City Research Online

City, University of London Institutional Repository

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: http://openaccess.city.ac.uk/15050/

Link to published version: http://dx.doi.org/10.1080/14697688.2016.1211798

City Research Online: http://openaccess.city.ac.uk/   publications@city.ac.uk

Jumps and stochastic volatility in crude oil prices and advances in average option pricing

Ioannis Kyriakou[*], Panos K. Pouliasis and Nikos C. Papapostolou

Cass Business School, City University London, 106 Bunhill Row, London EC1Y 8TZ, UK

(*July 21, 2016*)

Crude oil derivatives form an important part of the global derivatives market. In this paper, we focus on Asian options which are favoured by risk managers being effective and cost-saving hedging instruments. The paper has both empirical and theoretical contributions: we conduct an empirical analysis of the crude oil price dynamics and develop an accurate pricing setup for arithmetic Asian options with discrete and continuous monitoring featuring stochastic volatility and discontinuous underlying asset price movements. Our theoretical contribution is applicable to various commodities exhibiting similar stylized properties. We here estimate the stochastic volatility model with price jumps as well as the nested model with omitted jumps to NYMEX WTI futures vanilla options. We find that price jumps and stochastic volatility are necessary to fit options. Despite the averaging effect, we show that Asian options remain sensitive to jump risk and that ignoring the discontinuities can lead to substantial mispricings.

Keywords: Oil prices; Stochastic volatility; Jump diffusion; Arithmetic Asian options

JEL Classification: G12; G13; C63; C13

1. Introduction

The price of crude oil and oil products has taken centre stage in recent years among the leading indicators of the economy and is now always quoted when forecasting economic trends. This has occurred together with the growing acceptance of commodities as a mainstream financial and investment class, with the resulting growth in the volume and variety of financial products linked to them (see Campi and Galdenzi 2015). In general, the volume of exchanged-traded commodity derivatives has grown steadily since their introduction in the late 1980s. Focusing on the most recent years, as shown in Figure 1, the volume of crude oil futures and options has reached, respectively, up to 50% and 60% of the total energy contracts traded on NYMEX in 2015. This raises the importance of accurate pricing of commodity derivatives.

Risk managers use commodity options to hedge price risk. In particular, average (Asian) options, whose payoff is contingent on the average price of the underlying asset during a pre-specified time window, are popular for the risk management of commodity futures due to lower volatility of the average price compared to the price of the underlying asset (see Figure 2), hence are cheaper than plain vanilla options. Arithmetic averages prevent momentarily wild fluctuations from affecting transactions of large exchanged quantities, for example, of oil due to the time lapsed between the day a tanker leaves the production site until it reaches its destination (Geman 2005). For this, commodity Asian options are widely traded over-the-counter (OTC) and, in the recent years, options on futures, particularly on crude oil, have been introduced in exchanges worldwide and

---

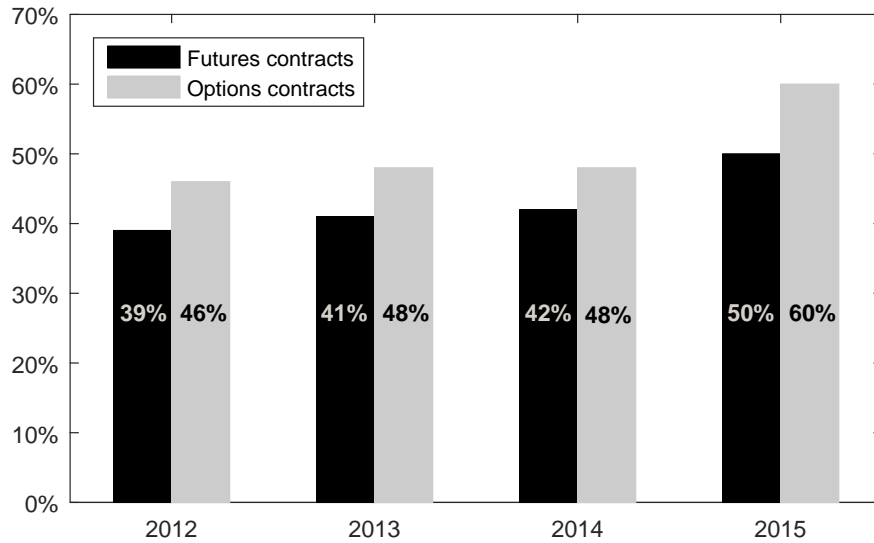[*]Corresponding author. Email: ioannis.kyriakou@city.ac.uk

Figure 1. Volumes of crude oil futures and options as % of total energy contracts traded on NYMEX. Source: CME Group (as of April 2015).
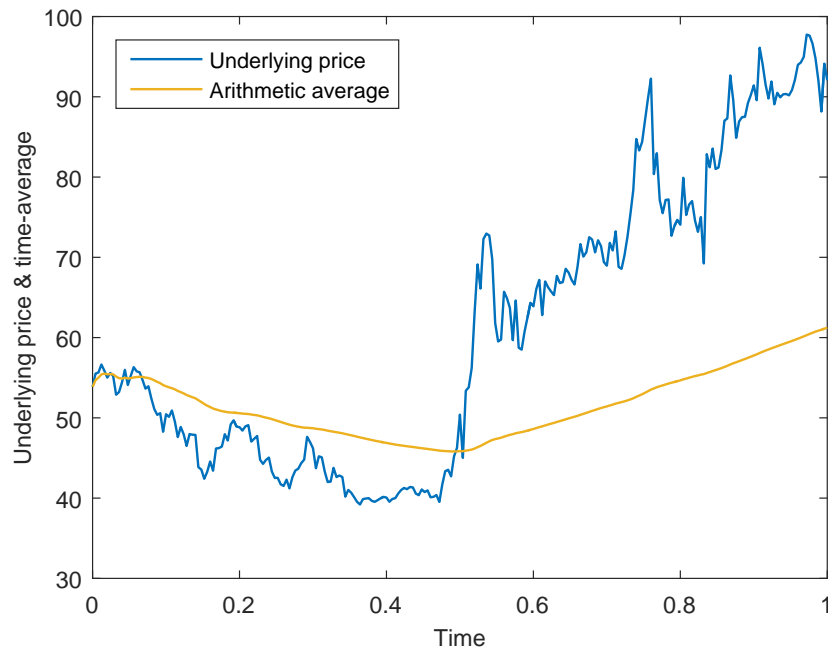

Figure 2. One-year-ahead simulated paths of crude oil price dynamics (daily) and of time-averaged prices (latter much less dispersed).

seen significant trading volume increases (see Alexander and Venkatramanan 2008). Asian options may have a fixed or a floating strike price. Also, averaging commencement may coincide with the inception of the contract or a forward date (forward start option). Finally, contracts which involve trades with different volumes over a period of time might use volume-weighted averages.

In light of the above, in the first part of the paper we perform a comprehensive empirical analysis of the crude oil prices. We consider daily WTI crude oil prices from April 11, 1983 to February 3, 2015. Our study concludes that there are systematic departures from the benchmark lognormal process which describe the crude oil price dynamics. More specifically, we find evidence of serial

dependence in the squared log-returns and that returns and their volatilities are negatively correlated in consistency with evidence by Trolle and Schwartz (2009), Larsson and Nossman (2011) and Geman and Shih (2009) in the post-2000 era. In addition, we find that mean-reversion is strongly rejected over the whole sample period and explain that seasonal patterns are absent from the term structure of futures prices. In light of these empirical facts about the crude oil prices, we adopt the stochastic volatility model of Heston (1993) and the Bates (1996) model overlaid with price jumps of random magnitude and arrival. We calibrate the two models to the prices of liquid WTI futures American vanilla options trading on the NYMEX division of the CME Group. We find that the fitted models are capable of mimicking volatility clustering, exhibit Samuelson's maturity effect, that is, a decaying term structure of volatility which is commonly observed in commodities markets, and negative correlation between volatility movements and log-returns, which is often interpreted in terms of the so-called leverage effect, i.e., the empirical observation that large downward moves in the price are associated with upward moves in volatility. When calibrating to batches of vanilla option prices with different strikes for given maturity, we find that the Bates model yields lower pricing error across different maturities implying the importance of modelling price discontinuities in order to flexibly describe abrupt changes in the WTI prices.

In the second part of the paper, we provide the development of our pricing methodology for arithmetic Asian options under stochastic volatility, in particular, the Heston and Bates models. Although a large volume of publications has been devoted to the pricing of Asian options, most of them rely on one-dimensional processes which may not always be able to reproduce certain stylized properties of asset prices, as in the case of commodities. Earlier contributions allowing for stochastic volatility are fewer and either represent price approximations for the arithmetic average option or are related to the less common geometric average option. We mention, for example, Fouque and Han (2003) who employ a PDE approach on a reduced state space in a purely diffusive model framework based on fast mean-reverting stochastic volatility asymptotic analysis. Based on the same principles, Wong and Cheung (2004) and Fouque and Han (2004) provide price approximations for continuously monitored geometric average options for further use as control variates in generating price estimates for the more prevalent arithmetic options. Tahani (2013) and Kim and Wee (2014) derive exact expressions for continuous geometric options under the square root variance model with/out mean-reverting log-asset price dynamics, whereas Hubalek *et al.* (2014) develop a pricing framework in a general affine stochastic volatility (ASV) model setup. Shiraya and Takahashi (2011) propose an approximation formula for pricing average options under the Heston and extended SABR stochastic volatility models. Using the principles set out in Curran (1994) and Rogers and Shi (1995), Dingeç *et al.* (2015) and, independently, Fusai and Kyriakou (2016) derive lower bounds under ASV models. Finally, Ewald *et al.* (2013) propose a solution by means of a PDE and a Monte Carlo simulation method, whereas Yamazaki (2014) a pricing formula based on the Gram–Charlier expansion.

Although we focus here on the Heston and Bates models, the method we propose is universally applicable to ASV models with known joint laws of the underlying (log) asset price and the variance via the associated characteristic functions. Exploiting this knowledge together with favourable option price convolution structure in the log-return dimension, we evaluate discrete arithmetic Asian options by means of Fourier transform, combined with numerical integration in the variance dimension. Our method is distinguished from previous contributions due to a number of appealing features, filling an important long-standing gap in the literature. First, it converges perfectly smoothly in the number of integration points, hence is suitable for further improvement through simple Richardson extrapolation leading to very high levels of precision. The convergence is also smooth in the number of monitoring dates, hence by the same means we can also obtain the price of a continuously monitored Asian option. Second, it can be applied to general ASV processes, such as time changed Lévy processes by integrated square root or Lévy-driven Ornstein–Uhlenbeck variance processes. Third, by direct differentiation of our pricing integrals, one can obtain exact representations of the option price sensitivities with respect to parameters of interest. Fourth, our method allows moving across monitoring dates without any intermediate time-discretization re-

quired by standard stochastic volatility Monte Carlo simulation schemes with significant impact on their order of convergence due to inborn bias, and, finally, it is free of any form of approximation error which can be hard to quantify.

