PRICING S&P 500 INDEX OPTIONS: A CONDITIONAL SEMI-NONPARAMETRIC APPROACH

MASSIMO GUIDOLIN and ERWIN HANSEN*

We price S&P 500 index options under the assumption that the conditional risk-neutral density function of the index follows a Semi-Nonparametric (SNP) process with GARCH variance. The model is estimated combining a set of option contracts written on the index and the daily index return time series in the period 1996–2011. The in-sample and out-sample performance of the model is compared with several benchmark models, beating most of them. We conclude that a pricing model dealing simultaneously with non-normalities and time-varying volatility helps to mitigate the observed S&P 500 index option biases. © 2015 Wiley Periodicals, Inc. Jrl Fut Mark 36:217–239, 2016

1. INTRODUCTION

The weak empirical performance of the Black–Scholes (BS, 1973) model is well documented in the literature. In particular, the observed prices of out-of-the-money puts and in-the money calls are too high when compared with the prices delivered by the BS model (Bakshi, Cao, & Chen, 1997; Dumas, Fleming, & Whaley, 1998). This empirical regularity, known as volatility smile/smirk, is the main challenge to be addressed by any model attempting to provide a realistic fit to prices in option markets. The restrictive assumptions of normality of log-returns and constant volatility explain the unsatisfactory empirical performance of the BS model. In this paper, we introduce and evaluate the empirical performance of an option pricing model that relaxes these two assumptions. In particular, we assume that the conditional distribution of log-returns can be approximated by a Semi-Nonparametric function (SNP) and that the log-returns' volatility belongs to the GARCH class. The SNP density function takes the form of an Hermite polynomial expansion whose leading term is the normal density and whose polynomial coefficients may vary with time as well.

The SNP density function, originally proposed by Gallant and Tauchen (1989), is attractive for at least three reasons: first, Gallant and Nychka (1987) have shown that, under mild conditions, the SNP density function can consistently approximate the true density function of the process; second, it provides a characterization of the conditional density function, and therefore, it accounts for the full dynamic of the underlying stock return; and

---

Massimo Guidolin is at Bocconi University, Department of Finance and CAREFIN, Via Roentgen, 1 (2nd floor), 20136, Milano, Italy. Erwin Hansen is at Facultad de Economía y Negocios, Universidad de Chile, Diagonal Paraguay 257, Santiago, Chile. Hansen acknowledges financial support by Program U-Apoya, University of Chile. Authors thank valuable comments to a previous version of this article to Stuart Hyde, Mike Bowe, Carrie Na, and seminar participants at Manchester Business School's Finance seminar.

*Correspondence author, Facultad de Economía y Negocios, Universidad de Chile, Diagonal Paraguay 257, Santiago, Chile. Tel: +56-2-29772125, Fax: +56-2-2220639, e-mail:ehansen@unegocios.cl.

Received June 2014; Accepted April 2015

The Journal of Futures Markets, Vol. 36, No. 3, 217–239 (2016)
© 2015 Wiley Periodicals, Inc.
Published online 27 May 2015 in Wiley Online Library (wileyonlinelibrary.com).
DOI: 10.1002/fut.21731

third, it always produces positive probabilities, which assures that the density function is properly defined.[1]

The SNP model has been studied in asset pricing applications modeling stock returns,[2] but much less used for the specific purpose of option pricing. In the context of option valuation, the SNP model has also been studied by León, Mencía, and Sentana (2009), and used as benchmark model by Lim, Martin, and Martin (2005). Our work, however, differs from them in several dimensions. First, in these studies, the SNP model is combined with a constant volatility assumption, whereas we explore a richer SNP model with time-varying GARCH volatility and time-varying polynomial coefficients. Christoffersen, Heston, and Jacobs (2006), Christoffersen et al. (2009) and Christoffersen et al. (2010) argue that in order to reduce observed biases in option prices, both assumptions are necessary: non-normal distribution and time-varying volatility. Second, as suggested by Gallant and Tauchen (2008), we use information criteria to identify the best model, instead of selecting a pre-determined specification. Finally, we use a more extended sample of options covering 16 years in a weekly basis, making it more suitable to identify the time-series dynamic of the model.[3]

We estimate the parameters of the model using a non-linear least square (NLS) method that minimizes the weighted squared distance between option prices observed in the market and those prices delivered by our SNP pricing model. After assuming that the conditional risk-neutral density function of the model belongs to the SNP class with GARCH variance, we can use this calibration approach to retrieve the risk-neutral parameters of the model and compute the information criteria needed to identify the best model. We estimate and test the performance of the model using daily S&P 500 index option data for the period 1996–2011. As required, we combine return data with the panel of option contracts and expect that the use of both sources of information improves the efficiency of our estimates.[4] Our dataset is built aiming to cover the longest possible time period, as in this way we are able to better capture the dynamic of the conditional risk-neutral density function, but keep the computational burden in the estimation of the models at manageable levels. Thus, our dataset contains 15 contracts for each of the 772 Wednesdays available in the period from 1996 to 2011. These 15 contracts are the closest to the 5 predetermined levels of moneyness (0.96, 0.98, 1, 1.02, 1.04) and 3 maturity levels. The trade-off between the time coverage versus the cross-sectional coverage in option dataset has been discussed by Ait-Sahalia and Lo (1998) and García, Ghysels, and Renault (2010).

The option pricing performance of the model is evaluated, in-sample and out-of-sample, against 5 benchmark models: the standard Black–Scholes model, the implied volatility model (IVF) proposed by Dumas et al. (1998), an extended IVF model proposed by Gonçalves and Guidolin (2006), the SNP model with constant volatility study by León et al. (2009), and the volatility model with the short-run and long-run component model of Christoffersen

[1]Jondeau and Rockinger (2001) and Jackwerth (2004) point out that alternative models based on polynomial expansions, may not necessarily preserve the integrity of the probability distribution, producing negative probabilities in some intervals of the domain.

[2]Among the asset pricing applications are Gallant and Tauchen (1989, 1992), Gallant, Rossi, and Tauchen (1992, 1993), Gallant, Hansen, and Tauchen (1990), Gallant, Hsieh, and Tauchen (1991), Tauchen, Zhang, and Liu (1996), Harrison and Zhang (1999) and Zhang (2000). More recently, in the context of the Efficient Method of Moments (EMM) of Gallant and Tauchen (1996), the SNP model has been used extensively as an auxiliary model. See a full list of references in George Tauchen's EMM website.

[3]For example, Lim et al. (2005) fit their model to a sample of options in a single day.

[4]The NLS approach has been used by Engle and Mustafa (1992), Bakshi et al. (1997), Chernov and Ghysels (2000), Pan (2002), Christoffersen and Jacobs (2004a, 2004b), Christoffersen et al. (2008), among others.

et al. (2008). The performance of the SNP option pricing model with GARCH volatility is competitive when compared with the benchmark models. In-sample, the SNP pricing model outperforms all the benchmark models, reducing the biases of the Black–Scholes model by 70%. The out-of-sample performance of the models is evaluated using a rolling window setup, and two forecasting horizons: 1-week and 4-weeks. As with the evaluation in-sample, the SNP model produces reasonable prices out-of-sample. For the 1-week horizon, the first-stage SNP models outperform all benchmark models, but the component volatility model of Christoffersen et al. (2008); whereas for the 4-week horizon, the first-stage SNP model again produces lower pricing errors than any of the benchmark models considered. Overall, we provide evidence that an option model accounting for deviation of normality and time-varying volatility, produces a better empirical performance than other competitive option models. Our results show that both features of the conditional distribution of the S&P 500 index must be taken into account in order to produce reasonable prices for the contracts written on index.

The rest of the paper is organized as follows. In Section 2, the SNP model is introduced. In Section 3, the closed-form option price formula of the model is presented. The dataset is described in Section 4. The SNP-GARCH model is estimated in Section 5. In Section 6, we report the in-sample and the out-sample performance of the model. Finally, we conclude in Section 7.

2. THE SNP CONDITIONAL DENSITY

Let $\{y_t\}_{t=-L}^{T}$ be a time-series realization of variable of interest, in our case the S&P 500 index log-returns. Let $x_t \equiv (y_{t-1}, \ldots, y_{t-L})$ be a vector of L lagged values of y_t. Let $z_t \equiv (y_t - \mu_x)/\sigma_x$ be a local-scaled transformation of y_t, where μ_x and σ_x are the location and the scale function, respectively. The location function, which depends on L_u lags of y_t, is defined as $\mu_{x_t} = b_0 + \sum_{i=1}^{L_u} b_i y_{t-i}$, whereas the scale function, σ_{x_t}, is characterized by a GARCH process, which depends on $L = max(L_r, L_g)$ lags of $y_t : \sigma_{x_t} = c_0 + \sum_{i=1}^{L_r} c_i \left|y_{t-i} - \mu_{x_{t-i-1}}\right| + \sum_{i=1}^{L_g} d_i \sigma_{x_{t-i}}$.

The inclusion of a GARCH structure in volatility is justified for at least for two reasons: first, GARCH option models have been successful in reducing index option biases as shown by Christoffersen et al. (2004, 2008, 2010); and second, it reduces the likelihood of suffering of overparametrization problem due to large polynomial expansion.[5]

We assume that the distribution of z_t, conditional on its history x_t, belongs to the SNP class proposed by Gallant and Tauchen (1989)

$$h(z_t|x_t, \theta) = \frac{[P(z_t, x_t)]^2\, \phi(z_t)}{\displaystyle\int [P(s, x_t)]^2\, \phi(s)\, \mathrm{d}s}, \tag{1}$$

where $P(z_t, x_t)$ is an Hermite polynomial of order (K_z, K_x). The polynomial term is squared to assure positiveness, and the whole expression is scaled by a proportionality constant, which assures that the conditional density integrates to one. $\phi()$ is the probability density function of a normal random variable. Given that the first term of the polynomial expansion is the identity matrix, the first term of the SNP conditional density corresponds to a normal distribution. Higher polynomial terms, then, capture deviation from normality.[6]

[5] See Liu and Zhang (1998) for a discussion on overparametrization in the context of SNP models.

