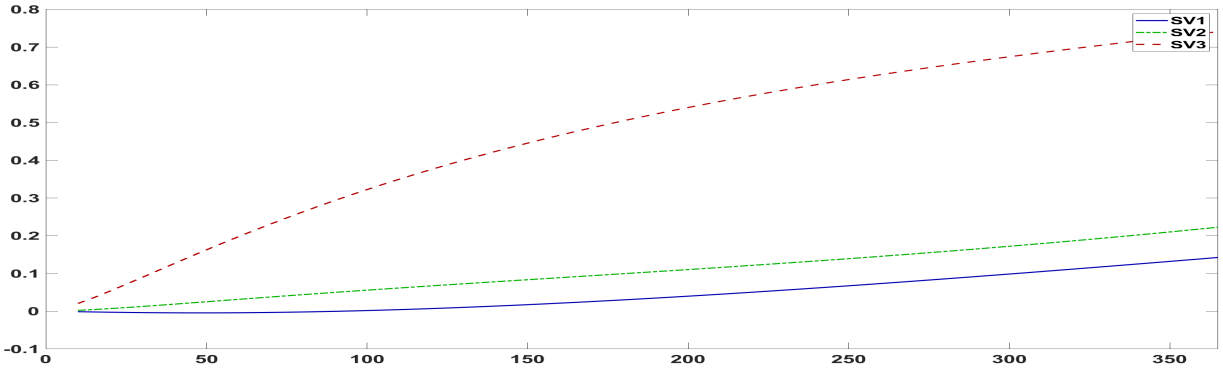

Panel (C): Kurtosis RP

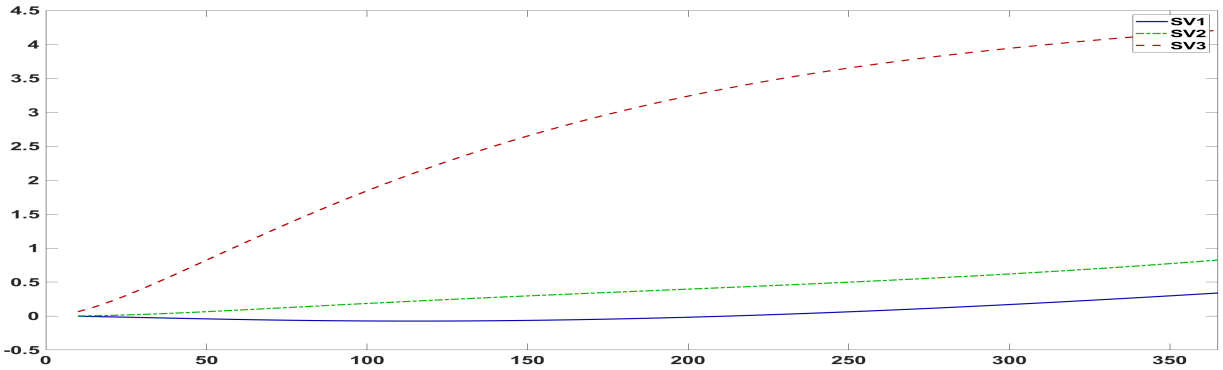

Notes: We plot the unconditional term structure (from one day up one year) of the variance, skewness and kurtosis risk premia of the SV models.

