




BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel 2017-52

Good Volatility, Bad Volatility and Option Pricing



by Bruno Feunou and Cédric Okou

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Bank of Canada Staff Working Paper 2017-52

December 2017

Good Volatility, Bad Volatility and Option Pricing

by

Bruno Feunou¹ and Cédric Okou²

¹ Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
feun@bankofcanada.ca

² ESG
UQAM
okou.cedric@uqam.ca

Acknowledgements

We thank Peter Christoffersen, Diego Amaya, Christian Dorion, Yoontae Jeon, and seminar participants at HEC Montréal for fruitful discussions. We gratefully acknowledge financial support from the Bank of Canada, the Université du Québec à Montréal (UQAM) research funds, and the IFSID. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

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 semi-variance dynamics driven by their model-free proxies. The new model outperforms common benchmarks, especially the alternative that splits the quadratic variation into diffusive and jump components.

Bank topics: Asset pricing; Econometric and statistical methods

JEL code: G12

Résumé

Grâce aux avancées dans le domaine de l'analyse des écarts, il est possible de diviser la variation quadratique totale d'un processus de diffusion à sauts en composantes à la hausse et à la baisse. Selon de récentes études, cette division améliore les prévisions de volatilité et fait ressortir que la différence entre les variances à la hausse et à la baisse constitue un facteur d'asymétrie dans les distributions des cours des actions. Pour estimer les gains économiques que procure cette division, nous concevons un nouveau modèle flexible d'évaluation des options dans lequel le prix de l'actif sous-jacent présente des dynamiques de variance à la hausse et à la baisse distinctes, déterminées par leurs équivalents non paramétriques. Le nouveau modèle surpasse les modèles de référence communs, surtout l'approche qui scinde la variation quadratique en composantes de diffusion et de sauts.

Sujets : Évaluation des actifs ; Méthodes économétriques et statistiques

Code JEL : G12

Non-Technical Summary

One of the most important recent developments in the financial econometrics literature is the use of intraday observations to precisely evaluate the variability of the price of any financial asset on a given day. That estimate is commonly known as the “realized variance.”

Using the realized variance, we can evaluate, for instance, the effect of a policy announcement or macroeconomic news on the uncertainty of a given asset’s price. Investors have asymmetric views of increases and decreases in asset prices. On the one hand, they like positive movements and would be willing to be exposed to them. On the other hand, they dislike negative movements and would ask to be paid a premium for taking on such an exposure. Different types of news and announcements affect these two movements differently.

Hence, it is of interest to academics and policy-makers to measure how much variability is attributable to positive movements in prices versus negative movements. Looking at intraday trading activities, researchers have provided a decomposition of realized variance as the sum of good variance (positive returns variability) and bad variance (negative returns variability).

This paper evaluates the economic significance of that decomposition by evaluating the mispricing of S&P 500 derivatives under two scenarios: ignoring or using the decomposition of the realized variance. We find that the split is very informative for option pricing as it reduces pricing errors significantly. This new model can be used to better measure downside risks embedded in asset prices. In addition, it can be used to understand compensation for downside risk. To be more specific, we disentangle the downside risk embedded in option prices in terms of an investor aversion component and a historical downside risk estimate.

1 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 et al. (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 paper develops an option valuation model in which the underlying asset price features specific upside (good) and downside (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 papers by Andersen et al. (2001a), Andersen et al. (2001b), and Andersen et al. (2003), to cite a few. Drawing on similar “infill asymptotics” arguments, Barndorff-Nielsen et al. (2010) show, in a model-free way, how to dissect the realized variance into upside and downside semi-variances 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 tradeoff (Guo et al., 2015), or explain the cross-section of stock returns (Bollerslev et al., 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 paper is related to a growing body of finance literature that aims at building option pricing models with empirically grounded properties. This strand of the option pricing literature is distinct from standard stochastic volatility option pricing models,¹ as it uses observed (realized) quantities to update factors: the factors are no longer latent. This modeling approach is not only practically appealing since we do not need a sophisticated filtering technique, but also bridges the gap between developments in the high-frequency econometrics and the option pricing literatures. A few studies

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

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 papers by Stentoft (2008), Corsi et al. (2013), and Christoffersen et al. (2014).

The aforementioned papers 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 et al. (2015) exploits an alternative dissection of the total quadratic variation into a diffusive volatility and a squared jump variation. Thus, the modeling approach in Christoffersen et al. (2015) can be viewed as the discrete time analog of a classical continuous time affine jump-diffusion specification such as Bates (2000), where the two factors capture the change in “normal” variation (diffusion) and the intensity of “extreme” moves (jumps), and are connected to their realized counterparts.

While modeling approaches based on these two decompositions (up-down variation versus 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, as we find that our preferred GSARV specification improves the option price fitting by a sizeable 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 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 paper is organized as follows. In Section 2, we present the theoretical and empirical arguments underpinning the construction and the use of good/bad (upside/downside) variation measures. Section 3 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 4, 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 5 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 6. In addition, we study the performance of our pricing kernel and document the determinants of the variance risk premium components. Section 7 concludes.