[6] As the polynomial order increases to infinity, the estimated conditional density function converges to the true conditional density function.

Following León et al. (2009), the polynomial term may take the form:

$$P(z_t, x_t) = \sum_{i=0}^{K_z} v_i(x_t) H_i(z_t). \quad (2)$$

Let $v(x_t) = (v_0(.)\ v_1(.) \ldots v_{K_z}(.))'$ be a time-varying vector such that $v(x_t)'v(x_t)$ defines the required proportionality constant $\int [P(s, z_t)]^2 \phi(s)ds$, and $H_i(z_t)$ is a normalized Hermite polynomial of order i.[7] Each element of the vector $v_i(x_t)$ takes the form:

$$v_i(x_t) = a_{0,i} + \sum_{k=1}^{Kx} \sum_{j=1}^{L_p} a_{k,j,i} y_{t-j}^k. \quad (3)$$

K_z and K_x are the two key parameters of the model: K_z defines to what extent the conditional density function deviates from normality, whereas K_x defines the extent to which this relation varies through time. If $K_x > 0$, $v(x_t)$ varies through time, and then, the full conditional distribution does. L_p is a parameter defining the number of lags to be included in $v(x_t)$.

As specified, the conditional SNP-GARCH density function has three possible channels of time dependency: the location function, μ_x, the scale function , σ_x, and the vector $v_i(x_t)$.

3. OPTION PRICING UNDER THE SNP CONDITIONAL DENSITY FUNCTION

3.1. Pricing

In this subsection, we derive a closed-form formula for pricing index option when log-returns are SNP distributed with GARCH volatility. The pricing formula derivation closely follows from León et al. (2009), so further details can be found in that article. We start by specifying the dynamic of log returns under the physical measure, then, an stochastic discount factor with exponential affine form is introduced such that the dynamic under the risk-neutral measure is recovered. Interestingly, the model form is the same under both measures.

Assume that, under the physical measure, P, the price of a stock at time $T > t$ can be written as a function of the current stock price S_t and the time to maturity $\tau = T - t$ as follows:

$$S_T = S_t \exp\left[\left(\mu_t - \frac{(\sigma_t)^2}{2}\right)\tau + \sigma_t\sqrt{\tau}z_t^*\right], \quad (4)$$

where μ_t and σ_t are the location and scale function defined above. The coefficients μ_t and σ_t represents the mean and the volatility of the log returns, $y_T \equiv log\left(\frac{S_T}{St}\right)$, conditional on the information set known at t. The random variable z_t^* is a standardized SNP variable with shape parameters v_t. It is assumed that $z_t = a(v_t) + b(v_t)x_t^P$, where x_t^P has an SNP distribution as well. The terms $a(v_t)$ and $b(v_t)$ are defined such that z_t has a zero mean and unit variance.[8] Under these assumptions, we have that $y_T = \delta_t + \lambda_t x^P$, where $\delta_t = \left(\mu_t - \frac{(\sigma_t)^2}{2}\right)\tau + \sigma_t\sqrt{\tau}a(v_t)$

[7] $H_i(z_t)$ is defined recursively as follows as $H_0(z_t) = 1$, $H_1(z_t) = z_t$ and $H_i(z_t) = \frac{z_t H_{i-1}(z_t) - \sqrt{i-1}H_{i-2}(z_t)}{\sqrt{i}}$ for $i \geqslant 2$.

[8] In particular, $a(v_t) \equiv \frac{-\mu_x'(1)}{\sqrt{\mu_x(2)}}$ and $b(v_t) \equiv \frac{1}{\sqrt{\mu_x(2)}}$, where $\mu_x'(n)$ and $\mu_x(n)$ are the uncentered and centered moments of order nth.

and $\lambda_t = \sigma_t b(v_t)$. Let $M_{t,T}$ be an exponential affine stochastic discount factor

$$M_{t,T} = \exp(\alpha_t y_T + \beta_t \tau). \tag{5}$$

León et al. (2009) establish the conditions under which the stochastic discount factor $M_{t,T}$ satisfies the required arbitrage-free conditions for pricing. Then, using standard results in option pricing based on the Radon–Nykodym derivative, the risk-neutral measure is defined as $f^{\mathbb{Q}}(y_T) = \exp(r_t\tau)M_{t,T} f^P(y_T)$.[9] If the asset price S_T is given by (4) under the real measure P, where the distribution of its log-return between t and T is a SNP of order K_Z with shape parameter v_t, the asset price under the risk-neutral measure $\mathbb{Q}$ is

$$S_T = S_t \exp\left[\left(\mu_t^{\mathbb{Q}} - \frac{\left(\sigma_t^{\mathbb{Q}}\right)^2}{2}\right)\tau + \sigma_t^{\mathbb{Q}}\sqrt{\tau}\kappa_t^*\right], \tag{6}$$

where

$$\begin{aligned} \mu^{\mathbb{Q}} &= \mu_t + \frac{\sigma_t^2}{2}\left(\frac{b(v_t)}{b(\theta_t)}\right)^2 - 1 + \frac{\sigma_t^2}{\sqrt{\tau}}\left[a(v_t) - a(\theta_t)\frac{b(v_t)}{b(\theta_t)}\right] + \alpha_t\sigma_t^2 b^2(v_t) \\ \sigma_t^{\mathbb{Q}} &= \sigma_t b(v_t)/b(\theta_t) \end{aligned}$$

and κ_t^* is a standardized SNP variable of order K_z with shape parameters $\theta_t = (\theta_{0t}, \ldots, \theta_{k_z t})'$ such that

$$\theta_{it} = \sum_{k=1}^{K_z} \frac{v_{kt}}{(k-i)!}\sqrt{\frac{k!}{i!}}(\alpha_t \lambda_{Pt})^{k-i}.$$

Finally, the option price is obtained using proposition 9 in León et al. (2009).

Proposition 1. The price at time t of a European call option with strike K written on the stock S_T defined by (6) under the risk neutral measure is

$$C_t^{SNP} = S_t P_{\mathbb{Q}_1}[x > d_t | I_t] - K\exp(-r_t\tau) P_{\mathbb{Q}}[x > d_t | I_t]$$

[9] If the conditional distribution of log-returns is SNP of order K_Z (5) satisfies the arbitrage-free conditions $E(M_{t,T}\exp(r_t\tau)) = 1$ and $E(M_{t,T}\exp(y_T)) = 1$ if and only if

$$\sum_{k=0}^{2K_Z} \frac{\gamma_k(v_t)(\alpha_t\lambda_t)^k}{\sqrt{k!}} = \exp\left[-\alpha_t\delta_t - \frac{1}{2}\alpha_t^2\lambda_t^2 - \beta_t\tau - r_t\tau\right]$$

and

$$\sum_{k=0}^{2K_Z} \frac{\gamma_k(v_t)(1+\alpha_t)^k\lambda_t^k}{\sqrt{k!}} = \exp\left[-(1+\alpha_t)\delta_t - \frac{1}{2}(1+\alpha_t)^2\lambda_t^2 - \beta_t\tau\right].$$

The unknowns α_t and β_t can be recovered solving these two equations.

where,

$$P_{\mathbb{Q}}[x > d_t | I_t] = \Phi(-d_t) + \phi(d_t) \sum_{k=1}^{2K_z} \frac{\gamma_k(\theta_t)}{\sqrt{k}} H_{t-1}(d_t),$$

$$P_{\mathbb{Q}_1}[x > d_t | I_t] = \exp(-r_t\tau + \delta_{\mathbb{Q}_t}) \sum_{k=0}^{2K_z} \gamma_k(\theta_t) I^*_{k,t},$$

$$I^*_{k,t} = \frac{1}{\sqrt{k}} \exp(\lambda_{\mathbb{Q}_t} d_t) H_{k-1}(d_t)\phi(d_t) + \frac{\lambda_{\mathbb{Q}_t}}{\sqrt{k}} I^*_{k-1,t},$$

$$I^*_{0,t} = \exp(\lambda^2_{\mathbb{Q}_t}/2)\Phi(\lambda_{\mathbb{Q}_t} - d_t),$$

$$\delta_{\mathbb{Q}_t} = \left(\mu_t^{\mathbb{Q}} - \frac{\left(\sigma_t^{\mathbb{Q}}\right)^2}{2}\right)\tau + a(\theta_t)\sigma_t^{\mathbb{Q}}\tau,$$

$$d_t = \frac{\log(K/S_t) - \delta_{\mathbb{Q}_t}}{\lambda_{\mathbb{Q}_t}}; \text{ and } \lambda_{\mathbb{Q}_t} = b(\theta_t)\sigma_t(\theta_t)\sqrt{\tau}.$$

$\phi()$ is the normal distribution function and $\Phi()$ is the cumulative normal distribution function. This closed-form price reduces to the Black–Scholes formula when the parameters K_z and K_x are set to zero. As long as K_z and K_x take positive values, the SNP conditional density captures complex underlying dynamics, including conditional heteroskedasticity, fat tails, skewness, among other empirical features.