To make concrete our study of the arithmetic Asian options, in the last part of the paper we apply to the case of crude oil options. As NYMEX WTI crude oil average options are primarily traded OTC, we evaluate these using the parameter estimates from our original calibration to WTI futures American vanilla options. We back up our claim for the importance of our work by showing that, although Asian options involve an averaging procedure which reduces the fluctuation of the underlying asset price process, they remain sensitive to jump risk and ignoring the discontinuities can lead to nontrivial mispricings.

We conclude this Introduction by noting that the use of our average option pricing framework is not limited to energy commodities but can be flexibly applied to other asset classes, from equities to foreign exchange to agriculturals and metals, exhibiting similar stylized features. We mention few examples to stress how the interest in average options is broader than just a purely challenging pricing task:

- In agricultural markets, where the underlying spot price is not easily identifiable, a "large player" may push the market up or down on a given day; this exercise of market power becomes more problematic in the case of average options (Geman 2005).
- Charterers operating in the freight market typically face freight rate exposure during a voyage. As the freight revenue process for a ship in the physical spot market is given by discretely sampled prices in this period, most freight derivatives are settled against an average of these spot freight rates published by the Baltic Exchange in response to concerns about potential market manipulation by large participants. Currently, liquidity in the freight options market is focused on the indices in the Capesize and Panamax drybulk sectors (see Nomikos *et al.* 2013).
- The London Metal Exchange offers Traded Average Price Options (TAPOs) giving the metal community a flexible way of hedging against fluctuations in the Monthly Average Settlement Price (MASP) for several metals. This is particularly useful as a large proportion of physical contracts are negotiated on the basis of the MASP (www.lme.com/trading/contract-types/tapos/).
- In foreign exchange markets, using weighted average options to hedge a stream of (received) payments (e.g., a USD average call can be bought to hedge the ongoing EUR revenues of a US-based company) can be much cheaper than using a stream of vanilla options (Wystup 2006).

The rest of this paper is organized as follows. In Section 2 we discuss the stylized facts of the crude oil market. Section 3 introduces the proposed stochastic volatility models for the crude oil price and presents their distributional properties. Section 4 describes the data and the model calibration. Section 5.1 presents the theoretical results of the paper and establishes our recursive-integral approach to valuing Asian options under stochastic volatility and price jumps, Section 5.2 the details of the numerical implementation, whereas Section 5.3 our application in pricing commodity options. Section 6 concludes.

2. Analysis of crude oil prices

In what follows, we examine the stylized properties of the spot series. The spot time series is obtained from Thomson Reuters Datastream for the period April 11, 1983 to February 3, 2015, i.e., 7,982 daily observations. First, we consider the whole sample of spot prices; in addition, we partition the entire set into two subsamples (Geman and Shih 2009): "pre-2000" period (April 11, 1983 to December 30, 1999, i.e., 4,200 observations) and "post-2000" period (January 3, 2000 to February 3, 2015, i.e., 3,782 observations).

Table 1. Descriptive statistics and tests of crude oil spot price returns.

	Whole sample	Subsample estimates		Rolling window estimates	
		Pre-2000	Post-2000	90% CI: 1,260 obs	
Panel A. Descriptive statistics of log-returns					
Mean (%)	1.605	-0.819	4.282	[-9.403, 21.05]	
Vol (%)	39.82	40.896	38.618	[28.21, 51.54]	
Skew	-0.971	-1.65	-0.076	[-2.874, 0.378]	
Exc. kurt.	23.35	36.25	5.197	[0.978, 49.10]	
JB stat.	$183,000^a$	$231,787^a$	$4,282^a$	[61.73, 128,279]	{0.00}
KS stat.	0.089^a	0.095^a	0.081^a	[0.044, 0.134]	{0.00}
$Q^2(1)$	356.6^a	201.9^a	85.67^a	[3.016, 96.97]	{4.20}
$Q^2(5)$	651.9^a	322.6^a	660.2^a	[13.83, 495.0]	{1.10}
$Q^2(10)$	$1,114^a$	557.1^a	$1,197^a$	[24.12, 787.0]	{1.40}
Panel B. GARCH model estimates					
a_0	0.015^a	0.013^a	0.027^a	[0.019, 1.207]	{0.70}
a_1	0.073^a	0.082^a	0.053^a	[0.037, 0.141]	{0.00}
a_2	0.926^a	0.917^a	0.943^a	[0.695, 0.944]	{0.00}
Residual diagnostics					
Skew	-0.381	-0.471	-0.288	[-0.834, 0.076]	
Exc. kurt.	3.397	4.866	1.964	[0.465, 8.737]	
JB stat.	$4,041^a$	$4,297^a$	663.5^a	[12.57, 4,155]	{1.51}
KS stat.	0.040^a	0.052^a	0.026^b	[0.023, 0.067]	{26.9}
$Q^2(1)$	2.167	0.404	4.292^c	[0.003, 5.864]	{86.9}
$Q^2(5)$	6.657	3.565	6.685	[0.785, 9.977]	{92.7}
$Q^2(10)$	21.98^b	13.3	17.25^c	[2.950, 18.17]	{90.8}
Panel C. Unit root tests on log-price levels					
ADF	-1.560	-3.029^b	-1.850	[-2.985, -0.774]	{15.9}
PP	-1.537	-3.093^b	-1.798	[-3.063, -0.581]	{10.7}

The table presents the results of our analysis over the whole sample period April 11, 1983 to February 3, 2015, i.e., 7,982 daily observations (obs.), and the two sub-periods: "pre-2000" (April 11, 1983 to December 30, 1999, i.e., 4,200 obs.) and "post-2000" (January 3, 2000 to February 3, 2015, i.e., 3,782 obs.). Panel A presents the descriptive statistics of the crude oil log-returns (vol: volatility, skew: skewness coefficient, exc. kurt.: excess kurtosis). Jarque–Bera (JB) tests the null hypothesis of a normal distribution for the sample log-return series. Kolmogorov–Smirnov (KS) tests the null hypothesis of equality of the historical log-return distribution to a normal distribution. $Q^2(n)$ is the Ljung–Box Q statistic for the n-th order sample autocorrelation of the squared log-return series (ARCH test of Engle 1982). Panel B presents the estimates of the fitted GARCH(1,1) model on the demeaned log-return series ε_t; the model is of the form $\sigma_t^2 = a_0 + a_1\varepsilon_{t-1}^2 + a_2\sigma_{t-1}^2$, where σ_t^2 is the variance process. Residual diagnostics are also reported. In Panel C, we test the null hypothesis that the crude oil log-price is non-stationary by employing two unit-root tests: augmented Dickey–Fuller (ADF) and Phillips–Perron (PP). Superscripts a, b, c indicate significance at the 1%, 5% and 10% levels, respectively. Rolling window estimates are calculated using rolling five-year 6,722 subsamples of 1,260 obs. each. Figures in [·] represent 90% confidence intervals (CI) of the estimates throughout the rolling period. Where relevant, figures in {·} represent the % numbers of times that the null hypothesis of the different tests cannot be rejected at the 10% significance level.

Panel A of Table 1 exhibits the descriptive statistics of the daily log-returns on the spot WTI crude oil at Cushing, Oklahoma. The sample mean of the log-returns is relatively small, whereas the volatility of the log-returns exceeds 38% per annum. Existing skewness and excess kurtosis suggest that the log-return distribution is not normal. Likewise, the Jarque–Bera (JB) statistic (Jarque and Bera 1980) rejects the null hypothesis of a normal distribution for the sample log-return series. In addition, we employ the Kolmogorov–Smirnov test for equality of the historical log-return distribution to a normal distribution; consistently with the JB statistic, this hypothesis is rejected at the 1% level. We assess the null hypothesis of a white-noise process for the sample squared log-returns by employing the Engle ARCH test (Engle 1982) based on the Ljung–Box Q statistic (Ljung and Box 1978); the calculated values for the 1st, 5th and 10th order sample autocorrelation suggest that there is significant evidence of serial dependence in the squared log-returns (heteroscedasticity).
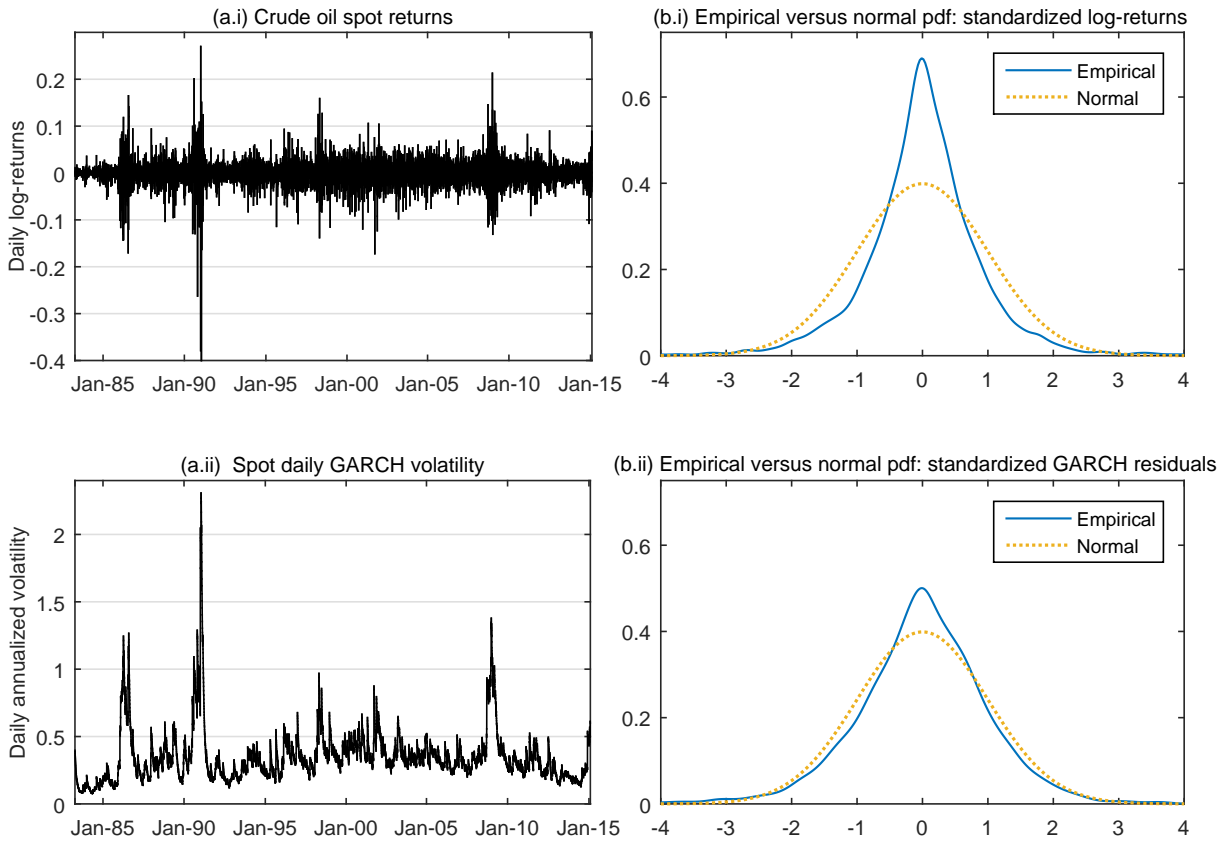
Figure 3. Crude oil log-returns (a.i) and annualized spot volatility (a.ii) from GARCH model from April 11, 1983 to February 3, 2015 using daily crude oil prices, coupled with probability densities (Matlab's kernel smoothing function estimates) of standardized log-returns (b.i) and standardized GARCH residuals (b.ii) versus standard normal distribution.