2 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 semi-variation 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.

2.1 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

$$ds_t = \mu_t dt + \sigma_t dW_t + \Delta s_t, \quad (1)$$

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_{t+1}^U), and the downside realized variance (denoted by RV_{t+1}^D) are defined as

$$\begin{aligned} RV_{t+1}^U &= \sum_{j=1}^{n_{t+1}} R_{t_j}^2 \mathbb{I}_{[R_{t_j} > 0]}, \\ RV_{t+1}^D &= \sum_{j=1}^{n_{t+1}} R_{t_j}^2 \mathbb{I}_{[R_{t_j} \leq 0]}, \end{aligned}$$

where $R_{t_j} \equiv s_{t_j} - s_{t_{j-1}}$ and $t < t_1 < \dots < t_{n_{t+1}} = t + 1$ are the time at which (trade or quote) prices are available. The total realized variance is $RV_{t+1} = RV_{t+1}^U + RV_{t+1}^D$.

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

$$\begin{aligned} RV_{t+1}^U &\xrightarrow{p} \frac{1}{2} \int_t^{t+1} \sigma_v^2 dv + \sum_{t \leq v \leq t+1} (\Delta s_v)^2 \mathbb{I}_{[\Delta s_v > 0]}, \\ RV_{t+1}^D &\xrightarrow{p} \frac{1}{2} \int_t^{t+1} \sigma_v^2 dv + \sum_{t \leq v \leq t+1} (\Delta s_v)^2 \mathbb{I}_{[\Delta s_v \leq 0]}. \end{aligned} \quad (2)$$

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 et al. (forthcoming) 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:

$$RSV_{t+1} = RV_{t+1}^U - RV_{t+1}^D.$$

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 can be found in Feunou et al. (2016). To provide more intuition on the behavior of RSV_{t+1} , we combine the previous asymptotic results to get

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

Assuming that jump sizes are *i.i.d.* and uncorrelated with the jump occurrence, the expectation of the realized skewness – a measure of the conditional skewness – is

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

Thus, $\mathbb{E}_t^{\mathbb{P}} [RSV_{t+1}]$ 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 their sizes, consistent with the volatility jump risk study of Bandi and Ren (2016). Bigger weights are assigned to larger jump sizes in the computation of RSV_{t+1} .

2.2 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:

$$\begin{aligned} RJV_{t+1} \equiv RV_{t+1} - BPV_{t+1} &\xrightarrow{p} \sum_{t \leq v \leq t+1} (\Delta s_v)^2, \\ BPV_{t+1} &\xrightarrow{p} \int_t^{t+1} \sigma_v^2 dv, \end{aligned} \quad (3)$$

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. While RJV_{t+1} estimates $\sum_{t \leq v \leq t+1} (\Delta s_v)^2$, 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 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 6.2.1, we also compare the empirical option pricing implications of these two decompositions.

2.3 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 five-minute interval to build a grid of five-minute equity index log-returns. Following Andersen et al. (2003, 2001a) 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

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

where $R_{j,t}^2$ is the j^{th} five-minute squared log-return, and n_t is the number of (five-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:

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

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

$$RV_t \equiv RV_t^U + RV_t^D. \quad (4)$$

2.4 Empirical dynamics of downside and upside volatilities

Figure 1 plots the daily time series of the return on S&P 500 index (Graph A), along with the square root of the daily realized variance (Graph C), the square root of its upside (Graph B) and the square root of its downside (Graph D) components from January 10, 1990 to August 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 of the realized variance and its components display a good level of synchronicity, as 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

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

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. Sizeable discrepancies between these two time series can arise from risk-neutral (rather than historical) expectations extracted from option data. Using a similar logic as Bollerslev et al. (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.

3 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 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.

3.1 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

$$E_t \left[RV_{t+1}^j \right] = h_{j,t}, \quad (5)$$

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

$$h_t \equiv E_t [RV_{t+1}] = h_{u,t} + h_{d,t}. \quad (6)$$

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 : (1) corresponds to the expected realized variance as in Equation 6, and (2) matches the conditional variance of $t + 1$ log-returns as

$$Var_t [R_{t+1}] = h_t. \quad (7)$$

As discussed in Section (2.1), 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.

3.2 Asset return process

Consider the following representation of the daily log-returns

$$R_{t+1} = E_t [R_{t+1}] + z_{t+1}, \quad (8)$$

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

$$E_t [\exp (R_{t+1})] = \exp (r_f + \lambda_u h_{u,t} + \lambda_d h_{d,t}), \quad (9)$$

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} .

3.2.1 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 innovations:

$$z_{t+1} = z_{u,t+1} - z_{d,t+1}. \quad (10)$$

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 $\tilde{z}_{d,t+1}$:

$$z_{j,t+1} = \tilde{z}_{j,t+1} - E_t[\tilde{z}_{j,t+1}], \text{ for } j = \{u, d\}. \quad (11)$$

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

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

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.