Figure A.8: Term Structure of Risk-Neutral Moments. SV2 and SV2JR Models

Panel (A): Variance

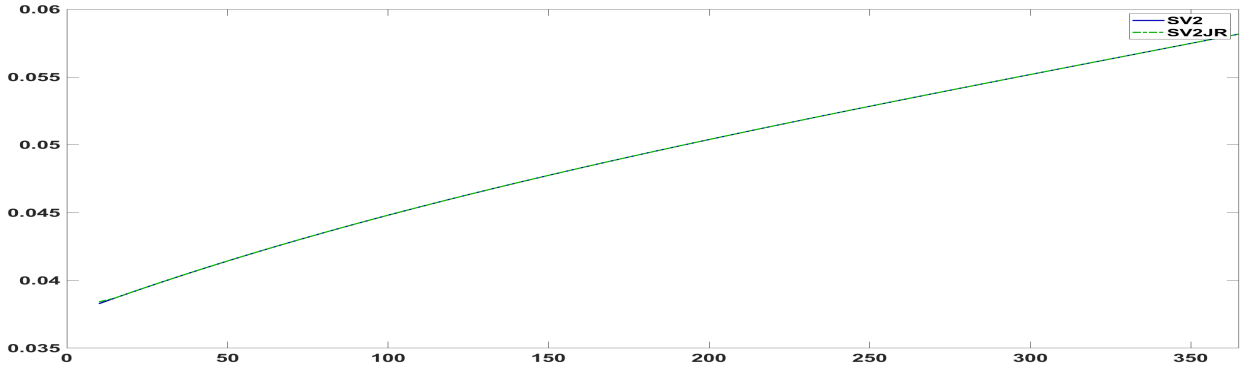

Panel (B): Skewness

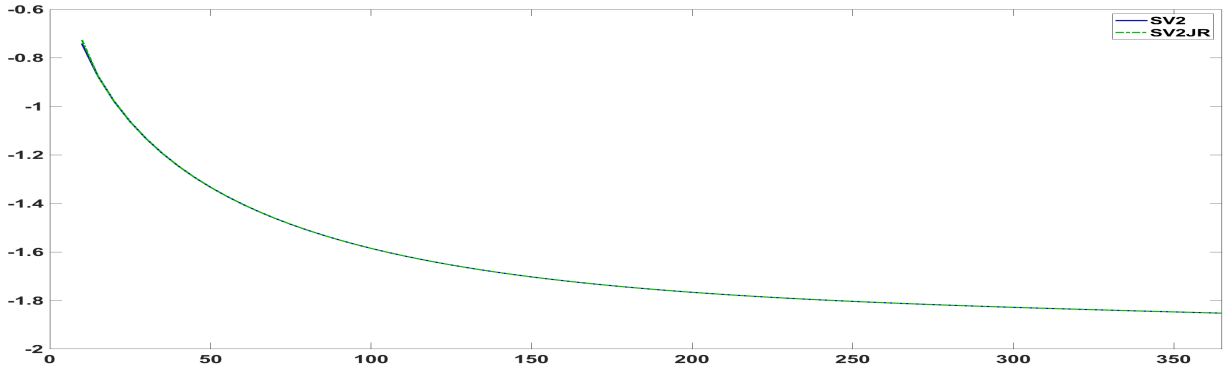

Panel (C): Kurtosis

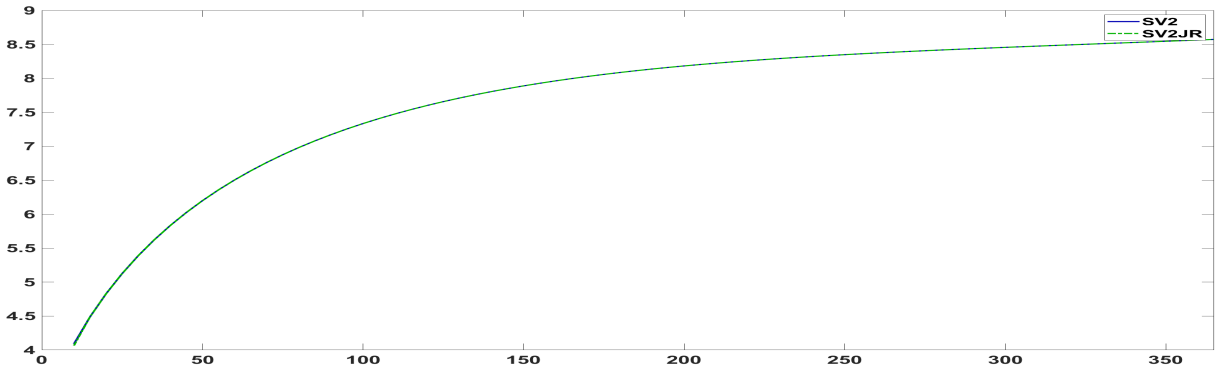

Notes: We plot the unconditional term structure (from one day up one year) of the riskneutral variance, skewness and kurtosis of the SV, SV2JR models.

Figure A.9: The Term Structure of the Variance Risk Premium.

