Institut de la finance structurée et des instruments dérivés de Montréal
Montreal Institute of Structured Finance and Derivatives

L'Institut bénéficie du soutien financier de l'Autorité des marchés financiers ainsi que du ministère des Finances du Québec

Document de recherche

DR 16-10

Good Volatility, Bad Volatility and Option Pricing

Novembre 2016

Ce document de recherche a été rédigé par :

Bruno Feunou	Banque du Canada
Cédric Okou	UQAM

Good Volatility, Bad Volatility and Option Pricing*

Bruno Feunou
Financial Markets Department
Bank of Canada

Cédric Okou
ESG
UQAM

November 30, 2016

Abstract

Advances in variance analysis permit to split the total quadratic variation of a jump-diffusion process into upside and downside components, commonly referred to as good and bad volatilities. This decomposition yields enhanced volatility predictions over standard approaches, as documented by many recent studies. To appraise the economic gain of the decomposition, we design a new and flexible option pricing model in which the underlying asset price exhibits distinct upside and downside semi-variances driven by their model-free empirical proxies and random innovations. This discrete-time option valuation framework outperforms common benchmark models in terms of fitting accuracy and likelihood improvements. The model also delivers realistic term structures of risk-neutral moments and variance risk premia.

JEL Classification: G12

Keywords: Dynamic Upside Volatility, Dynamic Downside Volatility, Dynamic Skewness, Realized Downside Volatility, Realized Upside Volatility.

*Please address correspondence to Bruno Feunou, who can be reached by phone at 613-782-8302 or by email at feun@bankofcanada.ca. Cédric Okou is at okou.cedric@uqam.ca. We gratefully acknowledge financial support from the Bank of Canada, UQAM research funds, and the IFSID. We are indebted to Peter Christoffersen, Diego Amaya, Christian Dorion, and Yoontae Jeong for helpful discussions. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

1 Introduction

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

Building on these insights, this paper develops an option valuation model in which the underlying asset price features specific upside and downside variance dynamics. In our modeling framework, good and bad volatilities are factors governing the return process, and evolve according to their nonparametric realized 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, Kinnebrock, and Shephard (2010) show, in a model-free way, how to dissect the realized variance in terms of 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), or predict the equity risk premium given the standard risk-return tradeoff (Guo, Wang, and Zhou, 2015). Our option pricing framework extends these studies by helping gauge the economic value-added of disentangling the upside semi-variance motion from its downside counterpart.

This article is related to a large body of the finance literature that aims at building option pricing models with empirically grounded properties. A brief review of the various developments in the valuation of option contracts is provided by Chernov et al. (2003). A few studies propose to model the joint dynamics of returns and realized variances in the context of option pricing. This class of option valuation models is shown to deliver superior pricing performance compared to models optimized only on returns. Recent developments include papers by Stentoft (2008), Corsi, Fusari, and Vecchia (2013), Christoffersen et al. (2014), and Christoffersen, Feunou, and

Jeon (2015).

The aforementioned papers focus exclusively on the total realized variation, or its diffusive and jump components, but 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 accounts for the asymmetry in the distribution of of the underlying asset. Moreover, the model is affine and cast in discrete time, which permits computing explicit pricing formulas, and entails a straightforward fitting procedure.

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 take the model to data and show that the newly proposed specification performs well in matching the historical as well as the risk-neutral distributions. Namely, the model improves significantly upon popular specifications with respect to several 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 in our specification matters for delivering empirically grounded term structures of risk-neutral moments as well as market premia for upside and downside risks. Allowing for distinct up/down variance dynamics in our pricing model is useful to track the time variation in the variance risk premium and its 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 empirical performance of our pricing kernel and document the determinants of variance risk premium components. Section 7 concludes.

2 Daily Returns and Realized Variation Measures

In this section, 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 (resp. bad) volatility are often constructed from intraday records exceeding (resp. falling below) a specified threshold.

2.1 Separating downside from upside volatility: Empirics

For our empirical analysis, we download intraday S&P500 cash index data from TickData.com. On a given day, we use the last record in each one-minute interval to build a grid of one-minute equity index values from which we compute five series of overlapping five minute log-returns. This allows us to compute five realized variance measures from the sum of squared five-minute returns for each observation day in our sample. Following Andersen et al. (2003, 2001a), we construct the realized variance of returns on any given trading day t as $RV_t = \sum_{j=1}^{n_t} r_{j,t}^2$, where $r_{j,t}^2$ is the j^{th} intraday squared log-return and n_t is the number of intraday returns recorded on that day. We add the squared overnight log-return (the difference in log price between when the market opens at t and when it closes at $t-1$), and we scale the RV_t series to ensure that the sample average realized variance equals the sample variance of daily log-returns. For a given threshold $\kappa = 0$, we decompose the realized variance into upside and downside realized variances following Barndorff-Nielsen, Kinnebrock, and Shephard (2010):

$$\begin{aligned} RV_t^U &= \sum_{j=1}^{n_t} r_{j,t}^2 \mathbb{I}_{[r_{j,t}>0]}, \\ RV_t^D &= \sum_{j=1}^{n_t} r_{j,t}^2 \mathbb{I}_{[r_{j,t}\leq 0]}. \end{aligned}$$

We add the squared overnight "positive" log-return (exceeding the threshold $\kappa = 0$) to the upside realized variance RV_t^U, and the squared overnight "negative" log-return (falling below the threshold $\kappa = 0$) to the downside realized variance RV_t^D. Because the daily realized variance sums

the upside and the downside realized variances, we apply the same scale to the two components of the realized variance. Specifically, we multiply both components by the ratio of the sample variance of daily log-returns over the sample average of the (pre-scaled) realized variance.

By construction, the cumulative realized variance adds up the cumulative realized upside and downside variances:

$$RV_t \equiv RV_t^U + RV_t^D. \tag{1}$$

2.2 Separating downside from upside volatility: Theory

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, Kinnebrock, and Shephard (2010), who assume that the stock price follows a jump-diffusion of the form

$$ds_t = \mu_t dt + \sigma_t dW_t + \Delta s_t,$$

where dW_t is an increment of standard Brownian motion and $\Delta s_t \equiv s_t - s_{t-}$ refers to the jump component. The instantaneous variance can be defined as $\widetilde{\sigma}_t^2 = \sigma_t^2 + (\Delta s_t)^2$. Under this general assumption on the instantaneous return process, Barndorff-Nielsen, Kinnebrock, and Shephard (2010) use infill asymptotics – convergence as the time separating two consecutive observations collapses to 0 – to prove that

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

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

Moreover, the difference between realized upside and downside variance can be perceived as a measure of (realized) skewness. This measure of asymmetry, denoted as RSV_t, is computed by

subtracting the downside semi-variance from the upside semi-variance:

$$RSV_{t+1} = RV^U_{t+1} - RV^D_{t+1}.$$

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, Jahan-Parvar, and Tédongap (2014). To provide more intuition on the behavior of RSV_{t+1}, we combine the previous asymptotic results to get

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

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

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

Thus, $\mathbb{E}^{\mathbb{P}}_t \left[RSV_{t+1}\right]$ captures – up to a multiplicative constant – the wedge between positive and negative jump intensities. Put differently, 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 Renò (2015). Obviously, bigger weights are assigned to larger jump sizes in the computation of RSV_{t+1}.

2.3 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 02, 1990 to April 26, 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 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, as compared to 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, as documented in Feunou, Jahan-Parvar, and Okou (2015). Using a similar logic as Bollerslev, Tauchen, and Zhou (2009), dynamic "variance-of-(semi-)variance" processes are needed to generate premia for second-order semi-moments. Specifically, the realized downside variance measure requires a dynamic specification of its own, one that is likely different from the dynamic specification of upside variation. Building a dynamic return model with such features is our next task.

3 A 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 our observable quantities. In addition, we focus on an empirical strategy for option pricing that can be implemented without resorting to Monte Carlo simulations.

3.1 The asset return process

Consider the following specification for daily log returns

$$R_{t+1} = \bar{r} + (\lambda_u - \xi_u)\, h_{u,t} + (\lambda_d - \xi_d)\, h_{d,t} + z_{u,t+1} - z_{d,t+1}, \tag{2}$$

where the innovations, $z_{u,t+1}$ and $z_{d,t+1}$, distinctly and respectively represent positive and negative shocks to the return dynamics.[1] These innovations are modeled as

$$z_{j,t+1} = \tilde{z}_{j,t+1} - E_t\left[\tilde{z}_{j,t+1}\right], \text{ for } j = \{u, d\},$$

where the underlying shocks $\tilde{z}_{j,t+1}$ have a positive support, namely, $\tilde{z}_{j,t+1} > 0$ for $j = \{u, d\}$. This specification ensures that the return shocks, $z_{u,t+1}$ and $z_{d,t+1}$, are conditionally demeaned. We further assume that the volatilities of returns shocks are time varying and driven by the state variables $h_{u,t}$ and $h_{d,t}$. Thus, the following equalities hold:

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

We therefore interpret $h_{u,t}$ and $h_{d,t}$ as good and bad stock market volatilities, that is, volatilities stemming from the right and left tail movements in stock returns.

