JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 54, No. 2, Apr. 2019, pp. 695–727
COPYRIGHT 2018, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S0022109018000777

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

Good Volatility, Bad Volatility, and Option Pricing

Bruno Feunou and Cédric Okou*

Abstract

Advances in variance analysis permit the splitting of the total quadratic variation of a jump-diffusion process into upside and downside components. Recent studies establish that this decomposition enhances volatility predictions and highlight the upside/downside variance spread as a driver of the asymmetry in stock price distributions. To appraise the economic gain of this decomposition, we design a new and flexible option pricing model in which the underlying asset price exhibits distinct upside and downside semivariance dynamics driven by the model-free proxies of the variances. The new model outperforms common benchmarks, especially the alternative that splits the quadratic variation into diffusive and jump components.

I. Introduction

The proper specification of underlying asset volatility dynamics is a key input for designing a successful option valuation framework. Volatility randomness, a persistent memory pattern, and substantial conditional tail thickness of the underlying distribution are a few empirical regularities often accounted for in valuation models to accurately fit the observed option prices. Heston and Nandi (2000), Bates (2000), Duffie, Pan, and Singleton (2000), and Huang and Wu (2004), among others, have made far-reaching contributions in this regard. Moreover, the asymmetric volatility response to positive versus negative shocks is a well-established stylized fact.

Building on these insights, this article develops an option valuation model in which the underlying asset price features specific upside (good) and downside

*Feunou (corresponding author), feun@bankofcanada.ca, Bank of Canada; Okou, okou.cedric@uqam.ca, Université du Québec à Montréal School of Management. We are indebted to Jennifer Conrad (the editor) and an anonymous referee for helpful comments that improved the article. We pay a special tribute to Peter Christoffersen, colleague and friend, who passed away on June 22, 2018, and whose guidance and support greatly shaped this research agenda. Our hearts and thoughts go out to his family. We also thank Diego Amaya, Christian Dorion, Yoontae Jeon, and seminar participants at HEC Montréal and Université Paris Nanterre for fruitful discussions. We gratefully acknowledge financial support from the Bank of Canada, the Université du Québec à Montréal (UQAM) research funds, and the Canadian Derivatives Institute (CDI). The views expressed in this article are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

(bad) variance dynamics. In our modeling framework, good and bad volatilities are factors governing the return process and are directly driven by model-free empirical measures. The theoretical and empirical justifications for constructing reliable realized variance measures using high-frequency observations are addressed in seminal articles by Andersen, Bollerslev, Diebold, and Ebens (2001), and Andersen, Bollerslev, Diebold, and Labys (2001), (2003), to cite a few. Drawing on similar "infill asymptotics" arguments, Barndorff-Nielsen, Kinnebrock, and Shephard (2010) show, in a model-free way, how to dissect the realized variance into upside and downside semivariances obtained by summing high-frequency positive and negative squared returns, respectively. This decomposition has been used to improve realized variance forecasts (Patton and Sheppard (2015)), predict the equity risk premium given the standard risk–return trade-off (Guo, Wang, and Zhou (2015)), or explain the cross-section of stock returns (Bollerslev, Li, and Zhao (2017)). We extend these studies by helping gauge the economic value added of disentangling the upside (semi-)variance motion from its downside counterpart in our option pricing framework.

This article is related to a growing body of finance literature that aims to build option pricing models with empirically grounded properties. This strand of the option pricing literature is distinct from standard stochastic volatility option pricing models because it uses observed (realized) quantities to update factors: The factors are no longer latent.[1] This modeling approach is not only practically appealing because we do not need a sophisticated filtering technique but also bridges the gap between developments in high-frequency econometrics and option pricing literature. A few studies propose to model the joint dynamics of returns and realized variances in the context of option pricing. This class of option valuation models is shown to deliver superior pricing performance compared with models optimized only on returns. Recent developments include articles by Stentoft (2008), Corsi, Fusari, and Vecchia (2013), and Christoffersen, Feunou, Jacobs, and Meddahi (2014).

The aforementioned articles focus exclusively on the total realized variation and do not incorporate the information pertaining to the direction of the variation. To the best of our knowledge, our framework is the first that explicitly prices options with distinct dynamics for observable upside and downside realized variations. By modeling the directional variations, our model successfully and explicitly accounts for the asymmetry in the distribution of the underlying asset. We refer to our specification as the generalized skew affine realized variance (GSARV) model. Moreover, the model is affine and cast in discrete-time, which permits computing explicit pricing formulas, and entails a straightforward fitting procedure.

The closely related bipower and jump variation option pricing model (BPJVM) developed in Christoffersen, Feunou, and Jeon (2015) exploits an alternative dissection of the total quadratic variation into a diffusive volatility and a squared jump variation. Thus, the modeling approach of Christoffersen et al. (2015) can be viewed as the discrete-time analog of a classical continuous-time affine jump-diffusion specification, such as that of Bates (2000), where the two

[1] See Andersen, Fusari, and Todorov (2015) for a review of the various developments in this literature.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

factors capture the change in "normal" variation (diffusion) and the intensity of "extreme" moves (jumps) and are connected to their realized counterparts.

Although modeling approaches based on these two decompositions (up–down variation vs. jump-diffusive variation) of the total quadratic variation account for the departure from a conditional normal distribution, their relative empirical performance remains an open question. In the continuous-time framework, a recent study by Andersen, Fusari, and Todorov (2015) underscores the importance of isolating negative expected jump variations from their positive counterparts. The authors demonstrate that accounting for directional jumps improves their option valuation model over the benchmark affine jump-diffusion specification of Bates (2000). Therefore, in a discrete-time setup, we expect our GSARV model to outperform the BPJVM. Our empirical investigation confirms this prediction; we find that our preferred GSARV specification improves the option price fitting by a sizable 10% over the BPJVM.

The contribution of this work is not, however, limited to constructing a novel option pricing framework. We also formally describe the major factors driving market compensations of good versus bad uncertainty. We estimate the model and show that it performs well in matching the historical as well as the risk-neutral distributions of the Standard & Poor's (S&P) 500 index returns. Namely, the model improves significantly upon popular specifications, with respect to various performance criteria, when optimized on a data set of S&P 500 index options, realized upside and downside variances, and returns. We find that the conditional asymmetry, mainly driven by the wedge between upside and downside volatilities in our specification, matters for delivering realistic market variance risk premia. Allowing for distinct up/down variance dynamics is useful to track the time variation in the variance risk premium and its up/down components at different horizons.

The article is organized as follows: In Section II, we present the theoretical and empirical arguments underpinning the construction and the use of good/bad (upside/downside) variation measures. Section III introduces a novel option pricing model that is general enough to accommodate specific upside and downside variance dynamics in the underlying return process while drawing relevant information from their empirical proxies. In Section IV, we describe the physical estimation strategy and discuss the different specifications that our option pricing framework encompasses. We also discuss the estimation findings based on historical observations. Section V investigates the empirical ability of the various nested models to fit the risk-neutral distribution embedded in option contracts. We implement a joint optimization procedure that combines historical information and option data in Section VI. In addition, we study the performance of our pricing kernel and document the determinants of the components of the variance risk premium. Section VII concludes.

II. Daily Returns and Realized Variation Measures

We outline the arguments supporting the decomposition of the total quadratic variation into its upside and downside components. Disentangling the upside realized semivariation from its downside counterpart can be achieved by exploiting

the fine structure of high-frequency observations. For instance, empirical measures of daily realized good (or bad) volatility are often constructed from intraday records exceeding (or falling below) a specified threshold.

A. Separating Downside from Upside Volatility: Theoretical Arguments

We briefly review the key theoretical results that allow us to separate daily positive from negative quadratic variation using intraday data. We mainly rely on Barndorff-Nielsen et al. (2010), who assume that the log stock price (denoted by s) follows a jump-diffusion motion of the form

$$ds_t = \mu_t dt + \sigma_t dW_t + \Delta s_t, \tag{1}$$

where dW_t is an increment of standard Brownian motion, and $\Delta s_t \equiv s_t - s_{t-}$ refers to the jump component.

For a given day of interest $t+1$, the upside realized variance (denoted by RV^U_{t+1}) and the downside realized variance (denoted by RV^D_{t+1}) are defined as

$$\mathrm{RV}^U_{t+1} = \sum_{j=1}^{n_{t+1}} R^2_{t_j} \mathbb{I}_{[R_{t_j}>0]},$$

$$\mathrm{RV}^D_{t+1} = \sum_{j=1}^{n_{t+1}} R^2_{t_j} \mathbb{I}_{[R_{t_j}\leq 0]},$$

where $R_{t_j} \equiv s_{t_j} - s_{t_{j-1}}$, $t < t_1 < \cdots < t_{n_{t+1}} = t+1$ are the time at which (trade or quote) prices are available, and $\mathbb{I}_{[*]}$ is the indicator function. The total realized variance is $\mathrm{RV}_{t+1} = \mathrm{RV}^U_{t+1} + \mathrm{RV}^D_{t+1}$.

Under this general assumption on the instantaneous return process, Barndorff-Nielsen et al. (2010) use infill asymptotics (i.e., convergence as the time separating two consecutive observations collapses to 0) to prove that

$$\mathrm{RV}^U_{t+1} \xrightarrow{p} \frac{1}{2}\int_t^{t+1} \sigma^2_\upsilon d\upsilon + \sum_{t\leq \upsilon \leq t+1} (\Delta s_\upsilon)^2 \mathbb{I}_{[\Delta s_\upsilon > 0]}, \tag{2}$$

$$\mathrm{RV}^D_{t+1} \xrightarrow{p} \frac{1}{2}\int_t^{t+1} \sigma^2_\upsilon d\upsilon + \sum_{t\leq \upsilon \leq t+1} (\Delta s_\upsilon)^2 \mathbb{I}_{[\Delta s_\upsilon \leq 0]}.$$

This decomposition of the realized variance into its up and down components is used by Patton and Sheppard (2015) to assess the information content and the predictive ability of signed squared jumps.

In a recent study, Feunou, Jahan-Parvar, and Okou (2018) investigate the asymmetric behavior of investors towards good versus bad uncertainty by analyzing the premia related to upside and downside realized variances. Moreover, the difference between realized upside and downside variance, also known as the signed jump variation, can be perceived as a measure of (realized) skewness. This measure of asymmetry, denoted as RSV_t, is obtained by subtracting the downside variance from the upside variance:

$$\mathrm{RSV}_{t+1} = \mathrm{RV}^U_{t+1} - \mathrm{RV}^D_{t+1}.$$

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

Thus, if $RSV_{t+1} < 0$, the distribution is left-skewed, and when $RSV_{t+1} > 0$, it is right-skewed. A theoretical justification for using RSV_{t+1} as a measure of skewness is given by Feunou, Jahan-Parvar, and Tédongap (2016). To provide more intuition on the behavior of RSV_{t+1}, we combine the previous asymptotic results to get

$$RSV_{t+1} \xrightarrow{p} \sum_{t<\upsilon\leq t+1} (\Delta s_\upsilon)^2 \left(\mathbb{I}_{[\Delta s_\upsilon > 0]} - \mathbb{I}_{[\Delta s_\upsilon \leq 0]}\right).$$

Assuming that jump sizes are independent and identically distributed (IID) and uncorrelated with the jump occurrence, the expectation of the realized skewness, a measure of the conditional skewness, is

$$\mathrm{E}_t^{\mathbb{P}}[RSV_{t+1}] \approx \mathrm{E}^{\mathbb{P}}\left[(\Delta s)^2\right] \sum_{t<\upsilon\leq t+1} \left(\mathbb{P}[\Delta s_\upsilon > 0] - \mathbb{P}[\Delta s_\upsilon \leq 0]\right).$$

Thus, $\mathrm{E}_t^{\mathbb{P}}[RSV_{t+1}]$ captures, up to a multiplicative constant, the wedge between positive and negative jump intensities. This realized skewness measure shows, in relative terms, how likely positive jumps are with respect to negative jumps. The realizations of directional jumps receive different weights according to their sizes, consistent with the study of volatility jump risk of Bandi and Renò (2016). Bigger weights are assigned to larger jump sizes in the computation of RSV_{t+1}.

B. Alternative Decomposition of the Total Quadratic Variation

Starting from the very general instantaneous return process in equation (1), Barndorff-Nielsen and Shephard (2004) show the following limiting result, as the sampling frequency goes to infinity:

$$(3) \qquad RJV_{t+1} \equiv RV_{t+1} - BPV_{t+1} \xrightarrow{p} \sum_{t\leq\upsilon\leq t+1} (\Delta s_\upsilon)^2,$$

$$BPV_{t+1} \xrightarrow{p} \int_t^{t+1} \sigma_\upsilon^2 d\upsilon,$$

where BPV_{t+1} denotes the bipower variation measuring the diffusive volatility. This alternative decomposition of the realized variation into smooth (diffusion) and rough (jumps) components is used by Christoffersen et al. (2015) in their option pricing framework. Although RJV_{t+1} estimates $\sum_{t\leq\upsilon\leq t+1}(\Delta s_\upsilon)^2$, it is "blind" to the asymmetric behavior of jumps. In continuous-time, Andersen, Fusari, and Todorov (2015) demonstrate the importance of disentangling negative jump dynamics from positive jump dynamics. Empirically, Andersen, Fusari, and Todorov's (2015) approach is shown to be superior to the alternative classic affine jump-diffusion approach of Bates (2000), where negative and positive jumps occur with the same intensity. In Section VI.B.1, we also compare the empirical option pricing implications of these two decompositions.

