HEC MONTRÉAL

Est-ce que les restrictions affines de volatilité sont encore coûteuses si le noyau de prix est quadratique?

par

Patrick Aoun

Sciences de la gestion
(Option Ingénierie financière)

Mémoire présenté en vue de l'obtention
du grade de maîtrise ès sciences en gestion
(M. Sc.)

Août 2018

Résumé

Inspiré par Christoffersen, Heston, et Jacobs (2013), je compare l'utilisation d'un noyau de prix exponentiellement quadratique avec le noyau de prix exponentiellement linéaire de Rubinstein (1976) pour la tarification d'options. Je l'utilise avec un modèle GARCH(1,1) affine, tel que dans CHJ (2013), et avec un modèle non-affine pour lequel je développe une mesure risque-neutre. Je produis ainsi deux axes de comparaison : affine versus non-affine, et noyau de prix exponentiellement linéaire versus quadratique. Je trouve que l'utilisation du noyau quadratique améliore les résultats pour les deux familles de modèles. J'obtiens de meilleurs résultats pour les modèles non-affine, avec et sans le noyau de prix quadratique et je ne trouve pas de diminution dans l'écart de performance entre le modèle affine et le non-affine. J'en conclus que l'utilisation du noyau de prix quadratique ne réduit pas significativement les coûts associés aux contraintes affines.

Mots clés : GARCH; noyaux de prix; volatilité stochastique.

Table des matières

Liste des tableaux et des figures

Remerciements

Je remercie tout d'abord mon directeur de thèse, Christian Dorion, qui m'a guidé dans mes recherches et la rédaction de ce mémoire, et qui a partagé généreusement son expertise. Je remercie également le Fonds de recherche du Québec – Nature et technologies pour avoir financé mes études.

Enfin, je tiens à remercier mes proches pour leur support, particulièrement mes parents, Michel et Raymonda, et ma copine, Pauline, qui m'ont comblé d'encouragements tout au long de mes études.

Introduction

En 1973, Black et Scholes ont développé une formule de tarification d'options qui persiste jusqu'à ce jour. Elle adopte une hypothèse de volatilité constante qui explique en grande partie sa faible performance empirique. En effet, loin d'être constante, on observe plutôt que la volatilité varie selon la période observée, et qu'elle se rassemble en périodes de forte ou faible volatilité. En combinaison avec l'hypothèse de distribution gaussienne des rendements, il en résulte le fameux « smirk » de volatilité implicite, provenant de la surestimation des prix des options lorsque le prix du sous-jacent est proche du prix d'exercice et de la sous-estimation des prix des options hors-jeu.

En réponse à ces faiblesses, Engle (1982) développe un modèle autorégressif conditionnel hétéroscédastique (ARCH ci-après), où la variance est fonction de l'innovation de la période précédente. Ce travail fut généralisé par son étudiant, Bollerslev (1986), produisant ainsi le modèle GARCH(p,q). Par la suite, Heston (1993) a développé un modèle à volatilité stochastique (SV ci-après). Son succès marqua la naissance d'une famille de modèles qui ont dominé la littérature des années 80 en tarification d'options et qui sont encore utilisés à ce jour. Similairement au sous-jacent, la variance suit un processus stochastique. Toutefois, elle ne peut être directement observée sur les marchés. Il en suit qu'il est nécessaire d'appliquer un filtre pour estimer sa valeur en tout temps. Finalement, motivé par cet inconvénient des modèles SV, Duan (1995) utilisa la relation de tarification localement risque-neutre pour adapter le modèle GARCH à la tarification d'options, propulsant cette famille de modèle à l'avant-garde de la littérature dans ce domaine.

Dans le cadre de ce mémoire, on ne s'attardera que sur les GARCH(1,1) avec discrétisation d'Euler:

$$y_t = x_t + \sqrt{h_t}\, z_t$$

$$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1} z_{t-1}^2$$

où $z \sim N(0,1)$

La variance ainsi définie est connue une période à l'avance. Sa dynamique se divise en trois composantes. La variance tend vers une moyenne long-terme, régie par le paramètre ω. Elle a une certaine persistance à travers le temps, c'est-à-dire que sa valeur dépend de celle à la période précédente. Cette composante, dirigée par β, permet de capturer l'effet de « clustering » de la volatilité observé dans les données et donne au modèle sa propriété d'auto-régression. Finalement, elle a une composante stochastique, gérée par α, qui modélise les nouvelles informations. Notons qu'elle est « indifférente » au signe de z, c'est-à-dire qu'elle réagit identiquement aux rendements positifs et négatifs. Cette propriété du modèle va directement à l'encontre de l'effet de levier soulevé par Black (1976), qui décrit la corrélation négative entre les rendements et les changements de volatilité.

Dans le but de remédier à cette faiblesse, Engle et Ng (1993) introduisent une version du modèle qui inclut un paramètre d'asymétrie. On nomme ce modèle le GARCH non-affine (NGARCH ci-après).

$$\ln\left(\frac{S_t}{S_{t-1}}\right) = R_t = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t} z_t$$

$$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1}(z_{t-1} - \gamma)^2$$

où S_t est la valeur d'un actif au temps t et r représente le taux sans risque.

Le paramètre γ, typiquement positif, introduit de l'asymétrie dans la réaction de la volatilité aux rendements. Un choc négatif aura un impact plus important sur la volatilité qu'un choc positif.

Alternativement, Heston et Nandi (2000) développent une variante affine du modèle (AGARCH ci-après).

$$R_t = r + \lambda h_t + \sqrt{h_t} z_t$$

$$h_t = \omega + \beta h_{t-1} + \alpha\left(z_{t-1} - \gamma\sqrt{h_{t-1}}\right)^2$$

Cette version leur permet de produire une solution semi-analytique pour le prix d'une option, permettant ainsi d'éviter les simulations Monte Carlo et le temps de calcul y étant attaché.

Cependant, les restrictions de volatilité affine imbriquées dans ce modèle s'avèrent couteuses. Plusieurs sources documentent leur infériorité par rapport à leurs contreparties non-affines. Entre autres, Hsieh et Richken (2005) trouvent qu'un modèle NGARCH est supérieur pour la réduction de biais dans les erreurs de tarification pour toutes maturités et degrés d'enjeu, particulièrement pour les options profondément hors-jeu. Supportant cette conclusion, Christoffersen, Dorion, Jacobs et Wang (2010, CDJW ci-après) trouvent que les modèles affines sous performent à la fois pour expliquer les rendements observés et pour la tarification d'options.

Étant donné que la comparaison affine versus non-affine constitue un des deux axes de comparaison étudié, il convient ici d'énoncer, discuter, et comparer certaines propriétés physiques de ces deux modèles. Débutant par la dynamique des rendements, on note que leurs espérances conditionnelles font en sorte que la dynamique risque-neutre conservera la forme non-affine ou affine. Pour le modèle affine, elle est linéaire en h_t, d'où son appellation.

Par définition, la variance des modèles GARCH est connue une période à l'avance. Deux périodes à l'avance, la variance du modèle non-affine est quadratique en h_t, contrairement au modèle affine. Cette caractéristique permet au NGARCH de mieux capter les périodes de forte incertitude. En effet, CDJW observent que le modèle NGARCH affiche une volatilité de la variance exceptionnellement élevée durant la première guerre du Golfe, la récession de 1990-1991, la chute de LTCM en 1998, et l'éclatement de la bulle dot-com de 2000-2002.

Pour les deux modèles, le paramètre d'asymétrie produit une corrélation négative entre les rendements et la variance lorsqu'il est positif, capturant l'effet de levier. Le modèle non-affine a une corrélation constante à travers le temps, contrairement au modèle affine où la corrélation est fonction de h_t. On note enfin qu'on s'attend à obtenir une valeur pour γ qui soit plusieurs ordres de magnitude plus élevée pour le modèle non-affine, de par la dynamique de la variance ainsi que les expressions de covariance et corrélations.

Pour la tarification d'options, les deux modèles utilisent le noyau de prix exponentiellement linéaire introduit par Rubinstein (1976). Plusieurs problèmes ont été soulevés par rapport à ce noyau de prix. Stein (1989) et Poteshman (2001) montrent que la variance implicite long-terme réagit de façon excessive aux changements de variance court-terme. De plus, il est amplement documenté, entre autres dans CDJW, que ces modèles ont de la difficulté à tarifer les options profondément hors-jeu. Enfin, plusieurs études ont cherché à réconcilier la distribution empirique des rendements avec la distribution risque-neutre implicite dans les prix d'options sans grand succès. Ensemble, ces problèmes sèment le doute sur la capacité du noyau de prix à expliquer les prix d'options et soulignent un important besoin pour un noyau plus général.

Suivant cette piste, Christoffersen, Heston, et Jacobs (2013, CHJ ci-après) proposent un nouveau noyau de prix incluant une prime pour la volatilité de la variance :

$$\frac{M_t}{M_{t-1}} = \left(\frac{S_t}{S_{t-1}}\right)^{\phi} \exp\left(\delta + \eta h_t + \xi(h_{t+1} - h_t)\right)$$

Ce noyau est fonction exponentiellement quadratique du rendement. Il imbrique le noyau de Rubinstein (1976), obtenu en posant $\xi = 0$. Les auteurs l'appliquent au modèle AGARCH(1,1) de Heston et Nandi (2000) et obtiennent d'excellents résultats. Le modèle ainsi obtenu (ACHJ ci-après) explique mieux les rendements et les prix d'options. Il produit des résultats comparables à un modèle « ad hoc » où les mesures physique et risque-neutre ne sont pas liées. Les auteurs concluent que ce nouveau modèle concilie les deux mesures avec succès et résout les problèmes du noyau exponentiellement linéaire.

Motivé par ces résultats, je cherche à appliquer le noyau de prix exponentiellement quadratique à un modèle NGARCH(1,1). J'obtiens ainsi deux axes de comparaison : affine versus non-affine, et noyau de prix linéaire versus quadratique. En somme, je compare 4 modèles GARCH(1,1) : Engle et Ng (1993), Heston et Nandi (2000), CHJ (2013), ainsi qu'un modèle non-affine avec noyau de prix quadratique que je développe en annexe (NCHJ ci-après). Spécifiquement, mon article cherche à répondre à la question titulaire : est-ce que les restrictions affines de volatilité sont encore coûteuses si le noyau de prix est quadratique?