3.2.2 Innovation conditional distribution

Since the objective is to have closed-form option price, we opt for the noncentral chi-squared distribution, which has been widely used in the discrete time option pricing literature (See Heston and Nandi (2000)). Shocks to returns are

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

with

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

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

$$z_{j,t+1} \xrightarrow[\omega_j \rightarrow 0]{\text{Distribution}} N(0, h_{j,t}),$$

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

The conditional skewness of log-returns is

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

where $h_t \equiv h_{u,t} + h_{d,t}$ is the total conditional variance of returns. Hence, two sources of conditional asymmetry emerge: (1) the difference between upside and downside variance $h_{u,t} - h_{d,t}$, and (2) the discrepancy between non-normalities 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.

3.3 Expected returns

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

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

where

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

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

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).

3.4 Incorporating realized upside and downside variation

Each day, realized upside and downside variations provide new information about the conditional semi-variances, $h_{u,t}$ and $h_{d,t}$. However, RV_{t+1}^U and RV_{t+1}^D are measured with error, and therefore, we specify the following measurement equation:

$$RV_{t+1}^j = h_{j,t} + z_{j,t+1}^{(2)}, \quad (16)$$

where $z_{j,t+1}^{(2)} = \tilde{z}_{j,t+1}^{(2)} - E_t [\tilde{z}_{j,t+1}^{(2)}]$ is a zero-mean innovation. By design, we have $E_t [RV_{t+1}^j] = h_{j,t}$.

To obtain closed-form option price formulas, we choose the noncentral chi-squared distribution

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

where

$$\varepsilon_{j,t+1}^{(2)} \stackrel{iid}{\sim} N(0, 1) \text{ for } j = \{u, d\},$$

and

$$\begin{aligned} E [\varepsilon_{u,t+1}^{(2)} \varepsilon_{d,t+1}^{(2)}] &= 0, \\ E [\varepsilon_{j,t+1}^{(1)} \varepsilon_{j,t+1}^{(2)}] &= \rho_j. \end{aligned}$$

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

3.5 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.

$$h_{j,t+1} - \omega_j = \tilde{\omega}_j + \tilde{\beta}_j (h_{j,t} - \omega_j) + \tilde{\alpha}_j (RV_{t+1}^j - \omega_j). \quad (18)$$

This is very similar to GARCH models, with the exception that less noisy innovations $RV_{t+1}^j - \omega_j$ are used to update conditional variances. The specification simply implies that $h_{u,t+1}$ and $h_{d,t+1}$

are both univariate $AR(1)$ processes. Thus, RV_{t+1}^U and RV_{t+1}^D belong to the class of univariate $ARMA(1,1)$ processes. Using equation (16), equation (18) can be rewritten as

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

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

3.6 Conditional second moments and moment generating function

From the model above, it is straightforward to derive the following one-day-ahead conditional second moments:

$$\begin{aligned} Var_t [R_{t+1}] &= h_{u,t} + h_{d,t}, \\ Var_t [RV_{t+1}^j] &= 2\sigma_j^2 (1 + 2\gamma_j^2 (h_{j,t} - \omega_j)), \\ Cov_t (R_{t+1}, RV_{t+1}^U) &= \sqrt{2\omega_u} \sigma_u \rho_u^2 + 2(h_{u,t} - \omega_u) \gamma_u \sigma_u \rho_u, \\ Cov_t (R_{t+1}, RV_{t+1}^D) &= -(\sqrt{2\omega_d} \sigma_d \rho_d^2 + 2(h_{d,t} - \omega_d) \gamma_d \sigma_d \rho_d), \\ Cov_t (R_{t+1}, h_{t+1}) &= \tilde{\alpha}_u Cov_t (R_{t+1}, RV_{t+1}^U) + \tilde{\alpha}_d Cov_t (R_{t+1}, RV_{t+1}^D). \end{aligned} \quad (20)$$

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 one-step-ahead conditional characteristic function of returns and factors is exponentially linear in those factors (see Darolles et al., 2006). The internal Appendix reports the expressions for the joint conditional moment-

generating function of returns, and expected upside and downside variations.

4 Physical Estimation

Up to this point in the paper, 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.