C. Separating Downside from Upside Volatility: Empirical Implementation

To empirically construct the upside and downside components of the realized variance, we download intraday S&P 500 cash index data from TickData.com. On a given day, we use the last record in each 5-minute interval to build a grid of 5-minute equity index log-returns. Following Andersen, Bollerslev,

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

Diebold, and Labys (2003), Andersen, Bollerslev, Diebold, and Ebens (2001), and Barndorff-Nielsen et al. (2010), we construct (pre-scaled) measures of total quadratic variation and their upside and downside components on any given trading day t as

$$\begin{aligned}\widetilde{\mathrm{RV}}_t &= \sum_{j=1}^{n_t} R_{j,t}^2,\\ \widetilde{\mathrm{RV}}_t^U &= \sum_{j=1}^{n_t} R_{j,t}^2 \mathbb{I}_{[R_{j,t}>0]},\\ \widetilde{\mathrm{RV}}_t^D &= \sum_{j=1}^{n_t} R_{j,t}^2 \mathbb{I}_{[R_{j,t}\leq 0]},\end{aligned}$$

where $R_{j,t}^2$ is the jth 5-minute squared log-return, and n_t is the number of (5-minute) intraday returns recorded on that day.[2] We add the squared overnight log-return (the difference in log-price between when the market opens at t and when it closes at $t-1$) to the total variation. We add the squared overnight log-return to the downside variation component when the overnight return is negative and to the upside variation component otherwise. To ensure that the sample average of daily realized variances equals the sample variance of daily log-returns, we perform the following scaling:

$$\begin{aligned}\mathrm{RV}_t &= \frac{\sum_{t=1}^{T}\left(R_t-\bar{R}_T\right)^2}{\sum_{t=1}^{T}\widetilde{\mathrm{RV}}_t}\widetilde{\mathrm{RV}}_t,\\ \mathrm{RV}_t^U &= \frac{\sum_{t=1}^{T}\left(R_t-\bar{R}_T\right)^2}{\sum_{t=1}^{T}\widetilde{\mathrm{RV}}_t}\widetilde{\mathrm{RV}}_t^U,\\ \mathrm{RV}_t^D &= \frac{\sum_{t=1}^{T}\left(R_t-\bar{R}_T\right)^2}{\sum_{t=1}^{T}\widetilde{\mathrm{RV}}_t}\widetilde{\mathrm{RV}}_t^D,\end{aligned}$$

where T is the sample size, and $\bar{R}_T$ is the sample average log-return. Different approaches to adjusting open-to-close realized variance measures are discussed by Hansen and Lunde (2006). By construction, the total realized variance adds up the realized upside and downside variances:

$$\mathrm{RV}_t \equiv \mathrm{RV}_t^U + \mathrm{RV}_t^D. \tag{4}$$

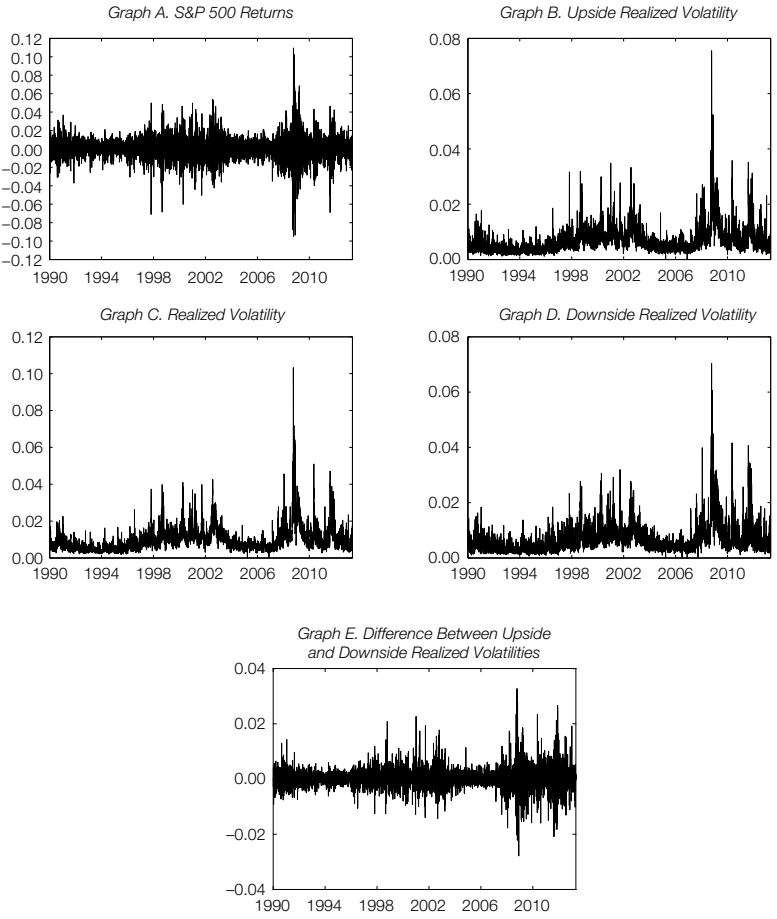
[2] On a typical trading day, we observe $n_t = 78$ 5-minute returns.

D. Empirical Dynamics of Downside and Upside Volatilities

Figure 1 plots the daily time series of the return on the S&P 500 index (Graph A), along with the square root of its upside (Graph B), the square root of the daily realized variance (Graph C), and the square root of its downside (Graph D) components, from Jan. 10, 1990, to Aug. 28, 2013. Periods of market instability, characterized by large swings in returns and a high level of volatility, are clearly apparent. These turbulent periods include the recent financial market meltdown of 2008–2009, the early 2000s, and the early 1990s. By contrast, the mid-1990s and the mid-2000s periods exhibit low levels of volatility. The time series

FIGURE 1
Daily Returns and Realized Variances

Figure 1 presents the daily returns on the Standard & Poor's (S&P) 500 index, R_t (Graph A); the daily upside realized volatilities, $\sqrt{RV_t^U}$ (Graph B); the daily realized volatilities, $\sqrt{RV_t}$ (Graph C); the daily downside realized volatilities, $\sqrt{RV_t^D}$ (Graph D); and the difference between the daily upside and downside realized volatilities, $\sqrt{RV_t^U} - \sqrt{RV_t^D}$ (Graph E). Daily realized volatility quantities are computed from 5-minute squared returns using a scaled realized variance (RV) estimator. The sample starts on Jan. 10, 1990, and ends on Aug. 28, 2013.



https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

of the realized variance and its components display a good level of synchronicity because they tend to rise and fall around the same periods. The difference between (square roots of) upside and downside variance (Graph E), which is an alternative measure of asymmetry, appears to be time varying, with important fluctuations during the 2008 financial crisis period.

Consistent with established empirical regularities, realized variance RV_t series exhibit markedly higher persistence, dispersion, positive asymmetry, and heavy-tailness compared with returns distribution. Moreover, we see that RV_t^U and RV_t^D have similar magnitudes. This observation is supported by the nearly identical average values of RV_t^U and RV_t^D presented in Table 1, along with other common summary statistics. Sizable discrepancies between these two time series can arise from risk-neutral (rather than historical) expectations extracted from option data.

Using similar logic as Bollerslev, Tauchen, and Zhou (2009), dynamic "variance-of-(semi-)variance" processes are needed to generate premia for second-order semi-moments. Specifically, the realized downside variance measure requires a dynamic specification of its own, one that is likely different from the dynamic specification of upside variation. Building a dynamic return model with such features is our next task.

III. New Dynamic Model for Asset Returns

This section builds a model for option valuation that incorporates the information in R_t, RV_t^U, and RV_t^D, computed at the end of any given day t. In our model, state variables are explicitly filtered from realized upside and downside variances in the spirit of generalized autoregressive conditional heteroscedasticity (GARCH) models. In addition, we focus on an empirical strategy for option pricing that can be implemented without resorting to Monte Carlo simulations. It is worth stressing the conceptual difference between this approach and standard continuous-time stochastic factor models. In this framework, there are no latent factors. Hence, we are not filtering or estimating factors from observed option prices. Instead, we update the factors in our model using their model-free ex post (realized) counterparts constructed from high-frequency historical returns on the underlying asset.

TABLE 1
Summary Statistics of Historical Series

Table 1 presents the summary statistics for the studied series. Mean, median, and standard deviation values are annualized and in percentages. AR(1) describes first-order autocorrelation coefficient values. The sample starts on Jan. 10, 1990, and ends on Aug. 28, 2013.

Series	Mean (%)	Median (%)	Std. Dev. (%)	Skewness	Kurtosis	AR(1)
Return	6.43	13.74	18.49	−0.22	11.41	−0.06
Volatility	15.32	12.66	10.34	3.25	23.17	0.81
Upside volatility	10.69	8.84	7.50	3.45	25.80	0.71
Downside volatility	10.51	8.52	7.79	3.16	21.18	0.70

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

A. Key Objectives

Our main objective is to build an option pricing model with two factors capturing the expected upside variation ($h_{u,t}$) and the expected downside variation ($h_{d,t}$). By design, we want to have

$$\mathrm{E}_t\left[\mathrm{RV}^j_{t+1}\right] = h_{j,t}, \tag{5}$$

with $j = u,d$. Thus, the sum of the factors (denoted by h_t) equals the expected total variation

$$h_t \equiv \mathrm{E}_t\left[\mathrm{RV}_{t+1}\right] = h_{u,t} + h_{d,t}. \tag{6}$$

We also want to embed these two factors into the log-return dynamic. Therefore, h_t plays a dual role in our return model. Specifically, we impose that h_t i) corresponds to the expected realized variance as in equation (6) and ii) matches the conditional variance of $t+1$ log-returns as

$$\mathrm{var}_t\left[R_{t+1}\right] = h_t. \tag{7}$$

As discussed in Section II.A, the difference between upside and downside variations is a measure of realized skewness. This implies a tight link between the difference between factors ($h_{u,t} - h_{d,t}$) and the conditional skewness of $t+1$ log-returns.

B. Asset Return Process

Consider the following representation of the daily log-returns:

$$R_{t+1} = \mathrm{E}_t\left[R_{t+1}\right] + z_{t+1}, \tag{8}$$

expressed as an expected component $\mathrm{E}_t\left[R_{t+1}\right]$ and an innovation term z_{t+1}. The expected return is set up to satisfy the following identity:

$$\mathrm{E}_t\left[\exp\left(R_{t+1}\right)\right] = \exp\left(r_f + \lambda_u h_{u,t} + \lambda_d h_{d,t}\right), \tag{9}$$

where r_f is the risk-free rate. Equation (9) enables us to interpret λ_u and λ_d as compensations for upside and downside volatility risk exposures, respectively. Next, we specify the conditional distribution of the total shock z_{t+1}.

1. Innovation Decomposition

To introduce the notion of positive and negative innovations in the return dynamic, we decompose z_{t+1} as the difference of two positive shocks:

$$z_{t+1} = z_{u,t+1} - z_{d,t+1}. \tag{10}$$

By construction, equation (10) implies that z_{t+1} increases with $z_{u,t+1}$ and decreases with $z_{d,t+1}$. To ensure that $z_{u,t+1}$ and $z_{d,t+1}$ are demeaned innovations, we will relate them to uncentered shocks $\tilde{z}_{u,t+1}$ and $\tilde{z}_{d,t+1}$:

$$z_{j,t+1} = \tilde{z}_{j,t+1} - \mathrm{E}_t\left[\tilde{z}_{j,t+1}\right], \quad \text{for } j = \{u,d\}. \tag{11}$$

Furthermore, we impose $\tilde{z}_{j,t+1} \geq 0$ for $j = \{u,d\}$ to ensure nonzero conditional skewness for the log-returns. We have

$$\begin{aligned} \text{var}_t\left[z_{u,t+1}\right] &= \text{var}_t\left[\tilde{z}_{u,t+1}\right] \equiv h_{u,t}, \\ \text{var}_t\left[z_{d,t+1}\right] &= \text{var}_t\left[\tilde{z}_{d,t+1}\right] \equiv h_{d,t}, \end{aligned}$$

which allows us to interpret $h_{u,t}$ and $h_{d,t}$ as stock market good and bad volatilities. Simply put, $h_{u,t}$ is the variance of positive shocks to returns, whereas $h_{d,t}$ is the variance of negative shocks to returns.

2. Innovation Conditional Distribution

Because the objective is to have a closed-form option price, we opt for the noncentral χ^2 distribution, which has been widely used in the discrete-time option pricing literature (see Heston and Nandi (2000)). Shocks to returns are

$$(12)\quad \tilde{z}_{j,t+1} = \sqrt{\frac{\omega_j}{2}}\left(\varepsilon^{(1)}_{j,t+1} - \sqrt{\frac{h_{j,t}-\omega_j}{2\omega_j}}\right)^2, \quad \text{with } \varepsilon^{(1)}_{j,t+1} \overset{\text{IID}}{\sim} \text{N}(0,1) \text{ for } j = \{u,d\},$$

with

$$\text{E}\left[\varepsilon^{(1)}_{u,t+1}\varepsilon^{(1)}_{d,t+1}\right] = 0.$$

By construction, we have $h_{j,t} = \text{var}_t\left[z_{j,t+1}\right]$. We also have

$$z_{j,t+1} \overset{\text{Distribution}}{\underset{\omega_j \to 0}{\longrightarrow}} \text{N}\left(0, h_{j,t}\right),$$

implying that ω_j drives the nonnormality in the conditional distribution of $z_{j,t+1}$.

The conditional skewness of log-returns is

$$(13)\quad \text{skew}_t\left[R_{t+1}\right] = \frac{h_t^{-3/2}}{\sqrt{2}}\left[3\sqrt{\omega_u}\left(h_{u,t}-h_{d,t}\right) + 3\left(\sqrt{\omega_u}-\sqrt{\omega_d}\right)h_{d,t} - \left(\omega_u^{3/2}-\omega_d^{3/2}\right)\right],$$

where $h_t \equiv h_{u,t} + h_{d,t}$ is the total conditional variance of returns. Hence, two sources of conditional asymmetry emerge: i) the difference between upside and downside variance $h_{u,t} - h_{d,t}$ and ii) the discrepancy between nonnormalities in good and bad shock distributions $\sqrt{\omega_u} - \sqrt{\omega_d}$. Hence, an interesting case is obtained by imposing $\omega_u = \omega_d$, implying that the wedge between the upside and downside volatility, $h_{u,t} - h_{d,t}$, is the unique channel for conditional skewness in the return dynamics. We will pay particular attention to that restriction in the empirical analysis.