La risque-neutralisation du modèle non-affine avec un noyau de prix identique à celui utilisé dans CHJ produit des paramètres dont la valeur varie dans le temps. Étant donné que les modèles non-affine exigent une tarification d'options par simulation Monte Carlo, avoir des paramètres qui nécessitent d'être recalculés à chaque pas de temps ajoute des coûts computationnels significatifs. Pour remédier à cette situation, j'adopte une légère variante du noyau de prix :

$$\frac{M_t}{M_{t-1}} = \left(\frac{S_t}{S_{t-1}}\right)^{\phi} \exp\left(\delta + \eta h_t + \xi\left(\frac{h_{t+1} - h_t}{h_t}\right)\right)$$

Cette modification, qui considère la différence relative de la variance plutôt que la différence absolue, ne produit qu'un seul paramètre dont la valeur varie stochastiquement.

Suivant CHJ, je calibre les modèles en maximisant une fonction de vraisemblance jointe. Cependant, je pondère la vraisemblance physique et risque-neutre de manière à ajuster pour le nombre d'observations de rendements et d'options, respectivement. Ce faisant, j'évite d'obtenir une vraisemblance physique négligeable par rapport à sa contrepartie risque-neutre, qui est construite à partir d'un jeu de données d'options considérablement plus grand que le nombre de rendements quotidiens. Ce choix représente une divergence par rapport à la méthodologie de CHJ. De plus, suivant CDJW, j'emploie la technique de ciblage de la variance pour la calibration du paramètre ω. Cette technique consiste à fixer la valeur du paramètre, qui régit l'espérance inconditionnelle, de façon à ce que cette dernière soit égale à la variance des rendements historiques.

Pour la tarification d'options, j'utilise la solution quasi-analytique de Heston et Nandi (2000) avec les modèles affines et la simulation Monte Carlo avec les modèles non-affine. Suivant la littérature, je choisi l'erreur quadratique moyenne relative sur la volatilité implicite (IV-RRMSE ci-après) comme mesure d'erreur de tarification. Utiliser une mesure relative évite d'accorder une importance disproportionnée aux options illiquides profondément hors-jeu, qui ont typiquement une volatilité implicite élevée.

Mon article est organisé comme suit. Dans la section 2, je présente les deux modèles GARCH(1,1) physiques et je discute brièvement de certaines de leurs caractéristiques. Dans la section 3, je dérive leurs mesures risque-neutre et je présente la méthodologie de tarification d'options. Dans la section 4, je calibre les modèles aux rendements et aux prix des options en utilisant une JMLE et je compare les résultats. Dans la section 5, je conclus et suggère des pistes pour de la recherche additionnelle.

[Page de garde]

Are Affine Volatility Restrictions Still Costly when the Pricing Kernel is Quadratic?

Patrick Aoun*

July 27, 2018

Abstract

I compare the quadratic kernel used by Christoffersen, Heston and Jacobs (2013) with the commonly used Rubinstein's (1976) power pricing kernel in terms of option valuation performance. I do so in both affine and nonaffine GARCH(1,1) models. I find that, in both cases, the quadratic kernel outperforms its linear counterpart. I find no evidence that the performance gap between the affine and the nonaffine models shrinks with the quadratic kernel.

Keywords: GARCH, pricing kernel, stochastic volatility

JEL Codes: G12

*HEC Montreal, Quebec, Canada

1 Introduction

Building on Engle (1982) and Bollerslev's (1986) work on GARCH(1,1) models, Engle and Ng (1993) introduced the widely used nonaffine GARCH(1,1) model. It was later used by Duan (1995) for option pricing. Following its success, Heston and Nandi (2000) proposed an affine variant. Its closed-form solution to option valuation allows for much easier pricing than the Monte Carlo alternative. However, it is now documented that the affine restrictions imposed by this family of models are quite costly both in terms of option pricing performance and asset return fit. Hsieh and Ritchken (2005) find that "a simple nonaffine NGARCH option model is superior in removing biases from pricing residuals for all moneyness and maturity categories especially for out-the-money contracts". Supporting these results, Christoffersen, Dorion, Jacobs and Wang (2010, henceforth CDJW) find that nonaffine models dominate affine models in terms of both fitting returns and option valuation.

On the other hand, Christoffersen, Heston and Jacobs (2013, henceforth CHJ) use an exponentially quadratic, U-shaped, pricing kernel. It includes an extra parameter which seeks to capture the risk of variance volatility in option valuation. They show that using the quadratic pricing kernel substantially improves pricing performance over the nested Rubinstein (1976) used in Heston and Nandi (2000).

Motivated by these findings, I seek to evaluate the performance gap between affine and nonaffine models using the quadratic kernel. I estimate four GARCH(1,1) models using a joint maximum likelihood estimation (henceforth JMLE) on both S&P500 returns and option contracts in order to benefit from the options' forward-looking nature. First, I evaluate the nonaffine and affine models with Rubinstein's exponentially linear pricing kernel. Then, I develop and estimate a nonaffine CHJ-like model (henceforth NCHJ) using a slight variant of their suggested pricing kernel. I compare it to the affine specification in CHJ (henceforth ACHJ). This results in 2 axes of comparison: affine versus nonaffine, and exponentially linear versus quadratic pricing kernel. I find that the quadratic pricing kernel significantly improves option fit in both models. However, the performance gap slightly widens, thus maintaining the nonaffine's superiority.

The article is organized as follows. In Section 2 I develop the two physical GARCH(1,1) asset models and I briefly discuss some of their characteristics. In Section 3 I derive their risk-neutral measures and present an option valuation methodology. In Section 4 I fit the models to asset returns and option prices using a JMLE and I compare the results. In Section 5 I conclude and suggest avenues for additional research.

2 The Physical Models

In this section I introduce both physical GARCH(1,1) models and discuss the affine versus nonaffine axis of comparison.

2.1 The Nonaffine GARCH(1,1)

Introduced by Engle and Ng (1993) and used for option pricing by Duan (1995), the nonaffine GARCH(1,1) model (henceforth NGARCH) has the following (Euler) discrete specification:

$$\begin{aligned} \ln(S_t) &= \ln(S_{t-1}) + r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}z_t \\ h_t &= \omega + \beta h_{t-1} + \alpha h_{t-1}(z_{t-1} - \gamma)^2, \end{aligned} \tag{1}$$

where S_t is the price of the underlying asset, r is the risk-free rate, λ is the variance risk premium, z_t is a N(0,1) shock, and h_t is the daily variance.

Let $R_t = \ln\left(\frac{S_t}{S_{t-1}}\right)$ denote the log returns. From the asset return dynamics in (1) we can derive the conditional return mean

$$\mathrm{E}_{t-1}\left[R_t\right] = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t. \tag{2}$$

This specification will produce a nonaffine dynamic for the conditional variance, under both physical and risk-neutral measures.

By design, the next-day variance is deterministic in GARCH models. However, the 2-days-ahead variance is stochastic. From (1) we can derive its conditional mean

$$\mathrm{E}_{t-1}\left[h_{t+1}\right] = \omega + \beta h_t + \alpha h_t \left(1+\gamma^2\right) \tag{3}$$

and conditional variance

$$\mathrm{Var}_{t-1}\left[h_{t+1}\right] = \alpha^2(2+4\gamma^2)h_t^2, \tag{4}$$

which we note is quadratic in h_t.

Also from (1), we derive the long-term unconditional expected value of the variance

$$\mathrm{E}\left[h_t\right] = \frac{\omega}{1-\beta-\alpha(1+\gamma^2)}. \tag{5}$$

The ω parameter can be fixed so that the unconditional expected value equals the historical variance of returns. This is true for both affine and nonaffine models. I will apply this optimization trick to every model considered in this article.

From (3), we see that the conditional variance h_t will mean-revert to its unconditional expected value with daily autocorrelation of

$$\pi = \beta + \alpha(1+\gamma^2). \tag{6}$$

Finally, using (1) to derive the conditional covariance between volatility and returns, as well as the ensuing correlation, yields

$$\mathrm{Cov}_{t-1}\left[R_t, h_{t+1}\right] = -2\alpha\gamma h_t^{\frac{3}{2}} \tag{7}$$

$$\mathrm{Corr}_{t-1}\left[R_t, h_{t+1}\right] = \frac{-2\gamma}{\sqrt{2+4\gamma^2}}. \tag{8}$$

The γ parameter, if larger than 0, captures the well-known "leverage effect" first documented by Black (1976). This stylized fact describes the negative correlation between returns and volatility. It has strong and well documented empirical support.

2.2 The Affine GARCH(1,1)

The affine specification (henceforth AGARCH) developed by Heston and Nandi (2000) has the following discrete form :

$$\begin{aligned} \ln(S_t) &= \ln(S_{t-1}) + r + \lambda h_t + \sqrt{h_t} z_t \\ h_t &= \omega + \beta h_{t-1} + \alpha (z_{t-1} - \gamma \sqrt{h_{t-1}})^2. \end{aligned} \tag{9}$$

Again, I derive the conditional return mean from the asset return dynamics,

$$\mathrm{E}_{t-1}\left[R_t\right] = r + \lambda h_{t+1}, \tag{10}$$

which is linear in h_t, hence the affine designation. From (9), the variance process' conditional mean and conditional variance are

$$\mathrm{E}_{t-1}\left[h_{t+1}\right] = \omega + \beta h_t + \alpha\left(1 + \gamma^2 h_t\right) \tag{11}$$

$$\mathrm{Var}_{t-1}\left[h_{t+1}\right] = 2\alpha^2 (1 + 2\gamma^2 h_t). \tag{12}$$

Note that the variance of variance is linear in h_t, unlike for the NGARCH. The unconditional expected value of h_t is

$$\mathrm{E}\left[h_t\right] = \frac{\omega + \alpha}{1 - \beta - \alpha\gamma^2}, \tag{13}$$

which requires $\omega + \alpha > 0$ to be positive. This allows ω to be negative. Eq (11) shows that the variance process has daily autocorrelation of

$$\pi = \beta + \alpha\gamma^2. \tag{14}$$

From (9), the conditional covariance and correlation are

$$\text{Cov}_{t-1}\left[R_t, h_{t+1}\right] = -2\alpha\gamma h_t \tag{15}$$

$$\text{Corr}_{t-1}\left[R_t, h_{t+1}\right] = \frac{-2\gamma\sqrt{h_t}}{\sqrt{2 + 4\gamma^2 h_t}}. \tag{16}$$

Note that the correlation is dependent on h_t, whereas the NGARCH conditional correlation is not time-varying.