To get closed-form expressions for equilibrium asset prices, we consider a return shock distribution for which the log moment-generating function is linear in the underlying variances $h_{u,t}$ and $h_{d,t}$. Namely,

$$\ln E_t\left[e^{\nu z_{j,t+1}}\right] = a_j(\nu) + b_j(\nu) h_{j,t},$$

where the functions $a_j(\nu)$ and $b_j(\nu)$ capture the shape of the moment-generating function of the underlying return shocks. Compound Poisson and Gamma laws are prominent examples of such distributions. Bekaert and Engstrom (2015) use a difference of demeaned Gamma shocks to study asset return dynamics in a "bad environment-good environment" framework. For this class of distributions, the function $a_j(\nu)$ and $b_j(\nu)$ are non-negative, convex, and asymmetric. That is, $a_j(\nu) > a_j(-\nu)$ and $b_j(\nu) > b_j(-\nu)$ for $\nu > 0$. In this paper, we opt for the Noncentral chi-squared distribution, which has been widely used in the discrete time option pricing literature (See Heston

[1] In our timing convention, $h_{j,t}$ for $j = \{u, d\}$ stands for the conditional semi-variance for day $t+1$ computed at the end of day t.

and Nandi (2000)). Thus, we cast the underlying shocks to return as

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

with

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

One can easily check that the corresponding $a_j(\nu)$ and $b_j(\nu)$ functions are

$$\begin{aligned} b_j(\nu) &= \frac{\nu^2}{2\left(1 - \sqrt{2\omega_j}\nu\right)}, \\ a_j(\nu) &= -\sqrt{\frac{\omega_j}{2}} - \omega_j b_j(\nu) - \frac{1}{2}\ln\left(1 - \sqrt{2\omega_j}\nu\right), \end{aligned} \tag{3}$$

and that by construction, we have

$$h_{j,t} = Var_t\left[z_{j,t+1}\right].$$

By setting $\omega_j = 0$ in functions $a_j(\nu)$ and $b_j(\nu)$, it is apparent that

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

Thus, the parameter ω_j drives the non-normalities, a property which is useful when modeling daily stock return series. To further analyze the sources of potential departure from normality, we compute the conditional skewness

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

where

$$h_t \equiv h_{u,t} + h_{d,t}$$

is the total conditional variance of returns. Two sources of non-normality emerge, that is, the difference between upside and downside variance $h_{u,t} - h_{d,t}$, and the discrepancy between non-normalities in good and bad shock distributions $\sqrt{\omega_u} - \sqrt{\omega_d}$. An interesting case is obtained by imposing $\omega_u = \omega_d$. This constraint implies that the wedge between the upside and downside volatil-

ity, $h_{u,t} - h_{d,t}$, offers a unique channel for inducing conditional skewness in the return dynamics.

3.2 Incorporating realized upside and downside variation

Each day, realized upside and downside variations provide new information about the 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

$$RV_{t+1}^j = h_{j,t} + \sigma_j \left[\left(\varepsilon_{j,t+1}^{(2)} - \gamma_j \sqrt{h_{j,t} - \omega_j} \right)^2 - \left(1 + \gamma_j^2 \left(h_{j,t} - \omega_j \right) \right) \right] \text{ for } j = \{u, d\}. \tag{4}$$

The measurement error variables introduced in the realized semi-variances dynamics,

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

feature the following correlation structure

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

where the $\varepsilon_{j,t+1}^{(1)}$s denote the return shocks.

The innovation term inside the brackets in equation (4) is constructed to have zero mean, ensuring that

$$E_t \left[RV_{t+1}^j \right] = h_{j,t}.$$

Note also that equation (4) allows for a nonlinear impact of $\varepsilon_{j,t+1}^{(2)}$ on RV_{t+1}^j via γ_j.

3.3 Upside and downside volatility dynamics

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

$$h_{j,t+1} - \omega_j = \tilde{\omega}_j + \tilde{\beta}_j \left(h_{j,t} - \omega_j \right) + \tilde{\alpha}_j \left(RV_{t+1}^j - \omega_j \right). \tag{5}$$

This is very similar to GARCH models, with the exception that less noisy measures are used to update conditional variances. This 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 (4), equation (5) can be rewritten as

$$h_{j,t+1} - \omega_j = \varpi_j + \beta_j \left(h_{j,t} - \omega_j\right) + \alpha_j \left(\varepsilon_{j,t+1}^{(2)} - \gamma_j \sqrt{h_{j,t} - \omega_j}\right)^2, \tag{6}$$

where

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

Equation (6) is very similar to the conditional variance dynamics in Heston and Nandi (2000). The only difference here is that the Gaussian shock $\varepsilon_{j,t+1}^{(2)}$ is not perfectly correlated to returns. We will refer to this general specification as the generalized skew affine realized variance (GSARV) model.

3.4 Expected returns and risk premiums

Equation (2) clearly entails that the one-day-ahead conditionally expected log returns in the model is simply

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

The coefficients ξ_u, ξ_d and $\bar{r}$, are functionally dependent on the deep parameters of the model, as follows

$$\bar{r} = r_f + \frac{1}{2}\left[\ln\left(1 - \sqrt{2\omega_u}\right) + \ln\left(1 + \sqrt{2\omega_d}\right) - \frac{\omega_u - \sqrt{2\omega_u}}{\left(1 - \sqrt{2\omega_u}\right)} - \frac{\omega_d + \sqrt{2\omega_d}}{\left(1 + \sqrt{2\omega_d}\right)}\right], \tag{7}$$

$$\xi_u = \frac{1}{2\left(1 - \sqrt{2\omega_u}\right)}, \quad \xi_d = \frac{1}{2\left(1 + \sqrt{2\omega_d}\right)}, \tag{8}$$

where r_f denotes the risk-free rate. These functional forms ensure that the conditionally expected total return is

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

which in turn allows to interpret λ_u and λ_d as compensations for upside and downside volatilities exposures, respectively. Substituting equation (2) into (9), taking expectations, and solving for ξ_u, ξ_d and $\bar{r}$ yields equations (7) and (8). Therefore, the parameters ξ_u, ξ_d and $\bar{r}$ are not estimated below, but instead, simply set to their values implied by the identities in equations (7) and (8).

3.5 Conditional second moments

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

$$\begin{aligned}
Var_t\left[R_{t+1}\right] &= h_{u,t} + h_{d,t} \\
Var_t\left[RV_{t+1}^j\right] &= 2\sigma_j^2\left(1 + 2\gamma_j^2\left(h_{j,t} - \omega_j\right)\right) \\
Cov_t\left(R_{t+1}, RV_{t+1}^U\right) &= \sqrt{2\omega_u}\sigma_u\rho_u^2 + 2\left(h_{u,t} - \omega_u\right)\gamma_u\sigma_u\rho_u \\
Cov_t\left(R_{t+1}, RV_{t+1}^D\right) &= -\left(\sqrt{2\omega_d}\sigma_d\rho_d^2 + 2\left(h_{d,t} - \omega_d\right)\gamma_d\sigma_d\rho_d\right) \\
Cov_t\left(RV_{t+1}^U, RV_{t+1}^D\right) &= 0.
\end{aligned} \tag{10}$$

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.

4 Physical Estimation

Up to here, 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 five 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

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

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

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

where $\Phi(.)$ is the cumulative distribution function of a standard normal, and the $Var_t\left[RV_{t+1}^j\right]$s are given in equation (10). 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, $f_{r,t}\left(R_{t+1}\right)$, 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

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

where functions $a_j(\nu)$ and $b_j(\nu)$ are given in equation (3), 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

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

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

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

and where $Cov_t\left(R_{t+1},RV_{t+1}^j\right)$, $Var_t\left[R_{t+1}\right]$, and $Var_t\left[RV_{t+1}^j\right]$ are defined in Equation (10).

Finally, the log-likelihood is calculated as

$$\ln L^P = \sum_{t=1}^{T-1} \ln \left(f_t \left(R_{t+1}, RV_{t+1}^U, RV_{t+1}^D \right) \right). \tag{11}$$

4.2 Nested specifications

Before turning to the estimation, we discuss five 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 the 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 Two factors affine realized variance (GARV) model

An alternative modeling choice is to conceive upside and downside variances as volatility components that play no role in the conditional skewness. By imposing $\omega_u = \omega_d = 0$, we get

$$R_{t+1} = r_f + \left(\lambda_u - \frac{1}{2} \right) h_{u,t} + \left(\lambda_d - \frac{1}{2} \right) h_{d,t} + z_{u,t+1} + z_{d,t+1},$$

with

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

where $\varepsilon_{j,t+1}^{(1)}$ are *i.i.d.* $N(0,1)$, and $\varepsilon_{u,t+1}^{(1)}$ and $\varepsilon_{d,t+1}^{(1)}$ are conditionally independent.