C. Expected Returns

We now discuss the specification of expected return dynamics. For equation (9) to hold, it must be the case that

$$\text{E}_t\left[R_{t+1}\right] = \bar{r} + \left(\lambda_u - \xi_u\right)h_{u,t} + \left(\lambda_d - \xi_d\right)h_{d,t},$$

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

where

$$\bar{r} = r_f + \frac{1}{2}\left[\ln\left(1-\sqrt{2\omega_u}\right)+\ln\left(1+\sqrt{2\omega_d}\right) - \frac{\omega_u-\sqrt{2\omega_u}}{\left(1-\sqrt{2\omega_u}\right)} - \frac{\omega_d+\sqrt{2\omega_d}}{\left(1+\sqrt{2\omega_d}\right)}\right], \tag{14}$$

$$\xi_u = \frac{1}{2\left(1-\sqrt{2\omega_u}\right)}, \quad \xi_d = \frac{1}{2\left(1+\sqrt{2\omega_d}\right)}. \tag{15}$$

Therefore, the parameters ξ_u, ξ_d, and $\bar{r}$ are not estimated but, instead, simply set to their values implied by the identities in equations (14) and (15).

D. Incorporating Realized Upside and Downside Variation

Each day, realized upside and downside variations provide new information about the conditional semivariances, $h_{u,t}$ and $h_{d,t}$. However, RV^U_{t+1} and RV^D_{t+1} are measured with error, and therefore, we specify the following measurement equation:

$$\mathrm{RV}^j_{t+1} = h_{j,t} + z^{(2)}_{j,t+1}, \tag{16}$$

where $z^{(2)}_{j,t+1} = \tilde{z}^{(2)}_{j,t+1} - \mathrm{E}_t\left[\tilde{z}^{(2)}_{j,t+1}\right]$ is a zero-mean innovation. By design, we have $\mathrm{E}_t\left[\mathrm{RV}^j_{t+1}\right] = h_{j,t}$.

To obtain closed-form option pricing formulas, we choose the noncentral χ^2 distribution:

$$\tilde{z}^{(2)}_{j,t+1} = \sigma_j\left(\varepsilon^{(2)}_{j,t+1} - \gamma_j\sqrt{h_{j,t}-\omega_j}\right)^2, \tag{17}$$

where

$$\varepsilon^{(2)}_{j,t+1} \overset{\text{IID}}{\sim} \mathrm{N}(0,1) \quad \text{for } j=\{u,d\},$$

and

$$\mathrm{E}\left[\varepsilon^{(2)}_{u,t+1}\varepsilon^{(2)}_{d,t+1}\right] = 0,$$
$$\mathrm{E}\left[\varepsilon^{(1)}_{j,t+1}\varepsilon^{(2)}_{j,t+1}\right] = \rho_j.$$

Note that equation (17) allows for a nonlinear impact of $\varepsilon^{(2)}_{j,t+1}$ on RV^j_{t+1} via γ_j.

E. Upside and Downside Volatility Dynamics

We are now ready to specify the dynamics of the expected upside and downside realized variances. We assume the following recursive dynamics, where new information on realized downside (or upside) variance at time $t+1$ is used to update our $t+2$ forecast of realized downside (or upside) variances:

$$h_{j,t+1} - \omega_j = \tilde{\omega}_j + \tilde{\beta}_j\left(h_{j,t}-\omega_j\right) + \tilde{\alpha}_j\left(\mathrm{RV}^j_{t+1}-\omega_j\right). \tag{18}$$

This is very similar to GARCH models, with the exception that less noisy innovations $\mathrm{RV}^j_{t+1}-\omega_j$ are used to update conditional variances. The specification

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

simply implies that $h_{u,t+1}$ and $h_{d,t+1}$ are both univariate first-order autoregressive (AR(1)) processes. Thus, RV^U_{t+1} and RV^D_{t+1} belong to the class of univariate autoregressive moving-average (ARMA(1, 1)) processes. Using equation (16), equation (18) can be rewritten as

$$h_{j,t+1} - \omega_j = \varpi_j + \beta_j \left(h_{j,t} - \omega_j\right) + \alpha_j \left(\varepsilon^{(2)}_{j,t+1} - \gamma_j \sqrt{h_{j,t} - \omega_j}\right)^2, \tag{19}$$

where

$$\tilde{\omega}_j = \varpi_j + \alpha_j, \quad \tilde{\beta}_j = \beta_j + \gamma_j^2 \alpha_j - \frac{\alpha_j}{\sigma_j}, \quad \tilde{\alpha}_j = \frac{\alpha_j}{\sigma_j}.$$

Equation (19) is similar to the conditional variance dynamics of Heston and Nandi (2000). The main difference here is that the Gaussian shock $\varepsilon^{(2)}_{j,t+1}$ is not perfectly correlated to returns. We refer to this general specification as the generalized skew affine realized variance (GSARV) model.

F. Conditional Second Moments and Moment-Generating Function

From the model described previously, it is straightforward to derive the following 1-day-ahead conditional second moments:

$$\begin{aligned}
\mathrm{var}_t\left[R_{t+1}\right] &= h_{u,t} + h_{d,t}, \\
\mathrm{var}_t\left[\mathrm{RV}^j_{t+1}\right] &= 2\sigma_j^2\left(1 + 2\gamma_j^2\left(h_{j,t} - \omega_j\right)\right), \\
\mathrm{cov}_t\left(R_{t+1}, \mathrm{RV}^U_{t+1}\right) &= \sqrt{2\omega_u}\sigma_u\rho_u^2 + 2\left(h_{u,t} - \omega_u\right)\gamma_u\sigma_u\rho_u, \\
\mathrm{cov}_t\left(R_{t+1}, \mathrm{RV}^D_{t+1}\right) &= -\left(\sqrt{2\omega_d}\sigma_d\rho_d^2 + 2\left(h_{d,t} - \omega_d\right)\gamma_d\sigma_d\rho_d\right), \\
\mathrm{cov}_t\left(R_{t+1}, h_{t+1}\right) &= \tilde{\alpha}_u \mathrm{cov}_t\left(R_{t+1}, \mathrm{RV}^U_{t+1}\right) + \tilde{\alpha}_d \mathrm{cov}_t\left(R_{t+1}, \mathrm{RV}^D_{t+1}\right).
\end{aligned} \tag{20}$$

Note that the model allows for two types of "leverage" effects: one via the return covariance with upside variation and another via the return covariance with downside variation.

Our modeling choices are mainly motivated by the need to derive option prices in closed form. To achieve this goal, a sufficient requirement is for the framework to be affine (see Duffie et al. (2000)). In discrete-time, it is sufficient to show that the joint 1-step-ahead conditional characteristic function of returns and factors is exponentially linear in those factors (see Darolles, Gourieroux, and Jasiak (2006)). The Appendix reports the expressions for the joint conditional moment-generating function of returns and the expected upside and downside variations.

IV. Physical Estimation

Up to this point, we have laid out a general framework for incorporating upside and downside realized variations when modeling the underlying asset dynamics. In this section we develop a likelihood-based method that enables us to estimate the physical parameters using daily observations on returns, as well as the good and bad realized variation measures. We also discuss special cases of the general specification that includes the Heston and Nandi (2000) benchmark GARCH model.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

A. Deriving the Likelihood Function

In deriving the conditional quasi-likelihood function, we compute the contribution of the day $t+1$ observation vector by multiplying the marginal densities with a Gaussian copula. Formally, we can write

$$\begin{aligned} f_t\left(R_{t+1}, \mathrm{RV}^U_{t+1}, \mathrm{RV}^D_{t+1}\right) &= f_{r,t}\left(R_{t+1}\right) f_{U,t}\left(\mathrm{RV}^U_{t+1}\right) f_{D,t}\left(\mathrm{RV}^D_{t+1}\right) \\ &\quad \times c_t\left(F_{r,t}\left(R_{t+1}\right), F_{U,t}\left(\mathrm{RV}^U_{t+1}\right), F_{D,t}\left(\mathrm{RV}^D_{t+1}\right)\right), \end{aligned}$$

where $f_{r,t}\left(R_{t+1}\right)$, $f_{U,t}\left(\mathrm{RV}^U_{t+1}\right)$, and $f_{D,t}\left(\mathrm{RV}^D_{t+1}\right)$ are the marginal conditional densities of returns, upside realized variance, and downside realized variance, respectively. Accordingly, $F_{r,t}\left(R_{t+1}\right)$, $F_{U,t}\left(\mathrm{RV}^U_{t+1}\right)$, and $F_{D,t}\left(\mathrm{RV}^D_{t+1}\right)$ are the marginal conditional cumulative distribution function (CDF) of returns, upside realized variance, and downside realized variance, respectively, and $c_t\left(\nu, \nu_U, \nu_D\right)$ denotes the density of the Gaussian copula. To characterize the randomness surrounding realized upside and downside variance measures, we employ Gaussian densities and CDFs. Thus,

$$\begin{aligned} f_{j,t}\left(\mathrm{RV}^j_{t+1}\right) &= \frac{\exp\left(-\frac{\left(\mathrm{RV}^j_{t+1}-h_{j,t}\right)^2}{2\mathrm{var}_t\left[\mathrm{RV}^j_{t+1}\right]}\right)}{\sqrt{2\pi\,\mathrm{var}_t\left[\mathrm{RV}^j_{t+1}\right]}}, \\ F_{j,t}\left(\mathrm{RV}^j_{t+1}\right) &= \Phi\left(\frac{\mathrm{RV}^j_{t+1}-h_{j,t}}{\sqrt{\mathrm{var}_t\left[\mathrm{RV}^j_{t+1}\right]}}\right), \quad \text{for } j=\{u,d\}, \end{aligned}$$

where $\Phi(\cdot)$ is the CDF of a standard normal, and the $\mathrm{var}_t\left[\mathrm{RV}^j_{t+1}\right]$ values are given in equation (20). Because we want to highlight the importance of $h_{u,t}-h_{d,t}$ in generating nonnormalities in the conditional distribution of returns, we need the exact marginal density function. Notice that the exact marginal density of returns, $f_{r,t}\left(R_{t+1}\right)$, is a convolution of two noncentral χ^2 densities, which does not have a known closed-form expression. Nonetheless, the conditional characteristic function of R_{t+1} is available in closed form as

$$\begin{aligned} \varphi_{r,t}\left(\nu\right) &\equiv \mathrm{E}_t\left[e^{i\nu R_{t+1}}\right] \\ &= \exp\left(\begin{array}{c} i\nu\left(\bar{r}+\left(\lambda_u-\xi_u\right)h_{u,t}+\left(\lambda_d-\xi_d\right)h_{d,t}\right) \\ +a_u\left(i\nu\right)+b_u\left(i\nu\right)h_{u,t} \\ +a_d\left(-i\nu\right)+b_d\left(-i\nu\right)h_{d,t} \end{array}\right), \end{aligned}$$

where the $a_j\left(\nu\right)$ and $b_j\left(\nu\right)$ functions are

$$\begin{aligned} b_j\left(\nu\right) &= \frac{\nu^2}{2\left(1-\sqrt{2\omega_j}\nu\right)}, \\ a_j\left(\nu\right) &= -\sqrt{\frac{\omega_j}{2}}-\omega_j b_j\left(\nu\right)-\frac{1}{2}\ln\left(1-\sqrt{2\omega_j}\nu\right), \end{aligned} \tag{21}$$

and i stands for the imaginary unit. Thus, we exploit Fourier inversion formulas to compute the quantities of interest:

$$\begin{aligned} F_{r,t}(R_{t+1}) &= \frac{1}{2} - \frac{1}{\pi}\int_0^\infty \frac{Im\left[e^{-i\nu R_{t+1}}\varphi_{r,t}(\nu)\right]}{\nu}\,d\nu, \\ f_{r,t}(R_{t+1}) &= \frac{1}{\pi}\int_0^\infty Re\left[e^{-i\nu R_{t+1}}\varphi_{r,t}(\nu)\right]d\nu. \end{aligned}$$

Moreover, the copula function is computed as

$$c_t(\nu,\nu_U,\nu_D) = \frac{1}{\sqrt{|\mathrm{CM}_t|}}\exp\left(-\frac{1}{2}(\nu,\nu_U,\nu_D)\left(\mathrm{CM}_t^{-1}-I_3\right)\begin{pmatrix}\nu\\ \nu_U\\ \nu_D\end{pmatrix}\right),$$

where CM_t is the conditional correlation matrix of $\left(R_{t+1},\mathrm{RV}_{t+1}^U,\mathrm{RV}_{t+1}^D\right)$, given by

$$\mathrm{CM}_t = \begin{bmatrix} 1 & \rho_t^U & \rho_t^D \\ \rho_t^U & 1 & 0 \\ \rho_t^D & 0 & 1 \end{bmatrix},$$

with

$$\rho_t^j = \frac{\mathrm{cov}_t\left(R_{t+1},\mathrm{RV}_{t+1}^j\right)}{\sqrt{\mathrm{var}_t\left[R_{t+1}\right]\mathrm{var}_t\left[\mathrm{RV}_{t+1}^j\right]}}, \quad \text{for } j=\{u,d\},$$

and where $\mathrm{cov}_t\left(R_{t+1},\mathrm{RV}_{t+1}^j\right)$, $\mathrm{var}_t\left[R_{t+1}\right]$, and $\mathrm{var}_t\left[\mathrm{RV}_{t+1}^j\right]$ are defined in equation (20).

Finally, the log likelihood is calculated as

$$\ln L^P = \sum_{t=1}^{T-1}\ln\left(f_t\left(R_{t+1},\mathrm{RV}_{t+1}^U,\mathrm{RV}_{t+1}^D\right)\right). \tag{22}$$

B. Nested Specifications

Before turning to the estimation, we discuss special cases of interest that are nested by the new option pricing model. In the following discussion, these restricted specifications are estimated along with the full model.