4.1 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

$$\begin{aligned} f_t(R_{t+1}, RV_{t+1}^U, RV_{t+1}^D) &= f_{r,t}(R_{t+1}) f_{U,t}(RV_{t+1}^U) f_{D,t}(RV_{t+1}^D) \\ &\quad \times c_t(F_{r,t}(R_{t+1}), F_{U,t}(RV_{t+1}^U), F_{D,t}(RV_{t+1}^D)), \end{aligned}$$

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

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

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal, and the $Var_t \left[RV_{t+1}^j \right]$ s are given in equation (20). Because we want to highlight the importance of $h_{u,t} - h_{d,t}$ in generating non-normalities in the conditional distribution of returns, we need the exact marginal density function. Notice that the exact marginal density of returns, $f_{r,t}(R_{t+1})$, is a convolution of two noncentral chi-squared densities, which does not have known closed-form expression. Nonetheless, the conditional characteristic function of R_{t+1} is available in closed-form as

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

where $a_j(\nu)$ and $b_j(\nu)$ functions are

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

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

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

where CM_t is the conditional correlation matrix of $(R_{t+1}, RV_{t+1}^U, RV_{t+1}^D)$ given by

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

with

$$\rho_t^j = \frac{Cov_t(R_{t+1}, RV_{t+1}^j)}{\sqrt{Var_t[R_{t+1}] Var_t[RV_{t+1}^j]}}, \quad \text{for } j = \{u, d\},$$

and where $Cov_t(R_{t+1}, RV_{t+1}^j)$, $Var_t[R_{t+1}]$, and $Var_t[RV_{t+1}^j]$ are defined in Equation (20).

Finally, the log-likelihood is calculated as

$$\ln L^P = \sum_{t=1}^{T-1} \ln(f_t(R_{t+1}, RV_{t+1}^U, RV_{t+1}^D)). \quad (22)$$

4.2 Nested specifications

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

4.2.1 The constrained generalized skew affine realized variance (CGSARV) model

One of the key objectives of this paper 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 above, 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.

4.2.2 The affine realized variance (ARV) 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 *i.i.d.* $N(0, 1)$.

Beginning with h_0 , we have

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

where

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

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

4.2.3 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.

4.3 Parameter estimates and model properties

Maximum likelihood estimation results from historical data are given in Table 2. As outlined above, 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 two-factor models, we estimate the downside market price 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 returns 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 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 RV , and the two-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 two-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 as 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 two-factor models are optimized on returns, RV_u , and RV_d , they yield slightly lower realized variance-implied log-likelihood values than the one-factor ARV model, which is fitted to returns, and RV . Nonetheless, the GSARV model delivers the highest realized variance-implied log-likelihood among two-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 one-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 about 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 about 1.12 for the other models. This confirms that two-factor models deliver a good fit for upside and

downside variances. Additional empirical properties of the various specifications are investigated in the external Appendix B.

5 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 study the performance of the various models relying on the implied volatility root-mean-squared-error.

5.1 Exploring option data

We use European-style options written on the S&P 500 index. The observations span the period January 10, 1996 through August 28, 2013.³ In line with the extant literature, we only include out-of-the-money (OTM) options with maturity 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 six 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 S&P 500 index option, we sort the data by moneyness, maturity, and market volatility index (VIX) level. Panel A of Table 4 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 Black-Scholes delta computed as

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

where $\Phi(*)$ stands for the normal cumulative distribution function (CDF), X is the strike price, r_f is the non annualized daily risk-free rate, M is the time-to-maturity expressed in days, and IV^{Mkt}

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

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.

5.2 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 model performance comparison tool 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 formula (BS). Thus, the model-based implied volatility can be formally extracted according to

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

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

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

Accordingly, the implied volatility error is computed as

$$e_j = IV_j^{Mkt} - IV_j^{Mod}.$$

It follows that the IVRMSE is obtained as

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

where N denotes the option sample size.

Finally, risk-neutral parameters are estimated by maximizing the Gaussian implied volatility error likelihood⁴

$$\ln L^O \propto -\frac{1}{2} \sum_{j=1}^N \{\ln (IVRMSE^2) + e_j^2/IVRMSE^2\}. \quad (23)$$

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 sizeable 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 IVs. As done in the historical estimation, we estimate $E^Q[h_j]$ and then back out the ϖ_j estimates from the theoretical unconditional risk-neutral variance formula.

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

5.3 Dissecting model fit

We now scrutinize the overall performance results reported at 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 put options seems more challenging. Interestingly, the ability of the various specifications to match the observed option-implied

⁴Below, 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.

volatility appears consistent across the term structure of the options, as 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 moneyness, maturity, and VIX level dimensions.

6 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 compensations of risks. Thus, we extend our investigation by jointly fitting the various models to historical returns, upside and downside realized variances, and options data.

The optimization is performed over a convolution of the quasi log-likelihood of returns and semi-realized variances denoted as $\ln L^P$ in Equation (22), and an option-based log-likelihood component $\ln L^O$ as per Equation (23). Accordingly, our joint maximization program writes

$$\max \ln L^P + \ln L^O. \tag{24}$$

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.

