



On practitioners closed-form GARCH option pricing

Sharif Mozumder^a, Bart Frijns^{b,*}, Bakhtear Talukdar^c, M. Humayun Kabir^d

^a Department of Mathematics, University of Dhaka, Bangladesh

^b Open Universiteit, the Netherlands

^c College of Business and Economics, University of Wisconsin Whitewater, USA

^d School of Economics & Finance, Massey University, Palmerston North, New Zealand

ARTICLE INFO

Keywords:

Heston-Nandi model
Market price of risk
Hedging
European-style options

ABSTRACT

This paper proposes a practitioner version of Heston and Nandi's (2000) (HN) model, which we term the Practitioner's Heston Nandi, or PHN model. We compare the option pricing and hedging performance of the PHN model vis-à-vis the HN model. Instead of using a one-period ahead volatility forecast for all options used in calibrations at any given time, the PHN model proposes using forward-looking *ad-hoc* volatilities (implied by market option prices) for each individual option and maturity in calibration and hedging. Since the proposed PHN model uses only option price data, it renders historical stock price data redundant, cutting the data requirement in derivative valuation. We employ options traded at CBOE for the period January 1, 2016 to December 31, 2018 and show that the proposed PHN model yields quick calibration and significantly improves pricing and hedging for European-style options.

1. Introduction

Following the work of Engle (1982) and Bollerslev (1986), numerous studies have developed Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models for volatility estimation and forecasting, as well as option pricing. Many studies have examined the empirical option pricing performance of specific GARCH models.¹ One of the benchmark models in many of these studies is that of Heston and Nandi (2000, HN hereafter). HN presents a closed-form option pricing model assuming normally distributed returns, a linear risk premium, and the same GARCH parameters for historical and pricing asset returns. The HN model provides good in- and short-term out-of-sample performance, reduces the required calibration time, and eliminates a counterfactual market price of risk interpretation. Using S&P 500 index options, HN shows that their GARCH model has smaller out-of-sample valuation errors than the *ad hoc* Black-Scholes Practitioner's Black-Scholes (PBS hereafter) model based on deterministic volatility functions (DVF)

proposed by Dumas, Fleming, and Whaley (1998). This is a remarkable result given that the PBS model is updated every period, while the HN-GARCH model parameters are held constant, and volatility is obtained from historical data. HN claims that any model outperforming the PBS model of Dumas et al. (1998) is compelling and rich.² However, it is unclear why a model with a joint estimation based on historical stock returns and prospective option price information would consistently outperform a model with purely implied measures based on prospective option prices alone. The HN implementation requires the calculation of volatility using historical return data in each iteration in the GARCH framework.³ However, this historical volatility estimate is backward-looking.

In this paper, we propose a Practitioner's version of the HN model (PHN hereafter), where using historical return data to compute volatility becomes redundant. The dynamics of the multi-period ahead density estimation in our model requires calculating a one-day ahead implied volatility using DVF mapping (see Dumas et al., 1998). Using implied

* Corresponding author.

E-mail addresses: sumozumder@du.ac.bd (S. Mozumder), bart.frijns@ou.nl (B. Frijns), talukdam@uww.edu (B. Talukdar), m.h.kabir@massey.ac.nz (M.H. Kabir).

¹ See Amin and Ng (1993), Engle and Mustafa (1992), Duan (1995, 1996), Hardle and Hafner (2000), Christoffersen and Jacobs (2004a, 2004b), Kim and Kim (2004), Christoffersen, Heston, and Jacobs (2006), Barone-Adesi, Engle, and Mancini (2008), Christoffersen, Jacobs, Ornthanalai, and Wang (2009), Huang and Hansen (2017), Augustyniak and Badescu (2021), Cao, Badescu, Cui, and Jayaraman (2021), Liang and Wang (2021), Liu, Jiao, & Guo (2022), Tong, Hansen, and Huang (2022), Wang, Cheng, Yin, and Yu (2022), Oh and Park (2023), Cheng, Chang, Lo, and Tsai (2023)

² In their terms a flexible but theoretically inconsistent model may appear to be better than their theoretically consistent GARCH model in terms of in-sample fit.

³ This is in line with HN's conjecture that it is possible that the index's history provides information about the future beyond the information contained in option prices.

volatility eliminates one parameter from being estimated, making historical return data redundant and the process more parsimonious. That parsimony makes our practitioner version more efficient than the traditional HN model.⁴

Using S&P 500 index options data from the Chicago Board Options Exchange (CBOE) from January 1, 2016, to December 31, 2018, we find that the PHN model calibrated with prospective market information in option prices significantly outperforms the HN and PBS models, both in- and out-of-sample. Moreover, the calibration of the PHN model is much faster than that of the HN model. Although the PBS model occasionally outperforms the unrestricted version of the HN model,⁵ this only occurs once for our PHN model. The source of the enhanced performance of the PHN model, even with fewer data points used in calibrations compared to the HN model, is that it incorporates option-wise multi-period ahead market volatilities in calibrations. In contrast, HN use a one-period ahead stock volatility for all options in calibrations. For the same reason, the hedging exercises of the PHN model also yield better Delta and Delta-Gamma hedging for options positions across moneyness and maturities.

The remainder of the paper is organized as follows. We briefly revisit the HN and PBS models and introduce our PHN model in Section 2. In Section 3, we present the data and discuss data cleaning and calibration issues. We also discuss in- and out-of-sample pricing results. Section 4 outlines our hedging exercises and discusses the hedging results. We conclude in Section 5.

2. Model

This section revisits the HN and PBS model structure, intuition, and pricing formulae. We subsequently introduce our PHN model and discuss its differences from the HN model.

2.1. The Heston and Nandi (2000) model

Following Heston and Nandi (2000), the log spot price, $s(t)$, follows a GARCH process, with time steps of length Δ ,

$$\log(s(t)) = \log(s(t-\Delta)) + r + \lambda h(t) + \sqrt{h(t)}z(t) \quad (1)$$

$$h(t) = \omega + \sum_{p=1}^P \beta_p h(t-p\Delta) + \sum_{q=1}^Q \alpha_q \left(z(t-q\Delta) - \gamma_q \sqrt{h(t-q\Delta)} \right)^2 \quad (2)$$

where r is the continuously compounded interest rate for time interval Δ , and $z(t)$ has a standard normal distribution. $h(t)$ is the conditional variance of the log return between $t-\Delta$ and t , based on the information set at $t-\Delta$ and appears in the mean equation as a return premium. This means the average spot return depends on the risk level, λ . The spot return is assumed to exceed the risk-free rate by an amount proportional to the variance, $h(t)$. Since volatility equals the square root of $h(t)$, this implies that the return premium per unit of risk is proportional to the square root of $h(t)$. As α_q and β_p approach zero, the variance becomes constant, and the model is equivalent to the Black-Scholes model

⁴ In the pricing model, HN assume a market price of risk parameter (λ) of -0.5 to characterize the time-varying moment-generating function of the stock's return distribution at maturity, while in the estimation, λ is given different values even though the GARCH dynamics assume a constant $h(t+\Delta)$. Empirically, λ has been found to vary widely. If we consider replacing GARCH forecast $h(t+\Delta)$ by the DVF of Dumas et al. (1998), we can avoid dual value assignment of λ to get a single model price. This also reduces the number of estimated parameters by one. Moreover, by replacing $h(t+\Delta)$ depending on maturity as well as strike prices with *ad-hoc* estimates, we can avoid the constancy of $h(t+\Delta)$ for all options and maturities.

⁵ The PBS model outperforms the unrestricted version of the HN model and non-updated version of HN often being outperformed by BS model.

observed at discrete time intervals.