1. The Constrained Generalized Skew Affine Realized Variance Model

One of the key objectives of this article is to highlight the importance of the difference between upside and downside variances $(h_{u,t}-h_{d,t})$ in generating the observed conditional skewness in returns and to assess its option pricing implications. As shown previously, the conditional skewness in the GSARV stems from two wedges, $\omega_u-\omega_d$ and $h_{u,t}-h_{d,t}$. Hence, we consider a special case where the first channel is shut down (i.e., $\omega\equiv\omega_u=\omega_d$), thus defining $h_{u,t}-h_{d,t}$ as the unique driver of the conditional skewness. We refer to this special case as the constrained generalized skew affine realized variance (CGSARV) model.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

2. The Affine Realized Variance Model

If we fix $\omega \equiv \omega_u = \omega_d = 0$ and $h_{u,t} = h_{d,t}$ in the GSARV specification, we obtain

$$R_{t+1} = r_f + \left(\lambda - \frac{1}{2}\right) h_t + z_{t+1},$$

with

$$z_{t+1} = \sqrt{h_t}\varepsilon_{t+1}^{(1)},$$

where $\varepsilon_{t+1}^{(1)}$ are IID N(0, 1).

Beginning with h_0, we have

$$\begin{aligned} h_{t+1} &= \varpi + \beta h_t + \alpha\left(\varepsilon_{t+1}^{(2)} - \gamma\sqrt{h_t}\right)^2, \\ \mathrm{RV}_{t+1} &= h_t + \sigma\left[\left(\varepsilon_{t+1}^{(2)} - \gamma\sqrt{h_t}\right)^2 - \left(1 + \gamma^2 h_t\right)\right], \end{aligned}$$

where

$$\mathrm{E}\left[\varepsilon_{t+1}^{(1)}\varepsilon_{t+1}^{(2)}\right] = \rho.$$

This single-factor specification is identical to the affine realized variance (ARV) model of Christoffersen et al. (2014).

3. The Heston and Nandi (2000) GARCH Model

Canonical GARCH-type option pricing models are obtained by setting $\rho = 1$ in the ARV model. Specifically, $\rho = 1$ implies that $\varepsilon_{t+1}^{(1)} = \varepsilon_{t+1}^{(2)}$, and therefore, the realized variance motion becomes irrelevant. We then get

$$\begin{aligned} h_{t+1} &= \varpi + \beta h_t + \alpha\left(\varepsilon_{t+1}^{(2)} - \gamma\sqrt{h_t}\right)^2, \\ &\equiv \varpi + \beta h_t + \alpha\left(\varepsilon_{t+1}^{(1)} - \gamma\sqrt{h_t}\right)^2, \end{aligned}$$

which is precisely the Heston and Nandi (2000) affine GARCH(1,1) model.

C. Parameter Estimates and Model Properties

Maximum likelihood estimation results from historical data are given in Table 2. As outlined previously, the estimated values are the empirical proxies of the parameters governing the return process under the physical distribution. A few remarks are in order regarding the estimation procedure. The parameters ϖ are inferred by targeting the unconditional sample variance, and therefore, they do not have standard errors. Moreover, for 2-factor models, we estimate the downside market price of risk λ_d, then back out its upside counterpart λ_u by exactly matching the observed average excess returns. Thus, we report the standard errors for the estimates of λ_d but not for λ_u.

For all models, the estimates of γ_u and γ_d are positive and statistically significant, thus pointing to the so-called leverage effect in the return dynamics. Moreover, the persistence in the conditional variance process is much lower for the GARCH model compared with the other specifications that exhibit variance persistence levels above 0.98.

The likelihood values allow us to appraise the performance of the alternative specifications under consideration. However, a straightforward comparison

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

TABLE 2

Estimation on Historical Returns and Realized Variance Components

Table 2 shows the maximum likelihood estimation results for four different models. We use daily historical returns, upside realized variances, and downside realized variances for the Standard & Poor's (S&P) 500 index from Jan. 10, 1990, through Aug. 28, 2013. We report the estimated parameters (Est.) with their corresponding standard errors (SE). For each model, we estimate the unconditional variance, then target the fitted value to back out the ϖ parameter, by exploiting the theoretical link between the former and the latter. The parameter λ_u is also inferred from the estimated value of λ_d, by exactly matching the observed (total) market price of risk. To allow for comparison among models, the last row gives the log likelihoods implied for the realized variances, and the second-to-last row indicates log-likelihood values when all models are estimated on returns only, including the benchmark Heston and Nandi (2000) generalized autoregressive conditional heteroscedasticity (GARCH) model. The third-to-last row shows the log likelihoods when the 1-factor affine realized variance (ARV) model is estimated on returns and realized variances. The fourth-to-last row reports log-likelihood values when the 2-factor constrained generalized skew affine realized variance (CGSARV), and generalized skew affine realized variance (GSARV) models are estimated on returns, upside realized variances, and downside realized variances.

	1-Factor Models				2-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
Parameters	Est.	SE	Est.	SE	Est.	SE	Est.	SE
λ_u	1.37E+00		1.37E+00		2.75E+00		2.75E+00	
ϖ_u	2.80E−13		1.51E−12		6.70E−09		4.90E−08	
ω_u					2.39E−07	4.51E−08	9.71E−08	3.31E−08
α_u	4.76E−06	3.12E−07	2.35E−06	6.60E−08	9.01E−07	1.27E−08	9.49E−07	6.61E−07
β_u	8.22E−01	1.25E−02	1.91E−08	2.41E−02	1.57E−05	4.08E−07	1.53E−02	6.73E−01
γ_u	1.73E+02	1.17E+01	6.47E+02	1.62E+01	1.05E+03	7.44E+00	1.01E+03	7.02E+02
σ_u			6.33E−06	1.57E−07	3.49E−06	2.71E−08	3.53E−06	2.44E−06
ρ_u			9.27E−02	8.28E−03	5.88E−01	1.13E−02	5.94E−01	1.21E−02
λ_d					1.93E−08	5.19E+00	4.66E−11	5.34E+00
ϖ_d					9.64E−09		7.32E−09	
ω_d					2.39E−07	4.51E−08	7.39E−06	7.73E−07
α_d					8.51E−07	1.76E−08	9.59E−07	6.15E−07
β_d					9.56E−06	4.45E−08	4.34E−02	6.06E−01
γ_d					1.08E+03	1.11E+01	9.90E+02	6.37E+02
σ_d					3.44E−06	4.02E−08	4.06E−06	2.61E−06
ρ_d					7.59E−01	7.33E−03	7.59E−01	1.00E−02
$E[h_u]$					6.77E−05		6.77E−05	
$E[h_d]$					6.79E−05		6.79E−05	
$E[h]=E[h_u+h_d]$	1.36E−04		1.36E−04		1.36E−04		1.36E−04	
Model Properties								
Avg. upside volatility					11.46		11.46	
Avg. downside volatility					11.40		11.46	
Avg. volatility	16.58		16.58		16.21		16.25	
Variance Persistence								
From RV_u					0.9865		0.9852	
From RV_d					0.9873		0.9840	
From RV			0.9827					
From returns	0.9649							
Log Likelihoods								
Returns, RV_u, and RV_d					120,683		121,036	
Returns and RV			66,565					
Marginalized on returns			19,191		19,237		19,265	
Maximized on returns	19,035		19,213		19,240		19,305	
Implied for RV			47,321		47,015		47,221	

of likelihoods among these models cannot be performed because the GARCH model is only fitted to returns, whereas the ARV model is fitted to returns and realized variance (RV), and the 2-factor models are fitted to returns, RV_u, and RV_d. We see in the bottom panel of Table 2 that the log likelihoods of 2-factor models are many times higher than those of single-factor models. To circumvent this challenge, we implement an additional estimation step for all models (except for the GARCH) by optimizing the likelihood only on returns. The second-to-last row of the log-likelihood panel gives the return-based log likelihoods. According to this metric, the GSARV model dominates the other specifications because it delivers the highest log-likelihood value. To further compare the different models, the last row contains the log likelihoods implied for the realized variance series. Because 2-factor models are optimized on returns, RV_u, and RV_d, they yield slightly lower realized-variance-implied log-likelihood values than the

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

1-factor ARV model, which is fitted to returns, and RV. Nonetheless, the GSARV model delivers the highest realized variance-implied log likelihood among 2-factor models, further confirming the superior performance of this specification.

We turn to Table 3 to assess the ability of the different models to forecast 1-day-ahead realized variance. Specifically, we run Mincer and Zarnowitz (1969) regressions of ex post (observed) on ex ante (model-based) variance components. The corresponding estimation fit is approximately 52% for the GARCH model and above 78% for the other specifications. Moreover, the slope coefficient of the model-implied quantities, which should be 1 for a perfect forecast, is 2.26 for the GARCH and approximately 1.12 for the other models. This confirms that 2-factor models deliver a good fit for upside and downside variances. Additional empirical properties of the various specifications are investigated in the Supplementary Material.

V. Risk-Neutral Estimation

We now estimate the different models by optimizing their fit on option data. This analysis aims at exploring the ability of each specification to properly match the risk-neutral distribution embedded in option contracts. We start by presenting the key features of the option data panel used in our empirical analysis, and then we study the performance of the various models relying on the implied volatility root-mean-squared error.

A. Exploring Option Data

We use European-style options written on the S&P 500 index. The observations span the period Jan. 10, 1996, through Aug. 28, 2013.[3] In line with the extant literature, we only include out-of-the-money (OTM) options with maturity

TABLE 3

Regressions of Model-Free on Model-Implied Variances

Table 3 shows the estimated coefficients (Est.) and the standard errors (SE) of the regression of model-free total (upside and downside) realized variance on the corresponding model-predicted total (upside and downside) variance from each of the four specifications defined in Table 2, in turn. We use the parameter estimates in Table 2 to generate model forecasts of the conditional variance. The sample period spans Jan. 10, 1990, through Aug. 28, 2013.

	1-Factor Models				2-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
Parameters	Est.	SE	Est.	SE	Est.	SE	Est.	SE
Variance								
Constant	−1.43E−04	4.41E−06	−1.63E−05	1.79E−06	−1.71E−05	2.09E−06	−1.70E−05	2.09E−06
Slope	2.26	0.03	1.12	0.01	1.13	0.01	1.12	0.01
R^2 (%)	52.24		84.11		78.60		78.57	
Upside Variance								
Constant					−9.72E−06	1.30E−06	−9.74E−06	1.30E−06
Slope					1.14	0.01	1.14	0.01
R^2 (%)					71.26		71.46	
Downside Variance								
Constant					−1.06E−05	1.20E−06	−1.06E−05	1.20E−06
Slope					1.16	0.01	1.16	0.01
R^2 (%)					75.21		75.00	

[3]Data are available through OptionMetrics, which supplies data for the U.S. option markets.

ranging between 15 and 180 days. This selection procedure is intended to guarantee that the contracts we use are liquid. We also filter out options that violate basic no-arbitrage criteria. For each maturity quoted on Wednesdays, we select only the 6 most liquid strike prices, which amounts to a data set of 21,283 option contracts. To ease calculation and interpretation, OTM put prices are converted into corresponding in-the-money call values by exploiting the call–put parity relationship.

Table 4 provides a crisp description of the option data. To highlight the main characteristics of the S&P 500 index option, we sort the data by moneyness, maturity, and market volatility index (VIX) level. Panel A groups the data by six moneyness buckets and shows the number of contracts, the average option price, the average Black and Scholes (1973) implied volatility, and the average bid–ask spread in dollars. Our measure of moneyness is based on the

TABLE 4
S&P 500 Index Option Data

Table 4 presents the characteristics of Standard & Poor's (S&P) 500 index option data by moneyness, maturity, and market volatility index (VIX) level. We use Wednesday closing out-of-the-money (OTM) call and put contracts from OptionMetrics for the period starting on Jan. 10, 1996, and ending on Aug. 28, 2013. The moneyness is measured by the Black–Scholes (1973) delta (see equation (23)). DTM denotes the number of calendar days to maturity. The average price and bid-ask spread are reported in dollars, whereas the average implied volatility is expressed in percentages.

	OTM Call			OTM Put			
	Delta < 0.3	0.3 ≤ Delta < 0.4	0.4 ≤ Delta < 0.5	0.5 ≤ Delta < 0.6	0.6 ≤ Delta < 0.7	Delta ≥ 0.7	All
Panel A. By Moneyness							
No. of contracts	3,788	1,391	1,781	2,846	2,746	8,731	21,283
Avg. price	7.85	20.94	32.28	45.30	65.93	132.41	74.35
Avg. implied volatility	16.72	18.40	19.31	20.40	21.71	25.09	21.62
Avg. bid–ask spread	1.046	1.674	1.955	2.018	1.834	1.228	1.470
	DTM < 30	30 ≤ DTM < 60	60 ≤ DTM < 90	90 ≤ DTM < 120	120 ≤ DTM < 150	DTM ≥ 150	All
Panel B. By Maturity							
No. of contracts	2,725	6,480	5,053	2,869	1,974	2,182	21,283
Avg. price	41.26	61.01	76.44	92.30	97.88	105.59	74.35
Avg. implied volatility	20.21	21.28	21.73	22.94	22.08	21.95	21.62
Avg. bid–ask spread	0.820	1.231	1.578	1.872	1.800	1.910	1.470
	VIX < 15	15 ≤ VIX < 20	20 ≤ VIX < 25	25 ≤ VIX < 30	30 ≤ VIX < 35	VIX ≥ 35	All
Panel C. By VIX Level							
No. of contracts	3,962	6,133	5,996	2,456	1,240	1,496	21,283
Avg. price	57.95	66.90	80.75	85.77	85.33	94.86	74.35
Avg. implied volatility	13.61	18.04	22.45	26.24	30.22	39.42	21.62
Avg. bid–ask spread	1.055	1.301	1.446	1.704	1.811	2.683	1.470

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

Black–Scholes delta computed as

$$\Delta = \Phi\left(\frac{\ln(S_t/X) + r_f M + 1/2\left(\mathrm{IV}^{\mathrm{Mkt}}\right)^2 M/365}{\mathrm{IV}^{\mathrm{Mkt}}\sqrt{M/365}}\right), \tag{23}$$

where $\Phi(\cdot)$ stands for the normal CDF, X is the strike price, r_f is the nonannualized daily risk-free rate, M is the time to maturity expressed in days, and $\mathrm{IV}^{\mathrm{Mkt}}$ denotes the annualized implied Black–Scholes volatility computed at the market price of the option. A few empirical regularities emerge at this point. We observe that deep OTM puts, which include the largest number of contracts with deltas exceeding 0.7, are the most expensive. This echoes the well-documented volatility smirk pattern in index options across moneyness.