3.2. Hedging

Even though our focus is on the pricing performance of the SNP model, a brief discussion about hedging is worth deserved, before moving forward.[10] It can be shown that the SNP option price (C_t^{SNP}) and the Black–Scholes option price (C^{BS}) are related through the equation

$$C_t^{SNP} = C_t^{BS} + \beta_{3t} sk_t + \beta_{4t}(ku_t - 3) + o(\sigma_{t,\tau}^{\mathbb{Q}\,2}),$$

where,

$$\beta_{3t} = \left(\frac{1}{3!}\right) S_t \sigma_{t,\tau}^{\mathbb{Q}} (\sigma_{t,\tau}^{\mathbb{Q}} - d_{1t})\phi(d_{1t}) + \left(\frac{1}{3!}\right) K \exp(-r_t\tau)\phi(d_{1t})\sigma_{t,\tau}^{\mathbb{Q}\,2},$$

$$\beta_{4t} = \left(\frac{1}{4!}\right) S_t \sigma_{t,\tau}^{\mathbb{Q}} (d_{1t}^2 - 3d_1\sigma_{t,\tau}^{\mathbb{Q}} - 1)\phi(d_{1t}),$$

$$d_{1t} = \frac{\log\left(\frac{S_t}{K}\right) + \left(r_t + \frac{\sigma_{t,\tau}^{\mathbb{Q}\,2}}{2}\right)}{\sigma_{t,\tau}^{\mathbb{Q}}\sqrt{\tau}}.$$

[10]We thanks a referee for suggesting us incorporating this point.

In this context, the delta hedging of the model is given by

$$\frac{\partial C_t^{SNP}}{\partial S_t} \simeq \frac{\partial C_t^{BS}}{\partial S_t} + \frac{\partial \beta_{3t}}{\partial S_t} sk_t + \frac{\partial \beta_{4t}}{\partial S_t}(ku_t - 3),$$

where sk_t is the conditional skewness and ku_t is conditional kurtosis.[11] This result shows that the SNP model's delta hedging converges to the BS model's delta hedging, $\Phi(d_{1t})$, when the underlying process is normal. Otherwise, given the magnitude of the derivatives, we would expect that the SNP model's delta hedge be slightly above BS model's delta hedge. The precise magnitude of the hedging should be clarified, however, in a more detailed empirical exercise that we left for future research.

4. DATA

We use a dataset of option contracts written on the S&P 500 index collected from Ivy Optionmetrics Database.[12,13] The data are available daily and covers the period from January 1996 to April 2011.

Following Ait-Sahalia and Lo (1998), and in order to cope with issues related to miss-recording and potential liquidity biases, we applied several filters on the raw data. In particular, we deleted contracts with: (i) implied volatilities higher than 70% (ii) option prices less than \$1/8, (iii) a time-to-maturity of less than 6 days and more than 365 days, and (iv) less than 10 transactions per day.[14] In order to deal with the unobservability of expected dividends paid by the index, we follow Ait-Sahalia and Lo (1998), which derived the implied expected dividend rate using both the spot-future parity and the put-call parity. We also drop those contracts violating standard arbitrage conditions in option markets.[15] The continuously compounded risk-free interest rate is proxied by a series of zero-coupon yield rates provided by Ivy Optionmetrics. We interpolate this series in order to match each of the maturities in the sample.

[11]The derivatives are

$$\begin{aligned}\frac{\partial \beta_{3t}}{\partial S_t} &= \left(\frac{1}{3!}\right)\sigma_{t,\tau}^{\mathbb{Q}\,2}\phi(d_{1t}) + \left(\frac{1}{3!}\right)S_t\sigma_{t,\tau}^{\mathbb{Q}\,2}\frac{\partial\phi(d_{1t})}{\partial d_{1t}}\frac{\partial d_{1t}}{\partial S} - \left(\frac{1}{3!}\right)\sigma_{t,\tau}^{\mathbb{Q}}d_{1t}\phi(d_{1t}) \\ &\quad - \left(\frac{1}{3!}\right)S_t\sigma_{t,\tau}^{\mathbb{Q}}\frac{\partial d_{1t}}{\partial S}\phi(d_{1t}) - \left(\frac{1}{3!}\right)S_t\sigma_{t,\tau}^{\mathbb{Q}}d_{1t}\frac{\partial\phi(d_{1t})}{\partial d_{1t}}\frac{\partial d_{1t}}{\partial S} \\ &\quad + \left(\frac{1}{3!}\right)K\exp(-r_t\tau)\frac{\partial\phi(d_{1t})}{\partial d_{1t}}\frac{\partial d_{1t}}{\partial S}\sigma_{t,\tau}^{\mathbb{Q}\,2},\end{aligned}$$

$$\frac{\partial \beta_{4t}}{\partial S_t} = \left(\frac{1}{4!}\right)\sigma_{t,\tau}^{\mathbb{Q}\,2}\left(\begin{array}{c}(d_{1t}^2 - 3d_1\sigma_{t,\tau}^{\mathbb{Q}} - 1)\phi(d_{1t}) + S_t\left(2d_{1t}\frac{\partial d_{1t}}{\partial S} - 3\sigma_{t,\tau}^{\mathbb{Q}}\frac{\partial d_{1t}}{\partial S}\right)\phi(d_{1t}) \\ +S_t(d_{1t}^2 - 3d_1\sigma_{t,\tau}^{\mathbb{Q}} - 1)\frac{\partial\phi(d_{1t})}{\partial d_{1t}}\frac{\partial d_{1t}}{\partial S}\end{array}\right),$$

and

$$\frac{\partial d_{1t}}{\partial S} = \frac{1}{S_t\sigma_{t,\tau}^{\mathbb{Q}}\sqrt{\tau}}.$$

[12]Ivy Optionmetrics database is a comprehensive dataset containing information for the entire US listed index and stocks options. The data are available online through the Wharton research data services (WRDS).

[13]S&P 500 index option data have been used by Baski et al. (1997), Dumas et al. (1998), Ait-Sahalia and Lo (1998), Heston and Nandi (2000), Ait-Sahalia, Wang, and Yared (2001), Christoffersen et. al. (2006, 2008, 2011), among others.

[14]There are many quotes in the original dataset with no transactions registered.

[15]In particular, we drop those contracts in which $c(S_t, X, \tau, r, \delta_{t,\tau}) \geq \max(S_t - K, S_t^* - Ke^{-r\tau})$ and $p(S_t, X, \tau, r, \delta_{t,\tau}) \geq \max(K - S_t, Ke^{-r\tau} - S_t^*)$, where S_t^* is the net of dividend implied spot index.

TABLE I
Summary Statistics for S&P 500 Index's Option Contracts

Variable	*Mean*	*Standard Deviation*	*Median*	*Min.*	*Max.*	*Skew.*	*Kurt.*
Option price (C, $)	36.5	24.7	33.5	0.1	184.5	0.7	3.3
Implied volatility (σ, %)	20.7	7.6	19.6	6.9	69.5	1.4	6.3
Days to maturity (τ)	57.8	44.1	51	7	338	2.0	9.8
S&P 500 index price (S, $)	1,140.4	210.0	1,152.2	592.1	1,561.7	−0.4	2.6
S&P 500 return (%)	0.1	4.1	0.3	−23.8	25.9	−0.3	8.4
Strike price (K, $)	1,149.7	216.4	1,160	600	1,675	−0.3	2.5
Interest rate (r, %)	3.5	2.2	3.9	0.2	7.0	−0.2	1.5
Number of observations	10,829						

Note. Our sample covers the period January, 1996–April, 2011. It includes, for each of the 772 Wednesdays in the sample, the 5 closest contracts to the pre-defined moneyness levels 0.96, 0.98, 1, 1.02, 1.04, for the three maturity levels with less missing values. Option prices are calculated as the average value of the highest closing bid price and the lowest closing ask price. S&P 500 index has been adjusted by discounting future expected dividends (see the details of the adjustment procedure in the main text). Interest rates are the continuously compounded zero-coupon yield rates provided by Optionmetrics, and interpolated to match the maturities in the dataset.

Despite being available at daily frequency, our analysis only considers contracts for one day of the week, Wednesday, reflecting the choice made in other studies. See for example Dumas et al. (1998) and Heston and Nandi (2000) and Christoffersen et al. (2008, 2010) and Christoffersen, Heston, and Jacobs (2011). Three reasons justify this choice. First, and most importantly, it reduces the number of contracts in the sample, facilitating computation. Second, Wednesday is the day of the week less likely to be a holiday. Third, it is less likely to observe biases associated with the day-of-the-week effect or the weekend effects on Wednesdays, as it would be the case in Mondays and Fridays, for example. For each Wednesday, we select the closest contracts to 5 predetermined levels of moneyness (0.96, 0.98, 1, 1.02, 1.04), defined as the ratio between the strike price of the contract and the future price of the S&P 500 index, and 3 levels of maturities; those with less missing values. Thus, our final sample contains 15 contracts in each of the 772 Wednesdays in the sample.

It is worth noting that our sample selection procedure aims to covering the longest time span in order to better capture the dynamic of conditional risk-neutral density function, and therefore, be able to produce accurate option price estimates. This choice, however, comes at the cost of losing cross-sectional representativeness. Ait-Sahalia and Lo (1998) and Garcí a et al. (2010) discuss about the trade-off between covering the time dimension versus the cross-sectional dimension in option prices dataset.

Table I reports summary statistics for our final sample.[16] Table II reports the average option price and the average implied volatility by moneyness and maturity levels in our sample.