When $P = Q = 1$, the process is stationary with finite mean and variance if $\beta_1 + \alpha_1 \gamma_1^2 < 1$. In that model, we can directly observe $h(t+\Delta)$ at time t , as a function of the spot price as follows:

$$h(t+\Delta) = \omega + \beta_1 h(t) + \alpha_1 \frac{(\log(s(t)) - \log(s(t-\Delta))) - r - \lambda h(t) + \gamma_1 (h(t))^2}{h(t)} \quad (3)$$

where, α_1 determines the kurtosis of the distribution. $\alpha_1 = 0$ implies a deterministic time-varying variance, and γ_1 captures the asymmetric impact of shocks, $z(t)$. In general, the conditional variance and spot return correlate as follows,

$$\text{Cov}_{t+\Delta} [h(t+\Delta), \log(s(t))] = -2\alpha_1 \gamma_1 h(t). \quad (4)$$

Since $\alpha_1 > 0$ and $\gamma_1 > 0$ in most estimations, the correlation between spot returns and variance is negative. Moreover, since γ_1 represents the skewness of the log return distribution, the distribution becomes symmetric if γ_1 and λ are equal to zero.

In a risk-neutral framework, we can write Eqs. (1) and (2) as

$$\log(s(t)) = \log(s(t-\Delta)) + r + \lambda h(t) + \sqrt{h(t)}z^*(t) \quad (5)$$

$$h(t) = \omega + \sum_{p=1}^P \beta_p h(t-p\Delta) + \sum_{q=1}^Q \alpha_q \left(z^*(t-q\Delta) - \gamma_q \sqrt{h(t-q\Delta)} \right)^2 + \alpha_1 \left(z^*(t-q\Delta) - \gamma_1^* \sqrt{h(t-q\Delta)} \right)^2 \quad (6)$$

with

$$z^*(t) = z(t) + \left(\lambda + \frac{1}{2} \right) \sqrt{h(t)} \quad (7)$$

$$\gamma_1^* = \gamma_1 + \lambda + \frac{1}{2} \quad (8)$$

where $z^*(t)$ is standard normally distributed under risk-neutral probabilities. For $z^*(t)$ to have a standard normal risk-neutral distribution, HN assume that an option's value with one period to expiration obeys the Black-Scholes-Rubinstein formula. We obtain the risk-neutral version in Eqs. (1) and (2) by replacing $\lambda = -1/2$ and γ_1 with $\gamma_1^* = \gamma_1 + \lambda + 1/2$. The conditional moment generating function of the asset price to produce option values, $f(\phi)$, for the GARCH process of Eqs. (1) and (2) takes the log-linear form,

$$f(\phi) = s(t)^\phi \exp \left\{ A(t; T, \phi) + \sum_{p=1}^P B_p(t; T, \phi) h(t+2\Delta-p\Delta) + \sum_{q=1}^{Q-1} C_q(t; T, \phi) \left(z(t+2\Delta-q\Delta) - \gamma_q \sqrt{h(t+2\Delta-q\Delta)} \right)^2 \right\} \quad (9)$$

where $A(t; T, \phi)$ and $B(t; T, \phi)$ are given by the recursive relations

$$A(t; T, \phi) = A(t+\Delta; T, \phi) + \phi r + B_1(t+\Delta; T, \phi) \omega - \frac{1}{2} \ln(1 - 2\alpha_1 B_1(t+\Delta; T, \phi)) \quad (10)$$

$$B(t; T, \phi) = \phi(\lambda + \gamma_1) - \frac{1}{2} \gamma_1^2 + \beta_1 B_1(t+\Delta; T, \phi) + \frac{1/2(\phi - \gamma_1)^2}{1 - 2\alpha_1 B_1(t+\Delta; T, \phi)} \quad (11)$$

For a single lag, $P = Q = 1$, the coefficients can be calculated recursively from the terminal conditions $A(T; T, \phi) = 0$ and $B(T; T, \phi) = 0$. We obtain the risk-neutral version by substituting $\lambda = -1/2$ and replacing γ_1 with $\gamma_1^* = \gamma_1 + \lambda + 1/2$. Since the moment-generating function of the spot price is the moment-generating function of the log spot price, $f(i\phi)$ is the characteristic function of the log spot price. To use the characteristic function, we note that ϕ in Eqs. (10) and (11) must be replaced by $i\phi$ everywhere. One can calculate risk-neutral probabilities

following Feller (1971) or Kendall and Stuart (1977) by inverting the characteristic function.

HN examine the empirical performance of S&P 500 index options for a GARCH model with a single lag ($P = Q = 1$). They set $\Delta = 1$ and use daily index returns to model volatility. There are two versions of the GARCH model. In the restricted/symmetric GARCH model where the skewness parameter, γ_1 , is constrained to zero, and the unrestricted/asymmetric GARCH model, where γ_1 is a free parameter. Rather than using maximum likelihood estimations (as in Bollerslev, 1986) to estimate model parameters and plugging the estimates into the options valuation formula to compute option values, HN use a non-linear least squares (NLS) procedure that minimizes the difference between model and observed option prices.⁶

2.2. Practitioner's Black-Scholes (PBS) model

Dumas et al. (1998) propose the Deterministic Volatility Function (DVF) framework inspired by the implied binomial trees of Derman and Kani (1994), Dupire (1994), and Rubinstein (1994), which considers the local volatilities of the underlying asset's returns as a deterministic function of asset price and time to expiration. Dumas et al. (1998) propose the following DVF:

$$\sigma^{iv} = a_0 + a_1 K + a_2 K^2 + a_3 T + a_4 T^2 + a_5 KT \quad (12)$$

where, σ^{iv} is the BS implied volatility, and $a_0, \dots, a_5$ are model parameters. Implied volatility is a function of the strike price, K , time to expiration, T , quadratic strike price, K^2 , quadratic time to expiration, T^2 and the interaction between the strike price and time to maturity, KT . They introduce a threshold of 0.01 to eliminate possible negative values of fitted volatility in the DVF model:

$$\widehat{\sigma}^{iv} = \max(0.01, a_0 + a_1 K + a_2 K^2 + a_3 T + a_4 T^2 + a_5 KT) \quad (13)$$

Minimizing the sum of squared errors between market and fitted option prices allows each model parameter to be easily estimated.

Dumas et al. (1998) show that the BS model leads to the largest valuation error compared to the last three DVF terms since volatility varies across moneyness and maturity. They find that the volatility function, including K, K^2, T, KT shows much better improvement over the BS model. There is a minor improvement when squared maturity, T^2 , is introduced. Practitioners generally smooth implied volatility across strike prices and value options using the smoothed volatility to consider the sneer patterns in BS implied volatilities. Dumas et al. (1998) label this procedure using the DVF option valuation model as the *ad hoc* approach. Christoffersen and Jacobs (2004b) also employ this model and refer to it as the Practitioner's Black-Scholes (PBS) model.

2.3. Practitioner's Heston Nandi (PHN) model

In our proposed Practitioner's HN (PHN) model, asset returns follow similar dynamics as in HN over the contract's life. The DVF coefficients do not change in the HN model, but the calibration process updates the structural parameters through optimization iterations. As a result, the DVF coefficients can be entered *ex-ante* into the calibration, similar to how time-to-maturity, strike price, and volatility are provided. The HN implementation requires the calculation of volatility using historical 252-day return data in each iteration for each week. In our practitioner's approach, the calculation of volatility using historical returns becomes redundant.⁷ The dynamics of the multi-period ahead density estimation

in our proposed model requires the calculation of the one-day ahead implied volatility by applying Dumas et al. (1998) DVF mapping. Using DVF to determine implied volatility reduces the number of parameters that need to be estimated by one. This is the main source of the reduced time requirement when implementing our proposed practitioner's version and gives the practitioner's version an efficiency edge over the traditional HN model.