Starting with $h_{j,0}$, we have

$$\begin{aligned} h_{j,t+1} &= \varpi_j + \beta_j h_{j,t} + \alpha_j \left(\varepsilon_{j,t}^{(2)} - \gamma_j \sqrt{h_{j,t}}\right)^2, \\ RV_{j,t+1} &= h_{j,t} + \sigma_j \left[\left(\varepsilon_{j,t+1}^{(2)} - \gamma_j \sqrt{h_{j,t}}\right)^2 - \left(1 + \gamma_j^2 h_{j,t}\right)\right], \end{aligned}$$

with

$$E\left[\varepsilon_{j,t+1}^{(1)} \varepsilon_{j,t+1}^{(2)}\right] = \rho_j,$$

and where $\varepsilon_{u,t+1}^{(2)}$ and $\varepsilon_{d,t+1}^{(2)}$ are conditionally independent random pairs.

Clearly, the GARV model is a particular case of the CGSARV, with the additional constraint $\omega = 0$.

4.2.3 The Skew affine realized variance (SARV) model

By setting $h_{d,t} = 0$ and $\omega_d = 0$, we have a single factor specification

$$R_{t+1} = \bar{r} + (\lambda - \xi)\, h_t + z_{t+1},$$

$$z_{t+1} = \sqrt{h_t}\eta_{t+1},$$

$$\eta_{t+1} = -\sqrt{\frac{\omega}{2h_t}} \left[\left(\varepsilon_{t+1}^{(1)} - \sqrt{\frac{h_t - \omega}{2\omega}}\right)^2 - \left(\frac{h_t + \omega}{2\omega}\right)\right],$$

where the $\varepsilon_{t+1}^{(1)}$s are *i.i.d.* $N(0,1)$.

Starting with $h_0 \geq \omega$, we have

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

where

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

$$\begin{aligned}\bar{r} &= r_f + \frac{1}{2}\left[\ln\left(1+\sqrt{2\omega}\right) - \frac{\omega+\sqrt{2\omega}}{\left(1+\sqrt{2\omega}\right)}\right], \\ \xi &= \frac{1}{2\left(1+\sqrt{2\omega}\right)}.\end{aligned}$$

The ARV of Christoffersen et al. (2014) is a special case of the SARV, obtained by letting $\omega = 0$. This model also bears close resemblance to the skew GARCH model in Christoffersen, Heston, and Jacobs (2006) with two notable differences: (i) we use realized variance to update the conditional variance and (ii) to model the departure from non-normality, we rely on the non-central chi-square distribution instead of the inverse gamma distribution.

4.2.4 The affine realized variance (ARV) model

If we fix $h_{u,t} = h_{d,t}$ in the GARV 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

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

where

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

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

4.2.5 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 processes 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.

In all six model specifications, 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, as compared to 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 and SARV models are 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 single factor models are roughly half (one-sixth in the case of the GARCH model) of that of the two-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 likelihood results. According to this metric, the generalized skew affine realized variance (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. Note that two-factor models yield lower realized variance-implied log-likelihood values than one-factor models (except for the GARCH specification), because the former class of specifications is optimized on returns, RV_u, and RV_d, but not on returns, and RV, as for the latter class of models. Among two-factor models, the GSARV model implies the highest log-likelihood for realized variance data, 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 on the model-implied variance, 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 used in our empirical analysis, and then study the performance of the various models relying on the implied volatility root-mean-square-error.

5.1 Exploring option data

We use European-style options written on the S&P500 index. The observations span the period January 10, 1996 through April 24, 2013[2]. 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

[2]Data are available through OptionMetrics, which supplies data for the US option markets.

the six most liquid strike prices, which amounts to a data set of 22,857 option contracts. To ease calculation and interpretation, out-of-the-money put prices are converted into corresponding in-the-money call values, by exploiting the call-put parity relation.

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

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

where $\Phi\left(*\right)$ stands for the normal cumulative distribution function (CDF) and IV^{Mkt} denotes the annualized implied Black-Scholes volatility computed at the market price of the option. A few empirical regularities emerge at this point. We observe that deep out-of-the-money 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 in favor of using the IVRMSE as a

proper model performance comparison tool in option pricing. Basically, the IVRMSE synthesizes the wedge 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 likelihood[3]

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

Table 5 contains the results of the option-based estimation. Clearly, our option fitting strategy yields accurate parameters 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

[3]Below, 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.

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 12% 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 in the bottom panel of Table 5. To this end, we dissect the IVRMSE by moneyness, maturity and VIX levels in Table 6. We use the same clusters as for the description of option data in Table 4.

We see that all models offer a satisfactory performance (low IVRMSEs) in matching at-the-money options contracts. By contrast, fitting deep out-of-the-money call and put options seems more challenging. Interestingly, the ability of the various specifications to match the observed implied 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.

5.4 Model fit for risk-neutral moments

We extend our analysis by gauging the ability of the different models to predict the term structure of risk-neutral volatility, skewness, and kurtosis. From the works of Bakshi and Madan (2000), Carr and Madan (2001), and Bakshi, Kapadia, and Madan (2003) we have a clear understanding of how to construct nonparametric risk-neutral moments. In Table 7, we run Mincer and Zarnowitz (1969) regressions of one-, three-, and six-month-ahead model-free (nonparametric) risk-neutral moments on corresponding model-impled series. Looking at regression coefficients (ideally, the intercept should be 0 and the slope should be 1), we notice that all specifications are well capable of mimicking the term structure of the risk neutral volatility. This observation is supported by the

top plot in Figure 2, where time series of 6-month model-implied risk-neutral volatilities track well the dynamics on the model-free series, including the prominent spike in volatility that occurred during the 2008 financial meltdown. Moreover, the GSARV model outperforms the other models in terms of fitting power measured by the R-squared. The different models also do a good job in mimicking the dynamic of risk-neutral skewness. Model-implied skewness values are fairly close the observed series, as the Mincer and Zarnowitz (1969) regressions slopes are close to 1. However, fitting the term structure of risk-neutral kurtosis is more challenging, possibly due to the fact that model-free estimates are noisy.

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 empirical performance of the chosen 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 11), and an option-based log likelihood component. Consistent with our previous methodology, we define Black-Scholes implied volatility errors as

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

and embed them in a Gaussian kernel to compute the option-based log likelihood

$$\ln L^O \propto -\frac{1}{2}\sum_{j=1}^{N}\left\{\ln\left(IVRMSE^2\right) + e_j^2/IVRMSE^2\right\}. \tag{12}$$

Our joint maximization program follows from Equations (11) and (12) as

$$\max \ln L^P + \ln L^O. \tag{13}$$

Solving the joint optimization problem should help elicit the different premia, including the total variance risk premium, the upside variance risk premium, the downside variance risk premium, and the skewness 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

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

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

where

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

Notice that our pricing kernel specification 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_2s are key elements that drive the wedge between the one-day-ahead conditional risk-neutral and physical variance (the variance spread). The magnitude of this spread may be non-trivial, as pointed out by Christoffersen, Heston, and Jacobs (2013), who also find an existing spread in the conditional Gaussian model. The External Appendix C contains the 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 is

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

where $\bar{r}^Q$ and ξ_j^Q are the risk-neutral analogs of $\bar{r}$ and ξ_j given in equations (7) and (8).

Under the risk-neutral measure, the risk-neutral shocks are noncentral chi-squared distributed according to

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

and

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

Note that the flexibility of the pricing kernel enables us to have both ω_j^Q and ω_j as independent[4] parameters to be estimated.

We now formally characterize the wedge between the latent conditional variance under the risk-neutral versus the physical distribution, a wedge which we refer to as variance spread (VSP). We show that this variance spread plays a non-trivial role in the determination of the variance risk premium (VRP) dynamics. We obtain the following expressions

$$VSP_t^j \equiv Var_t^Q\left[z_{j,t+1}^Q\right] - Var_t\left[z_{j,t+1}\right] = \vartheta_j + (\varsigma_j - 1)\, h_{jt}, \tag{14}$$

$$VRP_t^j \equiv E_t^Q\left[RV_{t+1}^j\right] - E_t^P\left[RV_{t+1}^j\right] = \theta_j + \varrho_j h_{jt} + \sigma_j^Q\left(\gamma_j^Q\right)^2 VSP_t^j, \tag{15}$$

with

$$\varsigma_d = \frac{\sqrt{2\omega_d^Q}\left(1+\sqrt{2\omega_d^Q}\right)}{\sqrt{2\omega_d}\left(1+\sqrt{2\omega_d}\right)} + 2\sqrt{2\omega_d^Q}\left(1+\sqrt{2\omega_d^Q}\right)\lambda_d,$$

$$\varsigma_u = \frac{\sqrt{2\omega_u^Q}\left(1-\sqrt{2\omega_u^Q}\right)}{\sqrt{2\omega_u}\left(1-\sqrt{2\omega_u}\right)} - 2\sqrt{2\omega_u^Q}\left(1-\sqrt{2\omega_u^Q}\right)\lambda_u,$$

[4]The equivalence between the two probability measure requires that $\omega_j^Q > 0$ if and only if $\omega_j > 0$.