To quantify the time-dependence in the volatility series, Panel B of Table 1 reports the maximum likelihood estimates of a simple GARCH(1,1) model; all coefficients are found significant at the 1% level. Figure 3 plots the daily log-return series and the daily annualized volatility series dynamics obtained from the GARCH model. Diagnostic tests on the standardized residuals indicate that by filtering the log-returns using a simple GARCH model, the ARCH effects wear away (rejected at the 1% significance level) and their distribution is closer to the normal; nevertheless, non-normality and, to some extent, heteroscedasticity are still present. The distance of the historical distribution from the normal is also illustrated in Figure 3 for the log-returns and the GARCH residuals.

As a robustness check, Table 1 reports the 90% confidence intervals of the abovementioned statistics by means of a rolling window estimation. More specifically, the statistics are calculated using rolling five-year subsamples[1] (1,260 observations each). For example, the first JB statistic is obtained using observations from the beginning of the sample period up to the 1,260th observation, the second test statistic using data from the 2nd up to the 1,261st, etc., up to the last observation. Where relevant, an associated rate is also shown, i.e., the % number of times that the p-value is greater than 0.1. Overall, our results from the rolling window estimation are found consistent with our original results.

---

[1]Ten-year rolling window estimates were also calculated and found qualitatively similar (available by the authors upon request).

Furthermore, our sample log-price series are tested for mean-reversion by employing the augmented Dickey and Fuller (1979, 1981) and Phillips and Perron (1988) unit-root tests (see Panel C of Table 1); both tests indicate that log-spot prices follow a unit-root process, i.e., mean-reversion is strongly rejected over the whole sample period and the 2000–2015 period, but not the 1983–1999 period. Our conclusion is consistent with recent evidence against mean-reversion (e.g., see Geman 2007, Geman and Shih 2009, Larsson and Nossman 2011). To discount the possibility that stationarity is sample-specific, we use rolling unit-root tests that explicitly take into account the possibility that a series may be more integrated during some periods or less so or not at all during other periods. The augmented Dickey–Fuller and Phillips–Perron tests reject the null of a unit root in only 15.9% and 10.7% of the total subsamples, respectively. Based on these results, in the ensuing analysis we do not consider mean-reversion in the log-spot prices.

Discontinuous price movements can flexibly accommodate implicit skewness and kurtosis in option prices at short time horizons (e.g., see Bakshi *et al.* 1997). Jumps are a salient feature of energy prices that is well-documented in the literature (e.g., see Li and Linetsky 2014 and various references therein) and can be attributed to temporary supply and demand imbalances, changes in market expectations, or even unanticipated macroeconomic developments (Hilliard and Reis 1998, Clewlow and Strickland 2000). Due to inelastic supply, even a relatively small change in demand can result in extreme sudden price movements.

Finally, we note that seasonal patterns are not common in the crude oil prices (although most petroleum products experience seasonality effects). For commodities prone to such fluctuations, the shape of the forward curve is determined by the expectations regarding the seasonal demand and supply dynamics; the crude oil market is a world market where seasonality is not observed in the term structure of futures prices (Borovkova and Geman 2006). For example, heating oil (gasoline) exhibits a relative upward pressure during the winter (summer) months and the storage capacity may not always be able to absorb the seasonal demand shocks, especially during the period of peak demand; hence, higher prices are anticipated in winter (summer), and lower prices during the inventory-accumulation period in summer (winter). As the demand for crude oil stems from the demand for its refined products, their distinct seasonality tends to balance out.

3. The models

Given our discussion in the previous section, in the ensuing analysis we consider the Bates (1996) model with price jumps and stochastic volatility as well as the nested Heston (1993) model without jumps, as candidates for describing the crude oil price dynamics.

Let $(\Omega, \mathcal{F}, \mathbb{F}, \hat{P})$ be a complete stochastic basis, i.e., the filtration $\mathbb{F} := \{\mathcal{F}_t\}_{t \in \mathbb{R}^+}$ satisfies the usual conditions. We interpret $\hat{P}$ as a risk neutral measure. In the Heston model, the stochastic variance V follows a Cox *et al.* (1985) (CIR) / Feller (1951) square root diffusion solving the SDE

$$dV_t = \alpha(\beta - V_t)dt + \gamma\sqrt{V_t}dW_t, \ V_0 = v_0, \tag{1}$$

where W is a standard Brownian motion and the parameters α, β, γ are assumed to satisfy the Feller condition, $d := 2\alpha\beta/\gamma^2 - 1 \geq 0$, to ensure that zero is an inaccessible boundary. The advantage of the CIR process in our contribution is in its analytical tractability; its transition probability density and the characteristic function of its time-integral conditional on the state of the process at time t are known in closed form. More specifically, the CIR transition density (see Cox *et al.* 1985) is

$$c(t, v|v_0) = \frac{2\alpha}{\gamma^2 \left(1 - e^{-\alpha t}\right)} e^{-\frac{2\alpha\left(v_0 e^{-\alpha t} + v\right)}{\gamma^2\left(1 - e^{-\alpha t}\right)}} \left(\frac{v}{v_0 e^{-\alpha t}}\right)^{\frac{d}{2}} I_d\left(\frac{4\alpha\sqrt{v_0 v e^{-\alpha t}}}{\gamma^2 \left(1 - e^{-\alpha t}\right)}\right), \tag{2}$$

where $I_d(\cdot)$ is the modified Bessel function of the first kind of order d. From Broadie and Kaya

(2006), $\zeta(t, u | v_0, v) := \hat{E}(\exp(iu \int_0^t V_s ds) | V_0 = v_0, V_t = v)$, that is, the characteristic function of the integrated process conditional on the state of the process at time t is

$$
\zeta(t, u | v_0, v) = \frac{\psi(u) \sqrt{e^{-(\psi(u) - \alpha)t}} \left(1 - e^{-\alpha t}\right)}{\alpha \left(1 - e^{-\psi(u)t}\right)}
$$

$$
\times \exp\left\{ \frac{v_0 + v}{\gamma^2} \left(\frac{\alpha \left(1 + e^{-\alpha t}\right)}{1 - e^{-\alpha t}} - \frac{\psi(u) \left(1 + e^{-\psi(u)t}\right)}{1 - e^{-\psi(u)t}} \right) \right\}
$$

$$
\times \frac{I_d \left(\frac{4\psi(u)\sqrt{v_0 v} e^{-\psi(u)t}}{\gamma^2 (1 - e^{-\psi(u)t})} \right)}{I_d \left(\frac{4\alpha\sqrt{v_0 v} e^{-\alpha t}}{\gamma^2 (1 - e^{-\alpha t})} \right)}, \tag{3}
$$

where $\psi(u) := \sqrt{\alpha^2 - 2iu\gamma^2}$. Result (3) will be revisited in Section 5.1 where its importance for the purposes of our pricing application will become obvious.

In the Heston model the log-asset price SDE is

$$
dX_t = \left(\mu - \frac{1}{2} V_t \right) dt + \sqrt{V_t} \left(\rho dW_t + \sqrt{1 - \rho^2} dB_t \right), \quad X_0 = \chi_0, \tag{4}
$$

where $\mu \in \mathbb{R}$ is constant, the standard Brownian motion B is independent of W, and ρ is the instantaneous correlation between the processes X and V. The Bates model is an extension of the Heston model to include jumps in the (log) asset price dynamics

$$
dX_t = \left(\mu - \lambda\kappa(-i) - \frac{1}{2} V_t \right) dt + \sqrt{V_t} \left(\rho dW_t + \sqrt{1 - \rho^2} dB_t \right) + dL_t, \tag{5}
$$

where L is an independent time-homogeneous compound Poisson process with intensity λ and normal jump sizes J with mean μ_J and standard deviation σ_J, and $\kappa(u) := \exp\left(iu\mu_J - \sigma_J^2 u^2 / 2 \right) - 1$.

4.   Data specification and calibration to option prices

We estimate models (1), (4) and (1), (5) on a panel data set of market quotes of NYMEX WTI futures American options trading on the NYMEX division of the CME Group. Both futures and options are quoted in US dollars and cents per barrel (bbl). They are traded for all deliveries (consecutive months) within the current year and the next five years. The June and December contracts are listed beyond the sixth year whereas additional months are added on an annual basis after the December contract expires, so that an additional June and a December contract are added nine years forward and the consecutive months in the sixth calendar year are filled in. Each futures contract is traded until the close of business on the third business day prior to the 25th calendar day of the month preceding the delivery month. If the 25th calendar day of the month is a non-business day, the trading ceases on the third business day prior to the business day preceding the 25th calendar day of the month preceding the delivery month. For each option contract, the trading ceases on the third business day prior to the business day preceding the termination of the trading in the underlying futures contract. One hundred different strike prices are traded in increments of $0.50 per bbl above and below the at-the-money (ATM) strike price, additional twenty strike prices in increments of $2.50 per bbl above and below the highest and lowest 50th increment, and ten strike prices in increments of $5.00 per bbl above and below the highest and lowest 250th increment.

The raw data set consists of settlement mid-point call and put option prices, open interest and daily volume for all available American vanilla options on WTI futures on February 3, 2015. The

raw options data set is obtained from the CME. We collect data for the first six nearby maturities, ranging from 42 days to maturity (April 2015) to 195 days to maturity (September 2015) with an average open interest and trading volume of 229,185 and 20,458, respectively. Beyond six months to expiry liquidity is concentrated on the contracts expiring in March, June, September and December, whereas beyond one year liquidity is concentrated on the contracts with December expiry. Thus, among the remaining contracts, we choose the first two contracts with expiration in December 2015 (287 days to maturity) and March 2016 (379 days to maturity) with an open interest of 16,556 and 67,822 and a trading volume of 402 and 300 contracts, respectively. Options on the remaining more distant maturity futures contracts are not considered, as the exposure to errors arising from the conversion of American to European option prices[1] and the assumption of constant interest rates becomes more serious for longer-term options.

When calibrating the models, we apply some standard filter rules on the option data. First, for each option maturity, we consider moneyness[1] ranging from 0.65 to 1.35, as very deep out-of-the-money (OTM) and in-the-money (ITM) options are plagued by illiquidity. Second, to minimize the effect of any errors due to the early-exercise approximation, we use only OTM and ATM put and call options (Trolle and Schwartz 2009); in addition, OTM options tend to be more liquid than ITM options as the latter are more expensive. Third, we consider only options that have an open interest in excess of 100 contracts. Finally, options with prices of less than 0.10 dollars are removed. After sorting the data, we are left with 322 options on the given trading day.