Formally, in the HN model, when $P = Q = 1$, the conditional moment generating function (9) becomes:

$$f(\varphi) = s_t^q \exp\{A(t; t+T, \varphi) + B(t; t+T, \varphi)h_{t+1}\} \quad (14)$$

and the price of a call option with strike price, K , and maturity, T , is obtained in closed form as:

$$c_{HN} = s_t \left(\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \text{Re} \left[\frac{K^{-i\varphi} f^*(i\varphi + 1)}{i\varphi f^*(1)} \right] d\varphi \right) - K e^{-rT} \left(\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \text{Re} \left[\frac{K^{-i\varphi} f^*(i)}{i\varphi} \right] d\varphi \right) \quad (15)$$

where f^* is the risk-neutral version of f , and $A(t; t+T, \varphi)$ and $B(t; t+T, \varphi)$ are obtained from the recursive relations (10) and (11) (see HN for details). Here, h_{t+1} is constant with respect to $f(\varphi)$ as in all GARCH models. In other words, h_{t+1} is a single number for all options used in calibration. In contrast, for PHN, h_{t+1} varies across options used in calibration following Eq. (13). Since h_{t+1} does not come from a GARCH process, the parameters α_1, β_1 , and γ_1 can be calibrated with forward-looking information contained in option prices alone. This makes the use of stock prices redundant.

The HN model uses Eqs. (1) and (2) to describe the complete dynamics of stock prices under the physical measure. However, when HN risk-neutralize their dynamics (Eqs. 5 and 6), they claim that λ in Eq. (1) can be replaced by $-1/2$, and γ_1 in Eq. (2) by $\gamma_1^* = \gamma_1 + \lambda + \frac{1}{2}$ while allowing λ in γ_1^* to vary freely (and not confining λ to $-1/2$ in Eq. 2, unlike what they imposed in Eq. 1). The argument HN present in favor of doing so is that it allows the asymmetry in risk-neutral dynamics to be different from the asymmetry in actual dynamics. First, empirically, λ varies widely and keeping $\lambda = -1/2$ in the pricing part is counterintuitive.⁸ Second, allowing λ in $\gamma_1^* = \gamma_1 + \lambda + \frac{1}{2}$ to be different from $-1/2$ implies that γ_1^* can be different from γ_1 . Third, this does not align with their model's continuous-time version, where risk-neutral and actual dynamics have the same asymmetry characterization (Heston, 1993).

Our PHN model is a version of HN where λ is replaced by $-1/2$ in Eqs. (1) and (2) to get the corresponding risk-neutral version.⁹ This removes any mathematical inconsistency, brings the discrete and continuous versions of the HN model in conformity, and reduces the number of parameters to be estimated, speeding up calibration. This also helps in interpreting $\gamma_1^* = \gamma_1$ more directly: $\gamma_1^* = \gamma_1 = 0$ yields a symmetric and $\gamma_1^* = \gamma_1 \neq 0$ yields an asymmetric version of the dynamics. Thus, the generating function given in eq. (14) is changed, and the function for PHN is given by:

$$f_{PHN}(\varphi) = s_t^q \exp\{A^{PHN}(t; t+T, \varphi) + B^{PHN}(t; t+T, \varphi)h_{t+1}^{DVF}\} \quad (16)$$

where A^{PHN} and B_1^{PHN} follow the recursions:

⁸ See the estimation using NLS and results in Table 2 (page 603) of HN.

⁹ That leaves us outside HN's evaluation that 'although constrained to use the variance from the history of asset prices, their GARCH model fits better because it matches the shape of the option prices better'. Their skewness parameter, γ_1^* , which in their model is influenced by γ_1 of log-return distribution is essentially the same as γ_1 in our PHN but not influenced by log-return and is solely estimated from forward looking option prices alone.

⁶ Mixing the historical information content of stock returns of 252 previous days with the forward-looking option information content of up to the longest maturity option used in calibration convolute both types of information.

⁷ Thus, there is no mixing of historical asset return information with forward looking option market information.

Table 1
Maximum likelihood estimation of HN and BS models

	α_1	β_1	γ_1	ω	λ	μ	σ	log-likelihood
3 Years								
HN-GARCH ($\gamma_1 \neq 0$)	7.75e-9 (1.06e-7)	0.999 (3.55e-4)	53.580 (6.12e-8)	2.22e-16 (4.79e-3)	2.235 (3.54e-3)			3648.6
BS						0.0857 (0.0945)	0.1637 (0.0035)	3658.3
5 Years								
HN-GARCH ($\gamma_1 \neq 0$)	9.73e-09 (0.41e-7)	0.999 (4.04e-05)	17.0235 (2.6e-8)	2.23e-16 (4.3e-3)	5.190 (3.1e-3)			6079.3
BS						0.1686 (0.0736)	0.1645 (0.0027)	6087.8
7 Years								
HN-GARCH ($\gamma_1 \neq 0$)	4.78e-8 (1.42e-8)	0.999 (2.63e-4)	-8.13 (8.88e-9)	2.22e-16 (0.0010)	-3.53 (0.0024)			8134.5
BS						0.1596 (0.0826)	0.2185 (0.0031)	7798.0
9 Years								
HN-GARCH ($\gamma_1 \neq 0$)	2.56e-8 (1.64e-8)	0.999 (0.50e-4)	-40.63 (6.09e-9)	2.22e-6 (0.0022)	-2.51 (0.0055)			10,156
BS						0.1011 (0.0771)	0.2314 (0.0029)	9836.8

Maximum Likelihood Estimates of the GARCH model with $P = Q = 1$ and $\Delta = 1$ (day) using the S&P 500 log-returns of unrestricted HN ($\gamma_1 \neq 0$) and BS model on December 31, 2018. All the estimations are carried out with the same initial values and optimization criteria. Standard errors, computed by numerically inverting the Hessian matrix, are reported below each estimate. Panels from top to bottom show estimations based on 3, 5, 7, and 9 years of log-return data.




Fig. 1. Density fits of BS and HN (unrestricted) models with 3-, 5-, 7-, and 9-years log-return using historical (actual) parameter estimates from Table 1.

Table 2

Non-linear least squares estimation of HN models (Non-updated on the first six months of each year).

	α_1	β_1	γ_1	ω	λ	θ	RMSE	Average Price	Observations	Time (Sec.)
2016										
BS							4.409			10.68
HN-non-updated ($\gamma_1^* = 0$)	3.98e-6 (6.22e-7)	0.926 (0.012)	-116.738 (16.554)	3.7e-19 (3.17e-9)	116.238 (16.554)	0.223	7.187	59.637	1994	661.70
Ad hoc BS (PBS)							1.721	59.637	1994	15.24
HN-non-updated	7.91e-7 (2.83e-8)	0.785 (0.003)	508.718 (12.993)	2.12e-9 (3.56e-8)	-14.393 (1.218)	0.138	3.530	59.637	1994	749.18
2017										
BS							2.967			15.81
HN-non-updated ($\gamma_1^* = 0$)	5.57e-19 (5.34e-6)	0.887 (0.546)	-182.229 (1343)	2.74-e6 (9.73e-6)	181.729 (1343)	0.078	3.567	58.790	2231	219.45
Ad hoc BS (PBS)							2.057	58.790	2231	21.32
HN-non-updated	7.52-e6 (6.39e-7)	0.233 (0.044)	249.848 (6.244)	2.48-e16 (2.18e-7)	4.709 (4.016)	0.080	2.722	58.790	2231	893.86
2018										
BS							6.766			16.822
HN-non-updated ($\gamma_1^* = 0$)	1.27e-6 (1.65e-7)	0.966 (0.007)	105.942 (12.991)	9.93e-20 (3.69e-11)	105.442 (12.991)	0.127	8.091	77.034	2758	777.67
Ad hoc BS (PBS)							2.423	77.034	2758	25.22
HN-non-updated	8.39e-7 (3.22e-11)	0.543 (3.5e-5)	721.637 (0.0003)	2.92e-7 (5.65e-11)	-18.443 (0.004)	0.119	4.402	77.034	2758	1607.57