and

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

Thus, the risk-neutral conditional variance ($h_{j,t+1}^Q$) and the dynamic of the realized variances under the risk-neutral measure are given by:

$$\begin{aligned}
h_{j,t+1}^Q &= h_{j,t}+VSP_t^j,\\
RV_{t+1}^j &= h_{j,t}+VRP_t^j+\sigma_j^Q z_{j,t+1}^{(*2)},
\end{aligned}$$

where

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

Interestingly, the dynamics of the risk-neutral factors $h_{u,t+1}^Q$ and $h_{d,t+1}^Q$ under the risk-neutral probability measure share a similar structure with that of $h_{u,t+1}$ and $h_{d,t+1}$ under the physical probability measure. That is

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

with

$$\begin{aligned}
\varpi_j^Q &= (1-\beta_j)\left(\vartheta_j + \varsigma_j\omega_j - \omega_j^Q\right) + \varsigma_j\varpi_j, \\
\alpha_j^Q &= \frac{\sigma_j^Q\varsigma_j}{\sigma_j}\alpha_j, \\
\rho_j^Q &= \frac{\rho_j\chi_j}{\sqrt{\left(1-\rho_j^2\right)\ \delta_j^2 + \rho_j^2\chi_j^2}}, \text{ where} \\
\chi_j &\equiv \left(\frac{\omega_j^Q}{\omega_j}\right)^{1/4}\left(\frac{\sigma_j^Q}{\sigma_j}\right)^{1/2}, \quad \delta_j^2 \equiv \frac{\left(1-\rho_j^2\right) + \sqrt{\left(1-\rho_j^2\right)^2 + 4\rho_j^2\chi_j^2}}{2}.
\end{aligned}$$

This general dynamic pricing framework is then applied to each of the nested models.

6.1.1 Analyzing variance spread and variance risk premium drivers

The variance spread expressions in equation (14) suggest that there are potentially three factors driving the wedge between risk-neutral and physical variances: the nonnormality ($\omega_j,\ \omega_j^Q \neq 0$), the distance between asymmetries ($\left\|\omega_j^Q - \omega_j\right\|$), and the market compensations for upside and downside variance exposures (λ_j). It is important to notice that our choice of a flexible stochastic discount factor induces a spread between risk-neutral and physical variances even in the Gaussian case ($\omega_j,\ \omega_j^Q = 0$). We refer the reader to the External Appendix C for technical details. In the non-Gaussian shock case, compensations for upside and downside variance 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.

In the option pricing literature, valuation models are expected to generate risk-neutral variances that are larger than physical variances in order to be 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 $\omega_j^Q > \omega_j$ imply that $VSP_t^j \geq 0$. The previous 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 (15), one can interpret ϱ_j as a difference between risk-neutral and physical variance of realized variance. That 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, Jahan-Parvar, and Okou (2015) have documented that the premium on the downside variance is positive $\left(VRP_t^d > 0\right)$, where as its upside counterpart is negative ($VRP_t^u < 0$). These empirical regularities are met within our model if $\varrho_d > 0$ and $\varrho_u < 0$, thus implying

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

6.2 Empirical results

Table 8 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 CGSARV and GSARV models, the inferred λ_u values are found to be negative. Thus, the joint estimation results for CGSARV and GSARV 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, Jahan-Parvar, and Tédongap (2013), among others. 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 8 also reports the corresponding joint log likelihood value along with its decomposition into the several components. To allow for comparison among models, we break down the global log likelihood into its 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.

Table 9 presents the Mincer and Zarnowitz (1969) regressions of model-free variance premium and its upside (resp. downside) component on their model-implied counterparts. The results clearly show that two-factor models deliver efficient predictions of variance risk premium values, as compared to single factor specifications. This evidence suggests that specifying distinct dynamics for upside and downside second-order variations helps significantly improve the model ability to track the market compensation of total variance risk, as illustrated by Figure 3. Looking at the variance risk premium components, we see that two-factor model forecasts for upside variance risk premium are clearly more accurate than the model-implied series for the downside variance risk premium. Predicting the downside component of the variance risk premium appears to be challenging, especially at short horizon, possibly due to its high variability. Namely, the predictive power of Mincer and Zarnowitz (1969) regressions at 1-month maturity is weak, as adjusted R-squared values are marginal. Interestingly, downside variance premium predictions from the GSARV model dominate that from the GARV specification. Thus, to generate realistic variance risk premium components, it is important to allow the difference in upside and downside variance dynamics drive the time variation in the conditional asymmetry.

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, yet empirically grounded, 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 empirically grounded risk-neutral moments and market compensations for upside and downside variance risks across various maturities.

References

Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens. 2001a. The distribution of realized stock return volatility. *Journal of Financial Economics* 61:43–76.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys. 2001b. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* 96:42–55.

———. 2003. Modeling and forecasting realized volatility. *Econometrica* 71:579–625.

Bakshi, G., N. Kapadia, and D. Madan. 2003. Stock return characteristics, skew laws and the differential pricing of individual equity options. *Review of Financial Studies* 16:101–43.

Bakshi, G., and D. Madan. 2000. Spanning and derivative security valuation. *Journal of Financial Economics* 55:205–38.

Bandi, F. M., and R. Renò. 2015. Price and volatility co-jumps. *Journal of Financial Economics, forthcoming* .

Barndorff-Nielsen, O. E., S. Kinnebrock, and N. Shephard. 2010. *Volatility and time series econometrics: Essays in honor of robert f. engle*, chap. Measuring downside risk: realised semivariance, 117–36. Oxford University Press.

Bates, D. 2000. Post-'87 crash fears in the s&p 500 futures option market. *Journal of Econometrics* 96:181–238.

Bekaert, G., and E. Engstrom. 2015. Asset return dynamics under bad environment good environment fundamentals. *Journal of Political Economy, forthcoming* .

Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81:637–54.

Bollerslev, T., G. Tauchen, and H. Zhou. 2009. Expected stock returns and variance risk premia. *Review of Financial Studies* 22:4463–92.

Carr, P., and D. Madan. 2001. Optimal positioning in derivative securities. *Quantitative Finance* 1:19–37.

Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen. 2003. Alternative models for stock price dynamics. *Journal of Econometrics* 116:225–57.

Christoffersen, P., R. Elkamhi, B. Feunou, and K. Jacobs. 2010. Option valuation with conditional heteroskedasticity and nonnormality. *Review of Financial Studies* 23:2139–83. doi:10.1093/rfs/hhp078.

Christoffersen, P., B. Feunou, K. Jacobs, and N. Meddahi. 2014. The economic value of realized volatility: Using high-frequency returns for option valuation. *Journal of Financial and Quantitative Analysis* 49:663–97. ISSN 1756-6916. doi:10.1017/S0022109014000428.

Christoffersen, P., B. Feunou, and Y. Jeon. 2015. Option valuation with observable volatility and jump dynamics. *Journal of Banking & Finance* 61, Supplement 2:S101 – S120. ISSN 0378-4266. doi:http://dx.doi.org/10.1016/j.jbankfin.2015.08.002. Recent Developments in Financial Econometrics and Applications.

Christoffersen, P., S. Heston, and K. Jacobs. 2006. Option valuation with conditional skewness. *Journal of Econometrics* 131:253 – 284. ISSN 0304-4076. doi:http://dx.doi.org/10.1016/j.jeconom.2005.01.010.

———. 2013. Capturing option anomalies with a variance-dependent pricing kernel. *Review of Financial Studies* 26:1963–2006. doi:10.1093/rfs/hht033.

Corsi, F., N. Fusari, and D. L. Vecchia. 2013. Realizing smiles: Options pricing with realized volatility. *Journal of Financial Economics* 107:284–304.

Duffie, D., J. Pan, and K. Singleton. 2000. Transform analysis and option pricing for affine jump-diffusions. *Econometrica* 68:1343–1377.

Feunou, B., M. R. Jahan-Parvar, and C. Okou. 2015. Downside Variance Premiun. *Working Paper, Bank of Canada, Federal Reserve Board, UQAM* .

Feunou, B., M. R. Jahan-Parvar, and R. Tédongap. 2013. Modeling Market Downside Volatility. *Review of Finance* 17:443–81.

———. 2014. Which parametric model for conditional skewness? *European Journal of Finance, forthcoming* .

Guo, H., K. Wang, and H. Zhou. 2015. Good jumps, bad jumps, and conditional equity premium. *Working Paper, University of Cincinnati, Xiamen University, and Tsinghua University* .

Heston, S. L., and S. Nandi. 2000. A closed-form garch option valuation model. *Review of Financial Studies* 13:585–625. doi:10.1093/rfs/13.3.585.

Huang, J.-Z., and L. Wu. 2004. Specification analysis of option pricing models based on time-changed lévy processes. *Journal of Finance* 59:1405–39.

Mincer, J., and V. Zarnowitz. 1969. *The evaluation of economic forecasts.* National Bureau of Economic Research.

Patton, A. J., and K. Sheppard. 2015. Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Review of Economics and Statistics* 97:683–97.

Renault, E. 1997. *Econometric models of option pricing errors, in advances in economics and econometrics, seventh world congress*, 223–78. Econometric Society Monographs, Cambridge University Press.

Stentoft, L. 2008. Option pricing using realized volatility. *CREATES Research Paper* 13.