Panel B of Table 4 sorts the data by maturity expressed in calendar days. Even though the term structure of volatility is nearly flat on average during the sample period, we notice that options with longer maturities are relatively more expensive.

Panel C of Table 4 classifies the data by the VIX level. It is immediately obvious that a large portion of the selected option contracts (75%) are quoted on days with VIX levels ranging between 15% and 35%.

Overall, a typical "median" contract features a delta above 0.6 and a time to expiry between 30 and 90 days and is quoted on "normal" days when the VIX lies within the [15%–25%] interval.

B. Fitting Options

We explore the performance of the different models by relying on the implied volatility root-mean-squared error (IVRMSE) metric. Renault (1997) advocates for using the IVRMSE as a proper tool for the comparison of model performance in option pricing. Basically, the IVRMSE synthesizes the discrepancy between model-based and market-based implied volatilities. To compute the IVRMSE, we invert the model-based option price C_j^{Mod} of each contract j using the Black–Scholes (1973) formula (BS). Thus, the model-based implied volatility can be formally extracted according to

$$\mathrm{IV}_j^{\mathrm{Mod}} = \mathrm{BS}^{-1}\left(C_j^{\mathrm{Mod}}\right).$$

Applying a similar procedure to the set of observed option contracts $\{C_j^{\mathrm{Mkt}}\}$ yields market-based implied volatilities

$$\mathrm{IV}_j^{\mathrm{Mkt}} = \mathrm{BS}^{-1}\left(C_j^{\mathrm{Mkt}}\right).$$

Accordingly, the implied volatility error is computed as

$$e_j = \mathrm{IV}_j^{\mathrm{Mkt}} - \mathrm{IV}_j^{\mathrm{Mod}}.$$

It follows that the IVRMSE is obtained as

$$\mathrm{IVRMSE} \equiv \sqrt{\frac{1}{N}\sum_{j=1}^{N} e_j^2},$$

where N denotes the option sample size.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

Finally, risk-neutral parameters are estimated by maximizing the Gaussian implied volatility error likelihood:[4]

$$\ln L^{O} \quad \propto \quad -\frac{1}{2}\sum_{j=1}^{N}\left\{\ln\left(\text{IVRMSE}^2\right)+e_j^2/\text{IVRMSE}^2\right\}. \tag{24}$$

Table 5 contains the results of the option-based estimation. Clearly, our option-fitting strategy yields accurate parameter estimates, as evidenced by fairly small standard errors and sizable model likelihoods. Because we are fitting the model only on options, the resulting estimates correspond to risk-neutral parameters. Thus, the two market prices of risk, λ_u and λ_d, are not estimated. Note that to ensure model consistency in the estimation step, we filter volatility on returns and RV while fitting option implied volatilities (IVs). As done in the historical estimation, we estimate $\text{E}^{Q}[h_j]$ and then back out the ϖ_j estimates from the theoretical unconditional risk-neutral variance formula.

TABLE 5
Estimation on Options

Table 5 shows the estimation results for the four different models defined in Table 2. We use Wednesday closing out-of-the-money (OTM) call and put contracts from OptionMetrics for the period starting on Jan. 10, 1996, and ending on Aug. 28, 2013. We report the estimated parameters (Est.) along with their corresponding standard errors (SE). For each model, we estimate the unconditional variance, then target the fitted value to back out the ϖ parameter, by exploiting the theoretical link between the former and the latter. The second-to-last row shows the implied volatility root-mean-squared errors (IVRMSEs) of all models. For comparison, the last row reports the IVRMSE ratio of each specification to the benchmark generalized autoregressive conditional heteroscedasticity (GARCH) model.

	1-Factor Models				2-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
Parameters	Est	SE	Est	SE	Est	SE	Est	SE
ϖ_u	8.15E−21		8.57E−09		3.91E−20		3.24E−20	
ω_u					9.62E−06	7.42E−07	9.03E−08	4.60E−07
α_u	8.17E−07	1.04E−08	7.88E−07	2.99E−08	1.67E−08	3.55E−09	1.22E−08	2.28E−09
β_u	8.89E−01	1.72E−03	2.62E−02	5.39E−04	1.03E−11	1.62E−04	6.10E−13	1.62E−04
γ_u	3.57E+02	4.20E+00	1.11E+03	2.08E+01	7.72E+03	1.45E+02	9.03E+03	8.05E+02
σ_u			4.30E−05	1.61E−06	9.70E−06	1.74E−06	5.82E−06	1.13E−06
ρ_u			4.64E−01	7.84E−03	9.88E−01	4.66E−07	9.88E−01	4.66E−07
ϖ_d					4.38E−08		1.09E−17	
ω_d					9.62E−06	5.01E−07	2.93E−05	4.92E−07
α_d					4.56E−07	1.59E−27	5.64E−07	2.94E−27
β_d					5.44E−11	7.42E−07	2.63E−11	4.60E−07
γ_d					1.48E+03	1.27E−08	1.33E+03	1.25E−08
σ_d					1.03E−05	1.80E+01	1.20E−05	1.04E+01
ρ_d					9.99E−01	2.69E−08	9.99E−01	2.69E−08
$E^{Q}[h_u]$					1.40E−05	5.24E−07	3.13E−06	5.36E−07
$E^{Q}[h_d]$					1.16E−04	2.67E−07	1.40E−04	2.13E−07
$E^{Q}[h]=E^{Q}[h_u+h_d]$	1.11E−04	4.91E−07	8.27E−05	6.21E−07	1.30E−04		1.44E−04	
Model Properties								
Log likelihoods	37,615		38,873		41,041		41,236	
Avg. upside model IV					9.31		8.47	
Avg. downside model IV					13.76		14.00	
Avg. model IV	16.62		17.76		16.73		16.45	
Variance Persistence								
From RV_u					0.9962		0.9960	
From RV_d					0.9953		0.9949	
From RV	0.9926		0.9904					
Option Errors								
IVRMSE	4.816		4.195		3.865		3.862	
Ratio to GARCH	1.000		0.871		0.803		0.802	

[4]Later in the article, the model optimization algorithm maximizes the joint likelihood on returns and options. Thus, to allow for comparison, we maximize the option likelihood at this point rather than minimizing the IVRMSE.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

The proposed GSARV model clearly outperforms the alternative specifications because it delivers the highest likelihood value and the smallest global IVRMSE. Specifically, the GSARV model offers approximately 10% and 20% improvement, respectively, in terms of log likelihood and IVRMSE over the benchmark GARCH model.

C. Dissecting Model Fit

We now scrutinize the overall performance results reported in the bottom panel of Table 5. To this end, we dissect the IVRMSE by moneyness, maturity, and VIX levels in Table 6. We use the same clusters as for the description of option data in Table 4.

We see that all models offer a satisfactory performance (low IVRMSEs) in matching at-the-money options contracts. By contrast, fitting deep OTM call and

TABLE 6

IVRMSE Option Error by Moneyness, Maturity, and VIX

Table 6 presents the implied volatility root-mean-squared error (IVRMSE) of the four models defined in Table 2, for contracts sorted by moneyness, maturity, and market volatility index (VIX) level. We use the parameter values estimated in Table 5 to fit our four models to Standard & Poor's (S&P) 500 index out-of-the-money (OTM) call and put option contracts from OptionMetrics. The sample starts on Jan. 10, 1996, and ends on Aug. 28, 2013. Panel A reports the IVRMSEs for contracts sorted by moneyness defined using the Black–Scholes (1973) delta (see equation (23)). Panel B reports the IVRMSEs for contracts sorted by days to maturity (DTM). Panel C reports the IVRMSEs for contracts sorted by the VIX level on the day corresponding to the option quote. The IVRMSEs are expressed in percentages.

	OTM Call			OTM Put		
Model	Delta $<$ 0.3	0.3 $\leq$ Delta $<$ 0.4	0.4 $\leq$ Delta $<$ 0.5	0.5 $\leq$ Delta $<$ 0.6	0.6 $\leq$ Delta $<$ 0.7	Delta $\geq$ 0.7
Panel A. IVRMSE by Moneyness						
GARCH	4.497	3.381	3.396	3.693	4.202	5.816
ARV	4.427	3.124	3.220	3.288	3.687	4.807
CGSARV	3.892	2.775	2.913	3.036	3.365	4.518
GSARV	3.985	2.757	2.910	3.034	3.375	4.490
Model	DTM $<$ 30	30 $\leq$ DTM $<$ 60	60 $\leq$ DTM $<$ 90	90 $\leq$ DTM $<$ 120	120 $\leq$ DTM $<$ 150	DTM $\geq$ 150
Panel B. IVRMSE by Maturity						
GARCH	4.833	4.888	4.647	4.644	5.218	4.806
ARV	4.125	4.068	3.944	4.058	4.880	4.699
CGSARV	3.793	3.916	3.706	3.657	4.300	3.984
GSARV	3.736	3.953	3.710	3.677	4.293	3.947
Model	VIX $<$ 15	15 $\leq$ VIX $<$ 20	20 $\leq$ VIX $<$ 25	25 $\leq$ VIX $<$ 30	30 $\leq$ VIX $<$ 35	VIX $\geq$ 35
Panel C. IVRMSE by VIX Level						
GARCH	3.233	3.089	5.020	5.921	6.034	8.629
ARV	2.950	2.469	4.245	5.465	6.050	7.128
CGSARV	2.598	2.369	3.935	5.126	5.636	6.310
GSARV	2.620	2.428	3.954	4.963	5.694	6.341

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

put options seems more challenging. Interestingly, the ability of the various specifications to match the observed option-implied volatility appears consistent across the term structure of the options because the IVRMSEs are of comparable magnitude. Moreover, the performance of these models tends to deteriorate nearly monotonically as a function of the VIX level. This observation suggests that the ability of the models to generate realistic option prices weakens in highly volatile times.

Across the board, the GSARV model dominates the other models along the dimensions of moneyness, maturity, and VIX level.

VI. Joint Estimation on Returns, RV^U, RV^D, and Options

The option-based optimization findings provide relevant information on the ability of each model to describe the risk-neutral dynamics embedded in option contracts. Nonetheless, these results offer a limited description of certain model features, namely, the performance of the pricing kernel. Recall that the pricing kernel is at the heart of the valuation of contingent claims and the analysis of market compensation of risks. Thus, we extend our investigation by jointly fitting the various models to historical returns, upside and downside realized variances, and option data.

The optimization is performed over a convolution of the quasi log likelihood of returns and semirealized variances, denoted as $\ln L^P$ in equation (22), and an option-based log-likelihood component $\ln L^O$ as per equation (24). Accordingly, our joint maximization program writes

$$\max \ln L^P + \ln L^O. \tag{25}$$

Solving the joint optimization problem should help elicit the different premia, including the total variance risk premium, the upside variance risk premium, and the downside variance risk premium. To this end, we need to specify the link between the physical and the risk-neutral probability measures.

A. Risk Neutralization and Implied Risk Premia

1. Pricing Kernel and Risk Neutralization

We now focus on the mapping between physical and risk-neutral probability measures using an exponential pricing kernel of the form

$$M_{t+1} = M_{t+1}^{(u)} M_{t+1}^{(d)}, \tag{26}$$

where

$$M_{t+1}^{(j)} = \frac{\exp\left(\nu_{1t}^{(1j)}\varepsilon_{j,t+1}^{(1)} + \nu_2^{(1j)}\left(\varepsilon_{j,t+1}^{(1)}\right)^2 + \nu_{1t}^{(2j)}\varepsilon_{j,t+1}^{(2)} + \nu_2^{(2j)}\left(\varepsilon_{j,t+1}^{(2)}\right)^2\right)}{\mathrm{E}_t\left[\exp\left(\nu_{1t}^{(1j)}\varepsilon_{j,t+1}^{(1)} + \nu_2^{(1j)}\left(\varepsilon_{j,t+1}^{(1j)}\right)^2 + \nu_{1t}^{(2j)}\varepsilon_{j,t+1}^{(2j)} + \nu_2^{(2j)}\left(\varepsilon_{j,t+1}^{(2j)}\right)^2\right)\right]}, \quad \text{for } j = u, d.$$