This table reports parameter estimates and in-sample valuation errors from minimizing the sum of squared errors between model option values and market option prices in the first half of each year. Asymptotic standard errors are in parentheses. Two versions of the GARCH model are estimated, one in which γ_1^* is unrestricted, and another in which $\gamma_1^* = 0$. BS is the Black-Scholes model in which a single implied volatility is estimated across all strikes and maturities on a given day, while the ad hoc BS (practitioner's BS) is an ad hoc version of the Black-Scholes model with strike and maturity-specific implied volatilities. Both BS and the ad hoc BS are calibrated weekly while the parameters of the HN (GARCH) model are held constant over the estimation period. θ is the annualized long-run standard deviation under the GARCH while $\gamma_1^* = \gamma_1 + \lambda + 1/2$ measures the skewness of the risk-neutral distribution. If $\gamma_1^* = 0$, then $\gamma_1 = -(\lambda + 1/2)$. RMSE is the root mean squared pricing error. Average price is the average option price in the sample.

$$A^{PHN}(t; T, \phi) = A^{PHN}(t + \Delta; T, \phi) + \phi r + B_1^{PHN}(t + \Delta; T, \phi)\omega - \frac{1}{2} \ln(1 - 2\alpha_1 B_1^{PHN}(t + \Delta; T, \phi)) \quad (17)$$

$$B_1^{PHN}(t; T, \phi) = \phi(-1/2 + \gamma_1) - \frac{1}{2}\gamma_1^2 + \beta_1 B_1^{PHN}(t + \Delta; T, \phi) + \frac{1/2(\phi - \gamma_1)^2}{1 - 2\alpha_1 B_1^{PHN}(t + \Delta; T, \phi)} \quad (18)$$

The call option price in the PHN model is obtained as follows:

$$c_{PHN} = S_t \left(\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \text{Re} \left[\frac{K^{-i\varphi} f_{PHN}^*(i\varphi + 1)}{i\varphi f_{PHN}^*(1)} \right] d\varphi \right) - K e^{-rT} \left(\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \text{Re} \left[\frac{K^{-i\varphi} f_{PHN}^*(i)}{i\varphi} \right] d\varphi \right) \quad (19)$$

where h_{t+1}^{DVF} in $f_{PHN}(\varphi)$ is obtained through Eq. (13) using all the cleanup options on day t to estimate the regression coefficients $a_0, \dots, a_5$. Once the regression coefficients on day t are estimated, the price of each option, c_{PHN} , is calculated using distinct h_{t+1}^{DVF} corresponding to the options' (K, T) capitalizing DVF relation in Eq. (13). This is the fundamental difference between our PHN with HN, as all options calibrated in c_{CFG} have a constant h_{t+1} .

3. Data and estimation results

This section discusses the data, in-sample calibration, and out-of-sample validation results for the BS, PBS, HN, and PHN models. The analysis of pricing performance is followed by a hedging analysis for the four models.

3.1. Data and sample

We use weekly data (sampled on Wednesdays) of S&P 500 index options traded on the CBOE from January 1, 2016 to December 31,

2018, where the in-sample period for non-updated calibrations is the first six months of each year, and for the updated calibration it is the last six months of each year. The out-of-sample validations resort to the subsequent Thursday option contracts of each week. We sample data on Wednesdays because Wednesdays are less likely to be a holiday or affected by a day-of-the-week effect. Two versions of the GARCH model are estimated, one for which γ_1^* is unrestricted, and another in which $\gamma_1^* = 0$.

Several filters are applied to avoid noise in parameter estimation. First, we exclude options with <10 days to expiration due to very small time-value and inaccuracy of implied volatilities due to sensitivity to market micro-structure problems and measurement errors (Hentschel & Kothari, 2001). Second, options with >180 days to maturity are excluded as the high trading premiums make trading unpopular. Third, we exclude options with moneyness (S/K) <0.90 and >1.10. Finally, to mitigate the impact of price discreteness on option valuation, price quotes lower than 3/8th of a dollar are excluded.

Although options data are weekly, the S&P 500 index returns are daily to calculate the conditional variance relevant to option values. We compute the in-sample variance estimate for the first six months of each year using daily logarithmic returns.

3.2. Return dynamics from the GARCH process

Following HN, the empirical analysis focuses on a GARCH model with a single lag. We set $\Delta = 1$ and use daily index returns to model the evolution of volatility. To illustrate the importance of the skewness parameter γ_1 , we estimate a restricted and unrestricted GARCH model. In Table 1, we report the volatility dynamics with parameter estimates from the history of index prices on December 31, 2018, using maximum likelihood estimation for the unrestricted HN-GARCH model. We perform all estimations with the same initial values and optimization criteria and numerically compute standard errors, reported below each estimate, by inverting the Hessian matrix. Panels from top to bottom

Table 3

Non-linear least squares estimation of the Practitioners' HN models (Non-updated on the first six months of each year)

	α_1	β_1	γ_1	ω	persistence	θ	RMSE	Average price	Observations	Time (Sec.)
2016										
PHN-non-updated ($\gamma_1^* = 0$)	3.9e-20 (1.26-e8)	0.999 (1.18e-6)		8.59e-9 (5.24e-9)	0.999	0.170	1.636	59.637	1994	210.73
PHN-non-updated	1.5e-7 (1.88e-8)	0.982 (1.1e-6)	332.264 (21.629)	1.4e-17 (1.55e-8)	0.998	0.144	1.558	59.637	1994	329.99
2017										
PHN-non-updated ($\gamma_1^* = 0$)	1.01e-6 (5.563-8)	0.962 (0.002)		2.1e-20 (9.53e-11)	0.962	0.082	2.683	58.790	2231	264.73
PHN-non-updated	3.16e-7 (3.69e-8)	0.935 (0.002)	412.806 (31.567)	1.57e-9 (2.58e-8)	0.989	0.086	1.624	58.790	2231	660.09
2018										
PHN-non-updated ($\gamma_1^* = 0$)	3.93e-7 (1.59e-10)	0.980 (7.61e-5)		5.37e-7 (1.42e-10)	0.980	0.108	4.582	77.034	2758	132.24
PHN-non-updated	9.97e-7 (2.08e-6)	0.921 (1.120)	243.802 (606.766)	4.1e-21 (1.1e-11)	0.980	0.112	3.600	77.034	2758	1253.27

This table reports the parameter estimates and in-sample valuation errors in the first half of each year (2016, 2017, and 2018) from minimizing the sum of squared errors between model option values and market option prices. Standard errors (asymptotic) appear in parentheses. Versions of GARCH models are: one in which γ_1^* is unrestricted and another in which $\gamma_1^* = 0$. PHN model is the practitioner's version of HN with strike and maturity specific implied volatilities across options for which DVF parameters (as in eq. 12) are calibrated every week (Wednesday) while the parameters of the GARCH model are held constant over the estimation period. θ is the annualized long run standard deviation under the GARCH. If $\gamma_1^* = 0$, then $\gamma_1 = -(\lambda + 1/2)$. RMSE is the root mean squared pricing error (in \$). Average price is the \$ average price of the options in the sample.

show estimations based on 3, 5, 7, and 9 years of log-return data. The skewness parameter γ_1 , is positive, indicating that shocks to returns and volatility are negatively correlated. In the restricted/symmetric model, the skewness parameter has an important effect on the behavior of the variance process. Thus, the high significance of the parameter estimates confirms that returns follow a GARCH process, and the GARCH model becomes superior over the BS model with an increase of information used in estimation as indicated by the log-likelihood ratio.¹⁰

3.3. In-sample valuation errors across models

Following DFW, HN construct a PBS model in which volatility enters the BS option formula as a function of strike price and time to maturity, as shown in Eqs. (12) and (13). The coefficients of the PBS model are estimated weekly on Wednesdays for the first six months of the year using OLS that minimizes the sum of squared errors between the BS implied volatilities across different strikes (and maturities) and the DFW model's functional form of the implied volatility. Following HN, we use NLS to estimate the parameters for the first six months of 2016, 2017, and 2018. The conditional variance, $h(t+1)$, is drawn from the daily evolution of index returns. We compute the variance for the first six months from daily returns in the non-updating estimation. In the updating estimation, the GARCH parameters are updated weekly.