Figure 1: Daily Returns and Realized Variances

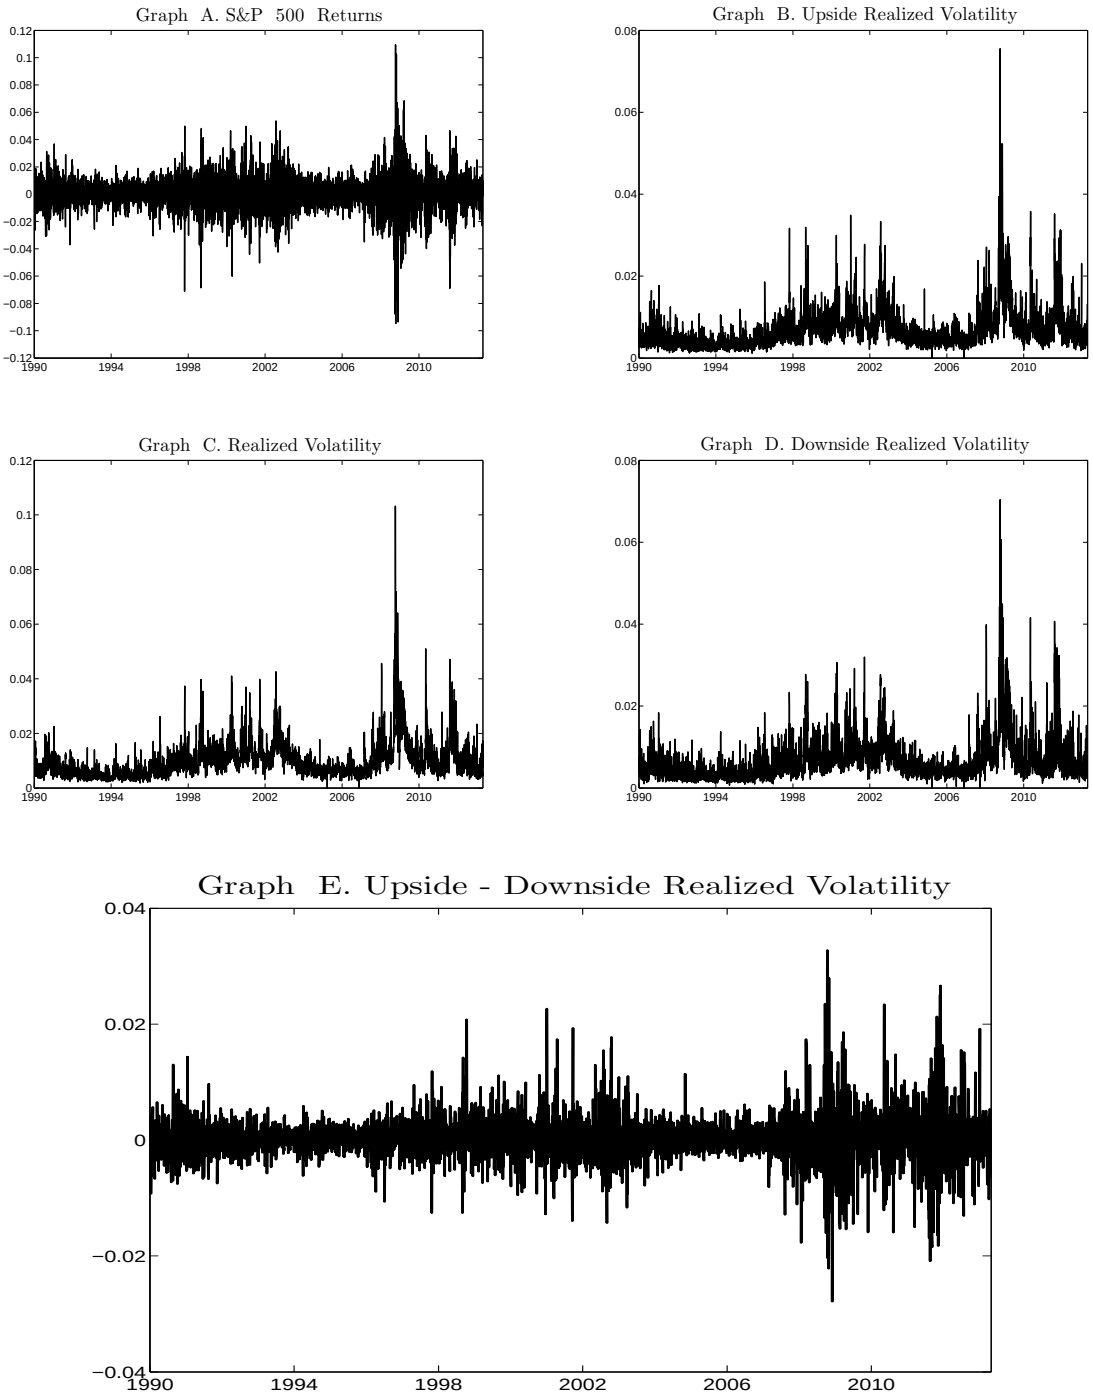

These figures present the daily returns on S&P 500 index, R_t (Graph A), the daily realized volatilities, $\sqrt{RV_t}$ (Graph C), the daily upside realized volatilities, $\sqrt{RV_t^U}$ (Graph B), the daily downside realized volatilities, $\sqrt{RV_t^D}$ (Graph D), and the difference between the daily upside and downside realized volatilities, $\sqrt{RV_t^U} - \sqrt{RV_t^D}$ (Graph E). Daily realized volatility quantities are computed from 5-minute squared returns, using an average RV estimator.

Figure 2: Model-Free and Model-Implied Risk-Neutral Moments

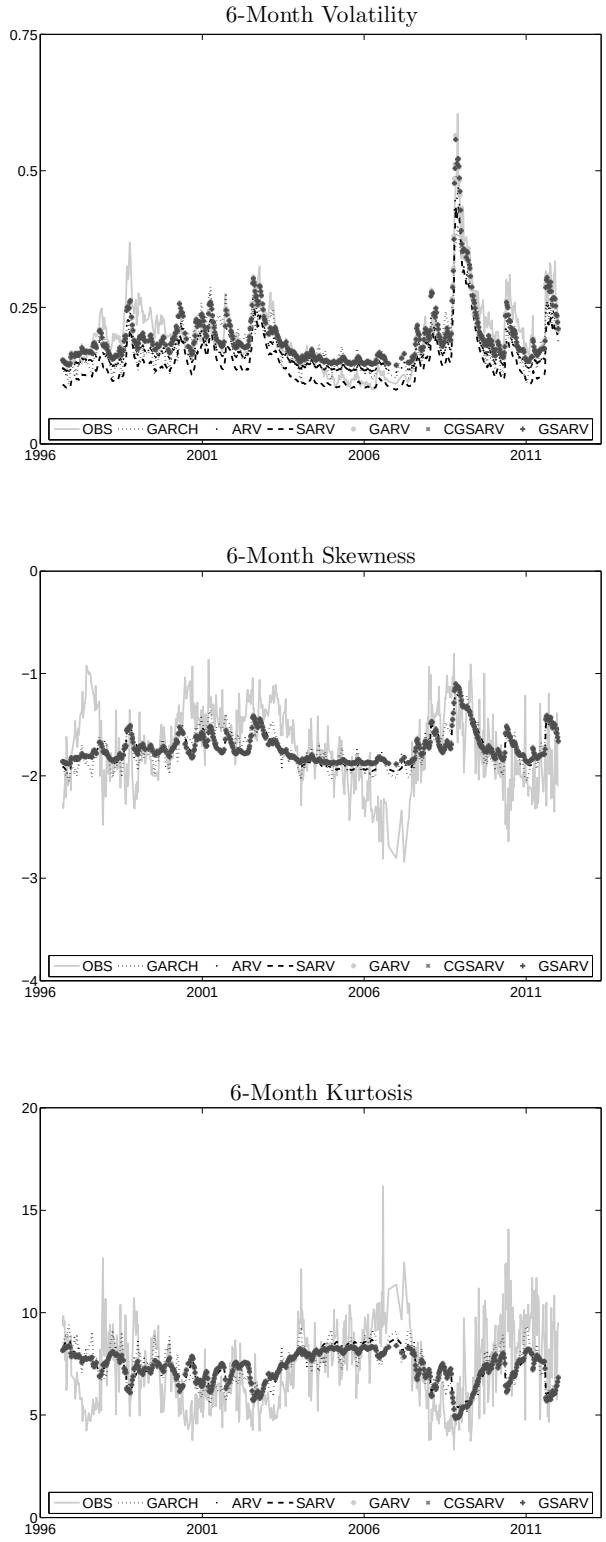

These figures plot the six-month (6M) model-free or observed (OBS) time series of risk-neutral volatility, skewness and kurtosis along with alternative model-implied series. We use the parameters estimated for each of the six models in Table 5. The sample runs from January 10, 1996 through April 24, 2013.