5. ESTIMATION OF THE SNP MODEL

In this section, we describe the non-linear least squared procedure to estimate the risk-neutral density function, we identify the best specification for the SNP model using the searching strategy proposed by Gallant and Tauchen (2008), and then, we report the estimates of our selected model.

[16]Missing values explain the difference between the theoretical and the effective number of observations in our sample.

TABLE II
Selected Statistics by Maturity and Moneyness

	0.96	*0.98*	*1*	*1.02*	*1.04*	*Total*
			Average option price ($)			
Short-term	57.8	39.7	25.2	15.2	8.9	29.4
Medium-term	75.2	57.9	45.1	36.4	27.0	48.9
Long-term	102.9	90.3	77.3	66.3	60.1	80.0
Total	64.3	46.4	32.4	22.6	15.4	36.5
			Average implied volatility (%)			
Short-term	23.5	21.4	19.9	19.0	18.3	20.4
Medium-term	22.9	21.3	20.5	21.1	20.4	21.3
Long-term	23.8	23.7	22.9	23.4	23.1	23.4
Total	23.3	21.4	21.1	19.7	19.0	20.7

Note. Short-term includes those contracts expiring up to 60 days, medium-term includes those expiring between 61 and 180 days, and long-term includes those expiring between 181 and 365 days. Moneyness is defined as the ratio between the strike price and the future index price.

5.1. Estimation of the Risk-Neutral Density Function

Under the assumption of absence of arbitrage opportunities in the market, the price of an option contract is the discounted value, under the risk-neutral density function, of the expected payoffs at maturity. Using this result, we can estimate the parameters of the risk-neutral conditional density function by minimizing the difference between a set of observed option prices in the market and the theoretical prices associated with the model. Thus, the estimated coefficients, $\widehat{\theta}$, are such that

$$\widehat{\theta}_1 = \arg\min_{\theta} \frac{1}{TN} \sum_{t=1}^{T} \sum_{i=1}^{N} \left(C_{i,t} - C_{i,t}^{SNP}(\theta) \right)^2 , \tag{7}$$

where T is the number of days in our sample and N is the total number of contracts per day (in our case, $T = 772$ and $N = 15$). $C_{i,j}$ is the option price observed in the market and $C_{i,j}^{SNP}$ is the theoretical option price delivered by the SNP model, for a contract with the same characteristics. The subscript 1 in $\widehat{\theta}_1$ stands for first stage estimates. Second-stage estimates, $\widehat{\theta}_2$, comes from minimizing the following expression

$$\widehat{\theta}_2 = \arg\min_{\theta} \frac{1}{TN} \sum_{t=1}^{T} \sum_{i=1}^{N} \left(C_{i,t} - C_{i,t}^{SNP}(\theta) \right) \widehat{\Omega}^{-1} \left(C_{i,t} - C_{i,t}^{SNP}(\theta) \right) , \tag{8}$$

where $\widehat{\Omega} \equiv e(\widehat{\theta}_1)e(\widehat{\theta}_1)'$ is the variance–covariance matrix of the first-stage pricing errors, $e(\widehat{\theta}_1)$.

This estimation procedure relies on two assumptions : the use of squared $ pricing errors in (7), and the use of $\widehat{\Omega}$ in (8). Regarding the first assumption, $ pricing errors have the advantage that they are computational stable, are easy to interpret and, due to our sample selection procedure, they do not suffer of the heterosckedasticity problem associated to the use of relatively wide range of option prices across maturity and moneyness described by Christoffersen and Jacobs (2004a, 2004b). On the other hand, the estimation is biased towards option contracts with high valuations (long-run contracts). This problem, however, is mitigated to some extent as more than 60% of the contracts in our data belong to the short-term category (less than 60-days).

Regarding the second assumption, the use of $\widehat{\Omega}$ is justified by the efficiency gains in the parameter estimation associated to the fact that more accurate priced contracts, in the first stage, are overweighted. Our procedure resembles the standard two-stage GMM estimation process. Huang and Wu (2004) argue that this procedure depends on the exact model being estimated, and therefore, it could be problematic to some extent. As an alternative, Huang and Wu (2004) propose to approximate the variance of the pricing error with the variance of the option price, as a percentage of the index level. However, this approximation is exact in processes in which the conditional return distribution does not vary over time. In our application, however, the SNP-GARCH distribution is time varying, and therefore, the use of this alternative approach may not be suitable. Overall, bear in mind that our estimates should be interpreted taking into consideration these limitations.[17]

5.2. Selection of the Best Specification

Gallant an Tauchen (2008) propose an identification strategy that combines a particular search path and the use of information criteria. The suggested search path first identifies the number of lags in the location equation, L_u, which is increased until the autoregressive process for the mean is obtained. Second, keeping fixed the optimal level of L_u from the first stage, the lags in the variance equation (L_r and L_g) are increased until the best (G)ARCH process is identified. Third, keeping fixed the optimal values of the previous two stages, the number of Hermite polynomials, K_z, is increased to find the best semi-nonparametric specification. Finally, K_x and L_p are increased, determining if a fully nonlinear specification is called for. Gallant and Tauchen (2008) emphasize that this model selection strategy does not necessarily produce the overall preferred model, according to a particular information criterion, however, it has been successful in finding reasonable approximations of the true density function in previous applications using the SNP method.

The best model specification is selected by minimizing the Bayesian Information Criterion (BIC), defined as:

$$\mathrm{BIC} = s(\widehat{\theta}) + \frac{1}{2}\frac{n_\theta}{NT}\log(NT),$$

where $s(\widehat{\theta})$ is the minimized squared option pricing errors that define the optimal $\widehat{\theta}$ in (7) and (8), respectively; n_θ is the number of model parameters and NT is the total number of observations. Note that we identify the best model using the estimated parameters retrieved from option contracts. In Table III, we reports our results. For the sake of completeness, we also report the Akaike information criterion, however, our selection is based on the BIC as it provides the more parsimonious model. In this way, we reduce a possible overfitting problem. We report first-stage and second-stage AIC and BIC.

The preferred SNP model is

$$L_u = 0, L_r = 1, L_g = 1, K_z = 4, K_x = 2, L_p = 1,$$

with first-stage and second-stage BIC of 1.8914 and 1.9151, respectively. The selected model exhibits a particularly rich dynamic, characterized by a 4th order Hermite polynomial expansion, and an additional time-varying channel given by $K_x = 2$.

[17] We thank a referee for bringing these points into the discussion.

TABLE III
Best SNP Model According to Information Criteria Implied by Options

L_u	L_r	L_g	K_z	K_x	L_p	θ	AIC (1)	BIC (1)	AIC (2)	BIC (2)
0	0	0	0	0	0	2	2.1803	2.1804	2.1849	2.1850
1	0	0	0	0	0	3	2.1806	2.1812	2.1852	2.1858
0	1	1	0	0	0	4	1.9222	1.9229	1.9308	1.9315
0	1	1	1	0	0	5	1.9028	1.9037	1.9309	1.9318
0	1	1	2	0	0	6	1.9013	1.9029	1.9361	1.9372
0	1	1	3	0	0	7	1.9007	1.9020	1.9381	1.9394
0	1	1	4	0	0	8	1.9006	1.9021	1.9378	1.9393
0	1	1	5	0	0	9	1.9008	1.9025	1.9388	1.9405
0	1	1	4	1	1	13	1.8906	1.8930	1.9238	1.9262
0	**1**	**1**	**4**	**2**	**1**	**18**	**1.8880**	**1.8914**	**1.9217**	**1.9151**
0	1	1	4	3	1	23	1.8891	1.8934	1.9216	1.9259
0	0	0	1	0	0	3	2.1823	2.1829	2.3924	2.3930
0	0	0	2	0	0	4	2.1787	2.1794	2.2464	2.2471
0	0	0	3	0	0	5	2.1786	2.1795	2.2389	2.2399
0	0	0	4	0	0	6	2.1805	2.1812	2.2362	2.2373
0	0	0	5	0	0	7	2.1788	2.1799	2.2178	2.2191
0	0	0	6	0	0	8	2.1784	2.1799	2.2240	2.2255

Note. L_u is the number of lags in the mean equation, L_r is the number of lags in the ARCH part of the variance, L_g is the number of lags in the GARCH part of the variance, K_z is the degree of the normalized Hermite polynomial in the SNP model, K_x is the degree of the polynomial defining the structure of the integration function v, and L_p is the number of lags in the integration function v. θ is the total number of parameters to be estimated. $AIC(1)$ and $AIC(2)$ are the Akaike information criterion obtained from the first and second stage SNP estimators, respectively. $BIC(1)$ and $BIC(2)$ are the Schwarz information criterion obtained from the first and second stage SNP estimators, respectively.

To identify the best model, Gallant and Tauchen (2008) advise to start the search with $K_z = 4$.This advice is adhered to and we do not stop the search at $K_z = 3$ for example, where the BIC is slightly lower (1.9020) than the BIC of the model with $K_z = 4$ (1.9021).

The model has a constant mean ($L_u = 0$) and the variance is characterized by a GARCH(1,1) structure. The gains associated to the inclusion of the GARCH(1,1) term are clear when we compared, for example, the models with and without this term, in the case of $L_u = 0$ and $K_z = 4$. The BICs are 1.9021 and 2.1812, respectively. It is important, from a empirical point of view, to validate the inclusion of the GARCH variance in our model in order to confirm the documented success of GARCH option pricing models in the literature. See for example, Heston and Nandi (2000) and Christoffersen et al. (2006, 2008, 2011). Because the SNP-GARCH model encompasses the simple GARCH (1,1) option model, we would expect further performance gains of this model.