6.1 Risk-neutralization and implied risk premia

6.1.1 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

$$M_{t+1} = M_{t+1}^{(u)} M_{t+1}^{(d)}, \quad (25)$$

where

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

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 et al. (2010) where only the $\varepsilon_{j,t+1}^{(i)}$ disturbances are priced. In our specification, the ν_2 s are key elements that drive the wedge between the one-day-ahead conditional risk-neutral and physical variance (termed as the variance spread). The magnitude of this variance spread may be non-trivial, as pointed out by Christoffersen et al. (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 external Appendix C also show that the risk-neutral motion of the underlying asset return is

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

where

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

The risk-neutral innovations in returns are noncentral chi-squared, distributed as

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

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

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

$$VSP_t^j \equiv h_{j,t}^Q - h_{j,t} = \vartheta_j + (\varsigma_j - 1) h_{j,t}, \quad (26)$$

with

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

and

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

We provide some analysis of this wedge in Section 6.1.2.

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

$$RV_{t+1}^j = h_{j,t} + VRP_t^j + z_{j,t+1}^{(*2)},$$

where

$$\begin{aligned} z_{j,t+1}^{(*2)} &= \tilde{z}_{j,t+1}^{(*2)} - E_t^Q \left[\tilde{z}_{j,t+1}^{(*2)} \right] \\ \tilde{z}_{j,t+1}^{(*2)} &= \sigma_j^Q \left(\varepsilon_{j,t+1}^{(*2)} - \gamma_j^Q \sqrt{h_{j,t}^Q - \omega_j^Q} \right)^2, \end{aligned}$$

$$\varepsilon_{j,t+1}^{(*2)} \stackrel{iid^Q}{\sim} N(0, 1), E^Q \left[\varepsilon_{u,t+1}^{(*2)} \varepsilon_{d,t+1}^{(*2)} \right] = 0, \text{ and } E_t^Q \left[\varepsilon_{j,t+1}^{(*1)} \varepsilon_{j,t+1}^{(*2)} \right] \equiv \rho_j^Q.$$

Moreover, there is a difference between the risk-neutral expectation of the realized variances $E_t^Q [RV_{t+1}^U], E_t^Q [RV_{t+1}^D]$ and their physical counterparts $h_{u,t}, h_{d,t}$. The difference is the variance risk premium, and is given by

$$VRP_t^j \equiv E_t^Q [RV_{t+1}^j] - E_t^P [RV_{t+1}^j] = \theta_j + \varrho_j h_{jt} + \sigma_j^Q \left(\gamma_j^Q \right)^2 VSP_t^j, \quad (27)$$

with

$$\begin{aligned} \varrho_j &= \sigma_j^Q \left(\gamma_j^Q \right)^2 - \sigma_j \gamma_j^2, \\ \theta_j &= \sigma_j^Q - \sigma_j + \sigma_j \gamma_j^2 \omega_j - \sigma_j^Q \left(\gamma_j^Q \right)^2 \omega_j^Q. \end{aligned}$$

It is important to point out the subtle difference between the variances spread and premium. The two notions are clearly related, as more spread induces a premium. But there can be a premium without spread; we will provide some analysis of these differences in Section 6.1.2.

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

$$h_{j,t+1}^Q - \omega_j^Q = \varpi_j^Q + \beta_j \left(h_{j,t}^Q - \omega_j^Q \right) + \alpha_j^Q \left(\varepsilon_{j,t+1}^{(*2)} - \gamma_j^Q \sqrt{h_{j,t}^Q - \omega_j^Q} \right)^2, \quad (28)$$

with $\varpi_j^Q = (1 - \beta_j) \left(\vartheta_j + \varsigma_j \omega_j - \omega_j^Q \right) + \varsigma_j \varpi_j$.

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

6.1.2 Analyzing variance spread and variance risk premium drivers

The variance spread expression in equation (26) reveals three potential factors driving the wedge between risk-neutral and physical variances: (1) the nonnormality ($\omega_j, \omega_j^Q \neq 0$), (2) the distance between asymmetries ($\|\omega_j^Q - \omega_j\|$) and (3) 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 ($\omega_j, \omega_j^Q = 0$); this is consistent with Christoffersen et al. (2013) findings. We refer the reader to the external Appendix C 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 $\omega_j^Q = \omega_j$, the risk-neutral variance becomes proportional to its historical counterpart. When $\omega_j^Q \neq \omega_j$, an increase in the gap between ω_j^Q and ω_j induces a larger variance spread.

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 (non negative). Moreover, combining these conditions with $\omega_j^Q > \omega_j$ implies that $VSP_t^j \geq 0$. 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 (27), 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 (as it increases the potential of substantial gains) but dislike bad uncertainty (as it increases the likelihood of severe losses), Feunou et al. (forthcoming) have documented that the premium on the downside variance is positive ($VRP_t^d > 0$), whereas its upside counterpart is negative ($VRP_t^u < 0$). These empirical regularities are met in our model if $\varrho_d > 0$ and $\varrho_u < 0$, thus implying economically sound

restrictions, i.e.,

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

6.2 Empirical results

6.2.1 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, while 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 two standard errors away from 0. Interestingly for GSARV and CGSARV models, the inferred λ_u values are found to be negative. Thus, the joint estimation results for GSARV and CGSARV specifications are consistent with the intuition that investors dislike downside uncertainty and demand a positive premium as a 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 towards good versus bad uncertainty is in line with the empirical regularities documented in recent works such as Feunou et al. (forthcoming), and Feunou et al. (2013), among others. Moreover, the joint estimation results for 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 that are able to generate distinct dynamics for market compensations 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 sizeable 10% improvement over 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).