The estimation results of the non-updated HN model for both actual and risk-neutral skewness (γ_1 and γ_1^*) assumptions where $\gamma_1^* = 0$ and $\gamma_1 = -(\lambda + 1/2)$ are reported in Table 2, along with those of the BS and PBS models. The RMSE of the BS model ranges from 2.967 in 2017 to 6.766 in 2018. The RMSE for the HN model (symmetric GARCH with $\gamma_1^* = 0$) is as high as 8.091 in 2018. In most cases, the symmetric HN model has a worse in-sample fit than the BS model. The RMSEs for the HN model with non-zero γ_1^* are 3.530, 2.722, and 4.402 in 2016, 2017, and 2018, respectively. This result confirms that the full GARCH model that allows for skewness effects provides a better in-sample fit than the

symmetric model. In comparison, the PBS model shows the best in-sample fit to options prices. The in-sample RMSEs for the PBS model are 1.721, 2.057, and 2.432 for 2016, 2017, and 2018, respectively.

The estimation results of practitioner non-updated HN models with and without asymmetry are presented in Table 3. The RMSEs for the non-updated PHN model with $\gamma_1^* = 0$ in 2016, 2017, and 2018 are 1.636, 2.683, and 4.582, respectively. The symmetric PHN model has a better in-sample fit than the symmetric HN model. The RMSEs for the non-updated model without restriction on γ_1^* are 1.558, 1.624, and 3.60 in 2016, 2017, and 2018, respectively. The non-updated PHN model without any restriction on γ_1^* performs better in-sample than the symmetric HN model. (The non-updated PHN model, with and without restriction on γ_1^* , performs significantly better in-sample than the non-updated HN model, with and without restriction on γ_1^* , respectively). Overall, the PHN model performs better than the PBS model in most cases.

To compare the PBS model with the model using the GARCH process, HN estimate an updated GARCH model by minimizing the RMSEs between the model and market option prices, allowing parameters to change weekly.¹¹ In the HN model, the parameters change weekly, but at time t , the conditional variance, $h(t+1)$, is drawn from the time series of returns from the previous 252 days. Instead of calculating volatility using historical return data in each iteration every week separately, our model calculates the one-day ahead implied volatility by applying Dumas et al. (1998) DVF mapping.

Table 4 presents the in-sample valuation errors from the weekly estimation using option prices in the second half of each year for the PBS, the updated HN-GARCH (restricted and unrestricted), and the updated PHN (restricted and unrestricted) models. Both the PBS and (P) HN models are estimated weekly using OLS and NLS, respectively. The parameter estimates are the averages of the weekly estimates. The RMSEs of the restricted PHN model are 0.929, 1.044, and 1.227 compared to 2.271, 2.161, and 3.80 for the restricted HN model in 2016, 2017, and 2018 respectively. Similarly, the RMSEs of the unrestricted PHN model (unrestricted HN model) are 0.558, 0.753, and 0.919 (0.631,

¹⁰ Likelihoods very similar to that of the BS model. BS likelihood remains better, for estimations with 3 and 5 years of log-return data (as shown in Table 1 and Fig. 1), than HN likelihoods which is not a big surprise as the BS model is often found outperforming the HN model. For example, for pricing options using OMXS30 (see Ekstorm et al. (2014)).

¹¹ HN (2000) argue that the improved in-sample fit of the PBS model might stems from a more flexible functional form of Dumas et al. (1998) or from the instability of the functional form of the GARCH process.

Table 4

In-sample comparison of the BS, PBS, the updated HN, and updated PHN models.

		RMSE	α_1	β_1	γ_1	ω	λ	Average Price	Observations	Time (Sec.)
2016	BS	2.300								9.47
	Ad hoc BS (PBS)	1.228								11.83
	HN-updated	2.271	7.5e-07	0.823	154.231	4.9e-06	−154.731	14.59	1415	24,791.82
	($\gamma_1^* = 0$)		(4.1e-08)	(0.008)	(0.617)	(3.82e-07)	(0.617)			
	HN-updated	0.631	1.3e-06	0.667	484.772	2.2e-07	−129.794	14.59	1415	59,129.71
			(2.3e-08)	(0.004)	(5.846)	(6.5e-09)	(8.167)			
	Practitioner HN-updated	0.929	7.8e-08	0.979	–	5.1e-07		14.59	1415	693.21
	($\gamma_1^* = 0$)		(4.7e-09)	(4.5e-05)		(5.6e-09)				
2017	BS	2.193								9.47
	Ad hoc BS (PBS)	1.162								13.06
	HN-updated	2.161	2.0e-06	0.700	147.693	3.2e-06	−148.194	12.88	1469	5659.32
	($\gamma_1^* = 0$)		(1.4e-07)	(0.010)	(0.382)	(1.6e-07)	(0.382)			
	HN-updated	0.826	7.2e-07	0.683	619.139	2.1e07	−328.013	12.88	1469	57,044.06
			(8.4e-09)	(0.004)	(6.114)	(5.2e-09)	(14.121)			
	Practitioner HN-updated	1.044	6.8e-08	0.979	–	2.7e-07		12.88	1469	688.80
	($\gamma_1^* = 0$)		(3.3e-09)	(0.4e-05)		(4.8e-09)				
2018	BS	3.534								9.79
	Ad hoc BS (PBS)	1.173								14.63
	HN-updated	3.800	3.1e-06	0.764	150.380	4.6e-06	−150.880	18.147	1643	33,629.79
	($\gamma_1^* = 0$)		(2.1e-07)	(0.011)	(0.441)	(2.2e-07)	(0.441)			
	HN-updated	0.870	2.6e-06	0.619	379.623	3.0e-07	−122.829	18.147	1643	68,174.42
			(8.7e-08)	(0.009)	(4.320)	(9.5e-09)	(7.074)			
	Practitioner HN-updated	1.227	8.2e-08	0.979	–	7.9e-07		18.147	1643	686.09
	($\gamma_1^* = 0$)		(6.1e-09)	(7.9e-05)		(2.0e-08)				
	Practitioner HN-updated	0.919	6.4e-07	0.876	399.719	7.2e-07		18.147	1643	7752.03
			(1.4e-08)	(0.003)	(4.242)	(2.0e-08)				

In-sample valuation errors from the weekly estimation using option prices in the second half of each year (2016, 2017, and 2018) for the PBS, the updated HN-GARCH (restricted and unrestricted), and the updated PHN-GARCH (restricted and unrestricted) models. The PBS and the GARCH models are estimated each week using OLS and NLS, respectively. For the GARCH model, $h(t+1)$ is computed from daily historical (last 252 days) S&P 500 returns. The parameter estimates are the average of the weekly estimates. Standard errors are in parenthesis.





Fig. 2. Variation of λ between HN and PHN models for updated weekly calibrations over July to December 2018 (every Wednesday) on the top. On the bottom, the corresponding weekly RMSE of HN and PHN (and PBS) models over the same period.

0.826, and 0.870) in 2016, 2017, and 2018, respectively. In other words, the updated PHN model shows the best in-sample fit to options prices compared to the HN model in most cases.

HN argue in favor of the empirical stability of the parameters, α_1 (the volatility of volatility) and γ_1 (the skewness of returns) compared to the other parameters, as the GARCH model fits the data well even with constant parameters. However, λ , the market price of risk, varies widely over time, as shown in Tables 2 and 3. Fig. 2 (top) also shows the

variation of λ between the HN and PHN models for updated weekly calibrations from July to December 2018 (every Wednesday). While λ is fixed at $-1/2$ in the PHN model, λ in the HN model fluctuates widely over time. We also plot (bottom) the corresponding weekly RMSEs of HN and PHN (and PBS) models for the same period. In the HN model, the GARCH dynamics assume a constant $h(t+\Delta)$, while λ has been assigned values different from $-1/2$ (even though $\lambda = -1/2$ in the pricing part). We replace the GARCH forecast $h(t+\Delta)$ by DVF mapping of Dumas

Table 5
Out-of-sample valuation errors.