3 Option Valuation Methodology

In this section I derive the risk-neutral measures for each model using the quadratic pricing kernel. Then, I suggest different pricing methods for the affine and nonaffine models and I show the convergence between these methods.

3.1 Risk Neutral Measures

The quadratic pricing kernel introduced by CHJ has the following discrete-time form:

$$\frac{M_t}{M_{t-1}} = \left(\frac{S_t}{S_{t-1}}\right)^{\phi} \exp\left(\delta + \eta h_t + \xi(h_{t+1} - h_t)\right) \tag{17}$$

where δ and η govern time-preference, ϕ the equity risk aversion, and ξ the variance risk aversion.

The novelty resides in the last part of the expression, $\xi(h_{t+1} - h_t)$, which seeks to capture the risk associated with changes in volatility. With $\xi = 0$ we find the nested Runbinstein

(1976) power pricing kernel.

Using it with the physical AGARCH(1,1) model described in (9) the authors get the following risk-neutral process:

$$\begin{aligned} \ln(S_t) &= \ln(S_{t-1}) + r - \frac{1}{2}h_t^* + \sqrt{h_t^*}z_t^* \\ h_t^* &= \omega^* + \beta h_{t-1}^* + \alpha^*(z_{t-1}^* - \gamma^*\sqrt{h_{t-1}^*})^2 \end{aligned} \tag{18}$$

where

$$\begin{aligned} z_t^* &\sim N(0,1) \\ h_t^* &= \frac{h_t}{1-2\xi\alpha} \\ \omega^* &= \frac{\omega}{1-2\xi\alpha} \\ \alpha^* &= \frac{\alpha}{(1-2\xi\alpha)^2} \\ \gamma^* &= (1-2\xi\alpha) \times \left(\lambda - \frac{1}{2} + \gamma\right) + \frac{1}{2}. \end{aligned}$$

Following a similar method, I derivate the risk-neutral process for the NCHJ model. However, using the pricing kernel in (17) leads to every risk-neutral parameter being time-varying. While this is theoretically acceptable, it will hinder Monte Carlo computation speed, as it requires recalculating parameter values at every time step. In an attempt to mitigate this inconveniance, I slightly modify the pricing kernel, leaving only γ_t^* to be time-varying:

$$\frac{M_t}{M_{t-1}} = \left(\frac{S_t}{S_{t-1}}\right)^{\phi} \exp\left(\delta + \eta h_t + \xi\left(\frac{h_{t+1} - h_t}{h_t}\right)\right). \tag{19}$$

I find the following risk-neutral process:

$$\begin{aligned} \ln(S_t) &= \ln(S_{t-1}) + r - \frac{1}{2}h_t^* + \sqrt{h_t^*}z_t^* \\ h_t^* &= \omega^* + \beta h_{t-1}^* + \alpha^* h_{t-1}^*\left(z_{t-1}^* - \gamma_{t-1}^*\right)^2 \end{aligned} \tag{20}$$

where

$$
\begin{aligned}
z_t^* &\sim N(0,1) \\
h_t^* &= \frac{h_t}{1-2\xi\alpha} \\
\omega^* &= \frac{\omega}{1-2\xi\alpha} \\
\alpha^* &= \frac{\alpha}{1-2\xi\alpha} \\
\gamma_t^* &= \sqrt{1-2\xi\alpha}\,(\lambda+\gamma) + \xi\alpha\sqrt{h_t^*}
\end{aligned}
$$

(*Details in Appendix A.*)

The risk-neutral variance dynamic is nonaffine and similar to the physical one, due to the chosen asset returns specification. From (18) and (20), it is evident that the ξ parameter has a major impact on the risk-neutral dynamics. It allows conditional variance, variance of variance, and persistence to differ from the physical measure. Setting $\xi = 0$ reveals the nested risk-neutral measures for the AGARCH and NGARCH models respectively.

3.2 Option Valuation

Under the nonaffine models, option pricing is done using a Monte Carlo approximation. I simulate 100,000 asset prices using the risk-neutral process described in (20) and compute the sample mean of the discounted payoffs.

For a european call option, this translates into:

$$
\begin{aligned}
C(t,T,K) &= \exp(-r(T-t))\,\mathrm{E}_t^{\mathbb{Q}}\left[\max(S_T-K,0)\right] \\
&\approx \exp(-r(T-t)) \\
&\quad \times \frac{1}{N}\sum_{i=1}^{N}\big[\max(S_{i,T}^*-K,0)\big],
\end{aligned}
\tag{21}
$$

where T is the maturity day, S^* the risk-neutral underlying asset price, K the strike price, and N the number of simulated paths.

For option pricing under the affine models, I use the closed-form option valuation formula developed in Heston and Nandi (2000) and presented in Appendix B. This method is considerably more efficient than Monte Carlo. Figure 1 shows the convergence between both pricing methods on synthetic call options with various maturities, spot volatilities, and moneyness, defined as the implied futures price divided by strike price.

4 Asset Return and Option Valuation Empirics

In this section, I fit the different models to return and option data using a joint likelihood function. I present the results of the JMLE in terms of fit and option pricing for each of the four models. Then, I examine the option pricing performance by moneyness, maturity, and volatility level. Finally, I make a short economic assessment of the option pricing.

4.1 Data

For my empirical analysis, I use daily S&P500 returns based on close price provided by the Center for Research in Security Prices. This covers the period from 1990 to 2013. Table 1 presents basic descriptive statistics on this series, as well as on a subset of the series covering the same period as the option contracts.

I also use out-of-the-money (OTM) S&P500 call and put options for the period from 1996 to 2013, the full sample period available from OptionMetrics at the time of writing. I apply the standard filters from Bakshi, Cao, and Chen (1997) and restrict the analysis to options traded on Wednesdays. This represents 39,571 option contracts with an average implied volatility (IV) of 20.96%. Table 1 presents descriptive statistics by moneyness and the number of days before maturity (DTM).

4.2 Joint Likelihood

Following CHJ, I estimate parameter values by maximizing the natural log of a two fold joint likelihood function, based on asset returns and option prices:

$$\max_{\Theta,\Theta^*} \left(\frac{N_{\text{ret}} + N_{\text{opt}}}{2}\right)\left(\frac{\ln L^R}{N_{\text{ret}}} + \frac{\ln L^Q}{N_{\text{opt}}}\right) \tag{22}$$

where N_{ret} and N_{opt} are the numbers of returns and options respectively, $\Theta = \{\omega, \alpha, \beta, \gamma, \lambda\}$ are the physical parameters, and Θ^* are the risk-neutral parameters, mapped according to (20). Weighting the physical and risk-neutral log likelihoods as such allows me to use the full options dataset without having the risk-neutral likelihood eclipse the physical one.

Using the conditional density function of daily returns,

$$f\left(R_t|h_t\right) = \frac{1}{\sqrt{2\pi h_t}} \exp\left(-\frac{(R_t - r - (\lambda + \frac{1}{2})h_t)^2}{2h_t}\right), \tag{23}$$

we get the following expression for the physical log likelihood:

$$\ln L^R = -\frac{1}{2}\sum_{i=1}^{N_{\text{ret}}}(\ln(2\pi h_i) + z_i^2) \tag{24}$$

where z_i and h_i are filtered from the time-series using (1) and (9) for, respectively, the nonaffine and affine models.

Following literature, I use relative IV pricing errors as the basis for the risk-neutral log likelihood,

$$\epsilon_i = \left(\frac{\text{IV}_i^{\text{MDL}}}{\text{IV}_i^{\text{MKT}}} - 1\right), \tag{25}$$

where IV_i^{MDL} and IV_i^{MKT} are the Black-Scholes volatility implied by the model option price and market price, respectively. Using a relative error measure avoids putting too much weight on illiquid deep-OTM options that have higher IV and are harder to price.

Following CHJ, we assume that these errors are i.i.d. and normally distributed. The risk-neutral log likelihood is

$$\ln L^Q = -\frac{1}{2} \sum_{i=1}^{N_{\text{opt}}} \left(\ln(2\pi\sigma_\epsilon^2) + \frac{\epsilon_i^2}{\sigma_\epsilon^2} \right), \tag{26}$$

where σ_ϵ^2 is the sample variance of ϵ_i.

4.3 Overall Results

Table 2 presents the estimated parameter values, the log likelihoods, and some physical properties for each of the 4 models. Standard errors, in parenthesis, are calculated as in Bollerslev and Wooldridge (1992). As in CHJ, the value of $(1-2\xi\alpha)^{-1}$ is shown in lieu of ξ because it represents the mapping between physical and risk-neutral variance, making it a more meaningful value.

For both affine and nonaffine models, I fix ω's value so that the long-term unconditional mean is equal to the daily return variance, i.e.

$$\mathrm{E}\left[h_t\right] = \widehat{\mathrm{Var}}\left[R_t\right] = 1.3377 \times 10^{-4}. \tag{27}$$

This represents a 18.36% annualized volatility. Because ω is fixed, no standard error is displayed for this parameter.

Figure 2 shows the annualized conditional volatility path of the models, i.e. $100\sqrt{252 h_t}$, covering the full period of the S&P500 time-series described in Section 4.1. The higher volatility of variance in nonaffine models is evident, especially during the 2008 financial crisis. The right column shows the relative difference between the conditional volatilities of the CHJ models and those of their non-CHJ counterparts. While there doesn't seem to be a strong pattern otherwise, the CHJ models do show consistently lower conditional volatility. This suggests that while the non-CHJ models need to over-estimate physical volatility for better option fit, the CHJ models use the ξ parameter to boost risk-neutral volatility as allowed by the mapping in (18) and (20). The higher physical likelihoods in the CHJ models support this hypothesis.

A quick review of the physical properties of the evaluated models reveals small but significant differences across both axes of comparison. Variance of variance, driven by α, goes down with the addition of the volatility premium. This finding was reported in CHJ and seems to apply to the NCHJ model as well. As in CDJW, nonaffine models exhibit variance of variance that is an order of magnitude higher than affine models. Variance persistence, driven by β, is very high across the board though slightly higher in the CHJ models. The leverage correlation, driven by γ, captures the "leverage effect" stylized fact mentioned in Section 2. Unsurprisingly, it is strongly negative in every model, and more so in the CHJ models.