Panel (A): SV1 model

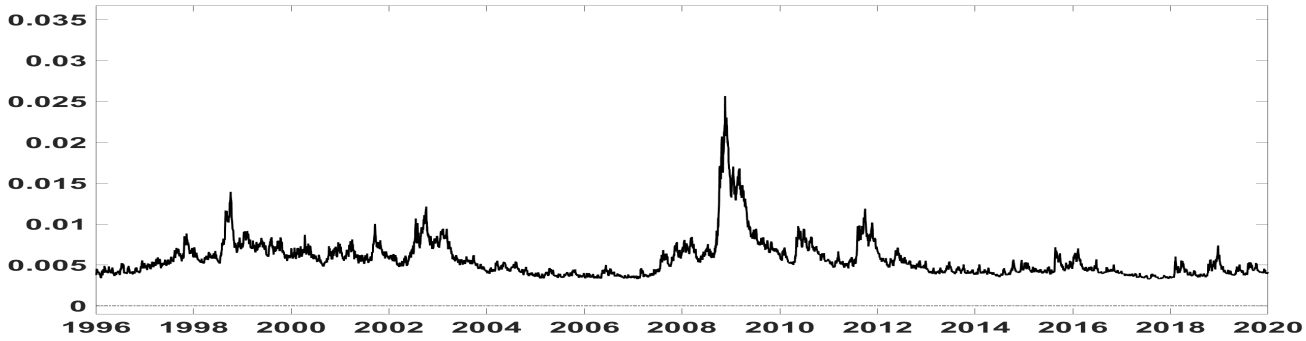

Panel (B): SV2 model

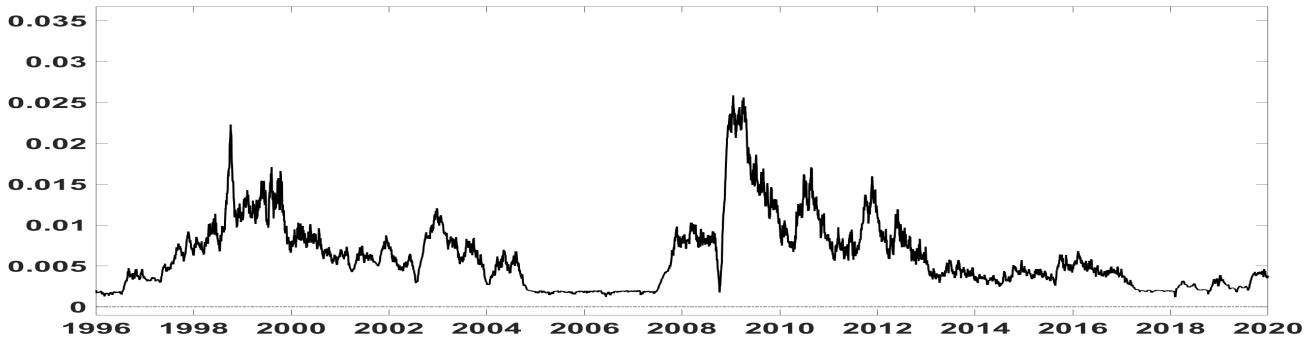

Panel (C): SV3 model

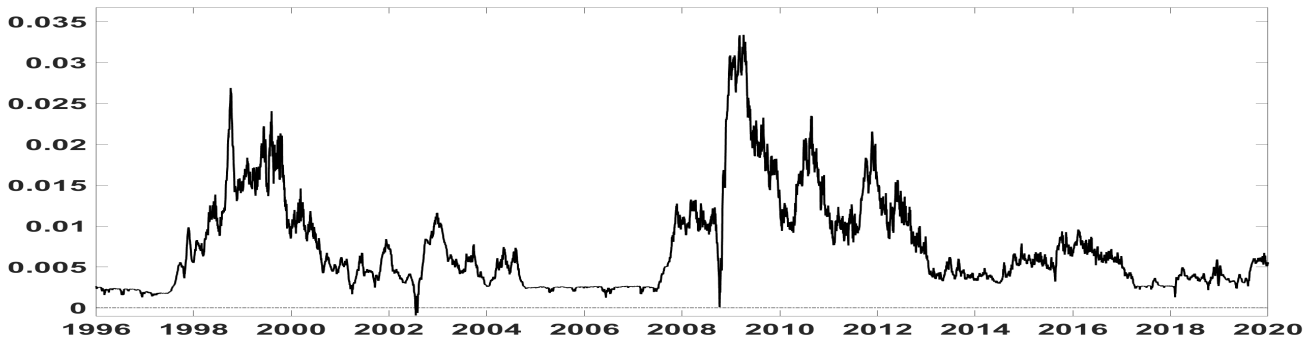

Notes: We plot the difference between the 365-day and 180-day variance risk premium. The sample period is from January 4, 1996 to December 31, 2019.