		RMSE (next Thursday of each weekly Wednesday in-sample calibration)	Average Price	Observations
2016	BS	8.161	14.34	1426
	PBS	1.762	14.34	1426
	HN-	2.762	14.34	1426
	updated ($\gamma_1^* = 0$)			
	HN-	1.837	14.34	1426
	updated			
	PHN-	1.651	14.34	1426
	updated ($\gamma_1^* = 0$)			
2017	PHN-	1.429	14.34	1426
	updated			
	BS	8.348	13.06	1478
	PBS	1.924	13.06	1478
	HN-	3.096	13.06	1478
	updated ($\gamma_1^* = 0$)			
	HN-	2.438	13.06	1478
	updated			
2018	PHN-	1.819	13.06	1478
	updated ($\gamma_1^* = 0$)			
	PHN-	1.678	13.06	1478
	updated			
	BS	10.431	18.08	1640
	PBS	3.298	18.08	1640
	HN-	11.605	18.08	1640
	updated ($\gamma_1^* = 0$)			
	HN-	5.005	18.08	1640
	updated			
	PHN-	3.380	18.08	1640
	updated ($\gamma_1^* = 0$)			
	PHN-	3.404	18.08	1640
	updated			

This table reports the average \$ out-of-sample valuation errors year by year (2016, 2017 and 2018). While the option values are computed every Thursday (not the immediate next day of every Wednesday updated calibration but the Thursday following the immediate next Thursday of each Wednesday (This is fundamentally different from HN who considered consecutive Wednesdays for in-sample and out-sample assessments. However most of the researchers in finance prefer using Wednesday's for in sample and the next day's (Thursday's) for out sample assessment. We avoid the immediate next Thursdays and consider the following Thursdays, exactly a week later, for our out-sample assessment; which is reflected in relatively higher RMSE's out-sample compared to those presented in [Heston and Nandi \(2000\)](#).) in the second half of each year, the GARCH (updated) model parameters are estimated in the previous week and variance is computed using historical daily S&P 500 levels. For BS model a single implied volatility is estimated across all strikes and maturities on every Wednesday, while PBS is an ad hoc version of the Black-Scholes model with strike and maturity specific implied volatilities. Same holds true for HN and our PHN models. Like BS and PBS models, HN and PHN models are estimated every week and then used to value options in the following week.

[et al. \(1998\)](#). This enables us to avoid dual value assignment of λ to get a single model price. Moreover, unlike keeping tomorrow's volatility forecast constant for all options over all maturities in the future, we replace $h(t + \Delta)$ by option-wise (depending on maturity and strike) *ad-hoc* estimates h_{t+1}^{DVF} following [Dumas et al. \(1998\)](#). This process provides more consistent pricing and reduces the calibration time significantly. Thus, the PHN model is parsimonious, as the table's last column shows.

3.4. Out-of-sample model comparison

Even though our proposed PHN model dominates in-sample, it is




Fig. 3. Simulated Delta-Gamma hedging performance of different models for a deep OTM call option. Models are calibrated on Wednesdays over July – Dec 2018; for hedging, the mean of the weekly estimates is used on January 3, 2019. The spot was 2447.9. For this illustration, 1000 simulations are used for dynamic delta-gamma hedging error calculation. As a secondary instrument, a call option with same maturity and $K_2 = K_1 + 1$ has been considered, where K_1 is the strike of the option being hedged.

important to examine the predictive power of the model for out-of-sample option prices. We examine the out-of-sample valuation performance of all updated models and the BS and PBS models for the next six months of each of the three years. The updated HN and PHN models, including the BS and PBS models, are estimated weekly in the second half of each year. For the PBS and updated HN models, the estimated parameters from the current week are used to value options for the following week. At each time, t , $h(t + 1)$ is drawn from the daily history of the index, including up to the first day in the index's history used for obtaining the in-sample estimates in the previous week. The BS, PBS, and updated HN models only predict one week ahead. For the PHN model, we obtain the aggregate out-of-sample pricing errors by aggregating each Thursday's out-of-sample valuation errors obtained using each Wednesday's in-sample estimates (obtained solely using options traded on that Wednesday).

[Table 5](#) reports the out-of-sample valuation errors for all models across three out-of-sample periods. Similar to the in-sample results, the BS formula fits worst out-of-sample. In contrast to HN, we find the PBS model outperforming both the restricted and unrestricted HN models. The RMSEs for the restricted HN model are 2.762, 3.096, and 11.605 in 2016, 2017, and 2018, respectively. The corresponding RMSEs for the restricted PHN model are 1.651, 1.891, and 3.380, which are substantially lower than the restricted HN model. The RMSEs for the unrestricted HN model are far lower than the restricted version, ranging from 1.837 in 2016 to 5.005 in 2018. However, the unrestricted PHN model substantially outperforms the HN model out-of-sample. The RMSEs for the unrestricted PHN model are 1.429, 1.678, and 3.404 in 2016, 2017, and 2018, respectively. Thus, our findings show that the PBS is superior to the HN models each year, but, out-of-sample, our proposed PHN model outperforms the PBS model.

We stress an obvious downside of the PHN model, unreported, to save space.¹² As a non-structural approach, the PHN model is not robust in its long-term out-of-sample forecast performance. Calibrating the most used benchmark models using Wednesday option (market) prices in-sample and assessing the out-of-sample performance next Wednesday

¹² These results are available on request.




Fig. 4. Delta and Delta-Gamma hedging performance of market options on January 3, 2019, of various strikes with increasingly out-of-the-money features. All options have 15 days to maturity. Models are calibrated on Wednesdays over July – Dec 2018; for hedging, the mean of the weekly estimates is used on January 3, 2019. The spot was 2447.9. For each option, we used 10,000 simulations in both dynamic Delta and dynamic Delta-Gamma hedging error calculation. In Gamma hedging, as a secondary instrument, a call option with the same maturity and $K_2 = K_1 + 1$ has been considered, where K_1 is the strike of the option being hedged.




Fig. 5. Delta and Delta-Gamma hedging performance of market options on January 3, 2019, of various strikes with increasingly out-of-the-money features. All the options have 43 days to maturity. Models are calibrated on Wednesdays over July – Dec 2018; for hedging, the mean of the weekly estimates is used on January 3, 2019. The spot was 2447.9. For each option, we used 10,000 simulations in both dynamic Delta and dynamic Delta-Gamma hedging error calculation. In Gamma hedging, as a secondary instrument, a call option with the same maturity and $K_2 = K_1 + 1$ has been considered, where K_1 is the strike of the option being hedged.

using Thursday option (market) prices show a clear dominance of PHN over HN and PBS. Similarly, calibrating over a six-month window of Wednesdays and then assessing the results over the next six-month window of Wednesdays shows the average performance of the non-structural PHN model compared to the HN model. Nonetheless, the out-of-sample performance of the purely non-structural PBS model assessed after a week is significantly improved when implemented with the HN approach (i.e., the PHN model). In that sense, we could see the PHN model as an improvement of the PBS model. Moreover, given its dominating in-sample and short-term out-of-sample performance over the HN and PBS models, the PHN model could become an attractive model updated daily to establish its appeal in actual market valuation. Indeed, people do not use even week-old estimates to determine the fair prices of derivatives in day-to-day business.