From the JMLE results in Table 2, we note that the superior joint likelihood of the nonaffine models is almost entirely due to better option fit. While physical likelihood is similar across the board, risk-neutral likelihood is what truly distinguishes both families of models. This uncontroversial result supports findings in CDJW that option RMSE is significantly lower in the nonaffine GARCH(1,1) model.

Additionally, we observe better overall likelihoods for the CHJ models. As expected, most of the improvement results from better option fit and higher risk-neutral likelihood. While the pricing kernel does not directly affect return fit, the addition of the ξ parameter reduces the duality of roles of other parameters, which are otherwise expected to fit both returns and options. Thus, it creates some "slack" in the JMLE, resulting in a better physical likelihood.

Interestingly, the addition of the volatility premium results in a greater improvement in the nonaffine models. In combination with CHJ's finding of "strong model-free evidence of a U-shaped pricing kernel", this suggests that the quadratic pricing kernel doesn't necessarily relax the affine restrictions, but rather outperforms Rubinstein's pricing kernel in a more general sense. These results answer the article's titular question: affine volatility restrictions are still costly when the pricing kernel is quadratic.

4.4 Option Valuation Results by Moneyness and Maturity

Table 3 presents relative IV RMSE (IV RRMSE) and relative bias by moneyness (strike price over forward price) and maturity for each model. The bins are the same as in Table 1. Figure 3 shows these results as a plot. The y-axis is the IV RRMSE in the top row, bias in the bottom row, and the x-axis is moneyness in the left column and days to maturity in the right column. The nonaffine models are marked with "+" and the affine models with "○". NCHJ and ACHJ are represented by dashed lines. Relative bias is the mean of the relative IV errors defined in Eq (25).

The kink where moneyness is above 1.1 is perplexing at first. Considering that I exclusively use OTM options, put-call parity leads us to expect equivalent pricing errors for a call at moneyness K/F and a put at F/K, everything else being equal. However, this visibly isn't the case for affine models. Further investigation reveals that options in that last bin disproportionately coincide with high spot conditional variance for the affine models only. This would explain the asymmetry, as high spot variance increases the chance of OTM options ending in the money. It seems that, as many authors argue, a good measure of moneyness should take spot variance into consideration. Incidentally, this also explains why, while the NGARCH clearly dominates the AGARCH in terms of IV RRMSE and bias magnitude in most bins, it presents a higher *overall* bias magnitude. Indeed, the AGARCH's high positive bias in these bins reduces overall bias magnitude, which is negative for all models.

As in CDJW, we see that affine models under perform in the pricing of OTM options, both in terms of error and bias, and that this gap shrinks at the money. The authors suggest that the affine models "perhaps [do not provide] sufficient nonnormality at the relevant horizons." This could also explain the gap for short maturity options, which show high IV RRMSE across the board that is notably higher in affine models. Negative bias for these options, while admittedly small, suggests indeed that fewer paths ended in the money than was required. Surprisingly, affine models seem to outperform for at-the-money and long-term options. This finding goes against CDJW, where nonaffine RMSE is consistently lower. This is surely because IV RRMSE gives smaller importance to cheap OTM options, which we already know affine models price poorly. It does however signify that affine models are not strictly inferior to nonaffine models.

Finally, we observe that the CHJ models' advantage is inconsistent and difficult to isolate. In fact, they slightly under perform in some bins. For instance, while the NCHJ model significantly outclasses the NGARCH for deep OTM option pricing, the ACHJ is weaker than the AGARCH in that same bin. In other words, the improvement obtained by the addition of the volatility premium does not seem to be strongly attributable to specific moneyness or maturity bins and a significant pattern does not emerge from the plots. Instead, we see an average improvement spread across every bin.

4.5 Economic Assessment of Option Valuation Performance

Table 4 presents the results of the regression of weekly IV RRMSE (Panel A) and weekly relative bias (Panel B) on the following economic variables and option characteristics: the VIX index, the S&P500 weekly return, the crude Brent Oil price, the 3-month T-bill rate, the credit spread, the term spread, as well as average option moneyness and maturity. Credit spread is defined as the yield on corporate bonds rated Baa minus the yield on Aaa bonds, as rated by Moody's. Term spread is defined as the difference between the yield on 10-year T-bond and the 3-month T-bill rate. In parenthesis, I present the t-statistics of each regressor, computing the standard deviation with White's robust variance matrix. Values greater than 2 in absolute terms are in bold type.

The results show that nonaffine models exhibit much lower sensitivity of pricing performance to market conditions. Panel A shows that the relationship between IV RRMSE and every economic variable is weaker for nonaffine models, except in the case of S&P500 weekly returns. Moreover, Panel B shows that the affine models' biases have a significant relationship with every economic variable while only the S&P500 returns correlate significantly with the nonaffine models' biases. Accordingly, note the much lower R-squared for the nonaffine models in both regressions.

After a negative weekly return on the S&P500, the IV RRMSE of nonaffine models significantly decreases and the bias *magnitude* of all models decreases. A return one standard deviation below the average weekly return of the index increases model biases by 0.84% (ACHJ) to 1.69% (NGARCH). Given that the average bias ranges from -3.39% to -3.91%, this means that the impact of weekly returns is of clear economic significance. At first sight, the fact the models are more precise after the market experiences a negative return might be puzzling. It also seems at odds with the sign of the coefficients on the level of the VIX. However, this could be explained by the negative loading on the volatility. Finally, we see that both IV RRMSE and bias magnitude are significantly negatively correlated with moneyness for all models. Additionally, while NGARCH IV RRMSE and bias are both significantly correlated with maturity, only AGARCH IV RRMSE is. These results support the findings in Table 3 discussed in the previous section.

On the other hand, using the quadratic pricing kernel makes no significant difference neither on the overall R-squared nor on specific regressors' significance. Apparently, while the quadratic pricing kernel improves the option fit of the models, it does not change the fact that pricing errors are statistically and economically significantly related to market conditions, especially for affine models. Dorion (2016) shows that accounting for business conditions can substantially improve the pricing performance of GARCH models. In his analysis, however, he does not consider a quadratic pricing kernel. In light of my results, an interesting avenue for future research could be to investigate the link between market conditions and the prices of risk in a quadratic pricing kernel context.

5 Conclusion

I have compared the asset returns and option pricing fits of 4 GARCH(1,1) models. Using a traditional exponentially linear pricing kernel, the NGARCH(1,1) model outperforms the AGARCH(1,1) as is well documented in the literature. Using a quadratic pricing kernel leads to better option pricing, and to a lesser extent better return fit, in both models. However, the improvement is stronger in the nonaffine model, thus widening the performance gap. I conclude that there is no evidence of relaxation in the affine volatility restrictions.

A few directions for future research arise from these results. First and foremost, I suggest attempting a similar study with the component versions of the GARCH(1,1) models I've been using in this article. First developed by Engle and Lee (1999), component GARCH models decompose volatility into short-term and long-term components. CDJW finds "strong evidence for the component structure in the affine GARCH models, but less so in the nonaffine models." Additionally, Babaoğlu, Christoffersen, Heston, and Jacobs (2017, henceforth BCHJ) develop an affine GARCH model with "multiple volatility factors, fat-tailed return innovations, and a variance dependent pricing kernel." Amongst these 3 features, they find that the U-shaped pricing kernel and the long-term volatility component are the most important ones, improving option fit by 17% and 9% respectively. Developing a nonaffine GARCH models with these features would fill a gap in the literature and bring the affine/nonaffine comparison to the forefront of current advancements.

Second, this article focused entirely on European option contracts. Applying these models to American-style option valuation would be a natural expansion to the results. Stentoft (2005) suggests a "new simulation technique" for pricing such options that easily accommodates GARCH models. It is worth noting that this would take away a significant advantage of the affine models; the closed-form option valuation formula does not apply to American options.

Finally, in a similar fashion, we could expand the study by using nonnormal shocks.

CDJW presents some evidence for a generalized error distribution. They found that it "improves the fit of all models to daily returns, but the improvement in option valuation is much less evident." We could also look to Christoffersen, Heston, and Jacobs (2006) and BCHJ (2016) for inverse-Gaussian distributions, or Duan, Ritchken, and Sun (2006) for jumps. We already know that affine models have a hard time providing sufficient nonnormality. We also know from BCHJ (2016) that the inverse-Gaussian shocks are complimentary to the U-shaped pricing kernel. Thus, it is reasonable to assume that if the nonaffine models do not profit from such complementarity the performance gap might shrink.

References

Babaoğlu, Kadir, Peter Christoffersen, Steven Heston, and Kris Jacobs (2017). "Option Valuation with Volatility Components, Fat Tails, and Nonmonotonic Pricing Kernels". In: *The Review of Asset Pricing Studies, forthcoming*.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen (1997). "Empirical performance of alternative option pricing models". In: *The Journal of Finance* 52, pp. 2003–2049.

Black, Fischer (1976). "Studies of stock price volatility changes". In: *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section*, pp. 177–181.

Black, Fischer and Myron Scholes (1973). "The Pricing of Options and Corporate Liabilities". In: *Journal of Political Economy* 81, pp. 637–659.

Bollerslev, Tim (1986). "Generalized autoregressive conditional heteroskedasticity". In: *Journal of econometrics* 31, pp. 307–327.

Bollerslev, Tim and Jeffrey M Wooldridge (1992). "Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances". In: *Econometric reviews* 11, pp. 143–172.

Christoffersen, Peter, Christian Dorion, Kris Jacobs, and Yintian Wang (2010). "Volatility components, affine restrictions, and nonnormal innovations". In: *Journal of Business & Economic Statistics* 28, pp. 483–502.

Christoffersen, Peter, Steve Heston, and Kris Jacobs (2006). "Option valuation with conditional skewness". In: *Journal of Econometrics* 131, pp. 253–284.

Christoffersen, Peter, Steven Heston, and Kris Jacobs (2013). "Capturing option anomalies with a variance-dependent pricing kernel". In: *The Review of Financial Studies* 26, pp. 1963–2006.

Dorion, Christian (2016). "Option Valuation with Macro-Finance Variables". In: *Journal of Financial and Quantitative Analysis* 51, pp. 1359–1389.

Duan, Jin-Chuan (1995). "The GARCH option pricing model". In: *Mathematical finance* 5, pp. 13–32.

Duan, Jin-Chuan, Peter Ritchken, and Zhiqiang Sun (2006). "Approximating GARCH-Jump Models, Jump-Diffusion Processes, and Option Pricing". In: *Mathematical Finance* 16, pp. 21–52.