We resort to standard practice for extracting risk neutral parameter estimates of the underlying price models from observed option prices. More specifically, let $C_{i,T}^M$ and $C_{i,T}^\theta$ be, respectively, the market option price and the theoretical option price under the underlying price model with parameter vector θ for an option with i-th strike price maturing at T. Theoretical option prices are computed with high accuracy using the Fourier-cosine series expansion of Fang and Oosterlee (2008). For the Bates model (1), (5), $\theta := \{V_0, \alpha, \beta, \gamma, \rho, \lambda, \mu_J, \sigma_J\}$; for the Heston model (1), (4) the parameter vector reduces to $\theta := \{V_0, \alpha, \beta, \gamma, \rho\}$. Note that $\mu = 0$ due to the underlying futures contracts. We estimate θ by minimizing the Euclidean distance between the observed and theoretical option prices for each maturity T

$$\theta_T^* := \arg\min_\theta \sum_{i=1}^m \left| C_{i,T}^M - C_{i,T}^\theta \right|^2, \tag{6}$$

where m is the number of strike prices. We obtain a proxy for the risk free interest rate used in our computations by fitting a Nelson and Siegel (1987) curve to the US treasury yield curve consisting of the 1-month, 3-month, 6-month, 12-month and 24-month rates.

4.1. *Results of the calibration*

Calibration results are reported in Table 2. Several remarks are in order. The long-run mean volatility $\sqrt{\beta}$ of the crude oil futures prices under the Heston model (Panel A) decreases from 90% to 43% with increasing maturity; the observed decaying term structure of volatility is consistent with the Samuelson (1965) effect which is common in commodities markets (short-end forward commodities curves are more sensitive to information flow). A similar pattern is observed for the volatility of variance γ. Our estimates for parameter α imply that the expected time for the model variances to return half-way toward level β increases from 4 to 26 weeks (half-life is given by $\ln 2^{1/\alpha}$). The correlation coefficient ρ is found negative, contrary to the inverse leverage effect (high prices

---

[1]For computational reasons, estimation is feasible only with European options necessitating a conversion of American to European option prices. To this end, we employ the procedure described in Trolle and Schwartz (2009, Appendix B). This requires an approximation of the early-exercise premium with the approximation error becoming more profound with increasing option maturity due to increasing early-exercise premium (as a percentage of the total option price).

[1]Moneyness is defined as the ratio of the strike price to the price of the underlying.

Table 2. Parameter estimates.

Maturities	Apr-15	May-15	Jun-15	Jul-15	Aug-15	Sep-15	Dec-15	Mar-16
S_0	53.86	54.95	56.11	57.27	58.35	59.26	61.28	62.83
r	0.00023	0.00047	0.00072	0.00104	0.00134	0.00170	0.00287	0.00423
Panel A. Heston model								
V_0	0.07600	0.05429	0.01033	0.07775	0.06321	0.07515	0.03042	0.04233
α	9.50206	6.31880	5.96210	3.83075	3.43548	2.68401	1.91593	1.41502
β	0.81483	0.65139	0.51239	0.40253	0.33722	0.29452	0.27397	0.18111
γ	3.93384	2.86685	2.23042	1.75488	1.52145	1.25605	1.02393	0.71546
ρ	-0.24516	-0.30362	-0.38460	-0.40354	-0.43811	-0.48035	-0.55634	-0.45537
Panel B. Bates model								
λ	2.06292	2.18428	1.21831	1.11805	0.76559	0.96278	0.80730	0.89852
μ_J	0.10440	0.16345	0.17049	0.15311	0.16134	0.14633	0.10748	0.07516
σ_J	0.11889	0.06909	0.08828	0.10017	0.10645	0.09321	0.09406	0.12787
V_0	0.02465	0.04522	0.01836	0.05299	0.03961	0.03738	0.01002	0.03397
α	15.72908	7.63344	6.12566	4.71097	4.61157	2.65080	2.44543	1.51012
β	0.55265	0.43761	0.40321	0.31188	0.25908	0.27714	0.22588	0.14665
γ	4.16949	2.58458	2.22234	1.71410	1.54578	1.21052	0.96353	0.66544
ρ	-0.46634	-0.70311	-0.66433	-0.72267	-0.73495	-0.81544	-0.80597	-0.79549

The table presents the model calibration outcome as on February 3, 2015. Annual parameter estimates for the Heston (Panel A) and Bates (Panel B) models are extracted from option prices with different strikes for each contract maturity (see Eq. 6): first six nearby maturities, ranging from 42 days to maturity (April 2015) to 195 days to maturity (September 2015) as well as December 2015 (287 days to maturity) and March 2016 (379 days to maturity).

associated with high volatility, e.g., see Pindyck 2004). However, our observation is consistent with recent evidence in Trolle and Schwartz (2009) and Larsson and Nossman (2011); Geman and Shih (2009) also find that the inverse leverage effect is not present in crude oil prices in the post-2000 era.

Incorporating price jumps in the dynamical structure of the log-return process (Bates model) naturally has a reduction effect on the long-run mean variance, the half-life, the volatility of variance and the correlation level as the inclusion of jumps reduces the need for the variance process to create large sudden movements (Panel B). For example, the long-run mean volatility $\sqrt{\beta}$ range reduces to 74%–38%. The estimated variance mean-reversion speeds translate to half-lives of approx. 2 to 24 weeks. The mean jump size μ_J lies within 7.5% to 17% with an average of 13.5% across maturities, whereas the standard deviation of the jump size σ_J lies within 6.9% to 12.8%. The average, across maturities, frequency of jumps λ is 1.25; the jump arrival rate is higher at nearby maturities and decays at more distant maturities, in consistency with the Samuelson (1965) effect. Finally, similarly to the Heston model, the correlation coefficient estimates are negative.

In addition, we test the performance of the two crude oil models in option pricing. To this end, we compute for each maturity T the following statistics which measure the tracking error from the market quotes

$$\text{MAE}_T := \frac{1}{m}\sum_{i=1}^{m}\left|C_{i,T}^M - C_{i,T}^{\theta_T^*}\right|, \qquad \text{MAPE}_T := \frac{1}{m}\sum_{i=1}^{m}\left|1 - \frac{C_{i,T}^{\theta_T^*}}{C_{i,T}^M}\right|, \tag{7}$$

$$\text{RMSE}_T := \sqrt{\frac{1}{m}\sum_{i=1}^{m}\left|C_{i,T}^M - C_{i,T}^{\theta_T^*}\right|^2}, \qquad \text{RMSPE}_T := \sqrt{\frac{1}{m}\sum_{i=1}^{m}\left|1 - \frac{C_{i,T}^{\theta_T^*}}{C_{i,T}^M}\right|^2}, \tag{8}$$

where θ_T^* is given by (6). Results are presented in Table 3.

A collective view of our calibration results suggests that both models generate low pricing errors. The aggregate, i.e., obtained across all maturities, MAPE (RMSPE) of 1.76% (2.23%) and 1.02% (1.69%), respectively, for the Heston and Bates models (Panels A–B) support their ability to capture

Table 3. Option pricing errors.

	Apr-15	May-15	Jun-15	Jul-15	Aug-15	Sep-15	Dec-15	Mar-16	Aggregate
Panel A. Heston model									
MAE	0.0094	0.0248	0.0554	0.0352	0.0276	0.0468	0.1143	0.0303	0.0464
MAPE	0.9380	1.3404	2.2882	1.3236	1.0873	1.6532	3.0435	0.7643	1.7550
RMSE	0.0125	0.0315	0.0661	0.0385	0.0334	0.0556	0.1334	0.0363	0.0684
RMSPE	1.2402	1.9105	2.5715	1.4737	1.2684	1.9507	3.6780	0.8903	2.2290
Panel B. Bates model									
MAE	0.0128	0.0163	0.0089	0.0075	0.0082	0.0156	0.1133	0.0182	0.0291
MAPE	0.8782	0.8669	0.4459	0.3220	0.2261	0.5283	2.9883	0.4470	1.0247
RMSE	0.0175	0.0206	0.0113	0.0089	0.0142	0.0200	0.1344	0.0219	0.0568
RMSPE	1.0899	1.1877	0.5607	0.4001	0.3120	0.6081	3.6795	0.5365	1.6920
Panel C. Pairwise t-statistics for model comparisons									
MAE	2.9051	-2.9141	-11.0159	-7.5163	-3.6091	-6.4501	-0.3031	-3.775	-9.8488
	[0.005]	[0.005]	[0.000]	[0.000]	[0.005]	[0.000]	[0.763]	[0.003]	[0.000]
	[0.002]	[0.997]	[1.000]	[1.000]	[0.998]	[1.000]	[0.619]	[0.998]	[1.000]
MAPE	-0.7862	-3.4698	-13.4752	-8.11	-3.8869	-6.4041	-1.6451	-4.8082	-11.92
	[0.434]	[0.001]	[0.000]	[0.000]	[0.003]	[0.000]	[0.106]	[0.001]	[0.000]
	[0.783]	[0.999]	[1.000]	[1.000]	[0.998]	[1.000]	[0.947]	[1.000]	[1.000]
RMSE	3.1533	-3.1205	-7.9804	-5.1265	-2.6974	-5.2083	2.0908	-2.7093	-1.4749
	[0.002]	[0.003]	[0.000]	[0.000]	[0.022]	[0.000]	[0.041]	[0.020]	[0.141]
	[0.001]	[0.998]	[1.000]	[1.000]	[0.989]	[1.000]	[0.021]	[0.990]	[0.929]
RMSPE	-1.8536	-2.7897	-7.7334	-3.8895	-2.8122	-3.9983	1.4142	-3.5058	-0.0774
	[0.068]	[0.007]	[0.000]	[0.001]	[0.018]	[0.000]	[0.163]	[0.005]	[0.938]
	[0.966]	[0.996]	[1.000]	[1.000]	[0.991]	[1.000]	[0.081]	[0.998]	[0.531]

Panels A and B of the table report, respectively, for the Heston and Bates models the average performance measures for each maturity as well as the aggregate average: mean absolute error (MAE), mean absolute percentage error (MAPE), root mean square error (RMSE) and root mean square percentage error (RMSPE) (see Eqs. 7–8). The theoretical option prices used in the computation of the measures are based on the parameter estimates in Table 2. Panel C presents pairwise t-statistics for each performance measure G (first entry of each cell), the p-values (in [·]) for the null H_0: $\bar{\text{G}}_{\text{Bates}} = \bar{\text{G}}_{\text{Heston}}$ (second entry) and for the null H_0: $\bar{\text{G}}_{\text{Bates}} < \bar{\text{G}}_{\text{Heston}}$ (third entry).

stylized properties of crude oil prices. In addition, the Bates model yields lower pricing error across maturities implying the importance of modelling price jumps in order to flexibly describe abrupt changes in the WTI prices. Based on the aggregate reports, for example, the improvement in the option price fitting performance achieved by the Bates model is 37%, 41%, 17% and 24% in MAE, MAPE, RMSE and RMSPE, respectively.

We conclude this analysis by referring to an additional robustness check in Panel C of Table 3, where we report pairwise comparisons of the two models to test whether the Bates model yields statistically significant improvements in the performance measures. Largely, our results suggest that ignoring the price jumps generates statistically higher errors. In particular, the hypothesis that the average performance measures across maturities are equal is rejected at conventional significance levels, with the exception of the Dec-15 MAE and aggregate RMSE and RMSPE. Yet, the hypothesis that the Bates model yields lower errors cannot be rejected at the 1% level in all cases.

5. Valuation of Asian options under stochastic volatility: a recursive-integral approach

5.1. Backward price recursion

Consider an Asian option with underlying S observed over the period $[0, T]$ at the equidistant times $t_0 = 0, t_1 = \Delta, t_2 = 2\Delta, \ldots, t_N = N\Delta = T$. Let $\{Z_j\}_{k=1}^{N}$ be a collection of random variables

Table 4. Coefficients $\{l_k\}_{k=0}^N$ for various types of Asian options.