6.2.2 Model fit for variance risk premia

From the works of Bakshi and Madan (2000), Carr and Madan (2001), Bakshi et al. (2003), 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 in Andersen et al. (2001a), Andersen et al. (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 model-free variance premium and its upside (downside) component on their model-implied counterparts. The results clearly show that the two-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-squared figures compared with one-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 variance risk premium components, we see that two-factor model forecasts 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. 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.

7 Conclusion

This study proposes a new and flexible option pricing model that can accommodate distinct upside and downside semi-variance 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 resorting 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 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.

References

- Andersen, Torben G., Bollerslev, Tim, Diebold, Francis X., and Ebens, Heiko (2001a), “The distribution of realized stock return volatility,” *Journal of Financial Economics*, 61, 43–76.
- Andersen, Torben G., Bollerslev, Tim, Diebold, Francis X., and Labys, Paul (2001b), “The distribution of realized exchange rate volatility,” *Journal of the American Statistical Association*, 96(453), 42–55.
- (2003), “Modeling and forecasting realized volatility,” *Econometrica*, 71(2), 579–625.
- Andersen, Torben G., Bondarenko, Oleg, and Gonzalez-Perez, Maria T. (2015), “Exploring return dynamics via corridor implied volatility,” *The Review of Financial Studies*, 28(10), 2902–2945.
- Andersen, Torben G., Fusari, Nicola, and Todorov, Viktor (2015), “Parametric inference and dynamic state recovery from option panels,” *Econometrica*, 83(3), 1081–1145.
- Bakshi, Gurdip, and Madan, Dilip (2000), “Spanning and derivative security valuation,” *Journal of Financial Economics*, 55(2), 205–238.
- Bakshi, Gurdip, Kapadia, N., and Madan, Dilip (2003), “Stock return characteristics, skew laws and the differential pricing of individual equity options,” *Review of Financial Studies*, 16(1), 101–143.
- Bandi, F.M., and Ren, R. (2016), “Price and volatility co-jumps,” *Journal of Financial Economics*, 119(1), 107 – 146.
- Barndorff-Nielsen, Ole E., and Shephard, Neil (2004), “Power and bipower variation with stochastic volatility and jumps,” *Journal of Financial Econometrics*, 2(1), 1–37.
- Barndorff-Nielsen, Ole E., Kinnebrock, Silja, and Shephard, Neil (2010) “Measuring downside risk: realised semivariance,” in *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, eds. T. Bollerslev, and J. Russell, and M. Watson, Oxford: Oxford University Press, pp. 117-136.
- Bates, David (2000), “Post-’87 crash fears in the S&P 500 futures option market,” *Journal of Econometrics*, 96(1–2), 181–238.
- Black, Fischer, and Scholes, Myron (1973), “The pricing of options and corporate liabilities,” *Journal of Political Economy*, 81(3), 637–654.
- Bollerslev, Tim, Li, Sophia, and Zhao, Bingzhi (2017), “Good volatility, bad volatility, and the cross-section of stock returns,” *Working Paper, Duke University*.
- Bollerslev, Tim, Tauchen, George, and Zhou, Hao (2009), “Expected stock returns and variance risk premia,” *Review of Financial Studies*, 22(11), 4463–4492.
- Carr, Peter, and Madan, Dilip (2001), “Optimal positioning in derivative securities,” *Quantitative Finance*, 1(1), 19–37.
- Christoffersen, Peter, Elkamhi, Redouane, Feunou, Bruno, and Jacobs, Kris (2010), “Option valuation with conditional heteroskedasticity and nonnormality,” *Review of Financial Studies*, 23(5), 2139–2183.