4. Hedging with HN and PHN

Hedging analysis often provides another benchmark to assess relative model stability. An option with a particular strike price and maturity can be hedged under different models using the model's Delta and Gamma, with the support of a secondary instrument, to obtain dynamic hedging errors under that particular model. Hedging various market options with increasingly out-of-the-money features for a particular maturity and repeating the process for several different maturities reveals the relative robustness of market-based hedging. This section




Fig. 6. Delta and Delta-Gamma hedging performance of market options on January 3, 2019, of various strikes with increasingly out-of-the-money features. All the options have 71 days to maturity. Models are calibrated on Wednesdays over July – Dec 2018; for hedging, the mean of the weekly estimates is used on January 3, 2019. The spot was 2447.9. For each option, we used 10,000 simulations in both dynamic Delta and dynamic Delta-Gamma hedging error calculation. In Gamma hedging, as a secondary instrument, a call option with the same maturity and $K_2 = K_1 + 1$ has been considered, where K_1 is the strike of the option being hedged.

conducts hedging experiments that confirm the PHN model's usefulness compared to the HN, BS, and PBS models.¹³

HN argue that the call and put Delta in their model are given by $\Delta_{call} = P_1$ and $\Delta_{put} = P_1 - 1$.

where P_1 is given by

$$P_1 = \frac{1}{2} + \frac{e^{-rT}}{S\pi} \int_0^\infty \text{Re} \left[\frac{K^{-i\phi} f^*(i\phi + 1)}{i\phi} \right] d\phi \quad (20)$$

The Gamma for both calls and puts can be obtained by differentiating Eq. (14) with respect to S :

$$\Gamma = \frac{dP_1}{dS} = \frac{e^{-rT}}{\pi} \int_0^\infty \text{Re} \left[\frac{K^{-i\phi} f^*(i\phi + 1)}{S^2} \right] d\phi \quad (21)$$

In the case of PHN, unlike keeping tomorrow's volatility forecast $h(t + \Delta)$ constant for all options over all maturities in the future (as in HN), we replace $h(t + \Delta)$ in $f(\varphi)$ option-wise (depending on maturity and strike) *ad-hoc* estimates of Dumas et al. (1998).¹⁴ The same applies in the case of PBS and BS hedge factor computations. It is worth noting that the practitioner's approach for PHN and PBS fundamentally differs from the original approach for HN and BS. In the practitioners' approach to hedging, every option considered for hedging on a given day will have its maturity- and strike-wise option implied volatility used in delta and gamma calculations to price it. However, in the non-practitioner's approach, all options considered for hedging have a common asset volatility used in Delta, Gamma computations, and pricing. This leaves the hedging performances of practitioner's and non-practitioner's approaches significantly different, despite applying the same evolution of risky or statistical dynamics for simulating underlying prices under the practitioner's and non-practitioner's approaches.¹⁵

The rationale for the practitioner's approach to hedging is that the day the hedging analysis takes place, it takes place after all derivative trades are completed. So $h(t + \Delta)$ is attainable following the method of Dumas et al. (1998) as $h_{t+\Delta}^{DVF}$, with all derivative contracts traded on day t . We further assume that the GARCH simulation of volatility, including the price evolution until the maturity day of the option being hedged,

starts with the initial $h(0) = h_{t+\Delta}^{DVF}$ in the PHN approach, and $h(0)$ is the daily volatility forecast (from Table 1) that corresponds to nine years of return data in HN hedging (i.e., in the non-practitioners approach). In other words, the simulation starts with the assumption that right from day one, HN $h(t)$ and PHN $h(t)$ are different. We assume the same for both BS and PBS models.

Our hedging experiments consider two sets of market calibrations. One set from Table 1 describes the price evolution of risky assets, the other is the risk-neutral one, as in Table 4, used in pricing derivatives and computing hedge factors. However, there are substantial differences in hedging experiments between HN and our PHN models. From the day of hedging up to the day of maturity of the option being hedged, HN hedging considers the daily evolution of stock price and volatility according to the risk dynamics described by the parameters in Table 1.¹⁶ Still, the rebalancing process through Delta-Gamma and option prices are based on the risk-neutral parameters in Table 4. However, in the case of PHN hedging, the stock prices evolve as in HN, but the rebalancing process through Delta-Gamma and option prices takes place with risk-neutral parameters in Table 4 together with $h(t)$, which specifically evolves for the options in hedging (with adjusted time remaining until maturity on each day) according to the Dumas et al. (1998) formula. This process uses regression coefficients from fitting all traded options on the hedging day.

First, for delta hedging, the hedging portfolio at time t composed of a short position in an option, long in $\Delta^{HO}(t)$ number of shares, and an amount in a risk-free bank account. The hedge is rebalanced at times $t + \Delta t$ by purchasing $(\Delta^{HO}(t + \Delta t) - \Delta^{HO}(t))$ shares. Thus, the dynamic delta hedging profit loss (PL) is calculated as (see Forsyth (2022), Armstrong (2017), Joshi (2003), Taleb (1997))¹⁷:

¹⁶ Risk neutralization in particular, and through the Girsanov transformation in general, is a mean correction and is assumed to leave volatilities under both measures same.

¹⁷ The hedging literature is very advanced, especially including the studies on option portfolios instead of individual option positions. However, as we focus on the practitioner's version of hedging, we need to demonstrate the difference by running the hedging experiments option-wise with various maturities and moneyness for which Dumas et al. (1998) volatilities can be extracted individually. Instead of designing an option portfolio hedging experiment, we hedge all market options (ATM, ITM and OTM) on a given day, 3RD Jan' 2019, separately for three different maturities (15, 43 and 71 DTM's) in the market.

¹³ It uses naive Monte Carlo simulations to hedge using options of various maturities and moneyness.

¹⁴ We then call it $f_{PHN}(\varphi)$ which is defined by Eq. 16.

¹⁵ With different starting values in simulations as explained in next paragraph.

$$\begin{aligned}
& PL(t + \Delta t) \\
& = PL(t)e^{r\Delta t} - \text{stock purchase cost for next hedging period of length } \Delta t \\
& = PL(t)e^{r\Delta t} - ((\Delta^{HO}(t + \Delta t) - (\Delta^{HO}(t))S(t + \Delta t))
\end{aligned} \quad (22)$$

where HO stands for hedged option. For each simulation or each final $PL(T)$ we obtain the market value of the delta-hedged portfolio at maturity T as:

$$\begin{aligned}
M(T) & = \text{Bank balance (i.e. } PL(T)) + \text{asset} - \text{liability} \\
& = P(T) + \Delta^{HO}(T - \Delta t)S(T) - \max\{S(T) - K, 0\}
\end{aligned} \quad (23)$$

Second, for delta-gamma hedging, we further need long position in a secondary derivative instrument (SI). This secondary instrument has different strike with the same maturity or higher than HO . The delta and gamma of HO and SI are denoted by Δ^{HO} , Γ^{HO} , and Γ^{SI} . Thus, the dynamic delta-gamma hedging profit loss is calculated as follows¹⁸:

$$\begin{aligned}
& PL(t + \Delta t) = PL(t)e^{r\Delta t} - \text{stock purchase cost for next hedging period of length } \Delta t \\
& \quad - \text{secondary instrument purchase cost for next hedging period of length } \Delta t = PL(t)e^{r\Delta t} \\
& - \left(\left(\Delta^{HO}(t + \Delta t) - \max \left(\min \left(\frac{\Gamma^{HO}(t + \Delta t)}{\Gamma^{SI}(t + \Delta t)}, 100 \right), -100 \right) \Delta^{SI}(t + \Delta t) \right) - \left(\Delta^{HO}(t) - \max \left(\min \left(\frac{\Gamma^{HO}(t)}{\Gamma^{SI}(t)}, 100 \right), -100 \right) \Delta^{SI}(t) \right) \right) S(t + \Delta t) \\
& \quad - \left(\max \left(\min \left(\frac{\Gamma^{HO}(t + \Delta t)}{\Gamma^{SI}(t + \Delta t)}, 100 \right), -100 \right) - \max \left(\min \left(\frac{\Gamma^{HO}(t)}{\Gamma^{SI}(t)}, 100 \right), -100 \right) \right) SI(t + \Delta t)
\end{aligned} \quad (24)$$

where SI stands for the secondary instrument used in hedging, which could be an option¹⁹ with a different strike with the same maturity or higher than HO . Now we are ready to simulate M possible profits or losses arising from hedging any option with maturity T , i.e., $\{PL^i(T); i = 1 \dots M\}$. For each simulation, i.e., for each final $PL(T)$, we obtain the market value of the delta-gamma hedged portfolio at maturity T as

$$\begin{aligned}
M(T) & = \text{Bank balance (i.e. } PL(T)) + \text{asset} - \text{liability} = P(T) + \left[\Delta^{HO}(T - \Delta t) - \max \left(\min \left(\frac{\Gamma^{HO}(T - \Delta t)}{\Gamma^{SI}(T - \Delta t)}, 100 \right), -100 \right) \Delta^{SI}(T - \Delta t) \right] S(T) \\
& \quad + \left[\max \left(\min \left(\frac{\Gamma^{HO}(T - \Delta t)}{\Gamma^{SI}(T - \Delta t)}, 100 \right), -100 \right) \right] SI(T) - \max\{S(T) - K, 0\}
\end{aligned} \quad (25)$$