Figure 3: Model-Free and Two-factor Model-Implied Variance Risk Premia

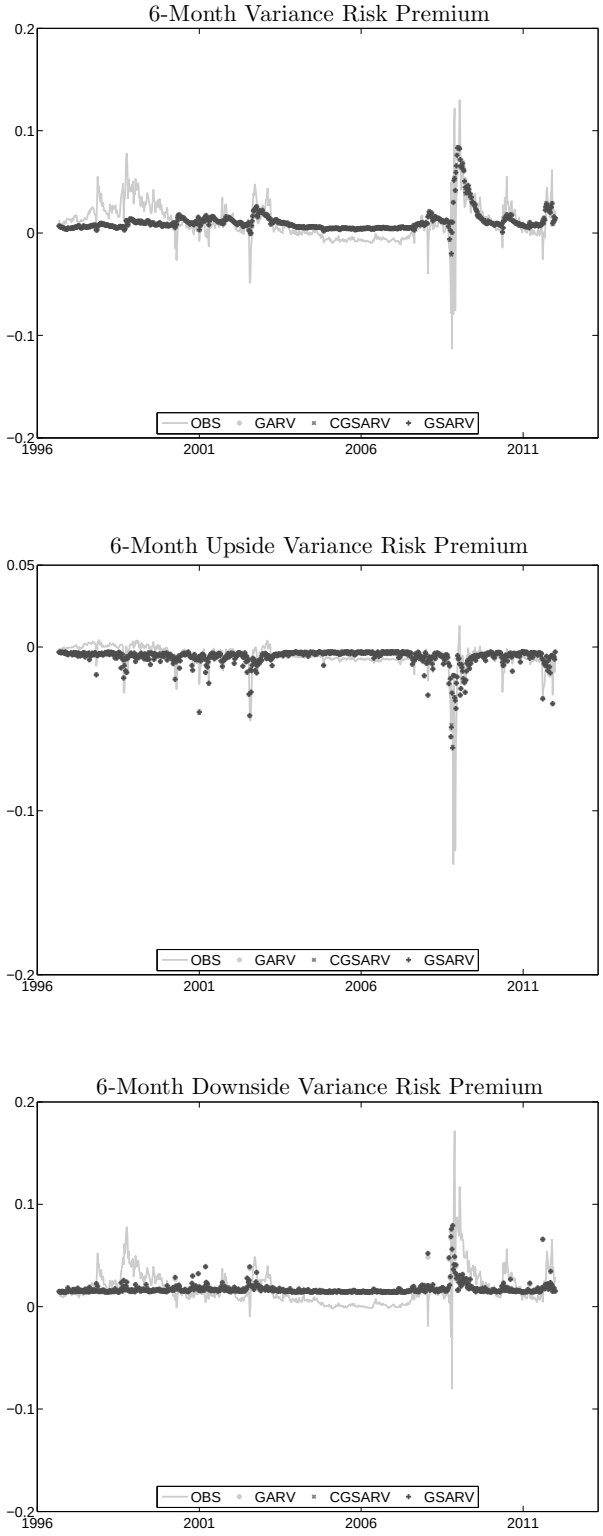

These figures plot the six-month (6M) model-free or observed (OBS) time series of (upside/downside) variance risk premium along with alternative two-factor model-implied series. We use the parameters estimated for each of the six models in Table 8. The sample runs from January 10, 1996 through April 24, 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 02, 1990 and ends on April 26, 2013.

Table 2: **Estimation on Historical Returns and Realized Variance Components**

	One-Factor Models						Tow-Factor Models					
	GARCH		ARV		SARV		GARV		CGSARV		GSARV	
Parameters	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
λ_u	1.37E+00		1.37E+00		1.37E+00		2.75E+00		2.75E+00		2.75E+00	
ϖ_u	2.80E-13		1.51E-12		1.65E-17		1.37E-10		6.70E-09		4.90E-08	
ω_u					3.90E-08	2.93E-05			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	2.35E-06	5.07E-07	8.85E-07	1.32E-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	9.31E-14	7.39E-12	1.09E-05	2.11E-03	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	6.46E+02	2.25E+02	1.06E+03	8.26E+00	1.05E+03	7.44E+00	1.01E+03	7.02E+02
σ_u			6.33E-06	1.57E-07	6.36E-06	1.37E-06	3.37E-06	2.86E-08	3.49E-06	2.71E-08	3.53E-06	2.44E-06
ρ_u			9.27E-02	8.28E-03	-1.67E-01	8.31E-03	-3.71E-01	1.93E-02	5.88E-01	1.13E-02	5.94E-01	1.21E-02
λ_d							3.45E-08	6.78E+00	1.93E-08	5.19E+00	4.66E-11	5.34E+00
ϖ_d							1.97E-10		9.64E-09		7.32E-09	
ω_d									2.39E-07	4.51E-08	7.39E-06	7.73E-07
α_d							9.99E-07	2.28E-08	8.51E-07	1.76E-08	9.59E-07	6.15E-07
β_d							8.79E-06	4.73E-03	9.56E-06	4.45E-08	4.34E-02	6.06E-01
γ_d							9.93E+02	1.18E+01	1.08E+03	1.11E+01	9.90E+02	6.37E+02
σ_d							3.79E-06	5.11E-08	3.44E-06	4.02E-08	4.06E-06	2.61E-06
ρ_d							7.38E-01	9.42E-03	7.59E-01	7.33E-03	7.59E-01	1.00E-02
$E[h_u]$							6.77E-05		6.77E-05		6.77E-05	
$E[h_d]$							6.79E-05		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		1.36E-04		1.36E-04	
Model Properties												
Avg. Upside Volatility							11.45		11.46		11.46	
Avg. Downside Volatility							11.40		11.40		11.46	
Avg. Volatility	16.58		16.58		16.58		16.21		16.21		16.25	
Variance Persistence												
From RV_u							0.9869		0.9865		0.9852	
From RV_d							0.9853		0.9873		0.9840	
From RV			0.9827		0.9826							
From Returns	0.9649											
Log Likelihoods												
Returns, RV_u, and RV_d							120,196		120,683		121,036	
Returns and RV			66,565		66,626							
Maximized on Returns	19,035		19,213		19,239		19,233		19,240		19,305	
Implied for RV			47,321		47,322		47,042		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 02, 1990 through April 26, 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**

	One-Factor Models						Tow-Factor Models					
	GARCH		ARV		SARV		GARV		CGSARV		GSARV	
Parameters	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
Variance												
Constant	-1.43E-04	4.41E-06	-1.63E-05	1.79E-06	-1.63E-05	1.79E-06	-1.67E-05	2.11E-06	-1.71E-05	2.09E-06	-1.70E-05	2.09E-06
Slope	2.26	0.03	1.12	0.01	1.12	0.01	1.12	0.01	1.13	0.01	1.12	0.01
R^2 (%)	52.24		84.11		84.06		78.29		78.60		78.57	
Upside Variance												
Constant							-9.74E-06	1.28E-06	-9.72E-06	1.30E-06	-9.74E-06	1.30E-06
Slope							1.14	0.01	1.14	0.01	1.14	0.01
R^2 (%)							72.03		71.26		71.46	
Downside Variance												
Constant							-1.04E-05	1.23E-06	-1.06E-05	1.20E-06	-1.06E-05	1.20E-06
Slope							1.15	0.01	1.16	0.01	1.16	0.01
R^2 (%)							74.04		75.21		75.00	

This table shows the estimated coefficients (Est) and the standard errors (SE) of the regression of model-free total (resp. upside, and downside) realized variance on the corresponding model-predicted total (resp. 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 02, 1990 through April 26, 2013.

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

	OTM Call			OTM Put			
	Delta < 0.3	$0.3 \leq$ Delta < 0.4	$0.4 \leq$ Delta < 0.5	$0.5 \leq$ Delta < 0.6	$0.6 \leq$ Delta < 0.7	Delta ≥ 0.7	All
Panel A: By Moneyness							
Number of contracts	4,509	1,424	2,330	2,728	2,746	9,120	22,857
Average price	8.28	25.09	39.31	54.22	75.92	171.49	91.22
Average implied volatility	17.20	19.23	20.46	21.27	22.27	27.45	22.84
Average bid-ask spread	1.061	1.802	2.078	1.971	1.769	1.000	1.380
	DTM < 30	$30 \leq$ DTM < 60	$60 \leq$ DTM < 90	$90 \leq$ DTM < 120	$120 \leq$ DTM < 150	DTM ≥ 150	All
Panel B: By Maturity							
Number of contracts	3,165	6,906	5,309	3,059	2,103	2,315	22,857
Average price	58.80	75.10	89.79	113.85	123.35	127.80	91.22
Average implied volatility	21.98	22.60	22.63	24.22	23.45	22.88	22.84
Average bid-ask spread	0.725	1.160	1.515	1.770	1.701	1.818	1.380
	VIX < 15	$15 \leq$ VIX < 20	$20 \leq$ VIX < 25	$25 \leq$ VIX < 30	$30 \leq$ VIX < 35	VIX ≥ 35	All
Panel C: By VIX Level							
Number of contracts	4,102	6,637	6,477	2,669	1,347	1,625	22,857
Average price	74.87	88.28	95.51	97.46	100.67	109.27	91.22
Average implied volatility	14.67	19.39	23.38	27.12	31.56	41.21	22.84
Average bid-ask spread	0.942	1.217	1.374	1.611	1.710	2.527	1.380

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 contrats from OptionMetrics for the period starting from January 10, 1996 and ending on April 24, 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 percentage.