Engle, Robert F (1982). "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation". In: *Econometrica: Journal of the Econometric Society* 50, pp. 987–1007.

Heston, Steven L (1993). "A closed-form solution for options with stochastic volatility with applications to bond and currency options". In: *The review of financial studies* 6, pp. 327–343.

Heston, Steven L and Saikat Nandi (2000). "A closed-form GARCH option valuation model". In: *The review of financial studies* 13, pp. 585–625.

Hsieh, KC and Peter Ritchken (2005). "An empirical comparison of GARCH option pricing models". In: *Review of derivatives research* 8, pp. 129–150.

Lee, Gary and Robert Engle (1999). "A permanent and transitory component model of stock return volatility". In: *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W.J. Granger*, pp. 475–497.

Ng, Victor and Robert Engle (1993). "Measuring and Testing the Impact of News on Volatility". In: *Journal of Finance* 48, pp. 1749–1778.

Poteshman, Allen M (2001). "Underreaction, overreaction, and increasing misreaction to information in the options market". In: *The Journal of Finance* 56, pp. 851–876.

Rubinstein, Mark (1976). "The valuation of uncertain income streams and the pricing of options". In: *Bell Journal of Economics* 7, pp. 407–425.

Stein, Jeremy (1989). "Overreactions in the options market". In: *The Journal of Finance* 44, pp. 1011–1023.

Stentoft, Lars (2005). "Pricing American options when the underlying asset follows GARCH processes". In: *Journal of Empirical Finance* 12, pp. 576–611.

Tables

Table 1 : Returns and Options Data

Panel A : Return Characteristics (Annualized)

	1990-2013	**1996-2013**
Mean	6.89%	6.26%
St. Dev.	18.36%	20.16%
Skewness	-0.236	-0.228
Kurtosis	11.62	10.50

Panel B : Option Data by Moneyness

K/F ∈	[0, 0.85]	]0.85, 0.90]	]0.90, 0.95]	]0.95, 1.00]	]1.00, 1.05]	]1.05, 1.10]	]1.10, ∞	All
Count	3700	4691	5653	6677	9330	5413	4107	39571
Avg IV	29,81%	26,47%	23,36%	20,23%	16,98%	17,11%	18,68%	20,96%
Avg Price	11,86	16,92	25,10	37,90	29,084	17,21	11,50	23,50
Avg Spread	1,26	1,42	1,61	1,82	1,64	1,43	1,35	1,54

Panel C : Option Data by Maturity

DTM ∈	[0, 30]	]30, 60]	]60, 90]	]90, 120]	]120, 180]	]180, 270]	]270, ∞	All
Count	5343	8279	6636	3837	4935	6002	4539	39571
Avg IV	20,96%	20,96%	20,99%	21,55%	20,78%	20,84%	20,52%	20,96%
Avg Price	5,73	10,70	17,20	22,92	28,09	39,45	51,39	23,50
Avg Spread	0,65	1,02	1,43	1,72	1,82	2,11	2,52	1,54

Table 2 : Parameter Estimation and Model Fit

Parameter values	**NGARCH**	**NCHJ**	**AGARCH**	**ACHJ**
ω	1.323E-06	1.261E-06	-9.765E-07	-8.919E-07
α	0.04635	0.04422	2.194E-06	1.831E-06
	(4.87E-04)	(5.06E-04)	(8.33E-09)	(8.51E-05)
β	0.8769	0.8721	0.8986	0.9058
	(0.0011)	(0.0011)	(4.13E-04)	(4.26E-04)
γ	1.2011	1.2961	205.15	218.21
	(0.0168)	(0.0173)	(3.57E-07)	(8.24E-07)
λ	0.05492	0.02521	3.930	2.047
	(0.0036)	(0.0025)	(5.91E-07)	(2.95E-05)
$(1-2\xi\alpha)^{-1}$	1	1.1079	1	1.0805
		(0.0056)		(0.0037)
Physical properties				
Annual volatility target	0.1836	0.1836	0.1836	0.1836
Daily autocorrelation	0.9901	0.9906	0.9909	0.9930
Volatility of variance	0.1463	0.1393	0.0433	0.0369
Correlation$[R_t, h_{t+1}]$	-0.921	-0.967	-0.931	-0.934
Empirical z kurtosis	4.5794	4.5819	5.2456	5.2643
Log Likelihoods				
Physical	19,749	19,765	19,612	19,622
Risk-neutral	12,235	12,377	10,685	10,713
IVRRMSE	17.75%	17.69%	18.46%	18.45%
Joint log-likelihood	81,527	81,666	80,117	80,168

Volatility of variance is the annual volatility of the annual variance, i.e. $\sqrt{h_t \times 252^3}$. Standard errors are given in parenthesis under the estimated value. Because variance targeting fixes the value of ω, no standard errors are given for that parameter. Where the physical properties are time-varying, their mean over the entire series is presented. Parameters and properties are daily unless specified otherwise.

Table 3 : IV RRMSE and Relative Bias by Moneyness and Maturity

Panel A : IV RRMSE (%) by Moneyness

K/F ∈	[0, 0.85]	]0.85, 0.90]	]0.90, 0.95]	]0.95, 1.00]	]1.00, 1.05]	]1.05, 1.10]	]1.10, ∞	All
NGARCH	18.65	17.95	16.93	15.77	**17.84**	18.87	19.06	17.75
NCHJ	**17.95**	**17.46**	**16.55**	15.69	18.21	19.06	19.18	**17.69**
AGARCH	23.14	21.36	18.31	14.80	18.22	18.55	15.99	18.46
ACHJ	23.43	21.55	18.40	**14.66**	18.04	**18.41**	**15.88**	18.45

Panel B : Relative Bias (%) by Moneyness

K/F ∈	[0, 0.85]	]0.85, 0.90]	]0.90, 0.95]	]0.95, 1.00]	]1.00, 1.05]	]1.05, 1.10]	]1.10, ∞	All
NGARCH	-11.77	-10.49	-8.11	-4.46	**0.85**	**1.56**	**-0.62**	-3.91
NCHJ	**-10.86**	**-9.75**	**-7.51**	-3.93	1.49	1.77	-0.68	**-3.39**
AGARCH	-18.69	-15.15	-10.06	**-3.21**	3.26	4.54	2.52	-3.87
ACHJ	-19.20	-15.61	-10.45	-3.41	3.44	5.00	3.28	-3.88

Panel C : IV RRMSE (%) by Maturity

DTM ∈	[0, 30]	]30, 60]	]60, 90]	]90, 120]	]120, 180]	]180, 270]	]270, ∞	All
NGARCH	**20.43**	18.78	17.17	16.53	16.56	16.35	17.27	17.75
NCHJ	20.69	**18.41**	**17.05**	16.51	16.61	16.40	17.18	**17.69**
AGARCH	24.76	20.80	17.16	**15.52**	15.49	15.32	15.99	18.46
ACHJ	24.78	20.82	17.18	15.53	**15.48**	**15.23**	**15.83**	18.45

Panel D : Relative Bias (%) by Maturity

DTM ∈	[0, 30]	]30, 60]	]60, 90]	]90, 120]	]120, 180]	]180, 270]	]270, ∞	All
NGARCH	-5.70	-5.91	-4.80	-5.81	-3.33	**-1.55**	1.01	-3.91
NCHJ	**-4.18**	**-4.36**	**-3.84**	**-5.58**	**-3.18**	-2.10	**-0.09**	**-3.39**
AGARCH	-5.43	-5.73	-4.84	-5.74	-3.41	-1.67	0.90	-3.87
ACHJ	-5.34	-5.66	-4.82	-5.67	-3.45	-1.79	0.75	-3.88

Relative bias is the mean of errors as defined in Eq (25). Moneyness is the ratio of Strike price on Forward price (K/F). Pricing done on the full option dataset (1996-2013) described in Section 4.1 using parameters from Table 2.

Table 4 : Pricing Errors Regression on Economic Variables

Panel A : Weekly IV RRMSE

	NGARCH	**NCHJ**	**AGARCH**	**ACHJ**
Average IV RRMSE	17.75%	17.69%	18.46%	18.45%
Coefficients and *t*-statistics				
Constant	**1.402**	**1.282**	**1.718**	**1.699**
	(4.079)	(3.628)	(5.623)	(5.686)
VIX	**0.0034**	**0.0035**	**0.0025**	**0.0026**
	(2.825)	(2.848)	(3.650)	(3.867)
S&P500 weekly return	**0.0079**	**0.0080**	0.0027	0.0025
	(2.782)	(2.807)	(1.533)	(1.419)
Crude Brent oil price	-0.000238	-0.000192	0.000073	0.000107
	(-1.055)	(-0.839)	(0.428)	(0.644)
3-month T-bill rate	0.0143	0.0129	**0.0247**	**0.0247**
	(1.623)	(1.455)	(4.140)	(4.254)
Credit spread	**-0.026**	**-0.024**	**-0.031**	**-0.030**
	(-2.420)	(-2.167)	(-4.468)	(-4.512)
Term spread	0.0095	0.0067	0.0161	0.0159
	(0.817)	(0.571)	(1.811)	(1.854)
Average moneyness (K/F)	**-1.361**	**-1.245**	**-1.689**	**-1.674**
	(-3.836)	(-3.408)	(-5.330)	(-5.399)
Average maturity (DTM)	**0.00074**	**0.00077**	**0.00067**	**0.00066**
	(2.690)	(2.763)	(2.857)	(2.862)
Adjusted R-squared	0.1071	0.0997	0.1919	0.1950

Panel B : Weekly relative bias

	NGARCH	**NCHJ**	**AGARCH**	**ACHJ**
Average relative bias	-3.91%	-3.39%	-3.87%	-3.88%
Coefficients and *t*-statistics				
Constant	**-3.327**	**-3.367**	**-3.823**	**-3.769**
	(-4.338)	(-4.345)	(-6.414)	(-6.471)
VIX	-0.0031	-0.0031	**-0.0119**	**-0.0120**
	(-1.109)	(-1.100)	(-9.139)	(-9.451)
S&P500 weekly return	**-0.0169**	**-0.0185**	**-0.0098**	**-0.0084**
	(-3.064)	(-3.303)	(-2.774)	(-2.451)
Crude Brent oil price	0.000939	0.000953	**0.000661**	**0.000635**
	(1.940)	(1.942)	(2.089)	(2.037)
3-month T-bill rate	-0.0133	-0.0115	**-0.0589**	**-0.0618**
	(-0.647)	(-0.552)	(-4.568)	(-4.882)
Credit spread	0.039	0.037	**0.099**	**0.102**
	(1.715)	(1.628)	(6.846)	(7.145)
Term spread	-0.0525	-0.0508	**-0.0916**	**-0.0936**
	(-1.966)	(-1.885)	(-5.154)	(-5.389)
Average moneyness (K/F)	**3.237**	**3.286**	**4.031**	**3.990**
	(4.064)	(4.081)	(6.609)	(6.687)
Average maturity (DTM)	**0.00152**	**0.00144**	0.00068	0.00063
	(2.590)	(2.435)	(1.436)	(1.354)
Adjusted R-squared	0.1733	0.1736	0.3230	0.3309

Relative bias is the mean of errors as defined in Eq (25). Credit spread is defined as the yield on corporate bonds rated Baa minus the yield on Aaa bonds, as rated by Moody's. Term spread is defined as the difference between the yield on 10-year T-bond and the 3-month T-bill rate. Moneyness is the ratio of Strike price on Forward price (K/F). S&P500 weekly returns are normalized. t-statistics, in parenthesis, are obtained using White's robust standard errors. Coefficients with an associated t-stat larger than two in absolute value are in bold.