- Christoffersen, Peter, Feunou, Bruno, and Jeon, Yoontae (2015), “Option valuation with observable volatility and jump dynamics,” *Journal of Banking & Finance*, 61, Supplement 2, S101 – S120.
- Christoffersen, Peter, Feunou, Bruno, Jacobs, Kris, and Meddahi, Nour (2014), “The economic value of realized volatility: Using high-frequency returns for option valuation,” *Journal of Financial and Quantitative Analysis*, 49, 663–697.
- Christoffersen, Peter, Heston, Steven, and Jacobs, Kris (2013), “Capturing option anomalies with a variance-dependent pricing kernel,” *Review of Financial Studies*, 26(8), 1963–2006.
- Corsi, Fluvio, Fusari, Nicola, and Vecchia, Davide La (2013), “Realizing smiles: Options pricing with realized volatility,” *Journal of Financial Economics*, 107, 284–304.
- Darolles, Serge, Gouriéroux, Christian, and Jasiak, Joann (2006), “Structural laplace transform and compound autoregressive models,” *Journal of Time Series Analysis*, 27(4), 477–503.
- Duffie, Darrell, Pan, Jun, and Singleton, Kenneth (2000), “Transform analysis and option pricing for affine jump-diffusions,” *Econometrica*, 68, 1343–1377.
- Feunou, Bruno, Jahan-Parvar, Mohammad R., and Okou, Cédric (forthcoming), “Downside variance risk premium,” *Journal of Financial Econometrics*.
- Feunou, Bruno, Jahan-Parvar, Mohammad R., and Tédongap, Roméo (2013), “Modeling Market Downside Volatility,” *Review of Finance*, 17(1), 443–481.
- Feunou, Bruno, Jahan-Parvar, Mohammad R., and Tédongap, Roméo (2016), “Which parametric model for conditional skewness?” *The European Journal of Finance*, 22(13), 1237–1271.
- Guo, Hui, Wang, Kent, and Zhou, Hao (2015), “Good jumps, bad jumps, and conditional equity premium,” *Working Paper, University of Cincinnati*.
- Hansen, Peter R., and Lunde, Asger (2006), “Realized variance and market microstructure noise,” *Journal of Business and Economic Statistics*, 24(2), 127–161.
- Heston, Steven L., and Nandi, Saikat (2000), “A closed-form garch option valuation model,” *Review of Financial Studies*, 13(3), 585–625.
- Huang, Jing-Zhi, and Wu, Liuren (2004), “Specification analysis of option pricing models based on time-changed lévy processes,” *Journal of Finance*, 59(3), 1405–1439.
- Mincer, Jacob, and Zarnowitz, Victor (1969) *The Evaluation of Economic Forecasts*, National Bureau of Economic Research.
- Patton, Andrew J., and Sheppard, Kevin (2015), “Good volatility, bad volatility: Signed jumps and the persistence of volatility,” *Review of Economics and Statistics*, 97(3), 683–697.
- Renault, Eric. (1997) “Econometric Models of Option Pricing Errors,” in *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*, eds. D.M. Kreps and K.F. Wallis, Cambridge: Cambridge University Press, pp. 223–278.
- Stentoft, Lars (2008), “Option pricing using realized volatility,” *CREATES Research Paper*.

Appendix: Moment-generating function and risk-neutral dynamic

A1-Moment-generating function

We show (see the external Appendix) that the model is affine, which implies that

$$E_t [\exp (\nu R_{t+1} + v_u h_{u,t+1} + v_d h_{d,t+1})] \equiv \exp (A_u (\nu, v_u) h_{u,t} + A_d (\nu, v_d) h_{d,t} + B (\nu, v_u, v_d)),$$

where

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

with

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

A2-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

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

We show that

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

By denoting,

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

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,

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

In our implementation, we choose to parameterize the model in terms of ω_j^Q and σ_j^Q , directly.

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

parameters according to

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


with

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

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

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

Figure 1: Daily Returns and Realized Variances



These figures present the daily returns on S&P 500 index, R_t (Graph A), the daily realized volatilities, $\sqrt{RV_t}$ (Graph C), the daily upside realized volatilities, $\sqrt{RV_t^U}$ (Graph B), 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 five-minute squared returns, using a scaled RV estimator. The sample starts from January 10, 1990 and ends on August 28, 2013.

Table 1: **Summary Statistics of Historical 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

This table presents the summary statistics for the studied series. Mean, median, and standard deviation values are annualized and in percentages. $AR(1)$ describes first autocorrelation coefficient values. The sample starts from January 10, 1990 and ends on August 28, 2013.

Table 2: Estimation on Historical Returns and Realized Variance Components

Parameters	One-Factor Models				Two-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
	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	

This table shows maximum likelihood estimation results for six different models. We use daily historical returns, upside, and downside realized variances for the S&P 500 index from January 10, 1990 through August 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 while the second-to-last row indicates log likelihood values when all models are estimated on returns only. The third-to-last row shows the log likelihoods when one-factor ARV and SARV models are estimated on returns and realized variances. The fourth-to-last row reports log likelihood values when two-factor GARV, CGSARV, and GSARV models are estimated on returns, upside realized variances, and downside realized variances.

Table 3: **Regressions of Model-Free on Model-Implied Variances**

Parameters	One-Factor Models				Two-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
	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	

This table 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 six specifications in turn. We use the parameter estimates in Table 2 to generate model forecasts of the conditional variance. The sample period spans January 10, 1990 through August 28, 2013.