Table 5: **Estimation on Options**

	One-Factor Models						Tow-Factor Models					
	GARCH		ARV		SARV		GARV		CGSARV		GSARV	
Parameters	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
ϖ_u	-8.47E-21		1.44E-20		-2.14E-20		3.92E-22		3.28E-15		8.57E-18	
ω_u					1.47E-05	1.47E-06			5.11E-06	9.69E-12	2.28E-06	9.69E-12
α_u	1.44E-06	1.13E-08	1.29E-06	1.74E-08	1.08E-06	2.48E-08	5.69E-22	9.13E-18	6.69E-15	1.80E-09	1.94E-17	1.80E-09
β_u	8.15E-01	1.64E-03	4.22E-20	0.00E+00	2.61E-20	0.00E+00	9.96E-01	8.82E-01	9.96E-01	1.62E-04	9.96E-01	1.62E-04
γ_u	3.49E+02	1.95E+00	8.77E+02	5.98E+00	9.57E+02	9.07E+00	7.14E+07	1.11E+13	1.74E+04	5.76E-01	2.05E+05	5.76E-01
σ_u			4.30E-05	5.26E-07	3.73E-05	7.84E-07	6.75E-20	1.08E-15	6.91E-13	1.95E-07	2.02E-15	1.95E-07
ρ_u			5.63E-01	3.58E-03	6.08E-01	8.97E-03	9.69E-01	3.76E+03	9.88E-01	4.66E-07	9.86E-01	4.66E-07
ϖ_d							2.03E-20		1.42E-20		-1.57E-20	
ω_d									5.11E-06	5.01E-07	5.55E-06	4.92E-07
α_d							1.87E-06	2.60E-08	1.88E-06	4.66E-07	1.98E-06	8.72E-07
β_d							3.43E-10	1.29E-20	3.01E-09	9.69E-12	6.98E-08	2.89E-12
γ_d							7.24E+02	5.16E+00	7.22E+02	4.23E-18	7.03E+02	2.97E-16
σ_d							1.63E-05	3.67E-07	1.76E-05	4.62E-08	1.68E-05	5.09E-08
ρ_d							1.00E+00	0.00E+00	1.00E+00	0.00E+00	1.00E+00	0.00E+00
$E^Q[h_u]$							2.48E-19	5.45E-17	5.11E-06	4.66E-07	2.28E-06	4.66E-07
$E^Q[h_d]$							1.07E-04	1.56E-06	1.08E-04	8.43E+00	1.03E-04	8.51E+00
$E^Q[h] = E^Q[h_u + h_d]$	1.48E-04	6.43E-07	1.17E-04	4.82E-07	1.20E-04	1.49E-06	1.07E-04		1.13E-04		1.05E-04	
Model Properties												
Log Likelihoods	30,308		32,145		32,242		32,975		33,075		33,089	
Avg. Upside Model IV							11.39		11.56		11.62	
Avg. Downside Model IV							13.74		13.83		13.74	
Avg. Model IV	17.81		18.73		18.81		17.97		18.13		18.11	
Variance Persistence												
From RV_u							0.9961		0.9956		0.9960	
From RV_d							0.9825		0.9817		0.9797	
From RV	0.9903		0.9890		0.9897							
Option Errors												
IVRMSE	6.425		5.937		5.913		5.720		5.696		5.692	
Ratio to GARCH	1.000		0.924		0.920		0.890		0.887		0.886	

This table shows estimation results for six different models. We use Wednesday closing out-of-the-money (OTM) call and put contrats from OptionMetrics for the period starting from January 10, 1996 and ending on April 24, 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. We also present the log likelihoods along with model-implied volatility components for all specifications. 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 6: **IVRMSE Option Error by Moneyness, Maturity, and VIX**

	OTM Call			OTM Put		
	Delta < 0.3	$0.3 \leq$ Delta < 0.4	$0.4 \leq$ Delta < 0.5	$0.5 \leq$ Delta < 0.6	$0.6 \leq$ Delta < 0.7	Delta ≥ 0.7
Panel A: IVRMSE By Moneyness						
Model						
GARCH	5.862	3.997	4.012	4.021	3.979	8.383
ARV	5.664	3.495	3.305	3.278	3.269	7.833
SARV	5.695	3.511	3.328	3.287	3.280	7.770
GARV	5.097	2.924	2.955	3.019	3.090	7.745
CGSARV	5.047	2.964	2.977	3.045	3.118	7.706
GSARV	5.018	2.927	2.962	3.054	3.130	7.709
	DTM < 30	$30 \leq$ DTM < 60	$60 \leq$ DTM < 90	$90 \leq$ DTM < 120	$120 \leq$ DTM < 150	DTM ≥ 150
Panel B: IVRMSE By Maturity						
Model						
GARCH	6.119	6.408	6.038	6.691	7.251	6.588
ARV	5.572	5.830	5.532	6.383	6.836	6.143
SARV	5.444	5.818	5.546	6.368	6.824	6.109
GARV	5.194	5.549	5.361	6.319	6.657	5.947
CGSARV	5.040	5.513	5.373	6.323	6.676	5.952
GSARV	5.036	5.510	5.361	6.317	6.681	5.951
	VIX < 15	$15 \leq$ VIX < 20	$20 \leq$ VIX < 25	$25 \leq$ VIX < 30	$30 \leq$ VIX < 35	VIX ≥ 35
Panel C: IVRMSE By VIX Level						
Model						
GARCH	5.835	5.965	6.544	6.449	7.587	7.920
ARV	5.456	5.359	5.771	6.040	7.656	7.880
SARV	5.510	5.378	5.695	5.986	7.563	7.842
GARV	5.241	5.382	5.521	5.642	7.363	7.337
CGSARV	5.405	5.364	5.431	5.578	7.237	7.283
GSARV	5.347	5.361	5.441	5.596	7.261	7.279

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&P500 index option contracts from OptionMetrics. The sample starts from January 10, 1996 and ends on April 24, 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 percentage.

Table 7: **Regressions of Model-Free on Model-Implied Risk-Neutral Moments**

	One-Factor Models						Tow-Factor Models					
	GARCH		ARV		SARV		GARV		CGSARV		GSARV	
Parameters	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
Volatility												
1-Month												
Constant	4.15E-03	4.91E-03	2.82E-03	3.30E-03	1.30E-02	3.43E-03	1.25E-03	3.03E-03	-1.64E-03	3.16E-03	-1.69E-03	3.16E-03
Slope	0.92	0.02	0.84	0.01	0.86	0.02	0.97	0.02	0.97	0.02	0.97	0.02
R^2 (%)	68.55		83.35		80.65		85.86		85.16		85.16	
3-Month												
Constant	7.40E-03	5.23E-03	4.31E-03	3.42E-03	2.33E-02	3.10E-03	3.62E-03	3.35E-03	2.57E-03	3.34E-03	2.53E-03	3.35E-03
Slope	1.00	0.03	0.98	0.02	1.00	0.02	1.01	0.02	1.01	0.02	1.01	0.02
R^2 (%)	68.75		84.35		84.54		85.02		85.18		85.18	
6-Month												
Constant	-2.63E-04	6.22E-03	-7.62E-03	4.50E-03	2.54E-02	3.59E-03	-4.36E-03	4.46E-03	-2.68E-03	4.34E-03	-2.71E-03	4.34E-03
Slope	1.08	0.03	1.16	0.02	1.15	0.02	1.03	0.02	1.03	0.02	1.03	0.02
R^2 (%)	63.65		78.39		80.62		78.18		78.88		78.89	
Skewness												
1-Month												
Constant	-1.22E+00	6.29E-02	-1.17E+00	7.58E-02	-1.21E+00	6.92E-02	-1.13E+00	7.95E-02	-1.08E+00	8.66E-02	-1.08E+00	8.66E-02
Slope	0.99	0.11	1.10	0.14	1.11	0.13	1.13	0.14	1.22	0.15	1.22	0.15
R^2 (%)	11.23		9.22		9.62		9.24		9.21		9.21	
3-Month												
Constant	-8.52E-01	6.92E-02	-8.16E-01	8.01E-02	-9.25E-01	6.56E-02	-7.81E-01	8.49E-02	-7.58E-01	8.88E-02	-7.58E-01	8.89E-02
Slope	1.10	0.08	1.16	0.10	1.23	0.09	1.19	0.10	1.22	0.11	1.22	0.11
R^2 (%)	22.54		18.81		21.70		18.09		17.37		17.37	
6-Month												
Constant	-6.05E-01	7.58E-02	-7.34E-01	8.80E-02	-9.17E-01	6.62E-02	-6.98E-01	9.30E-02	-6.92E-01	9.44E-02	-6.92E-01	9.44E-02
Slope	1.11	0.07	1.05	0.09	1.00	0.08	1.04	0.09	1.04	0.09	1.04	0.09
R^2 (%)	25.97		16.79		19.47		16.23		15.99		15.98	
Kurtosis												
1-Month												
Constant	7.41E-01	1.24E+00	5.05E-01	1.80E+00	-3.81E-02	1.93E+00	7.21E-01	1.70E+00	3.09E+00	1.21E+00	3.61E+00	1.28E+00
Slope	1.89	0.35	1.97	0.52	2.18	0.57	1.89	0.48	1.04	0.29	0.82	0.28
R^2 (%)	4.16		2.08		2.11		2.19		1.78		1.17	
3-Month												
Constant	1.31E+00	6.69E-01	1.41E+00	8.61E-01	8.76E-01	9.44E-01	1.26E+00	8.63E-01	1.18E+00	8.89E-01	1.18E+00	8.90E-01
Slope	1.56	0.16	1.54	0.21	1.82	0.25	1.56	0.21	1.59	0.22	1.59	0.22
R^2 (%)	12.46		7.59		7.44		7.90		7.65		7.64	
6-Month												
Constant	-1.31E+00	5.74E-01	-1.41E+00	7.61E-01	-1.25E+00	7.17E-01	-1.26E+00	7.39E-01	-1.21E+00	7.37E-01	-1.21E+00	7.37E-01
Slope	1.89	0.13	1.97	0.17	2.12	0.18	1.89	0.16	1.88	0.16	1.88	0.16
R^2 (%)	26.31		17.12		18.31		17.45		17.37		17.37	