Figures

Figure 1: Accuracy of Monte Carlo Approximation

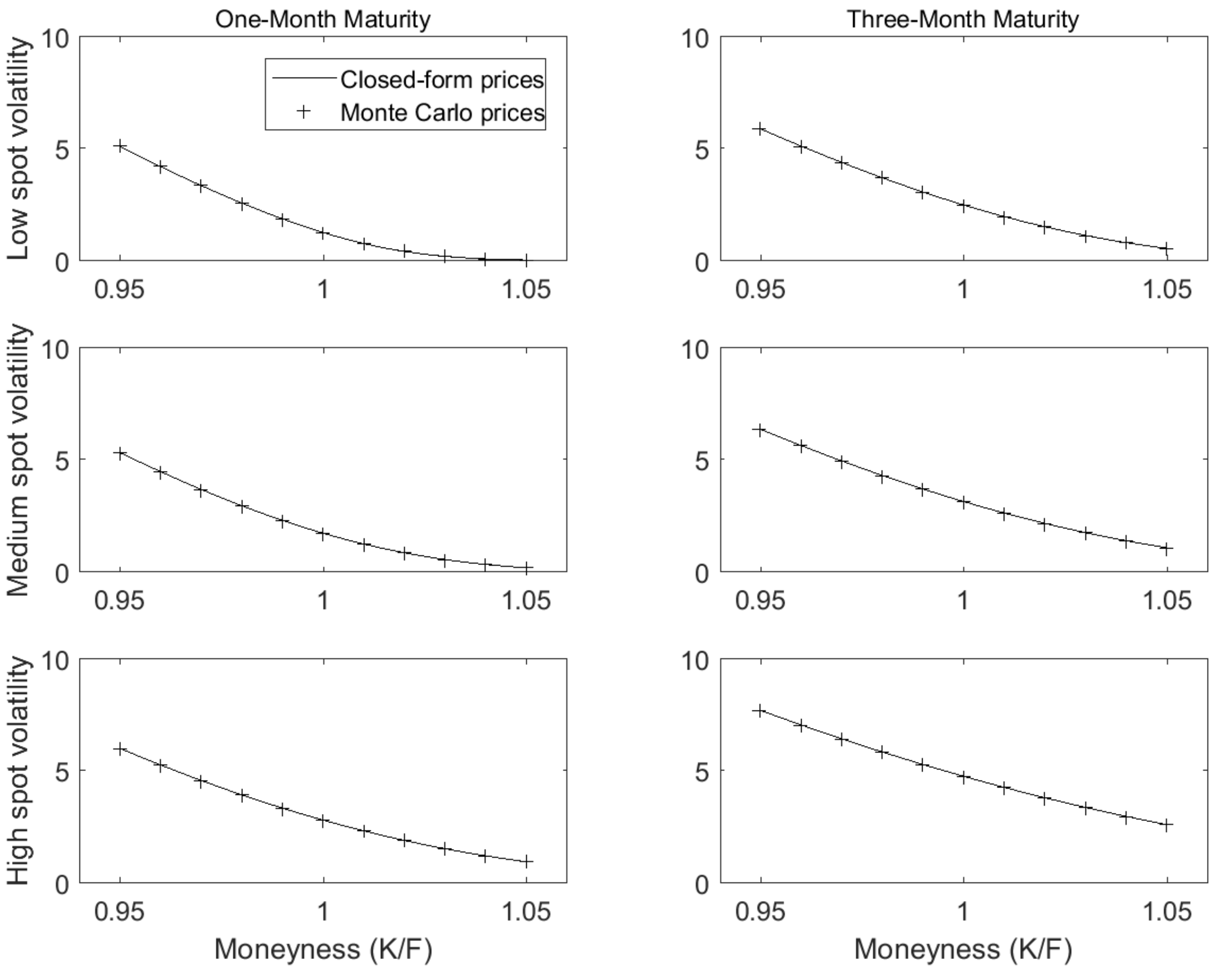

Using the AGARCH(1,1) model and the parameters from Table 2, I price call options using Monte Carlo approximation and the Heston-Nandi (2000) closed-form formula. The three spot volatility levels are 9.46%, 15.26%, and 25.52%, corresponding to the 5th, 50th, and 95th percentile of the model volatility on the returns dataset. The risk-free rate is 5%.

Figure 2: Conditional Volatility Paths

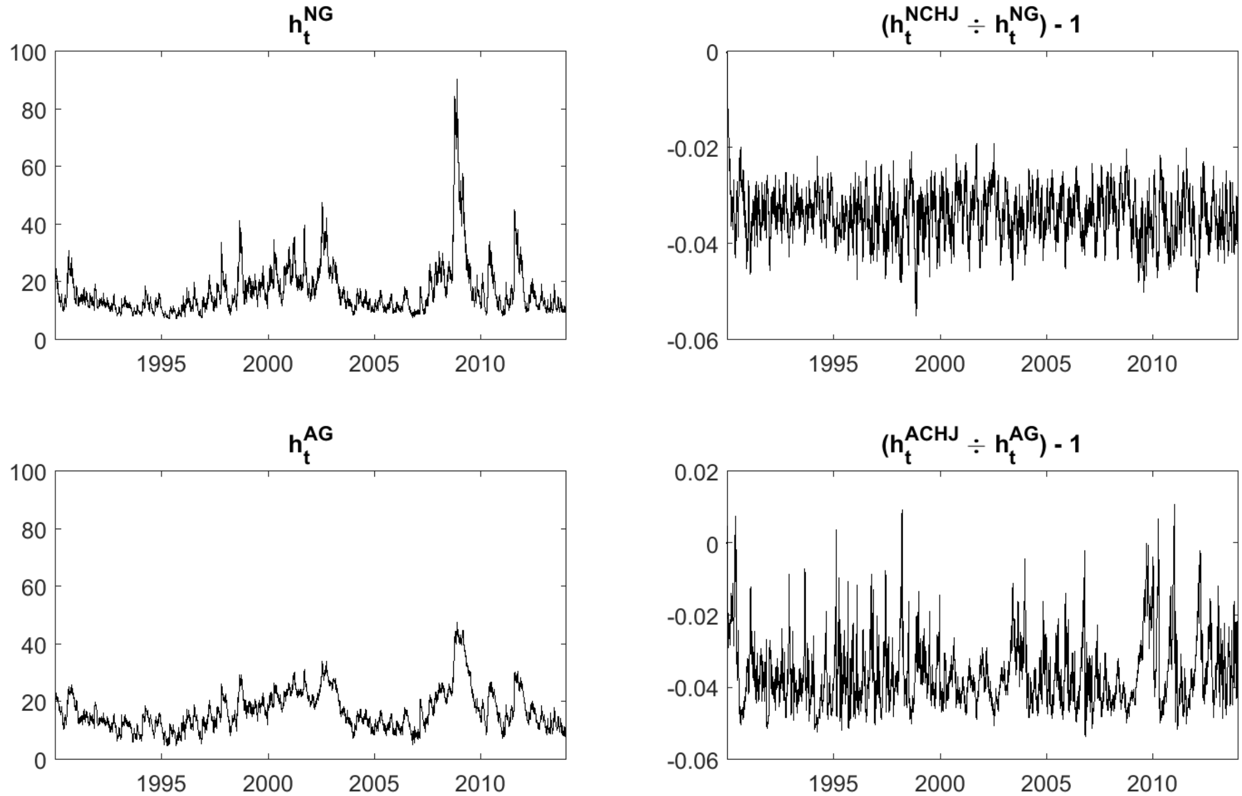

In the left column, I plot the physical annualized conditional volatility path of the NGARCH (h_t^{NG}) and AGARCH (h_t^{AG}) models, i.e. $100\sqrt{252 h_t}$. In the right column, I plot the relative difference with their CHJ counterparts. The plots cover the full period of the S&P500 time-series described in Section 4.1. The parameter values are taken from Table 2.