Our pricing kernel entails that both $\varepsilon_{j,t+1}^{(i)}$ shocks and their squares $\left(\varepsilon_{j,t+1}^{(i)}\right)^2$ are priced. This represents a significant departure from Christoffersen, Elkamhi, Feunou, and Jacobs (2010), where only the $\varepsilon_{j,t+1}^{(i)}$ disturbances are priced. In our specification, the ν_2s are key elements that drive the wedge between the 1-day-ahead conditional risk-neutral and physical variance (termed as the variance

spread).The magnitude of this variance spread may be nontrivial, as pointed out by Christoffersen, Heston, and Jacobs (2013) in their conditional Gaussian model. In the Appendix, we report explicit expressions of the pricing kernel parameters ($\nu_{1t}^{(1j)}$, $\nu_{2}^{(1j)}$, $\nu_{1t}^{(2j)}$, and $\nu_{2}^{(2j)}$) as functions of physical and risk-neutral parameters. The derivations in the Supplementary Material also show that the risk-neutral motion of the underlying asset return is

$$\begin{aligned} R_{t+1} &= \bar{r}^{\mathcal{Q}} - \xi_d^{\mathcal{Q}} h_{dt}^{\mathcal{Q}} - \xi_u^{\mathcal{Q}} h_{ut}^{\mathcal{Q}} + z_{u,t+1}^{\mathcal{Q}} - z_{d,t+1}^{\mathcal{Q}}, \\ z_{j,t+1}^{\mathcal{Q}} &= \tilde{z}_{j,t+1}^{\mathcal{Q}} - \mathrm{E}_t^{\mathcal{Q}}\left[\tilde{z}_{j,t+1}^{\mathcal{Q}}\right], \\ \mathrm{var}_t^{\mathcal{Q}}\left[z_{j,t+1}^{\mathcal{Q}}\right] &= \mathrm{var}_t^{\mathcal{Q}}\left[\tilde{z}_{j,t+1}^{\mathcal{Q}}\right] \equiv h_{j,t}^{\mathcal{Q}}, \quad \text{for } j=\{u,d\}, \end{aligned}$$

where

$$\begin{aligned} \bar{r}^{\mathcal{Q}} &= r_f + \frac{1}{2}\left[\ln\left(1+\sqrt{2\omega_d^{\mathcal{Q}}}\right) - \frac{\omega_d^{\mathcal{Q}}+\sqrt{2\omega_d^{\mathcal{Q}}}}{\left(1+\sqrt{2\omega_d^{\mathcal{Q}}}\right)}\right] + \frac{1}{2}\left[\ln\left(1-\sqrt{2\omega_u^{\mathcal{Q}}}\right) - \frac{\omega_u^{\mathcal{Q}}-\sqrt{2\omega_u^{\mathcal{Q}}}}{\left(1-\sqrt{2\omega_u^{\mathcal{Q}}}\right)}\right], \\ \xi_d^{\mathcal{Q}} &= \frac{1}{2\left(1+\sqrt{2\omega_d^{\mathcal{Q}}}\right)}, \quad \xi_u^{\mathcal{Q}} = \frac{1}{2\left(1-\sqrt{2\omega_u^{\mathcal{Q}}}\right)}. \end{aligned}$$

The risk-neutral innovations in returns are noncentral χ^2, distributed as

$$\tilde{z}_{j,t+1}^{\mathcal{Q}} = \sqrt{\frac{\omega_j^{\mathcal{Q}}}{2}}\left(\varepsilon_{j,t+1}^{(*1)} - \sqrt{\frac{h_{j,t}^{\mathcal{Q}}-\omega_j^{\mathcal{Q}}}{2\omega_j^{\mathcal{Q}}}}\right)^2, \quad \text{with } \varepsilon_{j,t+1}^{(*1)} \overset{\text{IID}^{\mathcal{Q}}}{\sim} \mathrm{N}(0,1) \text{ for } j=\{u,d\},$$

and $\mathrm{E}^{\mathcal{Q}}\left[\varepsilon_{u,t+1}^{(*1)}\varepsilon_{d,t+1}^{(*1)}\right]=0$.

Moreover, there is a spread between the 1-day-ahead risk-neutral variances $h_{u,t}^{\mathcal{Q}}$, $h_{d,t}^{\mathcal{Q}}$ and their physical counterparts $h_{u,t}$, $h_{d,t}$,

$$\text{(27)} \qquad \mathrm{VSP}_t^j \equiv h_{j,t}^{\mathcal{Q}} - h_{j,t} = \vartheta_j + (\varsigma_j - 1)h_{j,t},$$

with

$$\begin{aligned} \varsigma_d &= \frac{\sqrt{2\omega_d^{\mathcal{Q}}}\left(1+\sqrt{2\omega_d^{\mathcal{Q}}}\right)}{\sqrt{2\omega_d}\left(1+\sqrt{2\omega_d}\right)} + 2\sqrt{2\omega_d^{\mathcal{Q}}}\left(1+\sqrt{2\omega_d^{\mathcal{Q}}}\right)\lambda_d, \\ \varsigma_u &= \frac{\sqrt{2\omega_u^{\mathcal{Q}}}\left(1-\sqrt{2\omega_u^{\mathcal{Q}}}\right)}{\sqrt{2\omega_u}\left(1-\sqrt{2\omega_u}\right)} - 2\sqrt{2\omega_u^{\mathcal{Q}}}\left(1-\sqrt{2\omega_u^{\mathcal{Q}}}\right)\lambda_u, \end{aligned}$$

and

$$\begin{aligned} \vartheta_d &= \sqrt{2\omega_d^{\mathcal{Q}}}\left(1+\sqrt{2\omega_d^{\mathcal{Q}}}\right)\ln\left(\frac{1+\sqrt{2\omega_d}}{1+\sqrt{2\omega_d^{\mathcal{Q}}}}\right) + \frac{\sqrt{\omega_d^{\mathcal{Q}}}\left(\sqrt{\omega_d^{\mathcal{Q}}}-\sqrt{\omega_d}\right)}{1+\sqrt{2\omega_d}}, \\ \vartheta_u &= \sqrt{2\omega_u^{\mathcal{Q}}}\left(1-\sqrt{2\omega_u^{\mathcal{Q}}}\right)\ln\left(\frac{1-\sqrt{2\omega_u}}{1-\sqrt{2\omega_u^{\mathcal{Q}}}}\right) + \frac{\sqrt{\omega_u^{\mathcal{Q}}}\left(\sqrt{\omega_u^{\mathcal{Q}}}-\sqrt{\omega_u}\right)}{1-\sqrt{2\omega_u}}. \end{aligned}$$

We provide some analysis of this wedge in Section VI.A.2.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

The risk-neutral dynamics for upside and downside realized variances write

$$\mathrm{RV}^{j}_{t+1} = h_{j,t} + \mathrm{VRP}^{j}_{t} + z^{(*2)}_{j,t+1},$$

where

$$\begin{aligned} z^{(*2)}_{j,t+1} &= \tilde{z}^{(*2)}_{j,t+1} - \mathrm{E}^{Q}_{t}\left[\tilde{z}^{(*2)}_{j,t+1}\right] \\ \tilde{z}^{(*2)}_{j,t+1} &= \sigma^{Q}_{j}\left(\varepsilon^{*(2)}_{j,t+1} - \gamma^{Q}_{j}\sqrt{h^{Q}_{j,t} - \omega^{Q}_{j}}\right)^{2}, \end{aligned}$$

$\varepsilon^{(*2)}_{j,t+1} \overset{\mathrm{IID}^{Q}}{\sim} \mathrm{N}(0,1)$, $\mathrm{E}^{Q}\left[\varepsilon^{(*2)}_{u,t+1}\varepsilon^{(*2)}_{d,t+1}\right] = 0$, and $\mathrm{E}^{Q}_{t}\left[\varepsilon^{(*1)}_{j,t+1}\varepsilon^{(*2)}_{j,t+1}\right] \equiv \rho^{Q}_{j}$.

In addition, there is a difference between the risk-neutral expectation of the realized variances $\mathrm{E}^{Q}_{t}\left[\mathrm{RV}^{U}_{t+1}\right], \mathrm{E}^{Q}_{t}\left[\mathrm{RV}^{D}_{t+1}\right]$ and their physical counterparts $h_{u,t}, h_{d,t}$. The difference is the variance risk premium and is given by

$$\mathrm{VRP}^{j}_{t} \equiv \mathrm{E}^{Q}_{t}\left[\mathrm{RV}^{j}_{t+1}\right] - \mathrm{E}^{P}_{t}\left[\mathrm{RV}^{j}_{t+1}\right] = \theta_{j} + \varrho_{j} h_{jt} + \sigma^{Q}_{j}\left(\gamma^{Q}_{j}\right)^{2}\mathrm{VSP}^{j}_{t}, \tag{28}$$

with

$$\begin{aligned} \varrho_{j} &= \sigma^{Q}_{j}\left(\gamma^{Q}_{j}\right)^{2} - \sigma_{j}\gamma^{2}_{j}, \\ \theta_{j} &= \sigma^{Q}_{j} - \sigma_{j} + \sigma_{j}\gamma^{2}_{j}\omega_{j} - \sigma^{Q}_{j}\left(\gamma^{Q}_{j}\right)^{2}\omega^{Q}_{j}. \end{aligned}$$

It is important to point out the subtle difference between the variance spread and premium. The two notions are clearly related because more spread induces a premium. But there can be a premium without spread; we will provide some analysis of these differences in Section VI.A.2.

The dynamics of the risk-neutral factor $h^{Q}_{j,t+1}$ under the risk-neutral probability measure share a similar structure with their physical counterparts $h_{j,t+1}$. That is,

$$h^{Q}_{j,t+1} - \omega^{Q}_{j} = \varpi^{Q}_{j} + \beta_{j}\left(h^{Q}_{j,t} - \omega^{Q}_{j}\right) + \alpha^{Q}_{j}\left(\varepsilon^{(*2)}_{j,t+1} - \gamma^{Q}_{j}\sqrt{h^{Q}_{j,t} - \omega^{Q}_{j}}\right)^{2}, \tag{29}$$

with $\varpi^{Q}_{j} = \left(1-\beta_{j}\right)\left(\vartheta_{j} + \varsigma_{j}\omega_{j} - \omega^{Q}_{j}\right) + \varsigma_{j}\varpi_{j}$.

We provide the mapping between risk-neutral parameters, physical parameters, and pricing kernel in the Appendix.

2. Analyzing Variance Spread and Variance Risk Premium Drivers

The variance spread expression in equation (27) reveals three potential factors driving the wedge between risk-neutral and physical variances: i) the non-normality ($\omega_{j}, \omega^{Q}_{j} \neq 0$), ii) the distance between asymmetries ($\|\omega^{Q}_{j} - \omega_{j}\|$), and iii) the market compensations for upside and downside variance risk exposures (λ_{j}). Importantly, our choice of a flexible stochastic discount factor induces a spread between risk-neutral and physical variances even in the Gaussian case ($\omega_{j}, \omega^{Q}_{j} = 0$); this is consistent with the findings of Christoffersen et al. (2013). We refer the reader to the Supplementary Material for technical details. In the non-Gaussian shock case, compensations for upside and downside variance risk exposures (λ_{j}) affect variance spreads. Specifically, an increase in the compensation λ_{j} increases the corresponding variance spread. When $\omega^{Q}_{j} = \omega_{j}$, the risk-neutral variance becomes proportional to its historical counterpart. When $\omega^{Q}_{j} \neq \omega_{j}$, an increase in the gap between ω^{Q}_{j} and ω_{j} induces a larger variance spread.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

Option valuation models are expected to generate risk-neutral total variances that are larger than physical total variances in order to be economically and empirically relevant. Our framework entails that compensations for upside and downside variance exposures (λ_j) contribute to the dynamics of risk-neutral variances. Thus, $\lambda_d \geq 0$ and $\lambda_u \leq 0$ are sufficient conditions ensuring that risk-neutral variances are well defined (nonnegative). Moreover, combining these conditions with $\omega_j^Q > \omega_j$ implies that $\text{VSP}_t^j \geq 0$. These theoretical results clearly show that the variance spread is a key driver of the variance risk premium, with the two quantities moving in the same direction. Looking at equation (28), one can interpret ϱ_j as a difference between the risk-neutral and physical variance of realized variance. This difference plays a pivotal role in the determination of the variance risk premium. Building on the intuition that investors like good uncertainty (because it increases the potential of substantial gains) but dislike bad uncertainty (as it increases the likelihood of severe losses), Feunou et al. (2018) have documented that the premium on the downside variance is positive $\left(\text{VRP}_t^d > 0\right)$, whereas its upside counterpart is negative $\left(\text{VRP}_t^u < 0\right)$. These empirical regularities are met in our model if $\varrho_d > 0$ and $\varrho_u < 0$, thus implying economically sound restrictions, that is:

$$\begin{aligned} \sigma_d^Q \left(\gamma_d^Q\right)^2 &> \sigma_d \gamma_d^2, \\ \sigma_u^Q \left(\gamma_u^Q\right)^2 &< \sigma_u \gamma_u^2. \end{aligned}$$

B. Empirical Results

1. Model Fit

Table 7 presents the estimation results from the joint likelihood maximization on returns, RV^U, RV^D, and options. Our optimization procedure achieves an accurate fit of the entire set of parameters associated with the different models. Again, the ϖ parameter is calibrated by targeting the physical unconditional (total) variance. By contrast, the downside risk premium parameter λ_d is estimated as a free parameter, whereas its upside counterpart λ_u is inferred by exactly matching the observed (total) market price of risk.

The estimated values for λ_d are all positive and more than 2 standard errors away from 0. Interestingly for the GSARV and CGSARV models, the inferred λ_u values are found to be negative. Thus, the joint estimation results for the GSARV and CGSARV specifications are consistent with the intuition that investors dislike downside uncertainty and demand a positive premium as compensation for bearing that risk, whereas they find upside uncertainty desirable and are willing to pay (negative premium) for exposure to such risk. A typical investor's asymmetric behavior toward good versus bad uncertainty is in line with the empirical regularities documented in recent works, such as those by Feunou et al. (2018) and Feunou, Jahan-Parvar, and Tédongap (2013), among others. Moreover, the joint estimation results for the GSARV and CGSARV ($\omega \equiv \omega_u = \omega_d$) models are close, which suggests that the wedge between upside and downside conditional volatilities ($h_{u,t} - h_{d,t}$) is the main driver of the conditional skewness in our option valuation framework. Our findings underscore the importance of conditional asymmetry channels in designing empirically successful option valuation models

TABLE 7

Estimation on Historical Returns, Realized Variance Components, and Options

Table 7 shows the joint maximum likelihood estimation results for four different models defined in Table 2. We use daily historical returns, upside/downside realized variances, and options on the Standard & Poor's (S&P) 500 index from Jan. 10, 1996, through Aug. 28, 2013. We report the estimated parameters (Est.) with their corresponding standard errors (SE). For each model, we use physical unconditional variance targeting to back out the ϖ parameter. The parameter λ_u is also inferred from the estimated value of λ_d, by exactly matching the observed (total) market price of risk. We also present the joint log-likelihood value along with its decomposition into the several components. The second-to-last row shows the implied volatility root-mean-squared errors (IVRMSEs; in percentages) of all models. For comparison, the last row reports the IVRMSE ratio of each specification to the benchmark generalized autoregressive conditional heteroscedasticity (GARCH) model.