Option type	l_0	$l_1, \ldots, l_{N-1}$	l_N
Floating strike, put	$\frac{1}{N+1}$	$\frac{1}{N+1}$	$\frac{1}{N+1} - K$
Floating strike, call	$-\frac{1}{N+1}$	$-\frac{1}{N+1}$	$K - \frac{1}{N+1}$
Fixed strike, put	$\frac{K}{S_0} - \frac{1}{N+1}$	$-\frac{1}{N+1}$	$-\frac{1}{N+1}$
Fixed strike, call	$\frac{1}{N+1} - \frac{K}{S_0}$	$\frac{1}{N+1}$	$\frac{1}{N+1}$

representing the log-returns on S so that

$$S_k = S_0 \exp\left(\sum_{j=1}^{k} Z_j\right), \ k = 1, \ldots, N, \ S_0 > 0.$$

Under the unified framework of Večeř (2002), the payoff of an Asian option is generally given by

$$\left(\sum_{k=0}^{N} l_k S_k\right)^+,$$

where $(\cdot)^+$ denotes the positive part function and the coefficients $\{l_k\}_{k=0}^N$ are deterministic and take different values for different contract specifications (see Table 4). For example, we focus here on the case of a floating strike (also called average strike) put option with

$$l_k := \frac{1}{N+1}, \ 0 \le k < N, \ \text{and } l_N := \frac{1}{N+1} - K, \ K > 0, \tag{9}$$

from which we retrieve the payoff of the option

$$\left(\sum_{k=0}^{N} l_k S_k\right)^+ = \left(\frac{1}{N+1} \sum_{k=0}^{N} S_k - K S_N\right)^+. \tag{10}$$

The floating strike option is an important contract with a helpful structure in volatile or hardly predictable markets which justifies its high demand. Furthermore, beyond energy markets, as mentioned in the Introduction, fixed exchange rate linked quanto average strike options are, amongst others, particularly popular and actively traded over-the-counter as they can hedge both the foreign stock price and exchange risks for domestic investors (see Chang and Tsao 2011).

Consider the process

$$Y_k = \ln(e^{Y_{k-1}} + l_{k-1}) - Z_k, \ 1 < k \le N, \tag{11}$$
$$Y_1 = \ln l_0 - Z_1, \tag{12}$$

where $l_k > 0$ for $k = 0, \ldots, N-1$. By evaluating $\exp(Y_k)$ recursively using (11)–(12) it is straightforward to show that (10) can be written as

$$(S_N e^{Y_N} + S_N l_N)^+ = S_N(e^{Y_N} + l_N)^+, \ l_N \in \mathbb{R},$$

where the equality follows by homogeneity of the payoff function of degree 1. (This factorization approach has been contributed by Stewart Hodges and first appeared in Carverhill and Clewlow 1990.) The price of the option is given under the risk neutral measure $\hat{P}$ by

$$\hat{E}(S_N e^{-rt_N}(e^{Y_N} + l_N)^+) = \hat{E}(S_N e^{-rt_N})E((e^{Y_N} + l_N)^+),$$

where the equality follows from the numéraire change formula (Geman *et al.* 1995) with the second expectation being taken under a new measure P with the underlying asset representing the numéraire. This change of measure reduces effectively the pricing problem to a two-dimensional one. More specifically, the process Y is not Markov as its evolution is determined also by the level of volatility. To regain a Markov process one must consider the two-dimensional process (Y, V), then the expected value

$$E((e^{Y_N} + l_N)^+) \tag{13}$$

can be computed recursively under the measure P.

THEOREM 5.1 *Let c be the variance transition density given by (2) and f the log-return density conditional on the variance levels x_v and y_v at the endpoints of the time interval $[t_{k-1}, t_k]$. Consider constants $\{l_k\}_{k=0}^N$ given by (9). Define*

$$p_N(y) = (e^y + l_N)^+, \tag{14}$$

$$h_k(y) = \ln(e^y + l_k), \ 0 < k < N,$$

$$q_{k-1}(x, x_v) = \int_{\mathbb{R}^+} \int_{\mathbb{R}} p_k(x - z, y_v) g(\Delta, z, y_v | x_v) dz dy_v, \ 0 < k \le N, \tag{15}$$

$$= \int_{\mathbb{R}^+} \tilde{q}_{k-1}(x, x_v, y_v) c(\Delta, y_v | x_v) dy_v, \tag{16}$$

$$\tilde{q}_{k-1}(x, x_v, y_v) = \int_{\mathbb{R}} p_k(x - z, y_v) f(\Delta, z | x_v, y_v) dz, \ 0 < k \le N, \tag{17}$$

$$p_k(y, y_v) = q_k(h_k(y), y_v), \ 0 < k < N, \tag{18}$$

where q_{k-1} is the option value function at time t_{k-1} and

$$g(\Delta, z, y_v | x_v) = c(\Delta, y_v | x_v) f(\Delta, z | x_v, y_v)$$

is the joint density of the log-return and variance at t_k given the information at t_{k-1} from which (16)–(17) follow. (17) is the convolution of the function p_k with the log-return density f.

Then, the expected value (13) is given by

$$E((e^{Y_N} + l_N)^+) = q_0(\ln l_0, v_0).$$

Proof. We prove by induction on k that $E((e^{Y_N} + l_N)^+ | \mathcal{F}_k) = p_k(Y_k, V_k)$ for $k = 0, \ldots, N$. Trivially the result holds for $k = N$. Suppose $E((e^{Y_N} + l_N)^+ | \mathcal{F}_{k+1}) = p_{k+1}(Y_{k+1}, V_{k+1})$ holds for arbitrary $k < N - 1$. By iterated expectations,

$$E((e^{Y_N} + l_N)^+ | \mathcal{F}_k) = E[E((e^{Y_N} + l_N)^+ | \mathcal{F}_{k+1}) | \mathcal{F}_k] = E(p_{k+1}(Y_{k+1}, V_{k+1}) | \mathcal{F}_k)$$
$$= E(p_{k+1}(\ln(e^{Y_k} + l_k) - Z_{k+1}, V_{k+1}) | \mathcal{F}_k),$$

where the last equality follows from (11). Then

$$E(p_{k+1}(\ln(e^{Y_k} + l_k) - Z_{k+1}, V_{k+1}) | \mathcal{F}_k)$$
$$= \int_{\mathbb{R}^+} \int_{\mathbb{R}} p_{k+1}(\ln(e^{Y_k} + l_k) - z, y_v) g(\Delta, z, y_v | x_v) dz dy_v$$
$$= q_k(\ln(e^{Y_k} + l_k), V_k) = p_k(Y_k, V_k),$$

where the last two equalities follow from (15) and (18). Therefore, by induction,

$$E((e^{Y_N} + l_N)^+) = E((e^{Y_N} + l_N)^+ | \mathcal{F}_0) = q_0(\ln l_0, V_0).$$

$\square$

Few comments are in order. First, the inner integral in (15) is isolated and defined separately in (17) and (15) is re-expressed accordingly in (16) aiming to facilitate the numerical implementation described in Section 5.2. Second, the conditional log-return density f in (17) is not known explicitly, however through Proposition 5.1 we gain access to the associated characteristic function which serves efficiently when solving the problem numerically by Fourier transform as we describe in Section 5.2. Finally, it is worth noting that the valuation procedure developed in Theorem 5.1 reduces to the one-dimensional case considered in Černý and Kyriakou (2011) under Lévy log-increments of the underlying asset price. Given the price of the put option, the call option can also be priced via standard put-call parity.

PROPOSITION 5.1 *If processes V and X evolve according to (1) and (4), the characteristic function under the measure P of the log-return Z_k conditional on the states x_v and y_v of the variance process at the endpoints of the time interval $[t_{k-1}, t_k]$, i.e., $\phi(\Delta, u|x_v, y_v) = E(\exp(iuZ_k)|V_{k-1} = x_v, V_k = y_v)$, is given by*

$$\phi(\Delta, u|x_v, y_v) = \hat{E}\left(e^{Z_k}\right)^{-1} \exp\left\{i(u-i)\left(\left(\mu - \frac{\rho\alpha\beta}{\gamma}\right)\Delta + \frac{\rho}{\gamma}(y_v - x_v)\right)\right\}$$

$$\times \zeta\left(\Delta, (u-i)\left(\frac{\rho\alpha}{\gamma} - \frac{1}{2} + \frac{i(u-i)(1-\rho^2)}{2}\right)\Big| x_v, y_v\right), \qquad (19)$$

where ζ is given in (3).
If, instead, X follows (5), then

$$\phi(\Delta, u|x_v, y_v) = \hat{E}\left(e^{Z_k}\right)^{-1} \exp\left\{i(u-i)\left(\left(\mu - \lambda\kappa(-i) - \frac{\rho\alpha\beta}{\gamma}\right)\Delta + \frac{\rho}{\gamma}(y_v - x_v)\right) + \lambda\kappa(u-i)\Delta\right\}$$

$$\times \zeta\left(\Delta, (u-i)\left(\frac{\rho\alpha}{\gamma} - \frac{1}{2} + \frac{i(u-i)(1-\rho^2)}{2}\right)\Big| x_v, y_v\right). \qquad (20)$$

Proof. From the numéraire change formula

$$E\left(e^{iuZ_k}|V_{k-1} = x_v, V_k = y_v\right) = \hat{E}\left(e^{Z_k}\right)^{-1}\hat{E}\left(e^{i(u-i)Z_k}|V_{k-1} = x_v, V_k = y_v\right).$$

From the tower property of expectations,

$$\hat{E}\left(e^{i(u-i)Z_k}|V_{k-1} = x_v, V_k = y_v\right) = \hat{E}\left[\hat{E}\left(e^{i(u-i)Z_k}\Big| \int_{t_{k-1}}^{t_k} V_s ds\right)\Big| V_{k-1} = x_v, V_k = y_v\right]$$

$$= \exp\left\{i(u-i)\left(\left(\mu - \frac{\rho\alpha\beta}{\gamma}\right)\Delta + \frac{\rho}{\gamma}(y_v - x_v)\right)\right\}$$

$$\times \hat{E}\left[\exp\left\{i(u-i)\left(\frac{\rho\alpha}{\gamma} - \frac{1}{2} + \frac{i(u-i)(1-\rho^2)}{2}\right)\int_{t_{k-1}}^{t_k} V_s ds\right\}\Big| V_{k-1} = x_v, V_k = y_v\right], \qquad (21)$$

where the second equality follows from (4) and (1) and normality of Z conditional on the integral of the variance process. Then, (19) follows from (21) and (3). Extending to (20) is straightforward

14

by independence of the jump component in (5). $\qquad\square$

It is important to stress that our pricing setup can be adapted to general ASV models, including time changed Lévy processes (see Carr *et al.* 2003 and Carr and Wu 2004) using, for example, the integral of a CIR variance process, and the Barndorff–Nielsen–Shephard model of Barndorff-Nielsen and Shephard (2001) and Barndorff-Nielsen *et al.* (2002) with a Lévy-driven Ornstein–Uhlenbeck variance process. More specifically, from Bartlett (1938, p. 62–63), the required characteristic function of the log-return conditional on the variance state at time t (see Proposition 5.1) is generally given by

$$\phi(t, u|v_0, v) = \frac{\int e^{-iv\omega} \xi(t, \omega, u|v_0) d\omega}{\int e^{-iv\omega} \xi(t, \omega, 0|v_0) d\omega},$$

where $\xi(t, \omega, u|v_0) := E\left(\exp(i\omega V_t + iuZ_t)|V_0 = v_0\right)$ admits explicit representations under various ASV processes (e.g., see Fusai and Kyriakou 2016, Hubalek *et al.* 2014). In this paper, we consider the Heston log-asset price process, which forms a special case in this class being given by a CIR time changed arithmetic Brownian motion (e.g., see Kallsen and Pauwels 2011), as well as the Bates model with independent random jumps.