Figure 3: IV RRMSE and Relative Bias by Moneyness and Maturity

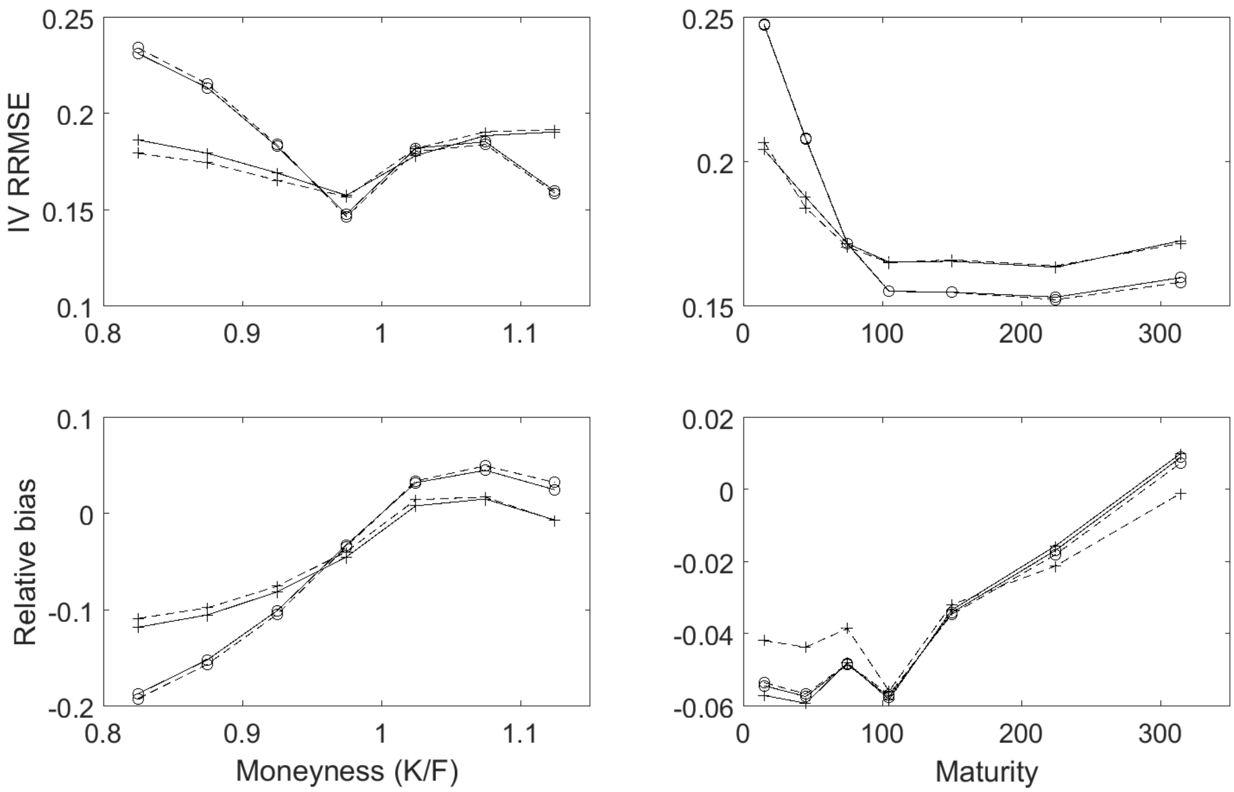

Graphical representation of the results in Table 3. Option IV RRMSE (top row) and relative bias (bottom row) by moneyness (left column) and maturity (right column). NGARCH(1,1) models are marked with "+" and AGARCH(1,1) with "○". CHJ models have dashed lines.

Appendix A. Risk Neutralization of the NCHJ model

Consider the following NGARCH(1,1) model

$$\begin{aligned}\ln(S_t) &= \ln(S_{t-1}) + r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}z_t \\ h_t &= \omega + \beta h_{t-1} + \alpha h_{t-1}(z_{t-1} - \gamma)^2 \end{aligned} \tag{28}$$

and this discrete-time pricing operator

$$\frac{M_t}{M_{t-1}} = \left(\frac{S_t}{S_{t-1}}\right)^{\phi} \exp\left(\delta + \eta h_t + \xi\left(\frac{h_{t+1} - h_t}{h_t}\right)\right). \tag{29}$$

From the NGARCH dynamic in (28) we can write

$$\begin{aligned}\frac{S_t}{S_{t-1}} &= \exp\left(r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}z_t\right) \\ h_{t+1} - h_t &= \omega + (\beta - 1)h_t + \alpha h_t(z_t - \gamma)^2. \end{aligned} \tag{30}$$

Substituting into (29) gives

$$\frac{M_t}{M_{t-1}} = \exp\left(\phi r + \phi\lambda\sqrt{h_t} - \phi\frac{1}{2}h_t + \phi\sqrt{h_t}z_t + \delta + \eta h_t + \frac{\xi\omega}{h_t} + \xi(\beta - 1) + \xi\alpha(z_t - \gamma)^2\right). \tag{31}$$

Expanding the square and collecting terms gives

$$\begin{aligned}\frac{M_t}{M_{t-1}} = \exp\Bigg(&\phi r + \delta + \xi(\beta - 1) + \xi\alpha\gamma^2 + \left[\eta + \frac{\xi\omega}{h_t^2} + \phi\left(\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2}\right)\right] h_t \\ &+\phi\sqrt{h_t}z_t + \xi\alpha z_t^2 - 2\xi\alpha\gamma z_t\Bigg). \end{aligned} \tag{32}$$

First, we use the fact that for any initial value h_t, the parameters must be consistent with the Euler equation for the riskless asset.

$$\begin{aligned}E_{t-1}\left[\frac{M_t}{M_{t-1}}\right] &= \exp(-r) \\ &= \exp\left(\phi r + \delta + \xi(\beta - 1) + \xi\alpha\gamma^2 + \left[\eta + \frac{\xi\omega}{h_t^2} + \phi\left(\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2}\right)\right] h_t\right) \\ &\quad \times E\left[\exp\left((\phi\sqrt{h_t} - 2\xi\alpha\gamma)z_t + \xi\alpha z_t^2\right)\right]. \end{aligned} \tag{33}$$

We need the following result

$$E\left[\exp(az^2 + 2abz)\right] = \exp\left(-\frac{1}{2}\ln(1-2a) + \frac{2a^2b^2}{1-2a}\right). \tag{34}$$

For our application we have

$$\begin{aligned} a &= \xi\alpha \\ b &= \frac{\phi\sqrt{h_t} - 2\xi\alpha\gamma}{2\xi\alpha} \\ 2a^2b^2 &= 2\xi^2\alpha^2\left(\frac{\phi\sqrt{h_t} - 2\xi\alpha\gamma}{2\xi\alpha}\right)^2 \\ &= \frac{1}{2}(\phi\sqrt{h_t} - 2\xi\alpha\gamma)^2. \end{aligned}$$

Therefore,

$$E\left[\exp\left((\phi\sqrt{h_t} - 2\xi\alpha\gamma)z_t + \xi\alpha z_t^2\right)\right] = \exp\left(-\frac{1}{2}\ln(1-2\xi\alpha) + \frac{(\phi\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2(1-2\xi\alpha)}\right) \tag{35}$$

and

$$\begin{aligned} E_{t-1}\left[\frac{M_t}{M_{t-1}}\right] = \exp\Bigg(&\phi r + \delta + \xi(\beta-1) + \xi\alpha\gamma^2 + \left[\eta + \frac{\xi\omega}{h_t^2} + \phi\left(\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2}\right)\right]h_t \\ &- \frac{1}{2}\ln(1-2\xi\alpha) + \frac{(\phi\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2(1-2\xi\alpha)}\Bigg). \end{aligned} \tag{36}$$

Rearranging and using (33) we get

$$\begin{aligned} 0 =& (\phi+1)r + \delta + \xi(\beta-1) + \xi\alpha\gamma^2 - \frac{1}{2}\ln(1-2\xi\alpha) \\ &+ \left[\eta + \frac{\xi\omega}{h_t^2} + \phi\left(\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2}\right) + \frac{(\phi\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2h_t(1-2\xi\alpha)}\right]h_t. \end{aligned} \tag{37}$$

Therefore we must have

$$\begin{aligned} \delta &= -(\phi+1)r - \xi(\beta-1) - \xi\alpha\gamma^2 + \frac{1}{2}\ln(1-2\xi\alpha) \\ \eta_t &= -\frac{\xi\omega}{h_t^2} - \phi\left(\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2}\right) - \frac{(\phi\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2h_t(1-2\xi\alpha)}. \end{aligned} \tag{38}$$

Now we use the Euler equation for the underlying index

$$E_{t-1}\left[\frac{S_t}{S_{t-1}}\frac{M_t}{M_{t-1}}\right] = 1. \tag{39}$$

First, note that $\frac{S_t}{S_{t-1}}\frac{M_t}{M_{t-1}}$ is equal to $\frac{M_t}{M_{t-1}}$ in (29) with ϕ replaced by $\phi + 1$, thus we can use the expression for $E_{t-1}\left[\frac{M_t}{M_{t-1}}\right]$ to write

$$\begin{aligned} E_{t-1}\left[\frac{S_t}{S_{t-1}}\frac{M_t}{M_{t-1}}\right] = \exp\Bigg(& (\phi+1)r + \delta + \xi(\beta - 1) + \xi\alpha\gamma^2 - \frac{1}{2}\ln(1-2\xi\alpha) \\ & + \left[\eta_t + \frac{\xi\omega}{h_t^2} + (\phi+1)\left(\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2}\right) + \frac{(\phi+1)\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2h_t(1-2\xi\alpha)}\right] h_t\Bigg). \end{aligned} \tag{40}$$

Taking logs, setting equal to zero and using the above solutions for δ_t and ϕ gives

$$0 = \left[\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2} + \frac{h_t + 2\phi h_t - 4\xi\alpha\gamma\sqrt{h_t}}{2h_t(1-2\xi\alpha)})\right] h_t. \tag{41}$$

Solving for ϕ yields

$$\phi_t = -\left(\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2}\right)(1-2\xi\alpha) + \frac{2\xi\alpha\gamma}{\sqrt{h_t}} - \frac{1}{2}. \tag{42}$$

To find the risk-neutral dynamic, note that the risk-neutral density is proportional to the physical density times the pricing kernel

$$f_{t-1}^{*}(S_t) = \frac{f_{t-1}(S_t)\, M_t}{E_{t-1}[M_t]}. \tag{43}$$

First, since $E_{t-1}\left[\frac{M(t)}{M(t-1)}\right] = e^{-r}$ and after inserting δ and η_t, we get:

$$\begin{aligned}
\frac{M(t)}{E_{t-1}[M(t)]} &= \frac{M(t)}{M(t-1)}\left(E_{t-1}\left[\frac{M(t)}{M(t-1)}\right]\right)^{-1} \\
&= \exp\left((\phi+1)r + \delta + \xi(\beta-1) + \xi\alpha\gamma^2 + \left[\eta_t + \frac{\xi\omega}{h_t^2} + \phi\left(\frac{\lambda}{\sqrt{h_t}} - \frac{1}{2}\right)\right]h_t \right. \\
&\qquad \left. + \phi\sqrt{h_t}z_t + \xi\alpha z_t^2 - 2\xi\alpha\gamma z_t\right) \\
&= \exp\left(\frac{1}{2}\ln(1-2\xi\alpha) - \frac{(\phi_t\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2(1-2\xi\alpha)} + \phi_t\sqrt{h_t}z_t + \xi\alpha z_t^2 - 2\xi\alpha\gamma z_t\right) \\
&= \sqrt{1-2\xi\alpha}\exp\left(-\frac{(\phi_t\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2(1-2\xi\alpha)} + \phi_t\sqrt{h_t}z_t + \xi\alpha z_t^2 - 2\xi\alpha\gamma z_t\right). \qquad (44)
\end{aligned}$$