	1-Factor Models				2-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
Parameters	Est.	SE	Est.	SE	Est.	SE	Est.	SE
λ_u	1.40E+01		1.51E+00		−6.67E−06		−5.70E−08	
ϖ_u	1.64E−08		7.00E−19		0.00E+00		0.00E+00	
ω_u					1.45E−06	1.04E−08	1.44E−06	2.61E−09
α_u	9.01E−07	1.86E−08	8.79E−08	6.93E−10	8.11E−07	3.01E−09	7.99E−07	3.23E−10
β_u	9.88E−01	6.09E−04	2.41E−01	2.31E−03	1.11E−16	4.22E−13	0.00E+00	1.72E−14
γ_u	6.22E+01	5.51E+00	2.94E+03	1.41E+01	1.10E+03	1.97E+00	1.11E+03	2.16E−01
σ_u			1.84E−06	8.03E−09	3.17E−06	6.95E−09	3.12E−06	4.77E−09
ρ_u			1.31E−01	2.28E−03	4.39E−01	4.40E−03	4.38E−01	4.85E−03
λ_d					3.01E+00	1.65E−08	3.01E+00	1.04E−06
ϖ_d					2.38E−14		7.41E−10	
ω_d					1.45E−06	1.04E−08	1.46E−06	2.37E−09
α_d					4.93E−08	2.05E−10	5.06E−08	2.87E−10
β_d					0.00E+00	1.60E−13	2.22E−16	4.39E−12
γ_d					4.50E+03	8.98E+00	4.44E+03	1.21E+01
σ_d					1.07E−06	2.31E−09	1.10E−06	3.51E−09
ρ_d					7.80E−01	5.41E−03	7.89E−01	4.64E−03
Pricing Kernel Parameters								
κ_1^u	1.05E+00	3.36E−03	1.57E+00	4.50E−03	6.10E+04	1.07E+03	6.18E+04	1.08E+03
κ_2^u			3.60E−02	3.34E−04	3.91E−06	6.88E−08	3.73E−06	6.54E−08
γ_u^Q	8.00E+01		6.96E+02	3.13E+00	5.38E+02	9.83E+00	5.13E+02	9.84E+00
κ_1^u					9.35E−01	3.27E−03	9.37E−01	3.30E−03
κ_2^u					1.30E−01	1.64E−03	1.34E−01	1.71E−03
γ_u^Q					1.50E+03	1.06E+01	1.57E+03	1.06E+01
Model Properties								
Avg. physical volatility	16.58		16.58		16.21		16.21	
Avg. model implied volatility	17.32		14.99		14.23		14.20	
Variance Persistence								
From RV_u					0.9889		0.9853	
From RV_d					0.9843		0.9867	
From RV			0.9835					
From returns	0.9893							
Log Likelihoods								
Returns, RV_u, RV_d, and options					163,429		163,432	
Returns, RV, and options			106,391					
Returns and options	52,770		59,294		61,070		61,098	
Returns	19,019		19,680		19,761		19,785	
Option Errors								
IVRMSE	5.740		4.243		4.031		4.031	
Ratio to GARCH	1.000		0.739		0.702		0.702	

that are able to generate distinct dynamics for the market compensation of upside and downside risks.

For each specification, Table 7 also reports the corresponding joint log-likelihood value along with its decomposition into the several components. To allow for comparison among models, we compute the log-likelihood component pertaining exclusively to returns. It is immediately clear that, in contrast to the GARCH model, the estimated CGSARV model delivers the highest log likelihood of returns.

In terms of fitting accuracy, the joint estimation results in Table 7 show a quite impressive 30% IVRMSE reduction for the GSARV model compared with the benchmark GARCH model. Furthermore, the global fit of the GSARV model (IVRMSE of 4.03%) reveals a sizable 10% improvement over

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

Christoffersen et al.'s (2015) BPJVM (with an IVRMSE of 4.77%), which features an alternative decomposition of the realized variance into diffusive and discontinuous components. This observation is further confirmed in Table 8, where the IVRMSEs reported at the bottom of Table 7 are disaggregated by intervals of moneyness, maturity, and VIX index level. It is immediately apparent that the overall improvement in option fit for the GSARV model is not driven by any particular subset of the data. Across the board, the GSARV specification delivers a better fit with respect to its nested variants as well as the BPJVM of Christoffersen et al. (2015).

TABLE 8

Joint Estimation of IVRMSE Option Error by Moneyness, Maturity, and VIX

Table 8 presents the implied volatility root-mean-squared errors (IVRMSEs) of the four models defined in Table 2, for contracts sorted by moneyness, maturity, and market volatility index (VIX) level. We use the parameter values estimated in Table 7 to fit our four models to Standard & Poor's (S&P) 500 index out-of-the-money (OTM) call and put option contracts from OptionMetrics. The sample starts on Jan. 10, 1996, and ends on Aug. 28, 2013. Panel A reports the IVRMSEs for contracts sorted by moneyness defined using the Black–Scholes (1973) delta (see equation (23)). Panel B reports the IVRMSEs for contracts sorted by days to maturity (DTM). Panel C reports the IVRMSEs for contracts sorted by the VIX level on the day corresponding to the option quote. The IVRMSEs are expressed in percentages. BPJVM denotes the bipower jump-variation option pricing model of Christoffersen et al. (2015).

	OTM Call			OTM Put		
Model	Delta < 0.3	0.3 ≤ Delta < 0.4	0.4 ≤ Delta < 0.5	0.5 ≤ Delta < 0.6	0.6 ≤ Delta < 0.7	Delta ≥ 0.7
Panel A. IVRMSE by Moneyness						
GARCH	5.338	4.001	3.823	3.896	4.228	8.112
ARV	4.530	3.214	3.321	3.422	3.703	4.697
CGSARV	4.355	3.002	3.032	3.110	3.347	4.645
GSARV	4.358	3.004	3.033	3.111	3.347	4.643
BPJVM	4.775	3.150	2.825	2.956	3.319	6.821
Model	DTM < 30	30 ≤ DTM < 60	60 ≤ DTM < 90	90 ≤ DTM < 120	120 ≤ DTM < 150	DTM ≥ 150
Panel B. IVRMSE by Maturity						
GARCH	5.259	5.700	5.640	5.852	6.497	5.834
ARV	4.514	4.206	4.029	3.938	4.486	4.157
CGSARV	4.433	4.113	3.801	3.683	4.314	3.950
GSARV	4.428	4.113	3.803	3.685	4.314	3.950
BPJVM	4.404	4.731	4.739	4.555	5.535	4.948
Model	VIX < 15	15 ≤ VIX < 20	20 ≤ VIX < 25	25 ≤ VIX < 30	30 ≤ VIX < 35	VIX ≥ 35
Panel C. IVRMSE by VIX Level						
GARCH	4.310	3.419	5.499	6.785	6.961	11.639
ARV	3.032	2.576	4.248	5.528	6.144	6.643
CGSARV	2.794	2.433	4.158	5.489	5.878	6.234
GSARV	2.796	2.431	4.157	5.488	5.882	6.235
BPJVM	3.335	3.202	5.358	5.868	5.888	7.231

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

2. Model Fit for Variance Risk Premia

From the works of Bakshi and Madan (2000), Carr and Madan (2001), Bakshi, Kapadia, and Madan (2003), and Andersen, Bondarenko, and Gonzalez-Perez (2015), we have a clear understanding of how to construct nonparametric option-implied moments. The corresponding model-free physical moments can also be obtained from high-frequency historical returns on the underlying asset, as discussed by Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Diebold, and Labys (2003), Barndorff-Nielsen et al. (2010), and Patton and Sheppard (2015), among others. Relying on these two strands of the literature, we compute nonparametric term structures of total, upside, and downside variance risk premia as spreads between risk-neutral and physical corresponding series at different maturities.

To further gauge the empirical performance of our pricing framework, we report in Table 9 the Mincer and Zarnowitz (1969) regressions of the model-free variance premium and its upside (downside) component on their model-implied counterparts. The results clearly show that the 2-factor GSARV model delivers efficient predictions of variance risk premia at different maturities, compared with single-factor specifications. Across maturities, the slope estimates for the GSARV model are virtually equal to 1, with higher adjusted R^2 figures compared with 1-factor models. This evidence suggests that specifying distinct dynamics for upside and downside second-order variations helps significantly improve the model's ability to track the term structures of market compensations for total, upside, and downside variance risks. Focusing on the components of the variance risk premium, we see that forecasts of the 2-factor model for upside variance risk premium are more accurate than the model-implied series for the downside variance risk premium. Predicting the downside component of the variance risk premium appears more challenging, especially at the short horizon, possibly due to its high empirical variability.

TABLE 9

Regressions of Model-Free on Model-Implied Variance Risk Premia

Table 9 shows the estimated coefficients (Est.) and the standard errors (SE) of the regressions of model-free variance risk premium (upside and downside variance risk premium) on corresponding model-predicted values at various maturities (1, 3, and 6 months) from each of the four specifications defined in Table 2, in turn. We use the parameter estimates in Table 7 to generate model forecasts. The sample period spans Jan. 10, 1990, through Aug. 28, 2013.

	1-Factor Models				2-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
Parameters	Est.	SE	Est.	SE	Est.	SE	Est.	SE
Panel A. VRP								
1-Month								
Constant	9.20E−04	9.13E−04	−7.40E−03	6.94E−04	−8.14E−03	5.48E−04	−3.74E−03	5.44E−04
Slope	−0.00	0.00	0.79	0.04	1.00	0.03	1.00	0.03
R^2 (%)	17.93		36.84		60.63		60.62	
3-Month								
Constant	4.44E−03	7.64E−04	6.09E−03	6.24E−04	−5.47E−03	8.20E−04	−2.83E−03	7.11E−04
Slope	0.00	0.00	0.09	0.01	1.00	0.06	1.00	0.06
R^2 (%)	0.47		4.61		27.76		27.84	
6-Month								
Constant	5.57E−03	8.41E−04	1.07E−02	7.09E−04	−3.91E−03	9.82E−04	−3.74E−03	9.74E−04
Slope	0.00	0.00	−0.01	0.01	1.00	0.05	1.00	0.05
R^2 (%)	10.94		0.00		30.27		30.30	

(continued on next page)

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

TABLE 9 (continued)
Regressions of Model-Free on Model-Implied Variance Risk Premia

	1-Factor Models				2-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
Parameters	Est.	SE	Est.	SE	Est.	SE	Est.	SE
Panel B. VRP^u								
1-Month								
Constant					1.24E−03	5.22E−04	1.44E−03	5.18E−04
Slope					1.07	0.03	1.07	0.03
R^2 (%)					61.54		67.54	
3-Month								
Constant					−3.23E−04	3.47E−04	−4.11E−04	3.45E−04
Slope					0.93	0.03	0.93	0.03
R^2 (%)					55.57		59.08	
6-Month								
Constant					−1.24E−03	3.12E−04	−1.44E−03	3.11E−04
Slope					1.07	0.05	1.07	0.05
R^2 (%)					37.42		39.75	
Panel C. VRP^d								
1-Month								
Constant					8.54E−03	1.44E−03	8.54E−03	1.05E−03
Slope					1.46	0.15	1.46	0.02
R^2 (%)					12.05		12.05	
3-Month								
Constant					1.30E−02	1.84E−03	1.30E−02	1.84E−03
Slope					1.46	0.18	1.46	0.18
R^2 (%)					12.17		12.17	
6-Month								
Constant					1.32E−02	2.10E−03	1.32E−02	2.10E−03
Slope					0.54	0.04	0.54	0.04
R^2 (%)					14.27		14.27	

Ultimately, the difference in upside and downside variance dynamics effectively induces the time variation in the conditional asymmetry and enhances the performance of our pricing model.

VII. Conclusion

This study proposes a new and flexible option pricing model that can accommodate distinct upside and downside semivariance dynamics in the underlying asset price process. Our approach takes advantage of the recent developments in high-frequency finance, allowing us to disentangle the upside from the downside quadratic variation. This decomposition offers an alternative and effective channel for modeling the asymmetry embedded in the distribution of stock prices. An important feature of our model is that the dynamics of upside and downside variances are governed by their nonparametric empirical proxies. Given that these proxies are constructed in discrete-time, our model is specified within the affine discrete-time family.

From a theoretical standpoint, the affine structure of the model enables us to derive closed-form valuation formulas nesting several option pricing specifications. This feature facilitates the estimation procedure, allows for a direct comparison among nested models, and avoids the need to resort to simulation techniques. We employ a general pricing kernel to characterize the mapping between physical and risk-neutral motions. In addition, we provide a formal description of the key factors driving the various market risk premia. Empirically, the new

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

specification performs well. It dominates standard benchmark models, in terms of fitting accuracy and likelihood, when fitted to S&P 500 index options, realized upside and downside variances, and returns. Moreover, the model implies a time-varying skewness in return, a feature that allows for delivering realistic market compensations for upside and downside variance risks across various maturities.