5.3. NLS Estimates of the Preferred SNP Model

In Table IV, we report the NLS estimates of the prefer model. In the first column, we report first-stage estimates, whereas in the second column we report second stage estimates. The preferred model contains 18 parameters. The estimated mean coefficient is close to zero and not significant.[18] In the GARCH(1,1) variance equation, the constant term is again close to

[18]The statistical significance of the parameters is verified using the inverse of the outer product of the gradient vector, evaluated at the minimized value, as standard errors.

TABLE IV
Estimated Coefficients of the Preferred SNP Model
$(L_u = 0, L_r = 1, L_g = 1, K_z = 4, K_x = 2, L_p = 1)$

		First stage	*Second stage*
Mean equation	b_0	0.0003	0.0003
Variance equation	c_0	0.0006	0.0007
	c_1	0.0521*	0.0701*
	d_1	0.8956*	0.8552*
Hermite polynomial	a_1	3.4348*	3.5525*
	a_2	−0.7931*	−0.8336*
	a_3	−0.0043*	−0.0057*
	a_4	−0.2559*	−0.2652*
	a_5	−1.4585*	−1.5203*
	a_6	−3.5*e*-05	0.4200
	a_7	0.1065*	0.1138*
	a_8	0.2347*	0.1068*
	a_9	−0.0052*	−0.0080*
	a_{10}	−0.3252*	−0.3392*
	a_{11}	−1.9484*	−1.9319*
	a_{12}	0.0058*	0.0072*
	a_{13}	0.2153*	0.2317*
	a_{14}	0.7406*	0.8078*

Note. The coefficients are estimated minimizing the squared option pricing errors produced by the preferred model. In the first stage, pricing errors are equally weighted in the minimization process, whereas in the second stage, the pricing errors are weighted by the inverse of variance–coveriance matrix estimated in the first stage. *Significant at 95%.

zero and insignificant, but the ARCH and the GARCH terms are both statistically significant. The variance is stationary as the sum of these two coefficients is less than 1. The coefficients of the Hermite polynomial are all significant, except for the sixth coefficient, indicating that the model is statistical valid.

Considering that the estimated coefficients of the polynomial expansion do not necessarily have an economic meaning on their own, a plot of the estimated conditional risk-neutral density function backed by these coefficients is more informative. In Figure 1 we report the 1-period ahead conditional density, using the estimates from the first-stage estimator. As a reference, we also plot a normal density function with the same mean and variance. The figure shows that the estimated conditional SNP density function has thicker tails than the normal distribution, indicating that the SNP model assigns higher probabilities to extreme events than the normal distribution does. Furthermore, we observe that the estimated density has negative skewness, indicating that negative events are more likely to occur than positive ones. Thus, the SNP model provides a more realistic estimate of the conditional density function of the S&P 500 log-returns than the normal case.

6. OPTION PRICING RESULTS

In this section, we study the pricing performance of the SNP model identified and estimated in the previous section. In particular, we compare the in-sample and the out-of-sample fit of the SNP option pricing model against several benchmark models. The out-of-sample evaluation is completed using a rolling window framework, for two forecasting horizons: 1 week and 4 weeks ahead.

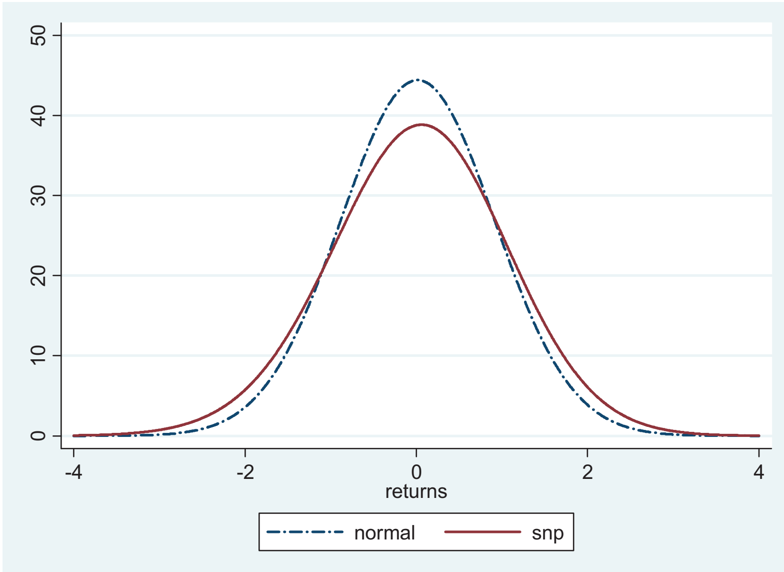

FIGURE 1
Estimated Conditional 1-Period Ahead SNP Distribution

Note. The figure shows the estimated conditional 1-period ahead distribution of the preferred SNP model ($L_u = 0, L_r = 1, L_g = 1, K_z = 4, K_x = 2, L_p = 1$). A normal distribution with the same mean and variance is superimposed as reference (dashed line). [Color figure can be viewed in the online issue, which is available wileyonlinelibrary.com.]

6.1. Benchmark Models

We study 5 benchmark models: the Black–Scholes model; the implied volatility function (IVF) model of Dumas et al. (1998), sometimes referred as practitioner Black-Sholes; an extended version of the IVF model proposed by Gonçalves and Guidolin (2006); a constant-volatility version of the SNP model studied by León et al. (2009); and finally, a conditional volatility model with short-run and long-run components, proposed by Christoffersen et al. (2008).

6.1.1. Implied volatility function model

The IVF model is an ad-hoc procedure that smooths BS implied volatilities across moneyness, maturities, and cross-product or squared values of these two terms. The procedure, originally introduced by Dumas et al. (1998), has shown remarkable success in fitting the cross-section of option prices; surpassing many alternatives specifications proposed in the literature. The model does not require the imposition of any additional assumptions either on the distribution of the stock index or on the preferences of the investors. Its ability to produce accurate option prices forecasts, in a relatively straightforward and simple way, is its main advantage.

The estimation procedure consists on running an OLS regression between the (log of) implied volatilities and moneyness (M_i), [19] time-to-maturity (τ_i), an additional terms

[19]Dumas et al. (1998) use a time-adjusted moneyness level, defined as $\frac{\left(\frac{K}{S^*}-1\right)}{\sqrt{T}}$, where K is the strike price, S^* is the dividend-adjusted spot price and T is the time to maturity.

involving these two variables. In particular, we use the following functional form:

$$\ln(\sigma_i) = \beta_0 + \beta_1 M_i + \beta_2 M_i^2 + \beta_3 \tau_i + \beta_4 (M_i \tau_i) + \varepsilon_i.$$

The estimated model is used to recover the fitted prices for each of the contracts in the dataset.

6.1.2. *Gonçalves and Guidolin (2006)'s model*

Gonçalves and Guidolin (2006) propose an extended IVF model, explicitly recognizing that the implied volatility surface is time-varying. Thus, for each day in the sample, the estimation procedure consists of building a time-series of daily OLS beta estimates, $\{\beta_t\}_{t=1}^{772} = \{(\beta_{0,t}, \beta_{1,t}, \beta_{2,t}, \beta_{3,t}, \beta_{4,t})'\}_{t=1}^{772}$. Then, a VAR model to the estimated betas is fitted:

$$\widehat{\beta}_t = \mu + \sum_{J=1}^{P} \Phi \widehat{\beta}_{t-j} + u_t, \tag{9}$$

where $u_t \sim i.i.d\ N(0, \Omega)$. The estimated dynamic model is used to forecast betas, which are used later to price options.

6.1.3. *SNP model with constant volatility*

León et al. (2009) study the performance of a SNP model with constant volatility to price S&P 500 index options. In particular, they identified the following model as the best performer model in their study: $L_u = 0, L_r = 0, L_g = 0, K_z = 4, K_x = 0, L_p = 0$. We study this model in order to verify if a richer specification produces more accurate option pricing estimates.[20]

6.1.4. *Volatility model with short-run and long-run components*

Christoffersen et al. (2008) introduce an option valuation model with long-run and short-run volatility components. This model constitutes a good benchmark because is a type of GARCH option pricing model that outperforms other well-known GARCH models, like the Heston and Nandi (2000)'s model. The model builds on Engle and Lee (1999) and Heston and Nandi (2000). The model, under the risk-neutral measure, is characterized by the following return dynamics in discrete time:

$$\begin{aligned} R_{t+1} &\equiv \ln(S_{t+1}) - \ln(S_t) = r - \frac{1}{2} h_{t+1} + \sqrt{h_{t+1}} z_{t+1}^* \qquad z_{t+1}^* \sim N(0, 1) \\ h_{t+1} &= q_{t+1} + \widetilde{\beta^*}(h_t - q_t) + \alpha v_{1,t}^* \\ q_{t+1} &= \omega + \rho^* q_t + \varphi v_{2,t}^*, \end{aligned} \tag{10}$$

where $(h_t - q_t)$ is the short run component, q_t is the long run component and $v_{i,t}^*$, for $i = \{1, 2\}$, are mean-zero innovations. This specification assures that the log-return of the asset equals the risk-free interest rate. Moreover, the variance of the log-returns under both

[20] León et al. (2009) also study SNP models with $K_z = 2$ and $K_z = 3$. Both of them produce a worst pricing performance than the SNP model with $K_z = 4$.

measures coincides. Under this structure, the price at time t of a European call option with strike price K that expires at time T is

$$C^{CJOW} = S_t \left(\frac{1}{2} + \frac{1}{\pi} \int_0^{\infty} \mathrm{Re} \left[\frac{K^{-i\phi} f^*(t, T; i\phi + 1)}{i\phi S_t e^{r(T-t)}} \right] d\phi \right)$$
$$-Ke^{-r(T-t)} \left(\frac{1}{2} + \frac{1}{\pi} \int_0^{\infty} \mathrm{Re} \left[\frac{K^{-i\phi} f^*(t, T; i\phi)}{i\phi} \right] d\phi \right), \quad (11)$$

where $f^*(t, T; i\phi)$ is the conditional characteristic function of the $\ln(S_t)$ under the risk-neutral measure. The function $Re(\cdot)$ takes the real value of the argument. Christoffersen et al. (2008) show that the characteristic function takes the form $f(t, T; \phi) = S_t^{\phi} \exp(A_t + B_{1,t}(h_{t+1} - q_{t+1}) + B_{2,t})$, where A_t, $B_{1,t}$ and $B_{2,t}$ are recursive coefficients[21].