This table shows the estimated coefficients (Est) and the standard errors (SE) of the regressions of model-free risk-neutral volatility (resp. skewness, and kurtosis) on the corresponding model-predicted risk-neutral values at various maturities (1,3, and 6-month) from each of the six specifications in turn. We use the parameter estimates in Table 5 to generate model forecasts. The sample period spans January 02, 1990 through April 26, 2013.

Table 8: **Estimation on Historical Returns, Realized Variance Components, and Options**

	One-Factor Models						Tow-Factor Models					
	GARCH		ARV		SARV		GARV		CGSARV		GSARV	
Parameters	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
λ_u	1.37E+00		1.37E+00		1.37E+00		2.75E+00		-3.98E-01		-7.34E-01	
ϖ_u	1.11E-13		1.00E-10		8.86E-13		2.65E-12		3.25E-12		4.10E-09	
ω_u					1.68E-08	1.94E-05			1.90E-09	8.72E-12	3.49E-08	1.04E-10
α_u	1.45E-06	5.96E-09	2.24E-06	3.28E-07	2.07E-06	2.95E-07	7.44E-07	1.07E-08	7.54E-07	3.27E-09	9.89E-07	2.19E-09
β_u	7.54E-01	3.29E-03	7.87E-04	2.45E-03	-2.22E-16	6.47E-14	2.12E-10	2.68E-13	1.73E-06	8.43E-11	3.59E-05	3.67E-09
γ_u	4.03E+02	3.21E+00	6.63E+02	1.62E+00	6.90E+02	3.48E+04	1.15E+03	8.39E+00	1.15E+03	2.51E+00	9.98E+02	1.12E+00
σ_u			6.77E-06	5.67E-06	6.86E-06	9.79E-07	3.20E-06	2.74E-08	2.84E-06	7.14E-09	3.57E-06	7.55E-09
ρ_u			2.55E-01	4.00E-03	3.55E-01	1.35E-03	-1.98E-01	1.90E-02	4.26E-01	8.44E-03	6.06E-01	7.46E-03
λ_d							5.26E-08	2.97E-06	3.13E+00	2.93E-13	3.47E+00	1.23E-10
ϖ_d							1.18E-11		1.76E-08		9.68E-09	
ω_d									1.90E-09	8.72E-12	7.80E-07	2.76E-10
α_d							8.95E-07	1.25E-08	1.05E-06	1.18E-09	8.81E-07	2.69E-09
β_d							2.20E-11	1.10E-11	4.81E-04	7.84E-09	4.50E-03	1.27E-03
γ_d							1.05E+03	7.46E+00	9.69E+02	5.55E-01	1.06E+03	2.26E+00
σ_d							3.87E-06	3.27E-08	3.90E-06	1.15E-08	3.54E-06	1.49E-08
ρ_d							7.26E-01	8.69E-03	7.80E-01	5.41E-03	7.89E-01	4.64E-03
Pricing Kernel Parameters												
κ_1^u	1.05E+00	3.36E-03	1.35E+00	9.15E-03	4.53E-01	3.35E-03	1.08E+00	1.82E-02	1.00E+00	4.20E-03	1.04E+00	1.65E-03
κ_2^u			6.70E-01	3.13E-03	6.69E-01	1.18E-03	8.34E-01	4.88E-02	7.85E-01	3.86E-03	9.51E-01	2.69E-03
γ_u^Q	4.23E+02		6.32E+02	3.40E+01	5.62E+02	2.51E-01	1.08E+03	3.32E+01	1.00E+03	8.66E-06	9.75E+02	8.10E-04
κ_1^u							1.11E+00	1.66E-02	1.41E+00	7.46E-03	9.97E-01	2.50E-04
κ_2^u							8.73E-01	1.23E-02	7.75E-01	9.78E-04	9.51E-01	4.24E-03
γ_u^Q							1.04E+03	8.45E+00	1.02E+03	1.95E-05	1.03E+03	9.91E-02
Model Properties												
Avg. Physical Volatility	16.58		16.58		16.58		16.21		16.21		16.21	
Avg. Model IV	17.31		15.81		27.34		17.66		17.06		18.28	
Variance Persistence												
From RV_u							0.9890		0.9889		0.9853	
From RV_d							0.9868		0.9843		0.9867	
From RV			0.9835		0.9847							
From Returns	0.9893											
Log Likelihoods												
Returns, RV_u, RV_d, and Options							150,103		150,431		150,702	
Returns, RV, and Options			96,523		96,930							
Returns and Options	49,690		49,368		50,495		49,253		49,166		49,270	
Returns	18,956		19,199		19,222		19,207		19,206		19,257	
Option Errors												
IVRMSE	6.299		5.855		5.849		5.837		5.829		5.827	
Ratio to GARCH	1.000		0.929		0.928		0.926		0.925		0.925	

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 April 24, 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 9: **Regressions of Model-Free on Model-Implied Variance Risk Premia**

	One-Factor Models						Tow-Factor Models					
	GARCH		ARV		SARV		GARV		CGSARV		GSARV	
Parameters	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
VRP												
1-Month												
Constant	9.20E-04	9.13E-04	-7.40E-03	6.94E-04	4.16E-04	6.27E-04	-1.96E-02	6.87E-04	-8.14E-03	5.48E-04	-3.74E-03	5.44E-04
Slope	-0.00	0.00	0.79	0.04	0.20	0.01	1.00	0.03	1.00	0.03	1.00	0.03
R^2 (%)	17.93		36.84		52.34		59.95		60.63		60.62	
3-Month												
Constant	4.44E-03	7.64E-04	6.09E-03	6.24E-04	6.75E-03	6.47E-04	-1.68E-02	1.40E-03	-5.47E-03	8.20E-04	-2.83E-03	7.11E-04
Slope	0.00	0.00	0.09	0.01	0.04	0.01	1.00	0.06	1.00	0.06	1.00	0.06
R^2 (%)	0.47		4.61		5.02		27.05		27.76		27.84	
6-Month												
Constant	5.57E-03	8.41E-04	1.07E-02	7.09E-04	1.04E-02	7.23E-04	-1.70E-02	1.60E-03	-3.91E-03	9.82E-04	-3.74E-03	9.74E-04
Slope	0.00	0.00	-0.01	0.01	-0.01	0.01	1.00	0.05	1.00	0.05	1.00	0.05
R^2 (%)	10.94		0.00		0.36		30.44		30.27		30.30	
VRP^u												
1-Month												
Constant							2.36E-04	5.14E-04	1.24E-03	5.22E-04	1.44E-03	5.18E-04
Slope							1.07	0.03	1.07	0.03	1.07	0.03
R^2 (%)							60.66		61.54		67.54	
3-Month												
Constant							-1.39E-03	3.44E-04	-3.23E-04	3.47E-04	-4.11E-04	3.45E-04
Slope							0.93	0.03	0.93	0.03	0.93	0.03
R^2 (%)							55.03		55.57		59.08	
6-Month												
Constant							-2.51E-03	3.10E-04	-1.24E-03	3.12E-04	-1.44E-03	3.11E-04
Slope							1.07	0.05	1.07	0.05	1.07	0.05
R^2 (%)							37.08		37.42		39.75	
VRP^d												
1-Month												
Constant							2.90E-03	1.67E-03	8.54E-03	1.44E-03	8.54E-03	1.05E-03
Slope							0.36	0.41	1.46	0.15	1.46	0.02
R^2 (%)							0.10		0.12		0.12	
3-Month												
Constant							-2.90E-03	1.88E-03	1.30E-02	1.84E-03	1.30E-02	1.84E-03
Slope							1.64	0.19	1.46	0.18	1.46	0.18
R^2 (%)							8.50		12.17		12.17	
6-Month												
Constant							2.56E-03	1.68E-03	1.32E-02	2.10E-03	1.32E-02	2.10E-03
Slope							0.36	0.04	0.54	0.04	0.54	0.04
R^2 (%)							9.79		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 (resp. 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 8 to generate model forecasts. The sample period spans January 02, 1990 through April 26, 2013.