Appendix. Moment-Generating Function and Risk-Neutral Dynamics

1. Moment-Generating Function

We show (see the Supplementary Material) that the model is affine, which implies that

$$\mathrm{E}_t\left[\exp\left(vR_{t+1}+v_u h_{u,t+1}+v_d h_{d,t+1}\right)\right]\equiv\exp\left(A_u(v,v_u)h_{u,t}+A_d(v,v_d)h_{d,t}+B(v,v_u,v_d)\right),$$

where

$$\begin{aligned}
A_u(v,v) &= \psi_u(-v,v)+u(\lambda_u-\xi_u)+v\beta_u,\\
A_d(v,v) &= \psi_d(v,v)+u(\lambda_d-\xi_d)+v\beta_d,\\
B(v,v_u,v_d) &= u\bar{r}+\zeta_u(-v,v_u)+\zeta_d(v,v_d)+v_u(\omega_u+\varpi_u-\omega_u\beta_u)\\
&\quad+v_d(\omega_d+\varpi_d-\omega_d\beta_d)-\omega_u\psi_u(-v,v_u)-\omega_d\psi_d(v,v_d),
\end{aligned}$$

with

$$\begin{aligned}
\zeta_j(u,v) &= -\frac{1}{2}\log\left(1-2v\alpha_j+\sqrt{2\omega_j}u\left(1-2v\alpha_j\left(1-\rho_j^2\right)\right)\right),\\
\psi_j(u,v) &= \frac{\frac{1}{2}\left(1-2v\alpha_j\left(1-\rho_j^2\right)\right)u^2+\gamma_j v\alpha_j\left(\gamma_j+2\left(\gamma_j\sqrt{\frac{\omega_j}{2}}-\rho_j\right)u\right)}{1-2v\alpha_j+\sqrt{2\omega_j}u\left(1-2v\alpha_j\left(1-\rho_j^2\right)\right)}.
\end{aligned}$$

2. Mapping between Risk-Neutral Parameters, Pricing Kernel Parameters, and Historical Parameters

We have

$$\rho_j^Q = \frac{\rho_j}{\sqrt{1-2(1-\rho^2)\nu_2^{(2j)}}\sqrt{1-2(1-\rho^2)\nu_2^{(1j)}}}.$$

Define

$$\begin{aligned}
\kappa_1^j &\equiv \frac{1-2\nu_2^{(1j)}-2\nu_2^{(2j)}+4\left(1-\rho_j^2\right)\nu_2^{(1j)}\nu_2^{(2j)}}{1-2\left(1-\rho_j^2\right)\nu_2^{(2j)}},\\
\kappa_2^j &\equiv \frac{1-2\nu_2^{(1j)}-2\nu_2^{(2j)}+4\left(1-\rho_j^2\right)\nu_2^{(1j)}\nu_2^{(2j)}}{1-2\left(1-\rho_j^2\right)\nu_2^{(1j)}}.
\end{aligned}$$

We show that

$$\omega_j^Q = \frac{\omega_j}{\left(\kappa_1^j\right)^2},\quad \sigma_j^Q = \frac{\sigma_j}{\kappa_2^j},\quad \alpha_j^Q = \frac{\varsigma_j\alpha_j}{\kappa_2^j}.$$

By denoting,

$$\chi_j \equiv \left(\frac{\omega_j^Q}{\omega_j}\right)^{1/4}\left(\frac{\sigma_j^Q}{\sigma_j}\right)^{1/2},\quad \delta_j^2 \equiv \frac{\left(1-\rho_j^2\right)+\sqrt{\left(1-\rho_j^2\right)^2+4\rho_j^2\chi_j^2}}{2},$$

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

we can express the pricing kernel parameters $\nu_2^{(1j)}$ and $\nu_2^{(2j)}$ as a function of physical and risk-neutral parameters. That is,

$$\begin{aligned}\nu_2^{(1j)} &= \frac{1}{2\left(1-\rho_j^2\right)}\left(1-\frac{\sigma_j^{\mathcal{Q}}}{\sigma_j\delta_j^2}\right)+\frac{1}{2}\left(\frac{\sigma_j^{\mathcal{Q}}}{\sigma_j\delta_j^2}-\sqrt{\frac{\omega_j}{\omega_j^{\mathcal{Q}}}}\right),\\ \nu_2^{(2j)} &= \frac{1}{2\left(1-\rho_j^2\right)}\left(1-\frac{\sigma_j}{\sigma_j^{\mathcal{Q}}}\delta_j^2\right).\end{aligned}$$

In our implementation, we choose to parameterize the model in terms of $\omega_j^{\mathcal{Q}}$ and $\sigma_j^{\mathcal{Q}}$, directly.

The prices of Gaussian innovations, $\nu_{1t}^{(1j)}$ and $\nu_{1t}^{(2j)}$, are time varying and related to both physical and risk-neutral parameters according to

$$\begin{aligned}\nu_{1t}^{(1j)} &= \pi_{1j}\sqrt{h_{j,t}-\omega_j}+\pi_{1j}^{\mathcal{Q}}\sqrt{h_{jt}^{\mathcal{Q}}-\omega_j^{\mathcal{Q}}},\\ \nu_{1t}^{(2j)} &= \pi_{2j}\sqrt{h_{j,t}-\omega_j}+\pi_{2j}^{\mathcal{Q}}\sqrt{h_{jt}^{\mathcal{Q}}-\omega_j^{\mathcal{Q}}},\end{aligned}$$

with

$$\begin{aligned}\pi_{1j} &= \frac{1}{1-\left(\rho_j^{\mathcal{Q}}\right)^2}\left(\frac{1}{\sqrt{2\omega_j^{\mathcal{Q}}}}-\gamma_j\frac{\rho_j}{\delta_j^2}\right),\\ \pi_{1j}^{\mathcal{Q}} &= \frac{1}{1-\left(\rho_j^{\mathcal{Q}}\right)^2}\left(\frac{1}{\sqrt{2\omega_j^{\mathcal{Q}}}}\left(\frac{\omega_j}{\omega_j^{\mathcal{Q}}}\right)^{1/4}-\gamma_j^{\mathcal{Q}}\frac{\rho_j}{\delta_j^2}\left(\frac{\sigma_j^{\mathcal{Q}}}{\sigma_j}\right)^{1/2}\right),\\ \pi_{2j} &= \frac{1}{1-\left(\rho_j^{\mathcal{Q}}\right)^2}\left(\gamma_j\frac{\sigma_j}{\sigma_j^{\mathcal{Q}}}-\frac{\sigma_j\delta_j^2\rho_j}{\left(1-\rho_j^2\right)\delta_j^2\sigma_j\sqrt{2\omega_j}+\rho_j^2\sigma_j^{\mathcal{Q}}\sqrt{2\omega_j^{\mathcal{Q}}}}\right),\\ \pi_{2j}^{\mathcal{Q}} &= \frac{1}{1-\left(\rho_j^{\mathcal{Q}}\right)^2}\left(\gamma_j^{\mathcal{Q}}\left(\frac{\sigma_j}{\sigma_j^{\mathcal{Q}}}\right)^{1/2}-\frac{\sigma_j\delta_j^2\rho_j}{\left(1-\rho_j^2\right)\delta_j^2\sigma_j\sqrt{2\omega_j}+\rho_j^2\sigma_j^{\mathcal{Q}}\sqrt{2\omega_j^{\mathcal{Q}}}}\left(\frac{\omega_j}{\omega_j^{\mathcal{Q}}}\right)^{1/4}\right).\end{aligned}$$

Given that $\pi_{1j}^{\mathcal{Q}}$ and $\pi_{2j}^{\mathcal{Q}}$ are linear in $\gamma_j^{\mathcal{Q}}$, we can parameterize the model in terms of either $\pi_{1j}^{\mathcal{Q}}$ or $\gamma_j^{\mathcal{Q}}$. We opt for the latter. Finally, the mapping between risk-neutral and physical innovation takes the form

$$\begin{aligned}\varepsilon_{j,t+1}^{(*1)} &= \sqrt{\kappa_1^j}\left(\varepsilon_{j,t+1}^{(1)}-\frac{\nu_{1t}^{(1j)}\left(1-2\left(1-\rho_j^2\right)\nu_2^{(2j)}\right)+\rho_j\nu_{1t}^{(2j)}}{\kappa_1^j\left(1-2\left(1-\rho_j^2\right)\nu_2^{(2j)}\right)}\right),\\ \varepsilon_{j,t+1}^{(*2)} &= \sqrt{\kappa_2^j}\left(\varepsilon_{j,t+1}^{(2)}-\frac{\nu_{1t}^{(2j)}\left(1-2\left(1-\rho_j^2\right)\nu_2^{(1j)}\right)+\rho_j\nu_{1t}^{(1j)}}{\kappa_2^j\left(1-2\left(1-\rho_j^2\right)\nu_2^{(1j)}\right)}\right).\end{aligned}$$

Supplementary Material

Supplementary Material for this article is available at https://doi.org/10.1017/S0022109018000777.

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press

References

Andersen, T. G.; T. Bollerslev; F. X. Diebold; and H. Ebens. "The Distribution of Realized Stock Return Volatility." *Journal of Financial Economics*, 61 (2001), 43–76.

Andersen, T. G.; T. Bollerslev; F. X. Diebold; and P. Labys. "The Distribution of Realized Exchange Rate Volatility." *Journal of the American Statistical Association*, 96 (2001), 42–55.

Andersen, T. G.; T. Bollerslev; F. X. Diebold; and P. Labys. "Modeling and Forecasting Realized Volatility." *Econometrica*, 71 (2003), 579–625.

Andersen, T. G.; O. Bondarenko; and M. T. Gonzalez-Perez. "Exploring Return Dynamics via Corridor Implied Volatility." *Review of Financial Studies*, 28 (2015), 2902–2945.

Andersen, T. G.; N. Fusari; and V. Todorov. "Parametric Inference and Dynamic State Recovery from Option Panels." *Econometrica*, 83 (2015), 1081–1145.

Bakshi, G.; N. Kapadia; and D. Madan. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies*, 16 (2003), 101–143.

Bakshi, G., and D. Madan. "Spanning and Derivative Security Valuation." *Journal of Financial Economics*, 55 (2000), 205–238.

Bandi, F., and R. Renò. "Price and Volatility Co-Jumps." *Journal of Financial Economics*, 119 (2016), 107–146.

Barndorff-Nielsen, O. E.; S. Kinnebrock; and N. Shephard. "Measuring Downside Risk: Realised Semivariance." In *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, T. Bollerslev, J. Russell, and M. Watson, eds. Oxford, UK: Oxford University Press (2010).

Barndorff-Nielsen, O. E., and N. Shephard. "Power and Bipower Variation with Stochastic Volatility and Jumps." *Journal of Financial Econometrics*, 2 (2004), 1–37.

Bates, D. "Post-'87 Crash Fears in the S&P 500 Futures Option Market." *Journal of Econometrics*, 96 (2000), 181–238.

Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), 637–654.

Bollerslev, T.; S. Li; and B. Zhao. "Good Volatility, Bad Volatility, and the Cross-Section of Stock Returns." Working Paper, Duke University (2017).

Bollerslev, T.; G. Tauchen; and H. Zhou. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies*, 22 (2009), 4463–4492.

Carr, P., and D. Madan. "Optimal Positioning in Derivative Securities." *Quantitative Finance*, 1 (2001), 19–37.

Christoffersen, P.; R. Elkamhi; B. Feunou; and K. Jacobs. "Option Valuation with Conditional Heteroskedasticity and Nonnormality." *Review of Financial Studies*, 23 (2010), 2139–2183.

Christoffersen, P.; B. Feunou; K. Jacobs; and N. Meddahi. "The Economic Value of Realized Volatility: Using High-Frequency Returns for Option Valuation." *Journal of Financial and Quantitative Analysis*, 49 (2014), 663–697.

Christoffersen, P.; B. Feunou; and Y. Jeon. "Option Valuation with Observable Volatility and Jump Dynamics." *Journal of Banking & Finance*, 61 (2015), S101–S120.

Christoffersen, P.; S. Heston; and K. Jacobs. "Capturing Option Anomalies with a Variance-Dependent Pricing Kernel." *Review of Financial Studies*, 26 (2013), 1963–2006.

Corsi, F.; N. Fusari; and D. L. Vecchia. "Realizing Smiles: Options Pricing with Realized Volatility." *Journal of Financial Economics*, 107 (2013), 284–304.

Darolles, S.; C. Gourieroux; and J. Jasiak. "Structural Laplace Transform and Compound Autoregressive Models." *Journal of Time Series Analysis*, 27 (2006), 477–503.

Duffie, D.; J. Pan; and K. Singleton. "Transform Analysis and Option Pricing for Affine Jump-Diffusions." *Econometrica*, 68 (2000), 1343–1377.

Feunou, B.; M. R. Jahan-Parvar; and C. Okou. "Downside Variance Risk Premium." *Journal of Financial Econometrics*, 16 (2018), 341–383.

Feunou, B.; M. R. Jahan-Parvar; and R. Tédongap. "Modeling Market Downside Volatility." *Review of Finance*, 17 (2013), 443–481.

Feunou, B.; M. R. Jahan-Parvar; and R. Tédongap. "Which Parametric Model for Conditional Skewness?" *European Journal of Finance*, 22 (2016), 1237–1271.

Guo, H.; K. Wang; and H. Zhou. "Good Jumps, Bad Jumps, and Conditional Equity Premium." Working Paper, University of Cincinnati (2015).

Hansen, P. R., and A. Lunde. "Realized Variance and Market Microstructure Noise." *Journal of Business and Economic Statistics*, 24 (2006), 127–161.

Heston, S. L., and S. Nandi. "A Closed-Form GARCH Option Valuation Model." *Review of Financial Studies*, 13 (2000), 585–625.

Huang, J.-Z., and L. Wu. "Specification Analysis of Option Pricing Models Based on Time-Changed Lévy Processes." *Journal of Finance*, 59 (2004), 1405–1439.

Mincer, J., and V. Zarnowitz. "The Evaluation of Economic Forecasts." In *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, J. Mincer, ed. Cambridge, MA: National Bureau of Economic Research (1969).

Patton, A. J., and K. Sheppard. "Good Volatility, Bad Volatility: Signed Jumps and the Persistence of Volatility." *Review of Economics and Statistics*, 97 (2015), 683–697.

Renault, E. "Econometric Models of Option Pricing Errors." In *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*, D. M. Kreps and K. F. Wallis, eds. Cambridge, UK: Cambridge University Press (1997).

Stentoft, L. "Option Pricing Using Realized Volatility." CREATES Research Paper (2008).

https://doi.org/10.1017/S0022109018000777 Published online by Cambridge University Press