6.2. Measures of Statistical Fit

To assess the forecasting ability of the fitted models (in-sample and out-of-sample), we compute 2 standard statistical measures: the root mean squared error (RMSE) and the mean absolute error (MAE). For a sample of option contracts for T days, containing N contracts per day, these measures are defined as

$$RMSE = \sqrt{\frac{1}{TN} \sum_{t=1}^{T} \sum_{i=1}^{N} (C_{t,i} - \widehat{C}_{t,i}^{Model}(\widehat{\theta}))^2}$$

$$MAE = \frac{1}{TN} \sum_{t=1}^{T} \sum_{i=1}^{N} \left| C_{t,i} - \widehat{C}_{t,i}^{Model}(\widehat{\theta}) \right|,$$

where $C_{t,i}$ is the market price of the option contract i at day t, and $\widehat{C}_{t,i}^{Model}$ is the price delivered by a particular option pricing model for a contract with the same characteristics (time-to-maturity, strike price, underlying stock price) as the one observed in the market.

6.3. In-Sample Fit Results

In Table V, we report the relative in-sample performance of the models, with respect to the Black–Scholes model. In particular, each value in the table corresponds to the ratio between the performance measure of a particular model and the performance of the BS model. A number smaller than 1 indicates that the model has a better performance than the BS model. In the limit, a ratio of 0, would indicate that the model produces a perfect fit. The first result is that both the first-stage and the second-stage SNP model considerably outperform the BS model. For example, in the case of the first-stage SNP model, the RMSE is 0.31, and the MAE is 0.25. These numbers indicate that the SNP model only produces 30% and 25% of the mispricing of the BS model. In other words, it reduces the relative in-sample biases by approximately 69% and 75%. Similar results are obtained for the second-stage SNP model, where the RMSE and MAE are 0.32 and 0.27, respectively.

[21]The exact definition of each of these recursive coefficients is given by Equation (25) in Christoffersen et al. (2008).

TABLE V
In-Sample Fit of the SNP Model and Benchmark Models

SNP (First stage)	RMSE	0.31
	MAE	0.25
SNP (Second stage)	RMSE	0.32
	MAE	0.27
Implied volatility function	RMSE	0.31
	MAE	0.34
Goncalves and Guidolin (2006)	RMSE	0.40
	MAE	0.38
León et al. (2009) SNP Kz = 4	RMSE	0.44
	MAE	0.39
Christoffersen et al. (2008)	RMSE	0.32
	MAE	0.27

Note. The table shows the relative in-sample performance of several option pricing models and the Black–Scholes (BS) model. In each cell, the ratio between the performance measures of the evaluated model and the BS model is reported. RMSE is the root mean squared error and MAE is the mean absolute error. SNP is the preferred Semi-nonparametric model with GARCH (1,1) variance introduced in this study. In particular, it corresponds to the case ($L_u = 0$, $L_r = 1$, $L_g = 1$, $K_z = 4$, $K_x = 2$, $L_p = 1$). The implied volatility function (IVF) model corresponds to the model introduced by Dumas et. al (1998), where in a first stage, the following regression is run, $\ln(\sigma_i) = \beta_0 + \beta_1 M_i + \beta_2 M_i^2 + \beta_3 \tau_i + \beta_4 (M_i \tau_i) + \varepsilon_i$, and in a second stage, options are priced using the predicted implied volatilities obtained from the model. In the above specification, σ_i is the implied volatility, M_i is a time-adjusted measure of moneyness, τ_i is the time to maturity in days, and ε_i is an error term. Goncalves and Guidolin (2006)'s model extends the Dumas et al. (1998) model, using a VAR model in the second stage to modeling the dynamic of the estimated beta coefficients, and then use this model to price options. León et al. (2009) is a SNP model, which is a simplified version of our preferred model. In particular, it corresponds to the case ($L_u = 0$, $L_r = 0$, $L_g = 0$, $K_z = 4$, $K_x = 0$, $L_p = 0$). Christoffersen et. al (2008) is an option pricing model with long-run and short-run volatility components. In particular, it corresponds to the persistent component model described in the paper.

A second result is that the first-stage SNP model is the best performing model among all the studied models. In terms of RMSE, only the IVF model has the same performance, 0.31. In terms of MAE, it delivers the higher value, 0.25. Other models are competitive as well, especially the Christoffersen et al. (2008) model. However, the SNP model is still slightly superior. It is worth noting that the SNP-GARCH model outperforms the constant variance SNP model of León et al. (2009). This result corroborates that the Hermite polynomial expansion must be combined with a time-varying volatility in order to produce accurate option prices estimates. Christoffersen et al. (2006) made the same point from a theoretical perspective. Finally, the performance of the IVF is reasonable as well, with RMSE of 0.31 and MAE of 0.34.

In Table VI, we report the in-sample performance estimates by moneyness and maturity levels. Our results show that both the first-stage and second-stage SNP models produce reasonable in-sample performance across all moneyness and maturity levels, when compared with the benchmark models. For example, for short-term contracts with moneyness of 0.96 (in-the money calls and out-of-the-money puts), the RMSE are 0.36 in both SNP models, 0.57 for the IVF model, 0.46 for the Gonçalves-Guidolin model, and 0.53 for the constant-volatility SNP model. Only the Christoffersen et al. (2008) model produces a better performance than the SNP model, with a RMSE of 0.28. Similar results are obtained for at-the money medium-term contracts. The SNP models produce RSME of 0.24 and 0.26, respectively. These numbers compared with values of 0.34, 0.35 and 0.27 for the Gonçalves and Guidolin (2006)'s model, the constant-volatility SNP model, and the Christoffersen et al. (2008) model, respectively. In this case, only the IVF model produces a superior performance, with a RMSE of 0.19.

TABLE VI
In-Sample Fit of the SNP Model and Benchmark Models by Maturity and Moneyness

			Moneyness				
			0.96	*0.98*	*1*	*1.02*	*1.04*
SNP (First stage)	Short-term	RMSE	0.36	0.29	0.26	0.39	0.51
		MAE	0.31	0.24	0.19	0.31	0.45
	Medium-term	RMSE	0.30	0.24	0.24	0.33	0.41
		MAE	0.26	0.21	0.19	0.24	0.32
	Long-term	RMSE	0.25	0.20	0.25	0.28	0.46
		MAE	0.21	0.17	0.18	0.19	0.29
SNP (Second stage)	Short-term	RMSE	0.36	0.30	0.27	0.40	0.53
		MAE	0.33	0.26	0.21	0.33	0.49
	Medium-term	RMSE	0.31	0.27	0.26	0.33	0.42
		MAE	0.27	0.23	0.21	0.26	0.35
	Long-term	RMSE	0.29	0.23	0.28	0.29	0.46
		MAE	0.25	0.19	0.21	0.22	0.32
Implied volatility function	Short-term	RMSE	0.57	0.42	0.32	0.44	0.64
		MAE	0.63	0.44	0.31	0.48	0.83
	Medium-term	RMSE	0.28	0.23	0.19	0.22	0.29
		MAE	0.28	0.22	0.19	0.22	0.31
	Long-term	RMSE	0.21	0.17	0.16	0.16	0.18
		MAE	0.19	0.16	0.14	0.15	0.18
Goncalves and Guidolin (2006)	Short-term	RMSE	0.46	0.38	0.31	0.39	0.51
		MAE	0.50	0.36	0.27	0.38	0.54
	Medium-term	RMSE	0.41	0.37	0.34	0.36	0.48
		MAE	0.39	0.35	0.31	0.35	0.48
	Long-term	RMSE	0.50	0.41	0.45	0.44	0.55
		MAE	0.46	0.37	0.41	0.43	0.56
León et al. (2009) SNP Kz = 4	Short-term	RMSE	0.53	0.44	0.37	0.53	0.67
		MAE	0.49	0.39	0.30	0.45	0.65
	Medium-term	RMSE	0.45	0.37	0.35	0.45	0.56
		MAE	0.39	0.33	0.30	0.39	0.50
	Long-term	RMSE	0.32	0.28	0.30	0.38	0.58
		MAE	0.27	0.24	0.24	0.30	0.44
Christoffersen et al. (2008)	Short-Term	RMSE	0.28	0.26	0.24	0.35	0.46
		MAE	0.24	0.22	0.19	0.29	0.43
	Medium-term	RMSE	0.28	0.28	0.27	0.34	0.42
		MAE	0.25	0.24	0.23	0.29	0.38
	Long-term	RMSE	0.42	0.36	0.37	0.40	0.48
		MAE	0.40	0.32	0.35	0.39	0.46

Note. The table presents the relative in-sample performance of several option pricing models for 3 maturity categories: short-term (<60 days), medium-term (61–180 days) and long-term (181–365 days), and 5 moneyness categories, where moneyness is defined as the ratio between the strike price of the contract and the future price of the S&P 500 index. A description of the models is given in Table V.