Inserting (44) in (43) , given that the physical innovation is standard normal, we get

$$\begin{aligned}
f_{t-1}^*(z_t) &= \sqrt{\frac{1}{2\pi}}\exp\left\{-\frac{1}{2}z_t^2\right\} \times \sqrt{1-2\xi\alpha}\exp\left\{-\frac{(\phi_t\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2(1-2\xi\alpha)} + \phi_t\sqrt{h_t}z_t + \xi\alpha z_t^2 - 2\xi\alpha\gamma z_t\right\} \\
&= \sqrt{\frac{1-2\xi\alpha}{2\pi}}\exp\left\{-\frac{(\phi_t\sqrt{h_t} - 2\xi\alpha\gamma)^2}{2(1-2\xi\alpha)} + \phi_t\sqrt{h_t}z_t + \xi\alpha z_t^2 - 2\xi\alpha\gamma z_t - \frac{1}{2}z_t^2\right\} \\
&= \sqrt{\frac{1-2\xi\alpha}{2\pi}}\exp\left\{-\frac{1}{2}\left(z_t^2 + \frac{(\phi_t\sqrt{h_t} - 2\xi\alpha\gamma)^2}{1-2\xi\alpha} - 2\phi_t\sqrt{h_t}z_t - 2\xi\alpha z_t^2 + 4\xi\alpha\gamma z_t\right)\right\} \\
&= \sqrt{\frac{1-2\xi\alpha}{2\pi}}\exp\left\{-\frac{1-2\xi\alpha}{2}\left(z_t^2 - \left(\frac{2\phi_t\sqrt{h_t} - 4\xi\alpha\gamma}{1-2\xi\alpha}\right)z_t + \left(\frac{\phi_t\sqrt{h_t} - 2\xi\alpha\gamma}{1-2\xi\alpha}\right)^2\right)\right\} \\
&= \sqrt{\frac{1-2\xi\alpha}{2\pi}}\exp\left\{-\frac{1-2\xi\alpha}{2}\left(z_t^2 - 2\left(\frac{\phi_t\sqrt{h_t} - 2\xi\alpha\gamma}{1-2\xi\alpha}\right)z_t + \left(\frac{\phi_t\sqrt{h_t} - 2\xi\alpha\gamma}{1-2\xi\alpha}\right)^2\right)\right\} \\
&= \sqrt{\frac{1-2\xi\alpha}{2\pi}}\exp\left\{-\frac{1-2\xi\alpha}{2}\left(z_t - \frac{\phi_t\sqrt{h_t} - 2\xi\alpha\gamma}{1-2\xi\alpha}\right)^2\right\}. \qquad (45)
\end{aligned}$$

It is therefore convenient to define a standardized risk-neutral innovation

$$\begin{aligned}
z_t^* &= \sqrt{1-2\xi\alpha}\left(z_t - \frac{\phi_t\sqrt{h_t} - 2\xi\alpha\gamma}{1-2\xi\alpha}\right) \\
&= \sqrt{1-2\xi\alpha}\left[z_t + \lambda + \left(\frac{\xi\alpha}{1-2\xi\alpha}\right)\sqrt{h_t}\right] \\
\Leftrightarrow z_t &= \frac{z_t^*}{\sqrt{1-2\xi\alpha}} - \lambda - \left(\frac{\xi\alpha}{1-2\xi\alpha}\right)\sqrt{h_t}. \qquad (46)
\end{aligned}$$

We substitute this in the physical GARCH process in (28):

$$\begin{aligned}\ln(S_t) &= \ln(S_{t-1}) + r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}\left(\frac{z_t^*}{\sqrt{1-2\xi\alpha}} - \lambda - \left(\frac{\xi\alpha}{1-2\xi\alpha}\right)\sqrt{h_t}\right)\\ &= \ln(S_{t-1}) + r - \frac{1}{2}\frac{h_t}{1-2\xi\alpha} + \sqrt{\frac{h_t}{1-2\xi\alpha}}z_t^*\end{aligned}$$

and

$$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1}\left(\left(\frac{z_t^*}{\sqrt{1-2\xi\alpha}} - \lambda - \left(\frac{\xi\alpha}{1-2\xi\alpha}\right)\sqrt{h_t}\right) - \gamma\right)^2. \tag{47}$$

Letting $h_t^* = \frac{h_t}{1-2\xi\alpha}$, we get

$$\begin{aligned}\ln(S_t) &= \ln(S_{t-1}) + r - \frac{1}{2}h_t^* + \sqrt{h_t^*}z_t^*\\ h_t^* &= \frac{1}{1-2\xi\alpha}\left[\omega + \beta h_{t-1} + \alpha h_{t-1}\left(\left(\frac{z_{t-1}^*}{\sqrt{1-2\xi\alpha}} - \lambda - \left(\frac{\xi\alpha}{1-2\xi\alpha}\right)\sqrt{h_{t-1}}\right) - \gamma\right)^2\right].\end{aligned} \tag{48}$$

Letting $\gamma_t^* = \sqrt{1-2\xi\alpha}\,(\lambda+\gamma) + \xi\alpha\sqrt{h_t^*}$, we get

$$\begin{aligned}h_t^* &= \frac{1}{1-2\xi\alpha}\left[\omega + \beta h_{t-1} + \alpha h_{t-1}\left(\frac{z_{t-1}^*}{\sqrt{1-2\xi\alpha}} - \frac{\gamma_{t-1}^*}{\sqrt{1-2\xi\alpha}}\right)^2\right]\\ &= \frac{1}{1-2\xi\alpha}\left[\omega + \beta h_{t-1} + \alpha h_{t-1}^*\left(z_{t-1}^* - \gamma_{t-1}^*\right)^2\right].\end{aligned} \tag{49}$$

With $\omega^* = \frac{\omega}{1-2\xi\alpha}$ and $\alpha^* = \frac{\alpha}{1-2\xi\alpha}$ we can simplify to

$$h_t^* = \omega^* + \beta h_{t-1}^* + \alpha^* h_{t-1}^*\left(z_{t-1}^* - \gamma_{t-1}^*\right)^2. \tag{50}$$

Summing up, we have:

$$\begin{aligned}\ln(S_t) &= \ln(S_{t-1}) + r - \frac{1}{2}h_t^* + \sqrt{h_t^*}z_t^*\\ h_t^* &= \omega^* + \beta h_{t-1}^* + \alpha^* h_{t-1}^*\left(z_{t-1}^* - \gamma_{t-1}^*\right)^2\end{aligned} \tag{51}$$

Where

$$
\begin{aligned}
\omega^* &= \frac{\omega}{1-2\xi\alpha} \\
\alpha^* &= \frac{\alpha}{1-2\xi\alpha} \\
\gamma_t^* &= \sqrt{1-2\xi\alpha}\,(\lambda+\gamma) + \xi\alpha\sqrt{h_t^*}.
\end{aligned}
$$

Note that

$$
h_0^* = \frac{h_0}{1-2\xi\alpha}.
$$

Appendix B. Heston-Nandi Option Valuation

From the risk-neutral dynamic in (18),

$$\begin{aligned} \ln(S_t) &= \ln(S_{t-1}) + r - \frac{1}{2}h_t^* + \sqrt{h_t^*}z_t^* \\ h_t^* &= \omega^* + \beta h_{t-1}^* + \alpha^*\left(z_{t-1}^* - \gamma^*\sqrt{h_{t-1}^*}\right)^2, \end{aligned} \tag{52}$$

we get the conditional moment generating function:

$$g_{t,T}^*(\Phi) \equiv \mathrm{E}_t^{\mathbb{Q}}\left[\exp\left(\Phi\ln(S_T)\right)\right] = \exp\left(\Phi\ln(S_t) + A_{t,T}(\Phi) + B_{t,T}(\Phi)h_{t+1}^*\right) \tag{53}$$

where

$$\begin{aligned} A_{t,T}(\Phi) &= A_{t+1,T}(\Phi) + \Phi r + B_{t+1,T}(\Phi)\omega^* - \frac{1}{2}\ln\left(1 - 2B_{t+1,T}(\Phi)\alpha^*\right) \\ B_{t,T}(\Phi) &= -\frac{1}{2}\Phi + B_{t+1,T}(\Phi)\left(\beta + \alpha^*\gamma^{*2}\right) + \frac{\Phi^2 + 4B_{t+1,T}(\Phi)\alpha^*\gamma^*(B_{t+1,T}(\Phi)\alpha^*\gamma^* - \Phi)}{2 - 4B_{t+1,T}(\Phi)\alpha^*}. \end{aligned}$$

It is computed recursively from the terminal condition that stems from the fact that S_T is known at maturity:

$$A_{T,T}(\Phi) = B_{T,T}(\Phi) = 0. \tag{54}$$

Call option prices can be computed using

$$C(t,T,K) = S_t P_{1,t} - K\exp(-r(T-t))P_{2,t} \tag{55}$$

where

$$\begin{aligned} P_{1,t} &= \left(\frac{1}{2} + \frac{\exp(-r(T-t))}{\pi}\int_0^\infty \mathrm{Re}\left[\frac{K^{-i\Phi}g_{t,T}^*(i\Phi+1)}{i\Phi S_t}\right]d\Phi\right) \\ P_{2,t} &= \left(\frac{1}{2} + \frac{1}{\pi}\int_0^\infty \mathrm{Re}\left[\frac{K^{-i\Phi}g_{t,T}^*(i\Phi)}{i\Phi}\right]d\Phi\right). \end{aligned}$$

Put options can be priced using put-call parity.

Conclusion

Je compare les vraisemblances et les erreurs de tarification d'options de 4 modèles GARCH(1,1). Utilisant un noyau de prix exponentiellement linéaire, le modèle NGARCH(1,1) performe mieux que le AGARCH(1,1), tel qu'il est largement documenté dans la littérature. Pour les deux familles de modèles, utiliser un noyau de prix exponentiellement quadratique produit de plus petites erreurs de tarification, ainsi qu'une vraisemblance physique légèrement plus élevée. Par contre, cette amélioration est supérieure dans les modèles non-affines, élargissant ainsi l'écart de performance. Je conclus que l'utilisation du noyau de prix quadratique ne réduit pas les coûts associés aux contraintes affines.