Table A.1: Out-Of-Sample Option Fit. 2020-2021 Option Sample

	Maturity (days)					
Moneyness (K/S)	5-30	30-60	60-90	90-180	180-365	All
SV1						
0.85-0.90	9.252	15.109	18.569	20.958	21.003	16.978
0.90-0.95	11.537	14.818	16.204	16.720	16.587	15.173
0.95-1.00	13.619	15.801	14.021	12.007	11.295	13.349
1.00-1.05	14.711	17.191	15.403	11.929	12.398	14.326
1.05-1.10	7.260	10.643	12.337	12.325	15.065	11.526
1.10-1.15	3.704	4.182	7.312	10.392	16.315	8.381
All	10.014	12.957	13.974	14.055	15.444	13.289
SV2/SV1						
0.85-0.90	0.755	0.756	0.846	0.930	0.949	0.847
0.90-0.95	0.717	0.705	0.787	0.911	0.942	0.812
0.95-1.00	0.624	0.537	0.539	0.738	0.862	0.660
1.00-1.05	0.704	0.674	0.692	0.757	0.646	0.695
1.05-1.10	0.817	0.867	0.844	0.813	0.680	0.804
1.10-1.15	1.147	1.505	1.143	0.943	0.668	1.081
All	0.794	0.841	0.809	0.849	0.791	0.817
SV3/SV1						
0.85-0.90	0.625	0.533	0.597	0.683	0.649	0.617
0.90-0.95	0.590	0.482	0.546	0.686	0.706	0.602
0.95-1.00	0.618	0.552	0.512	0.704	0.840	0.645
1.00-1.05	0.738	0.675	0.651	0.656	0.648	0.674
1.05-1.10	0.764	0.746	0.678	0.623	0.507	0.664
1.10-1.15	1.099	1.261	0.947	0.789	0.563	0.932
All	0.739	0.708	0.655	0.690	0.652	0.689
SV1JR/SV1						
0.85-0.90	0.965	0.965	0.976	0.999	1.090	0.999
0.90-0.95	0.992	0.998	0.999	0.993	1.067	1.010
0.95-1.00	0.990	0.993	1.019	1.005	1.032	1.008
1.00-1.05	0.917	0.932	0.955	0.986	1.070	0.972
1.05-1.10	0.957	0.922	0.933	0.974	1.084	0.974
1.10-1.15	1.035	0.951	0.960	0.992	1.072	1.002
All	0.976	0.960	0.974	0.992	1.069	0.994
SV2JR/SV1						
0.85-0.90	0.754	0.708	0.790	0.914	0.997	0.832
0.90-0.95	0.710	0.649	0.736	0.910	1.004	0.802
0.95-1.00	0.682	0.606	0.562	0.764	0.890	0.701
1.00-1.05	0.801	0.760	0.747	0.763	0.652	0.744
1.05-1.10	0.899	0.933	0.897	0.833	0.669	0.846
1.10-1.15	1.099	1.442	1.124	0.901	0.651	1.043
All	0.824	0.850	0.809	0.848	0.810	0.828
SV3JR/SV1						
0.85-0.90	0.696	0.628	0.702	0.817	0.839	0.736
0.90-0.95	0.653	0.569	0.650	0.819	0.871	0.712
0.95-1.00	0.652	0.572	0.523	0.721	0.880	0.670
1.00-1.05	0.768	0.704	0.667	0.666	0.642	0.689
1.05-1.10	0.839	0.825	0.747	0.656	0.567	0.727
1.10-1.15	1.079	1.279	0.957	0.744	0.545	0.921
All	0.781	0.763	0.708	0.737	0.724	0.743

Notes: The table reports the out-of-sample root mean squared option pricing error within each moneyness-maturity range of the SV1 model, using all options from 2020 and 2021. For the other models, the table shows the relative gain in out-of-sample root mean squared option pricing error with respect to the SV1 model.