Again, we observe that the SNP model with GARCH(1,1) volatility produces better results than the constant volatility SNP model, across all the moneyness and maturity levels considered. When we compare the performance of the IVF model and the Gonçalves and Guidolin (2006)'s model, we observe that the Gonçalves and Guidolin (2006)'s model produces better in-sample results than the IVF model for short-term contracts. This result is reversed in medium-term contracts, however. Finally, these results confirm that the Christoffersen et al. (2008) model is the most challenging benchmark. This model out-

TABLE VII
Out-of-Sample Fit of the SNP Model and Benchmark Models, 1 Week and 4 Weeks Ahead

		1 week	*4 weeks*
SNP (First stage)	RMSE	0.62	0.66
	MAE	0.58	0.63
SNP (Second stage)	RMSE	0.72	0.77
	MAE	0.68	0.74
Implied volatility function	RMSE	1.00	1.11
	MAE	0.92	0.99
Goncalves and Guidolin (2006)	RMSE	0.92	0.99
	MAE	0.80	0.96
León et al. (2009) SNP Kz = 4.	RMSE	0.98	0.98
	MAE	0.98	0.99
Christoffersen et al. (2008)	RMSE	0.51	0.89
	MAE	0.47	0.69

Note. The table reports the relative out-of-sample performance of several option pricing models with respect to the BS model, 1 week and 4 weeks ahead. The out-of-sample performance measures of the models is computed using a 5-year rolling window setup. A description of the models is given in Table V.

performs the SNP model in the case of short-term contracts. However, the SNP produces smaller RMSE and MAE in some medium-term categories.

6.4. Out-of-Sample Fit Results

In the previous subsection, we have shown that the SNP option pricing model has good in-sample properties. However, we aim at identifying a model also able to produce good option price forecasts, out-of-sample. A model producing reasonable forecasts could be used for investment and hedging purposes. We turn to that evaluation here.

We evaluate the models using a 5-year rolling window framework. The first window, containing 260 Wednesdays, starts in January 1996 and ends in December 2000. We estimate the models in this window, and then we compute the 1-week and 4-week forecasts. Thereafter, we move the window 1 period forward, and we re-estimate the models and compute the forecasts again. We continue moving the window forward until the end of our sample. We repeat this process 508 times. Finally, the performance measures are computed as averages across windows.

In Table VII, we report our out-of-sample results. For the 1-week forecasting horizon, the SNP model produces a RMSE of 0.62 and a MAE of 0.58. The SNP model outperforms all the benchmark models but Christoffersen et al. (2008) model. For example, the IVF model produces RMSE and MAE of 1.00 and 0.92, respectively. Because this evaluation is relative to the performance of the BS model, these numbers indicate that the IVF model performance is quite similar to the BS performance, out-of-sample. The Gonçalves-Guidolin (2006) model produces slightly better results than the IVF model, but worse than the SNP model, with RMSE and MAE of 0.92 and 0.80, respectively. The constant-volatility SNP model also produces poor results when compared with our SNP model. The model with long-run and short-run volatility components of Christoffersen et al. (2008) shows the best performance at the 1-week forecasting horizon. Its RMSE is 0.51 and its MAE is 0.47, well below the numbers of the SNP model.

TABLE VIII
Out-of-Sample Fit of the SNP Model and Benchmark Models by Maturity and Moneyness

(1 week ahead)

			Moneyness				
			0.96	*0.98*	*1*	*1.02*	*1.04*
SNP (First stage)	Short-term	RMSE	0.61	0.59	0.59	0.62	0.65
		MAE	0.69	0.58	0.53	0.54	0.55
	Medium-term	RMSE	0.63	0.63	0.58	0.60	0.65
		MAE	0.69	0.64	0.55	0.52	0.54
	Long-term	RMSE	0.68	0.62	0.68	0.64	0.67
		MAE	0.74	0.62	0.61	0.62	0.63
SNP (Second stage)	Short-term	RMSE	0.68	0.68	0.68	0.71	0.73
		MAE	0.79	0.68	0.67	0.68	0.69
	Medium-term	RMSE	0.69	0.71	0.65	0.67	0.71
		MAE	0.78	0.73	0.63	0.62	0.64
	Long-term	RMSE	0.72	0.69	0.73	0.75	0.78
		MAE	0.79	0.67	0.71	0.75	0.80
Implied volatility function	Short-term	RMSE	0.89	0.80	0.77	0.70	0.78
		MAE	1.05	0.88	0.78	0.71	0.76
	Medium-term	RMSE	1.00	0.98	0.96	0.92	1.13
		MAE	1.08	1.01	0.95	0.88	1.00
	Long-term	RMSE	1.92	1.27	1.65	1.16	2.30
		MAE	1.62	1.27	1.48	1.14	1.90
Goncalves and Guidolin (2006)	Short-term	RMSE	0.72	0.74	0.68	0.60	0.62
		MAE	0.93	0.82	0.73	0.65	0.64
	Medium-term	RMSE	0.73	0.83	0.78	0.74	0.76
		MAE	0.87	0.91	0.84	0.80	0.82
	Long-term	RMSE	0.82	0.89	1.03	1.14	1.03
		MAE	0.97	1.04	1.17	1.30	1.36
León et al. (2009) SNP Kz = 4	Short-term	RMSE	0.97	1.00	1.00	1.00	1.02
		MAE	1.07	1.03	0.97	0.96	1.00
	Medium-term	RMSE	0.92	0.97	0.97	1.00	1.03
		MAE	0.95	0.97	0.96	0.99	1.02
	Long-term	RMSE	0.74	0.79	0.84	0.92	0.98
		MAE	0.84	0.83	0.85	0.97	1.01
Christoffersen et al. (2008)	Short-term	RMSE	0.23	0.26	0.26	0.28	0.35
		MAE	0.25	0.26	0.27	0.29	0.38
	Medium-term	RMSE	0.25	0.26	0.34	0.31	0.33
		MAE	0.27	0.27	0.33	0.30	0.34
	Long-term	RMSE	0.38	0.40	0.51	0.47	0.49
		MAE	0.41	0.40	0.50	0.45	0.50

Note. The table presents the relative out-of-sample performance, 1-week ahead, of several option pricing models for 3 maturity categories: short-term (<60 days), medium-term (61–180 days) and long-term (181–365 days), and 5 moneyness categories, where moneyness is defined as the ratio between the strike price of the contract and the future price of the S&P 500 index. The out-of-sample performance measure of the models is computed using a 5-year rolling window setup. A description of the models is given in Table V.

The forecasting evaluation at the 4-week horizon confirms the ability of the SNP model in producing reasonable out-of-sample forecasts. A first point to note is that, as expected, the forecasting ability of the models is reduced at longer horizons: consistently across models, the 4-week forecast performance is worse than the 1-week forecasting performance. In relative terms, we observe that the SNP model is quite competitive when compared with the

TABLE IX
Out-of-Sample Fit of the SNP Model and Benchmark Models by Maturity and Moneyness

(4 weeks ahead)

			Moneyness				
			0.96	*0.98*	*1*	*1.02*	*1.04*
SNP (First stage)	Short-term	RMSE	0.70	0.67	0.68	0.72	0.73
		MAE	0.72	0.62	0.58	0.58	0.58
	Medium-term	RMSE	0.69	0.69	0.65	0.66	0.72
		MAE	0.74	0.69	0.60	0.57	0.59
	Long-term	RMSE	0.68	0.62	0.64	0.68	0.77
		MAE	0.76	0.64	0.61	0.65	0.67
SNP (Second stage)	Short-Term	RMSE	0.77	0.76	0.77	0.80	0.80
		MAE	0.84	0.73	0.73	0.74	0.73
	Medium-term	RMSE	0.75	0.77	0.74	0.74	0.80
		MAE	0.83	0.79	0.71	0.69	0.72
	Long-term	RMSE	0.70	0.65	0.69	0.75	0.82
		MAE	0.77	0.60	0.62	0.70	0.78
Implied volatility function	Short-term	RMSE	1.01	0.92	0.86	0.82	0.95
		MAE	1.11	0.94	0.86	0.80	0.87
	Medium-term	RMSE	0.97	0.99	0.97	1.01	1.40
		MAE	1.07	1.02	0.98	0.94	1.07
	Long-term	RMSE	2.38	2.15	1.68	2.78	2.93
		MAE	2.13	1.78	1.57	2.02	2.18
Goncalves and Guidolin (2006)	Short-term	RMSE	0.79	0.80	0.74	0.73	0.74
		MAE	1.01	0.90	0.81	0.76	0.75
	Medium-term	RMSE	0.77	0.86	0.83	0.80	0.81
		MAE	0.91	0.94	0.90	0.85	0.86
	Long-term	RMSE	0.96	0.95	1.00	1.20	1.04
		MAE	1.06	1.05	1.07	1.32	1.36
León et al. (2009) SNP Kz = 4	Short-term	RMSE	0.98	1.00	1.00	1.00	1.02
		MAE	1.07	1.03	0.97	0.96	1.00
	Medium-term	RMSE	0.92	0.96	0.97	1.00	1.04
		MAE	0.95	0.97	0.96	0.99	1.02
	Long-term	RMSE	0.75	0.79	0.82	0.93	0.99
		MAE	0.85	0.84	0.83	0.97	1.03
Christoffersen et al. (2008)	Short-term	RMSE	0.40	0.47	0.46	0.68	0.68
		MAE	0.48	0.50	0.49	0.63	0.67
	Medium-term	RMSE	0.98	0.78	1.27	0.88	0.83
		MAE	1.08	0.89	1.27	0.91	0.93
	Long-term	RMSE	0.66	0.69	0.68	0.74	0.52
		MAE	0.89	0.87	0.84	0.91	0.76

Note. The table presents the relative out-of-sample performance, 4-week ahead, of several option pricing models for 3 maturity categories: short-term (between the strike price of the contract and the future price of the S&P 500 index. The out-of-sample performance measure of the models is computed using a 5-year rolling window setup. A description of the models is given in Table V.

benchmark models. Actually, at this forecast horizon, the performance of the SNP model is superior to the one of Christoffersen et al. (2008) model. The SNP model produces RMSE and MAE of 0.66 and 0.63, whereas the Christoffersen et al. (2008) model produces values of 0.89 and 0.69. Regarding the other models, we obtain similar results to those for the 1-week forecasting horizon. The SNP model outperforms the IVF model, the Gonçalves-Guidolin (2006) model, and the constant-volatility SNP model. Also, we observed that the first-stage

SNP estimates yield better out-of-sample results than second-stage SNP estimates at both forecasting horizons.