Finally, another advantage of our method is that the option price sensitivities (also known as the Greeks) can be obtained assuming that we can differentiate equations (16)–(17) under the integral sign (a usual assumption in option pricing via Fourier transform) with respect to parameters of interest, e.g., the initial variance v_0, etc. (see Ballotta *et al.* 2016 for the case of general Lévy log-increments of the underlying asset price).

5.2. *Implementation of backward recursion by combined Fourier transform-quadrature method*

Equations (14)–(18) define the problem we want to solve numerically. We compute by backward recursion at each monitoring date the functions $\tilde{q}$, q and p starting from maturity to obtain ultimately the option prices on a grid of variance values at time t_0. To ensure fast computation of the integrals (16)–(17), we employ a numerical method that combines standard trapezoidal quadrature rule and discrete Fourier transform. Our numerical algorithm proceeds as follows.

Step 0 Preliminary computations: grid construction, density function and characteristic function valuations.

  a) Select equally spaced grids for the variance and the value of the underlying asset on the log-scale, respectively, $\ln \boldsymbol{x_v} = \{\ln x_{v,0} + m_v \delta_v\}_{m_v=0}^{n_v-1}$ and $\boldsymbol{x} = \{x_0 + m\delta\}_{m=0}^{n-1}$ on which q_{k-1} in (16) is evaluated (see *Step 2*). Based on the rules set out in Fang and Oosterlee (2011), choose $\ln x_{v,0} = \ln E(V_\Delta | V_0 = v_0) - \varrho(1+d)^{-1} = \ln\left(v_0 e^{-\alpha\Delta} + \beta(1 - e^{-\alpha\Delta})\right) - \varrho(1+d)^{-1}$, where $d = 2\alpha\beta\gamma^{-2} - 1$ and ϱ is a user-defined proportionality constant, and $x_0 = c_1 - \varrho\sqrt{c_2}$ where c_j is the j-th cumulant of the log-asset price distribution (this can be computed explicitly by straightforward differentiation of its log-characteristic function using any symbolic computation package such as Mathematica). For given grid sizes n_v and n and assuming a symmetric grid, compute, respectively, grid spacings $\delta_v = (\ln x_{v,n_v-1} - \ln x_{v,0})n_v^{-1} = 2\varrho(1+d)^{-1}n_v^{-1}$ and $\delta = (x_{n-1} - x_0)n^{-1} = 2\varrho\sqrt{c_2}n^{-1}$.
  b) Select grid $\boldsymbol{y_v}$ (for convenience, $\boldsymbol{y_v} = \boldsymbol{x_v}$) and evaluate the CIR variance transition density given in (2). Let $\boldsymbol{c} = \{c(\Delta, \boldsymbol{y_{v,j_v}}|\boldsymbol{x_{v,m_v}})\}_{j_v=0,m_v=0}^{n_v-1,n_v-1}$ be the density function values on the grid $(\boldsymbol{y_v}, \boldsymbol{x_v})$; store for use for all k in *Step 2*.

  REMARK 5.1 *The Bessel function $I_d(z)$ (see Eq. 2) can be computed, for example, in Matlab using the built-in function* `besseli(d,z)` *for real order d and complex argument z. To avoid computational problems due to extreme Bessel function values, it*

15

is advisable to compute first the scaled version $\exp(-|\mathrm{Re}\,z|)\,I_d(z)$ based on the robust package of Amos (1985, 1986) available as `besseli(d,z,1)` in Matlab, and rescale to retrieve $I_d(z)$. In addition, when the Feller condition is not satisfied, i.e., $d < 0$, the process can reach zero when started from $v_0 > 0$, and zero is an instantaneously reflecting boundary. Numerical experiments in Fang and Oosterlee (2011) show that decreasing d affects the behaviour of the CIR density (does not tail off on both sides) on the left (slow decay, constant tail or drastic increase in value), which can be avoided by transforming to the log-variance domain, for which reason we construct the variance grid on the log-scale in the above step.

    c) Select grid $\boldsymbol{y}$ (for convenience, $\boldsymbol{y} = \boldsymbol{x}$) on which p_k in (18) is evaluated (see *Step 3*).

    d) Select uniform, symmetric grid $\boldsymbol{u} = \{(m_u - n/2)\delta_u\}_{m_u=0}^{n-1}$ of size n and spacing δ_u for use next and in *Step 1*. The range of values of the grid $\boldsymbol{u}$ is chosen so that the tail of the absolute value of the characteristic function (19) or (20) is sufficiently captured at the tails, i.e., $|\phi| \leq 10^{-\varrho'}$ where $\varrho' \in \mathbb{N}$ is guided by the desired precision.

    e) Let $\boldsymbol{\phi} = \{\phi(\Delta, \boldsymbol{u}_{m_u}|\boldsymbol{x}_{\boldsymbol{v},m_v}, \boldsymbol{y}_{\boldsymbol{v},j_v})\}_{m_u=0,m_v=0,j_v=0}^{n-1,n_v-1,n_v-1}$ be the values of the characteristic function ϕ on the grid $(\boldsymbol{u}, \boldsymbol{y}_{\boldsymbol{v}}, \boldsymbol{x}_{\boldsymbol{v}})$; store for use for all k in *Step 1*.

Step 1 Consider $\tilde{q}_{k-1}$ in (17). For $j_v = 0, \ldots, n_v - 1$ and $m_v = 0, \ldots, n_v - 1$, compute, conditional on the variance levels $\boldsymbol{x}_{\boldsymbol{v},m_v}$ and $\boldsymbol{y}_{\boldsymbol{v},j_v}$ at the endpoints of the time interval $[t_{k-1}, t_k]$, the values of $\tilde{q}_{k-1}$ on the grid of log-asset values $\boldsymbol{x}$ by means of the inverse discrete Fourier transform:

$$\tilde{\boldsymbol{q}}_{k-1,\cdot,j_v,m_v} = \frac{\delta_u}{2\pi} e^{i\frac{n}{2}\delta_u(\boldsymbol{x}^{\mathrm{T}}-x_0)} \circ \sum_{m_u=0}^{n-1} e^{-i\frac{2\pi}{n}m_u m} \left(e^{-i\boldsymbol{u}^{\mathrm{T}}x_0} \circ \boldsymbol{P}_{k,\cdot,j_v} \circ \boldsymbol{\phi}_{\cdot,j_v,m_v} \right)_{m_u}, \qquad (22)$$

where $\circ$ denotes the Hadamard element-wise product and, for $j_v = 0, \ldots, n_v - 1$,

$$\boldsymbol{P}_{k,\cdot,j_v} = \delta e^{i(\boldsymbol{u}+\frac{n}{2}\delta_u)y_0} \circ \sum_{j=0}^{n-1} e^{i\frac{2\pi}{n}m_u j} \left(e^{-i\frac{n}{2}\delta_u \boldsymbol{y}} \circ \boldsymbol{p}_{k,\cdot,j_v}{}^{\mathrm{T}} \circ \boldsymbol{w} \right)_j \qquad (23)$$

are the values of the discrete Fourier transform of p_k on the grid $\boldsymbol{u}$ for given variance level $\boldsymbol{y}_{\boldsymbol{v},j_v}$ at time t_k. $\boldsymbol{w} = [1/2, 1, 1, \cdots, 1, 1, 1/2]$ is the vector of trapezoid weights. In (23), $\boldsymbol{p}_k$ on the grid $(\boldsymbol{y}, \boldsymbol{y}_{\boldsymbol{v}})$ is given by (14) for $k = N$; see *Step 3* for $k < N$.

REMARK 5.2 *The sums in (22) and (23) can be computed using the fast Fourier transform (FFT) which is readily available in Matlab as `fft` and `ifft`. As required for a FFT implementation, check that the Nyquist relation $\delta_u \delta = 2\pi/n$ holds; if not, adjust the original grids accordingly.*

Step 2 Consider q_{k-1} in (16). For $m_v = 0, \ldots, n_v - 1$, compute for given variance level $\boldsymbol{x}_{\boldsymbol{v},m_v}$ at time t_{k-1} the values of q_{k-1} on the grid of log-asset values $\boldsymbol{x}$:

$$\boldsymbol{q}_{k-1,\cdot,m_v} = \tilde{\boldsymbol{q}}_{k-1,\cdot,\cdot,m_v} \left(\boldsymbol{c}_{\cdot,m_v} \circ \boldsymbol{w}^{\mathrm{T}} \right). \qquad (24)$$

REMARK 5.3 *As the Feller condition is satisfied in our application, there is no noticeable impact on the accuracy of the trapezoid rule. Violation of the Feller condition would necessitate the use of a larger grid, hence of a higher-order quadrature rule.*

Step 3 Consider p_{k-1} in (18). The values of p_{k-1} on the grid $(\boldsymbol{y}, \boldsymbol{y}_{\boldsymbol{v}})$ are given by $p_{k-1}(\boldsymbol{y}, \boldsymbol{y}_{\boldsymbol{v}}) = q_{k-1}(h_{k-1}(\boldsymbol{y}), \boldsymbol{y}_{\boldsymbol{v}})$: fit a cubic interpolating spline to the nodes $(\boldsymbol{x}, \boldsymbol{q}_{k-1,\cdot,m_v})$, for $m_v = 0, \ldots, n_v - 1$, with not-a-knot endpoint conditions to evaluate q_{k-1} at $h_{k-1}(\boldsymbol{y}) \subseteq \boldsymbol{x}$; for

Table 5.   Asian option prices.