Table A.2: Out-Of-Sample Option Fit. Options not used in estimation

	Maturity (days)					
Moneyness (K/S)	5-30	30-60	60-90	90-180	180-365	All
SV1						
0.85-0.90	2.248	2.861	3.406	3.978	3.734	3.245
0.90-0.95	2.894	3.678	4.030	4.194	4.067	3.773
0.95-1.00	3.405	3.350	2.972	2.724	4.462	3.383
1.00-1.05	3.387	3.444	3.132	2.786	5.257	3.601
1.05-1.10	2.637	2.504	2.557	2.949	4.974	3.124
1.10-1.15	2.516	2.496	2.637	2.776	4.163	2.918
All	2.848	3.056	3.122	3.235	4.443	3.341
SV2/SV1						
0.85-0.90	0.870	0.775	0.740	0.754	0.941	0.816
0.90-0.95	0.749	0.638	0.590	0.662	0.814	0.691
0.95-1.00	0.679	0.656	0.697	0.798	0.638	0.694
1.00-1.05	0.745	0.801	0.871	0.920	0.599	0.787
1.05-1.10	0.877	0.884	0.885	0.878	0.676	0.840
1.10-1.15	0.989	0.988	0.969	0.933	0.819	0.939
All	0.818	0.790	0.792	0.824	0.748	0.794
SV3/SV1						
0.85-0.90	0.742	0.547	0.479	0.458	0.579	0.561
0.90-0.95	0.601	0.464	0.432	0.471	0.596	0.513
0.95-1.00	0.612	0.584	0.612	0.731	0.568	0.621
1.00-1.05	0.689	0.675	0.666	0.709	0.508	0.649
1.05-1.10	0.874	0.769	0.671	0.630	0.538	0.696
1.10-1.15	0.961	0.863	0.807	0.761	0.652	0.809
All	0.746	0.650	0.611	0.627	0.574	0.642
SV1JR/SV1						
0.90-0.95	0.813	0.778	0.730	0.737	0.996	0.811
0.85-0.90	0.819	0.778	0.788	0.894	1.228	0.901
0.95-1.00	0.981	0.990	0.991	0.940	0.923	0.965
1.00-1.05	0.936	0.917	0.980	1.063	0.944	0.968
1.05-1.10	0.978	0.949	0.924	0.907	0.973	0.946
1.10-1.15	0.981	0.957	0.945	0.946	0.948	0.956
All	0.918	0.895	0.893	0.914	1.002	0.925
SV2JR/SV1						
0.85-0.90	0.850	0.719	0.676	0.707	0.906	0.772
0.90-0.95	0.712	0.564	0.534	0.657	0.834	0.660
0.95-1.00	0.693	0.674	0.706	0.816	0.650	0.708
1.00-1.05	0.820	0.862	0.886	0.899	0.565	0.806
1.05-1.10	0.896	0.909	0.851	0.787	0.626	0.814
1.10-1.15	0.988	0.977	0.946	0.877	0.760	0.910
All	0.827	0.784	0.767	0.790	0.724	0.778
SV3JR/SV1						
0.85-0.90	0.794	0.614	0.549	0.560	0.691	0.642
0.90-0.95	0.634	0.458	0.427	0.527	0.673	0.544
0.95-1.00	0.653	0.655	0.692	0.770	0.594	0.673
1.00-1.05	0.767	0.779	0.780	0.799	0.525	0.730
1.05-1.10	0.879	0.817	0.715	0.666	0.539	0.723
1.10-1.15	0.974	0.916	0.854	0.762	0.635	0.828
All	0.783	0.706	0.670	0.681	0.610	0.690

Notes: The table reports the out-of-sample root mean squared option pricing error within each moneyness-maturity range of the SV1 model, using all options that have not been used for parameter estimation of the models . For the other models, the table shows the relative gain in out-of-sample root mean squared option pricing error with respect to the SV1 model.

Appendix References

Andrieu, Christophe, Arnaud Doucet, and Roman Holenstein, 2010, Particle Markov Chain Monte Carlo Methods, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72, 269–342.

Atchadé, Y., and J. Rosenthal, 2005, On Adaptive Markov Chain Monte Carlo Algorithms, *Bernoulli* 11(5), 815–828.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options, *Review of Financial Studies* 16, 101–143.

Bakshi, Gurdip, and Dilip Madan, 2000, Spanning and Derivative-security Valuation, *Journal of Financial Economics* 55, 205–238.

Carr, Peter, and Dilip B. Madan, 1998, Option Valuation Using the Fast Fourier Transform, *Journal of Computational Finance* 2, 61–73.

Eraker, Bjørn, 2004, Do Stock Prices and Volatility Jump? Evidence From Spot and Option Prices, *Journal of Finance* 59, 1367–1403.

Laloy, Eric, and Jasper A Vrugt, 2012, High-dimensional posterior exploration of hydrologic models using multiple-try DREAM (ZS) and high-performance computing, *Water Resources Research* 48.

Lord, Roger, Remmert Koekkoek, and Dick Van Dijk, 2010, A comparison of biased simulation schemes for stochastic volatility models, *Quantitative Finance* 10, 177–194.

Vrugt, Jasper A, CJF Ter Braak, CGH Diks, Bruce A Robinson, James M Hyman, and Dave Higdon, 2009, Accelerating Markov chain Monte Carlo simulation by differential evolution with self-adaptive randomized subspace sampling, *International journal of nonlinear sciences and numerical simulation* 10, 273–290.