Table 4: **S&P 500 Index Option Data**

	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							
Number of contracts	3,788	1,391	1,781	2,846	2,746	8,731	21,283
Average price	7.85	20.94	32.28	45.30	65.93	132.41	74.35
Average implied volatility	16.72	18.40	19.31	20.40	21.71	25.09	21.62
Average 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							
Number of contracts	2,725	6,480	5,053	2,869	1,974	2,182	21,283
Average price	41.26	61.01	76.44	92.30	97.88	105.59	74.35
Average implied volatility	20.21	21.28	21.73	22.94	22.08	21.95	21.62
Average 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							
Number of contracts	3,962	6,133	5,996	2,456	1,240	1,496	21,283
Average price	57.95	66.90	80.75	85.77	85.33	94.86	74.35
Average implied volatility	13.61	18.04	22.45	26.24	30.22	39.42	21.62
Average bid-ask spread	1.055	1.301	1.446	1.704	1.811	2.683	1.470

This table presents the characteristics of S&P 500 index option data by moneyness, maturity, and VIX level. We use Wednesday closing out-of-the-money (OTM) call and put contracts from OptionMetrics for the period starting from January 10, 1996 and ending on August 28, 2013. The moneyness is measured by the Black-Scholes delta. DTM denotes the number of calendar days to maturity. The average price is reported in dollars and the average implied volatility is expressed in percentages.

Table 5: **Estimation on Options**

Parameters	One-Factor Models				Two-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
	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	

This table shows estimation results for six different models. We use Wednesday closing out-of-the-money (OTM) call and put contracts from OptionMetrics for the period starting from January 10, 1996 and ending on August 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 GARCH model.

Table 6: **IVRMSE Option Error by Moneyness, Maturity, and VIX**

	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
Panel A: IVRMSE By Moneyness						
<u>Model</u>						
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
	DTM < 30	30 < DTM < 60	60 < DTM < 90	90 < DTM < 120	120 < DTM < 150	DTM ≥ 150
Panel B: IVRMSE By Maturity						
<u>Model</u>						
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
	VIX < 15	15 < VIX < 20	20 < VIX < 25	25 < VIX < 30	30 < VIX < 35	VIX ≥ 35
Panel C: IVRMSE By VIX Level						
<u>Model</u>						
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

This table presents the implied volatility root mean squared error (IVRMSE) of the six models for contracts sorted by moneyness, maturity, and VIX level. We use the parameter values estimated in Table 5 to fit our six models to S&P 500 index option contracts from OptionMetrics. The sample starts from January 10, 1996 and ends on August 28, 2013. The first panel (Panel A) reports IVRMSE for contracts sorted by moneyness defined using the BlackScholes delta. The second (Panel B) reports IVRMSE for contracts sorted by days to maturity (DTM). The third panel (Panel C) reports the IVRMSE for contract sorted by the VIX level on the day corresponding to the option quote. The IVRMSE is expressed in percentages.

Table 7: Estimation on Historical Returns, Realized Variance Components, and Options

Parameters	One-Factor Models				Two-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
	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_Q^u	8.00E+01		6.96E+02	3.13E+00	5.38E+02	9.83E+00	5.13E+02	9.84E+00
κ_1^d					9.35E-01	3.27E-03	9.37E-01	3.30E-03
κ_2^d					1.30E-01	1.64E-03	1.34E-01	1.71E-03
γ_Q^d					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 IV	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	

This table shows the joint maximum likelihood estimation results for six different models. We use daily historical returns, upside/downside realized variances, and options on S&P 500 index from January 10, 1996 through August 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 GARCH model.

Table 8: Joint Estimation IVRMSE Option Error by Moneyness, Maturity, and VIX

	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
Panel A: IVRMSE By Moneyness						
<u>Model</u>						
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
	DTM < 30	30 ≤ DTM < 60	60 ≤ DTM < 90	90 ≤ DTM < 120	120 ≤ DTM < 150	DTM ≥ 150
Panel B: IVRMSE By Maturity						
<u>Model</u>						
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
	VIX < 15	15 ≤ VIX < 20	20 ≤ VIX < 25	25 ≤ VIX < 30	30 ≤ VIX < 35	VIX ≥ 35
Panel C: IVRMSE By VIX Level						
<u>Model</u>						
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

This table presents the implied volatility root mean squared error (IVRMSE) of the six models for contracts sorted by moneyness, maturity, and VIX level. We use the parameter values estimated in Table 7 to fit our six models to S&P 500 index option contracts from OptionMetrics. The sample starts from January 10, 1996 and ends on August 28, 2013. The first panel (Panel A) reports IVRMSE for contracts sorted by moneyness defined using the Black-Scholes delta. The second (Panel B) reports IVRMSE for contracts sorted by days to maturity (DTM). The third panel (Panel C) reports the IVRMSE for contracts sorted by the VIX level on the day corresponding to the option quote. The IVRMSE is expressed in percentages. BPJVM denotes the bipower jump variation option pricing model of Christoffersen et al. (2015).

Table 9: **Regressions of Model-Free on Model-Implied Variance Risk Premia**

Parameters	One-Factor Models				Two-Factor Models			
	GARCH		ARV		CGSARV		GSARV	
	Est	SE	Est	SE	Est	SE	Est	SE
<i>VRP</i>								
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	
<i>VRP^u</i>								
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	
<i>VRP^d</i>								
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	

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