Back. recursion	Opt. LB	CVMC	Std. err. $\times 10^{-5}$	95% CI
(a) Heston model				
10.1464609	10.128	10.14651	11.413	[10.1463, 10.1467]
9.0447223	9.032	9.04481	7.254	[9.0447, 9.0449]
8.2269388	8.219	8.22694	4.407	[8.2269, 8.2270]
7.1199900	7.114	7.11999	3.283	[7.1199, 7.1201]
6.5601931	6.556	6.56019	2.513	[6.5601, 6.5602]
6.0491417	6.046	6.04914	2.115	[6.0491, 6.0492]
5.5730025	5.570	5.57300	1.635	[5.5730, 5.5730]
4.5498436	4.548	4.54983	1.089	[4.5498, 4.5498]
(b) Bates model				
8.9372259	8.925	8.93726	6.213	[8.9371, 8.9373]
8.0169212	8.009	8.01688	3.997	[8.0168, 8.0169]
7.5854509	7.579	7.58549	3.418	[7.5854, 7.5855]
6.7104526	6.705	6.71046	2.624	[6.7104, 6.7105]
6.1970137	6.193	6.19700	2.134	[6.1970, 6.1970]
6.0241464	6.020	6.02410	2.108	[6.0241, 6.0241]
5.513648	5.511	5.51365	1.528	[5.5136, 5.5137]
4.4666072	4.465	4.46661	1.160	[4.4666, 4.4666]

The table presents the results from our backward recursive-integral approach (back. recursion), the optimized option price lower bounds of Fusai and Kyriakou (2016) (opt. LB) and the Monte Carlo price estimates obtained using opt. LB as control variate (CVMC) (std. err.: standard error of CVMC price estimates, CI: confidence interval for 10^7 runs). Model parameters: see Table 2 and $\mu = 0$ (due to the underlying futures contracts). Other parameters: $K = 1$ (see Eq. 10), $T = 1$, $N = 12$.

$h_{k-1}(\boldsymbol{y}) \nsubseteq \boldsymbol{x}$ use linear extrapolation in e^x (for example, use `interp1` available in Matlab). Continue with *Step 1* until $k = 1$.

REMARK 5.4   *The outcome of the numerical scheme when $k = 1$ comprises option values on the grid of n_v initial variance values v_0 (see Eq. 24).*

5.3.   *Application to the valuation of commodity options*

5.3.1.   Discrete average options. In order to test our backward recursive approach, we perform numerical experiments under the Heston and Bates model parameterizations obtained in Section 4. We note that our parameter sets (see Table 2) correspond to various levels of volatility, skewness coefficient and excess kurtosis of the log-returns. We compare our results with benchmark prices generated by a control variate Monte Carlo (CVMC) strategy which uses as control variate the optimized option price lower bound of Fusai and Kyriakou (2016). To this end, we employ standard CVMC setup with the CV coefficient estimated in a pilot run (e.g., see Glasserman 2004). In general, stochastic volatility models are hard to simulate accurately and various approaches have been proposed (e.g., see discussion in Tse and Wan 2013). We choose a widely used and efficient one in the literature, that is, the quadratic-exponential method of Andersen (2008) to simulate the square root variance process with central discretization of the integrated variance.

We produce numerical results for a floating strike Asian option with monthly monitoring frequency[1] ($N = 12$). In Figure 4 we present error patterns of computed option prices in the Heston and Bates models with increasing number of integration points n (see Section 5.2). We observe that smoothly diminishing error patterns are preserved under both models and all parameter sets, ensuring convergence to the desired number of decimal places with increasing n. To quantify the

---

[1]Higher monitoring frequencies have also been considered (for example in Figure 5 we use $N = 16, 32, 64$) with an equally good performance. Additional results are omitted for brevity, however can be made available by the authors upon request.
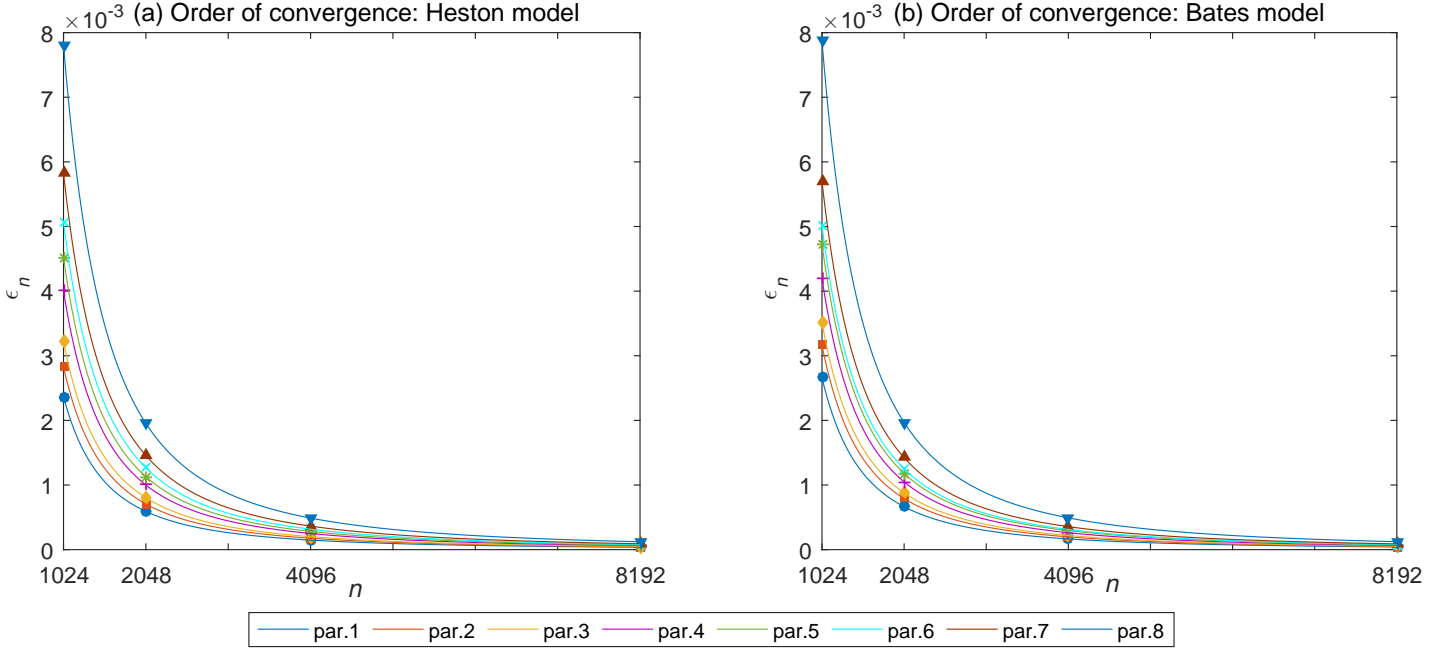
Figure 4. Error patterns of our backward recursive-integral pricing method in number of integration points n under Heston and Bates models for each set of parameters in Table 2. price_n denotes the Asian option price calculated using the numerical scheme in Section 5.2 with n integration points. $\epsilon_n = |\text{price}_n - \text{price}_{n/2}|$ are the absolute price differences for $n = 2^{10}, 2^{11}, 2^{12}, 2^{13}$.

convergence, we fit regression lines. We define the order of convergence as b where the absolute error is $\epsilon_n = an^{-b}$; thus, the log-absolute error $\ln \epsilon_n$ is linear in $\ln n$ and has slope $-b$. In all cases we find that the slope coefficient is $\hat{b} = 2$ and the coefficient of determination is $R^2 = 1.0$ implying order 2 convergence, in consistency with the error analyses of Andricopoulos *et al.* (2003), Lord *et al.* (2008) and Černý and Kyriakou (2011) for the case of one-dimensional pricing problems under Lévy log-returns and continuous value functions. We exploit the order 2 convergence to achieve very high precision faster by extrapolation via a simple Richardson-type procedure; in Table 5 we report results with accuracy of seven decimals achieved in $30s$.[2] Even with a very effective control variate, the Monte Carlo method is not competitive as it takes $200s$ in the case of the Heston model to achieve four decimal places of accuracy (at the 95% confidence level) using 10^7 simulation runs; the simulation of the Bates model is substantially slower due to the additional simulation of jumps. Contrary to our method, it can be hard for users to gauge the precision of Monte Carlo simulation as its convergence may be damped by the existence of discretization bias (see discussion above) requiring the use of a very large number of simulations to evaluate this precisely (such task is beyond our scope), hence rendering Monte Carlo slow; also, it yields varying standard error depending on the choice of model parameters.

5.3.2. **Continuous average options.**
Next, we assess the performance of our pricing methodology in the case of the continuous average. Inherently, our construction refers to a discrete average. Nevertheless, as shown in Figure 5, we achieve smoothly diminishing error patterns with increasing number of monitoring dates N.[1] We gauge the order of convergence by calculating the absolute

---

[2]All CPU time reports correspond to a Matlab R2015b implementation on a PC with an Intel Core i7-4870HQ CPU @ 2.5GHz and 16GB of RAM.

[1]In Figure 5 we present results based on the first set of parameters in Table 2; additional results can be made available by the authors upon request.
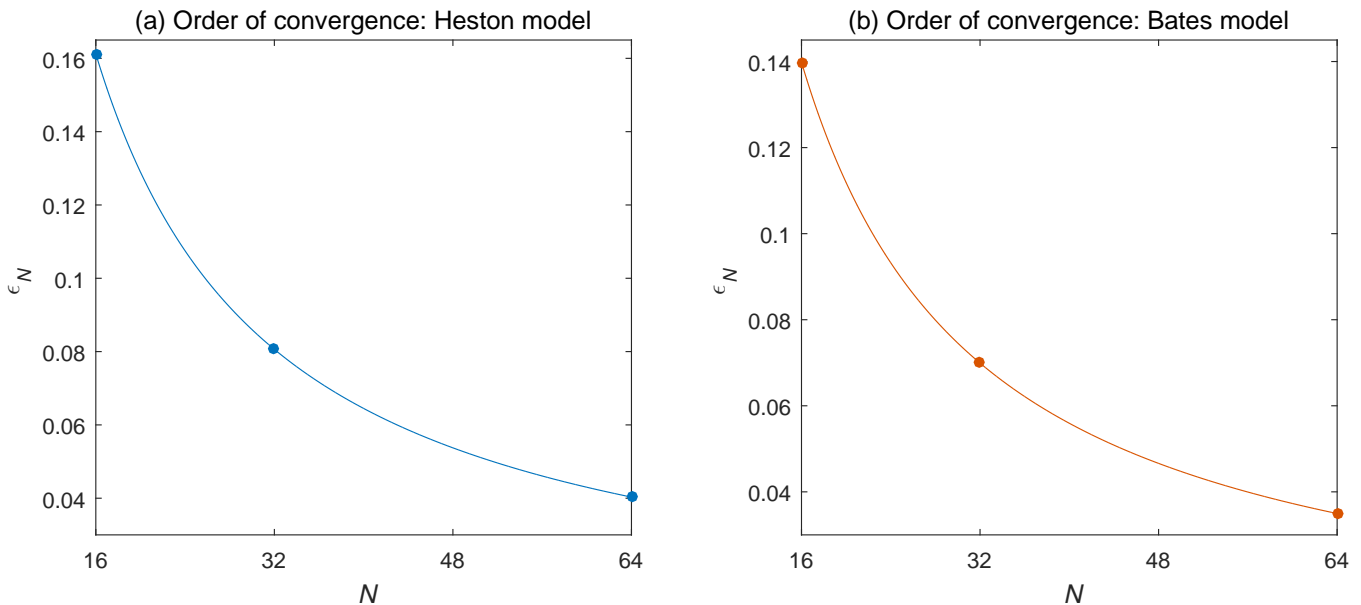
Figure 5. Error patterns of our backward recursive-integral pricing method in number of monitoring dates N under Heston and Bates models for parameter set 1 in Table 2. price_N denotes the price of the Asian option with N monitoring dates calculated using the numerical scheme in Section 5.2. $\epsilon_N = |\mathrm{price}_N - \mathrm{price}_{N/2}|$ are the absolute price differences for $N = 2^4, 2^5, 2^6$.

error $\epsilon_N = aN^{-b}$. By fitting regression lines, we find that the slope coefficient (log-log scale) is $\hat{b} = 1$ and the coefficient of determination is $R^2 = 1.0$ implying order 1 convergence of the discretely monitored Asian option to the continuous solution, which is consistent with the error analyses of Andricopoulos *et al.* (2003) and Lord *et al.* (2008) for high frequency observations and continuous American features. Therefore, the discretely monitored Asian option prices are amenable to Richardson extrapolation in the time dimension speeding up convergence to the price of the continuously monitored Asian option ($N \to \infty$), or an Asian option with a large finite number of monitoring dates ($0 \ll N < \infty$) whose pricing by direct implementation of our method could be computationally demanding. For the sake of exemplification, we use the parameter set 1 and compute the price of the continuously Asian option under the Heston and Bates models: with accuracy of four decimals we obtain, respectively, 10.3546 and 9.1162. (Higher precision can be achieved by extrapolating discrete Asian option prices with $N > 2^6$—for illustration purposes in Figure 5 we consider up to $N = 2^6$.) Existing simulation error in the Monte Carlo price estimates affects the smooth convergence, hence prohibits the use of extrapolation techniques in this case, therefore our approach remains competitive with increasing number of monitoring dates.