Similar to the in-sample evaluation, we evaluate the out-of-sample performance of the models across moneyness and maturity categories. Our results are reported in Tables VIII and IX. For the 1-week horizon, we find that the first-stage SNP model outperforms, across moneyness and maturity levels, the IVF model, the Gonçalves-Guidolin (2006) model, and the constant volatility SNP model. With regards to this last model, the evidence favors a more general specification in order to obtain reliable option prices forecasts. As before, we also find that the first-stage SNP model is superior to the second-stage SNP model. The performance results of Christoffersen et al. (2008) model are particularly good and outperform the SNP model. The best results of this model are obtained in the short-term category. For the 4-week forecasting horizon, we obtain similar results to those of the 1-week horizon, but now, the SNP model outperforms the Christoffersen et al. (2008) model for medium-term and long-term contracts. For short-term contracts, we observe the same pattern reported above, where the SNP model is outperformed by the Christoffersen et al. (2008) model. With regards to the other models, we again obtain that first-stage SNP estimates outperform the second-stage SNP model, the IVF model, Gonçalves-Guidolin (2006) model, and the constant-volatility model, across moneyness and maturity levels.

7. CONCLUSIONS

This study evaluates the empirical performance of a semi-nonparametric model with time-varying volatility in pricing S&P 500 index option contracts. In particular, we assume that the conditional risk-neutral density function of the index is characterized by an Hermite polynomial expansion and by a GARCH structure that accounts for time-varying variance of log returns. Furthermore, the model allows for time-varying polynomial coefficients, incorporating an extra channel of time dynamic.

The empirical performance of the model is evaluated in-sample and out-of-sample against 5 benchmark models in the period 1996–2011. In-sample, the first-stage SNP model produces the best results among all the studied models, reducing the Black–Scholes biases by 70%, approximately. Out-of-sample, the results are promising as well. For the 1-week horizon, the first-stage SNP model outperforms all the benchmark models except the long-run and short-run volatility model of Christoffersen et al. (2008). For the 4-week horizon, again, the SNP model outperformed all the competing models. This evidence highlights the importance of incorporating both non-normalities and time-varying volatilities simultaneously in pricing index options as suggested by Christoffersen et al. (2006).

For future research, it would be interesting to evaluate the SNP-GARCH model using alternative measures of performance such as hedging measures, to evaluate the performance of investment strategies with options, using stock options instead of index options, using alternative benchmark models like option GARCH with jumps models, among others.

REFERENCES

Aït-Sahalia, Y., & Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. The Journal of Finance, 53, 499–547.

Ait-Sahalia, Y., Wang, Y., & Yared, F. (2001). Do option markets correctly price the probabilities of movement of the underlying asset? Journal of Econometrics, 102, 67–110.

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. The Journal of Finance, 52, 2003–2049.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. The Journal of Political Economy, 81, 637–654.
Chernov, M., & Ghysels, E. (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. Journal of Financial Economics, 56, 407–458.
Christoffersen, P., Dorion, C., Jacobs, K., & Wang, Y. (2010). Volatility components, affine restrictions, and non-normal innovations. Journal of Business and Economic Statistics, 28, 483–502.
Christoffersen, P., Elkamhi, R., Feunou, B., & Jacobs, K. (2009). Option valuation with conditional heteroskedasticity and nonnormality. Review of Financial Studies, 23, 2139–2183.
Christoffersen, P., Heston, S., & Jacobs, K. (2006). Option valuation with conditional skewness. Journal of Econometrics, 131, 253–284.
Christoffersen, P., Heston, S., & Jacobs, K. (2011). A GARCH option model with variance-dependent pricing kernel. Working paper SSRN 1538394. Available from: http://ssrn.com/abstract=1538394.
Christoffersen, P., & Jacobs, K. (2004a). The importance of the loss function in option valuation. Journal of Financial Economics, 72, 291–318.
Christoffersen, P., & Jacobs, K. (2004b). Which GARCH model for option valuation? Management Science, 50, 1204–1221.
Christoffersen, P., Jacobs, K., Ornthanalai, C., & Wang, Y. (2008). Option valuation with long-run and short-run volatility components. Journal of Financial Economics, 90, 272–297.
Dumas, B., Fleming, J., & Whaley, R. (1998). Implied volatility functions: Empirical tests. Journal of Finance, 53, 2059–2106.
Engle, R., & Mustafa, C. (1992). Implied ARCH models from options prices. Journal of Econometrics, 52, 289–311.
Gallant, A., & Nychka, D. (1987). Semi-nonparametric maximum likelihood estimation. Econometrica, 55, 363–390.
Gallant, A., Rossi, P., & Tauchen, G. (1992). Stock prices and volume. Review of Financial Studies, 5, 199–242.
Gallant, A., Rossi, P., & Tauchen, G. (1993). Nonlinear dynamic structures. Econometrica, 61, 871–907.
Gallant, A., & Tauchen, G. (1989). Seminonparametric estimation of conditionally constrained heterogeneous processes: Asset pricing applications. Econometrica, 57, 1091–1120.
Gallant, A., & Tauchen, G. (1996). Which moments to match? Econometric Theory, 12, 657–681.
Gallant, A., & Tauchen, G. (2008). SNP: A program for nonparametric time series analysis. User's guide. Working paper, Duke Economic Department. Available from: http://public.econ.duke.edu/pub/arg/compecon/src/cases/snp/snpman/snpguide.pdf.
Gallant, A. R., Hansen, L. P., & Tauchen, G. (1990). Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. Journal of Econometrics, 45, 141–179.
Gallant, A. R., & Tauchen, G. (1992). A nonparametric approach to nonlinear time series analysis: Estimation and simulation. New Directions in Time Series Analysis, 46, 71–92.
Gallant, R., Hsieh, D., & Tauchen, G. (1991). On fitting a recalcitrant series: The pound/dollar exchange rate, 1974–1983. In W. Barnett, J. Powell, and G. Tauchen (Eds.), Nonparametric and semiparametric methods in econometrics and statistics (pp. 199–240). Cambridge: Cambridge University Press.
García, R., Ghysels, E., & Renault, E. (2010). The econometrics of option pricing. Handbook of Financial Econometrics, 1, 479–552.
Gonçalves, S., & Guidolin, M. (2006). Predictable dynamics in the S&P 500 index options implied volatility surface. Journal of Business, 79, 1591–1635.
Harrison, P., & Zhang, H. (1999). An investigation of the risk and return relation at long horizons. Review of Economics and Statistics, 81, 399–408.
Heston, S., & Nandi, S. (2000). A closed-form Garch option valuation model. Review of Financial Studies, 13, 585–625.
Huang, J.-Z., & Wu, L. (2004). Specification analysis of option pricing models based on time-changed lévy processes. The Journal of Finance, 59, 1405–1440.
Jackwerth, J. (2004). Option-implied risk-neutral distributions and risk aversion. CFA Institute Research Foundation of AIMR Publications, 1, 1–86.
Jondeau, E., & Rockinger, M. (2001). Gram-Charlier densities. Journal of Economic Dynamics and Control, 25, 1457–1483.
Lee, E., & Engle, R. (1999). A permanent and transitory component model of stock return volatility. In R. Engle and H. White (Eds.), Cointegration causality and forecasting a festschrift in honor of Clive W.J. Granger (pp. 475–497). Oxford University Press.
León, Á., Mencía, J., & Sentana, E. (2009). Parametric properties of semi-nonparametric distributions, with applications to option valuation. Journal of Business and Economic Statistics, 27, 176–192.
Lim, G., Martin, G., & Martin, V. (2005). Parametric pricing of higher order moments in S&P 500 options. Journal of Applied Econometrics, 20, 377–404.

Liu, M., & Zhang, H. (1998). Overparameterization in the seminonparametric density estimation. Economics Letters, 60, 11–18.

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

Tauchen, G., Zhang, H., & Liu, M. (1996). Volume, volatility, and leverage: A dynamic analysis. Journal of Econometrics, 74, 177–208.

Zhang, H. (2000). Explaining bond returns in heterogeneous agent models: The importance of higher-order moments. Journal of Economic Dynamics and Control, 24, 1381–1404.