We compute the relative hedging error (RHE) for each simulation as follows:

$$RHE = \frac{e^{-rT}M(T)}{\text{market price of } HO} \quad (26)$$

The RMSE of hedging is then obtained as:

$$RMSE = \sqrt{\frac{\sum_{i=1}^M RHE_i^2}{M}} \quad (27)$$

We use 10,000 ($M = 10,000$) simulations in the RMSE calculation for each option. In the case of delta-gamma hedging as secondary hedging

instruments, we use options with similar characteristics as the primary options being hedged (HO). Still, the strike of the secondary instrument option (K2) is considered one more than the strike of the primary option being hedged (K1), i.e., $K_2 = K_1 + 1$. The hedged option, HO , has maturity T , and the secondary instrument used in hedging, SI , is the call option with the same maturity, T . Fig. 3 shows the relative hedging errors of delta-gamma hedging for 1000 simulations corresponding to an OTM option. Models are calibrated on Wednesdays over July – Dec 2018. For hedging, the mean of the weekly estimates is used on January 3, 2019. The spot was 2447.9. We find excellent hedging performance of the PHN model in simulation.

Figures 4–6 show the RMSE of delta and delta-gamma hedging errors for various market options on January 3, 2019, corresponding to options with 15 (Fig. 4), 43 (Fig. 5), and 71 (Fig. 6) days to maturity. Models are calibrated on Wednesdays over July – Dec 2018. For hedging, the mean of the weekly estimates is used on January 3, 2019. The spot was 2447.9.

Hedging errors increase as the contracts become increasingly out-of-the-money (OTM) in all three cases with three different maturities. Overall, delta hedging is much worse than delta-gamma hedging in all cases. The wide variations of λ (the market price of risk) can be so large that the average estimates for updated unrestricted versions of HN offer poor forecasting performance, as evidenced by the hedging graphs (Figs. 3, 4, 5, and 6); and PHN provides a solution.

5. Conclusions

The Practitioner's Black and Scholes model of Dumas et al. (1998) and the closed-form GARCH option pricing model of Heston and Nandi (2000) remain two popular approaches in the derivatives literature. Their complementary strengths are combined in this article to offer another model for obtaining more accurate prices. Better prices for the PHN model are obtained much faster than with the HN model, and PHN reduces the data requirement of implementing the HN model by half. PHN most often outperforms both HN and PBS models, both in- and out-of-sample. The main reason behind the improved performance for PHN is its flexibility to incorporate option-wise future multi-period ahead market volatilities into current pricing instead of using a single one-period ahead asset volatility for all options. Such option-wise volatility yields better hedging performance for the PHN model, as hedging market options show that the PHN hedge factors 'delta' and 'gamma' yield lower delta and delta-gamma hedging errors for options positions irrespective of moneyness and maturity. The PHN model is essentially

¹⁸ Use of 'max' and 'min' operators are simply to avoid possible blow ups due to possible small gamma values.

¹⁹ Or possibly other instruments.

the practitioner's version of the HN model, but it also offers a way to improve the performance of the PBS model. Since any model that outperforms PBS sets a new benchmark, the PHN model offers a new challenging benchmark for future researchers.

Data availability

Data will be made available on request.

References

- Amin, K., & Ng, V. (1993). *ARCH processes and option valuation*. Working paper. Ann Arbor, MI: University of Michigan.
- Armstrong, J. (2017). *C++ for financial mathematics*. CRC Press.
- Augustyniak, M., & Badescu, A. (2021). On the computation of hedging strategies in affine GARCH models. *Journal of Futures Markets*, 41, 710–735.
- Barone-Adesi, G., Engle, R. F., & Mancini, L. (2008). A GARCH option pricing model with filtered historical simulation. *Review of Financial Studies*, 21, 1223–1258.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31, 307–327.
- Cao, H., Badescu, A., Cui, Z., & Jayaraman, S. K. (2021). Valuation of VIX and target volatility options with affine GARCH models. *Journal of Futures Markets*, 40, 1880–1917.
- Cheng, H. W., Chang, L. H., Lo, C. L., & Tsai, J. T. (2023). Empirical performance of component GARCH models in pricing VIX term structure and VIX futures. *Journal of Empirical Finance*, 72, 122–142.
- Christoffersen, P., Heston, S., & Jacobs, K. (2006). Option valuation with conditional skewness. *Journal of Econometrics*, 131, 253–284.
- 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., Ornathanalai, C., & Wang, Y. (2009). Option valuation with long-run and short-run volatility components. *Journal of Financial Economics*, 90, 272–297.
- Derman, E., & Kani, I. (1994). Riding on the smile. *Risk*, 7, 32–39.
- Duan, J.-C. (1995). The GARCH option pricing model. *Mathematical Finance*, 5, 13–32.
- Duan, J.-C. (1996). Cracking the smile. *Risk*, 9, 55–59.
- Dumas, B., Fleming, J., & Whaley, R. (1998). Implied volatility functions: Empirical tests. *Journal of Finance*, 53, 2059–2106.
- Dupire, B. (1994). Pricing with a smile. *Risk*, 7, 18–20.
- Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation. *Econometrica*, 50, 987–1008.
- Engle, R., & Mustafa, C. (1992). Implied ARCH models from options prices. *Journal of Econometrics*, 52, 289–311.
- Feller, W. (1971). *An introduction to probability theory and its applications*. New York: Wiley.
- Forsyth, P. (2022). *An introduction to computational finance without agonizing pain*. Working Book Draft: University of Waterloo.
- Hardle, W., & Hafner, C. (2000). Discrete time option pricing with flexible volatility estimation. *Finance and Stochastics*, 4, 189–207.
- Hentschel, L., & Kothari, S. P. (2001). Are Corporations Reducing or Taking Risks with Derivatives? *Journal of Financial and Quantitative Analysis*, 36(1), 93–118.
- Heston, S., & Nandi, S. (2000). A closed-form GARCH option pricing model. *Review of Financial Studies*, 13, 585–626.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2), 327–343.
- Huang, Z., Wang, T., & Hansen, P. R. (2017). Option pricing with the realized GARCH model: An analytical approximation approach. *Journal of Futures Markets*, 37, 328–358.
- Joshi, M. S. (2003). *The concepts and practice of mathematical finance* (vol. 1). Cambridge University Press.
- Kendall, M., & Stuart, A. (1977). *The advanced theory of statistics: Distribution theory*. London: Griffin.
- Kim, I. J., & Kim, S. (2004). Empirical comparison of alternative stochastic volatility option pricing models: Evidence from Korean KOSPI 200 index options market. *Pacific-Basin Finance Journal*, 12(2), 117–142.
- Liang, G., & Wang, X. (2021). Pricing vulnerable options in a hybrid credit risk model driven by Heston–Nandi GARCH processes. *Review of Derivatives Research*, 24, 1–30.
- Liu, G., Jiao, Y., & Guo, S. (2022). GARCH pricing and hedging of VIX options. *Journal of Futures Markets*, 42(6), 1039–1066.
- Oh, D. H., & Park, Y.-H. (2023). GARCH option pricing with volatility derivatives. *Journal of Banking & Finance*, 146, Article 106718.
- Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance*, 49, 771–818.
- Taleb, N. (1997). *Dynamic hedging*. Wiley.
- Tong, C., Hansen, P. R., & Huang, Z. (2022). Option pricing with state-dependent pricing kernel. *Journal of Futures Markets*, 42(8), 1409–1433.
- Wang, T., Cheng, S., Yin, F., & Yu, M. (2022). Overnight volatility, realized volatility, and option pricing. *Journal of Futures Markets*, 42(7), 1264–1283.