5.3.3. Heston versus Bates model. Comparing the option prices in the left panel of Figure 6 obtained under the Heston and Bates models, we see that the Heston prices are higher. The same remark applies in the case of the fitted vanilla option prices from the calibration (6) in Section 4. In our Asian option application, large discrepancies emerge for "shorter-term" calibrations, more specifically, the percentage discrepancy ranges from 13.5% in the case of the Apr-15 calibration (parameter set 1) to 1%−2% in the case of the Mar-16 calibration (parameter set 8). This behaviour stems from the joint effect of the negative skewness, excess kurtosis and standard deviation of the risk neutral distribution of the difference $\frac{1}{N+1}\sum_{k=0}^{N} S_k - S_N$ (see right panel of Figure 6) which results in higher prices under the Heston model. From the right panel of Figure 6 we observe that the Heston model yields larger standard deviation, kurtosis and (absolute) skewness. We also observe decreasing cumulant levels and discrepancies between the two models with increasing time scale; the effect of jumps becomes more visible at shorter time scales where the need to allow for
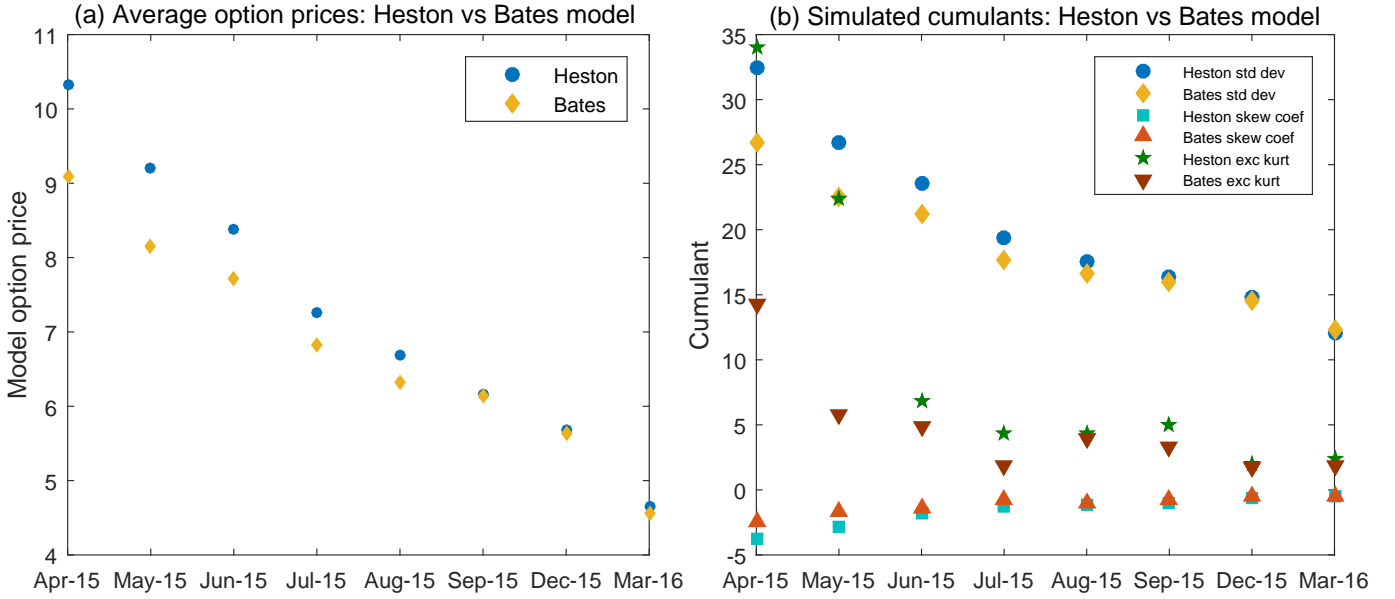
Figure 6. Floating strike Asian put option price patterns across parameter sets in Table 2 corresponding to model calibrations based on WTI futures American vanilla options with maturities Apr-15, May-15, Jun-15, Jul-15, Aug-15, Sep-15, Dec-15 and Mar-16 (left panel), alongside simulated cumulants of the difference $\frac{1}{N+1}\sum_{k=0}^{N} S_k - S_N$ (right panel).

them becomes more stringent, hence the Heston model on itself becomes less sufficient in explaining the option prices leading to a substantial mispricing. This adds to our earlier observation of the Heston model being unable to capture extreme price movements, such as during the Gulf War (1990–91) and/or the events during the 2008 financial turmoil (see Figure 3), but also of the Bates model's superior performance in our vanilla option price fitting exercise in Section 4 as is evident from Table 3. As seen in Figure 6, the pricing error reduces at longer time scales with the two models performing similarly and the cumulants lying at the same levels.

6. Conclusions

We have developed a pricing framework for arithmetic Asian options in the presence of stochastic volatility and price jumps. Our model setup accounts for volatility clustering, price discontinuities, exhibits Samuelson's maturity effect and negative correlation between volatility movements and log-returns, which are ubiquitous features of crude oil prices. In this paper, we have used data from the crude oil derivatives market as this is the most important and liquid one. Nevertheless, our option pricing architecture is generally suited to other energy commodities, metals, agriculturals and freight rates exhibiting similar stylized features. We have estimated the stochastic volatility model with discontinuous price movements as well as its nested version without jumps to NYMEX WTI futures vanilla options and found that both price jumps and stochastic volatility are necessary to fit options. We also show that ignoring the jumps results in nontrivial pricing errors for Asian options especially at short time horizons. Hence, model risk is less severe when both ingredients are included.

In the current application we focus on the Heston and Bates models, nonetheless our methodology is applicable to general ASV models with known joint laws of the underlying (log) asset price and the stochastic variance via the associated characteristic functions. We further plan investigating possible application in time changed Markov jump processes (see subordinate Ornstein–Uhlenbeck (subOU) and time changed subOU models with stochastic volatility in Li and Linetsky 2014)

with state-dependent jumps, i.e., with the jump direction and amplitude dependent on the current states of the driving processes; this remains an open, interesting problem with various mathematical and computational challenges. Currently, our pricing scheme converges smoothly in the number of integration points, hence can reach high levels of precision by means of Richardson extrapolation. In future work we aim to introduce parallelism in the pricing algorithm (e.g., see Corsaro *et al.* 2015, Fusai *et al.* 2010) to further speed up the evaluation procedure while preserving accuracy. Our method is currently suited to both discrete and continuous average options, is free of any form of bias or approximation error that can be hard to quantify, and is adaptable to the computation of the price sensitivities.

Acknowledgements

References

Alexander, C. and Venkatramanan, A., Commodity options. In *The Handbook of Commodity Investing*, The Frank J. Fabozzi Series, edited by F.J. Fabozzi, R. Füss and D.G. Kaiser, pp. 570–595, 2008 (John Wiley & Sons, Inc.: New Jersey).

Amos, D.E., A subroutine package for Bessel functions of a complex argument and nonnegative order. *Sandia National Laboratory Report SAND85-1018, Sandia National Laboratories, Albuquerque, N.M.*, 1985.

Amos, D.E., Algorithm 644: A portable package for Bessel functions of a complex argument and nonnegative order. *ACM Transactions on Mathematical Software*, 1986, **12**, 265–273.

Andersen, L., Simple and efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance*, 2008, **11**, 1–42.

Andricopoulos, A.D., Widdicks, M., Duck, P.W. and Newton, D.P., Universal option valuation using quadrature methods. *Journal of Financial Economics*, 2003, **67**, 447–471.

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

Ballotta, L., Gerrard, R. and Kyriakou, I., Hedging of Asian options under exponential Lévy models: computation and performance. *The European Journal of Finance*, 2016 forthcoming.

Barndorff-Nielsen, O.E., Nicolato, E. and Shephard, N., Some recent developments in stochastic volatility modelling. *Quantitative Finance*, 2002, **2**, 11–23.

Barndorff-Nielsen, O.E. and Shephard, N., Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2001, **63**, 167–241.

Bartlett, M.S., The characteristic function of a conditional statistic. *Journal of the London Mathematical Society*, 1938, **13**, 62–67.

Bates, D.S., Jumps and stochastic volatility: exchange rate processes implicit in deutsche mark options. *Review of Financial Studies*, 1996, **9**, 69–107.

Borovkova, S. and Geman, H., Seasonal and stochastic effects in commodity forward curves. *Review of Derivatives Research*, 2006, **9**, 167–186.

Broadie, M. and Kaya, O., Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research*, 2006, **54**, 217–231.

Campi, C. and Galdenzi, F., Oil markets and products. In *Handbook of Multi-Commodity Markets and Products: Structuring, Trading and Risk Management*, The Wiley Finance Series, edited by A. Roncoroni, G. Fusai and M. Cummins, pp. 3–66, 2015 (John Wiley & Sons, Ltd: Chichester).

Carr, P., Geman, H., Madan, D.B. and Yor, M., Stochastic volatility for Lévy processes. *Mathematical Finance*, 2003, **13**, 345–382.

Carr, P. and Wu, L., Time-changed Lévy processes and option pricing. *Journal of Financial Economics*, 2004, **71**, 113–141.

Carverhill, A. and Clewlow, L., Flexible convolution. *Risk*, 1990, **3**, 25–29.

Černý, A. and Kyriakou, I., An improved convolution algorithm for discretely sampled Asian options. *Quantitative Finance*, 2011, **11**, 381–389.

Chang, C.C. and Tsao, C.Y., Efficient and accurate quadratic approximation methods for pricing Asian strike options. *Quantitative Finance*, 2011, **11**, 729–748.

Clewlow, L. and Strickland, C., *Energy Derivatives: Pricing and Risk Management*, 2000 (Lacima Publications: London).

Corsaro, S., Marazzina, D. and Marino, Z., A parallel wavelet-based pricing procedure for Asian options. *Quantitative Finance*, 2015, **15**, 101–113.

Cox, J.C., Ingersoll, J. E., J. and Ross, S.A., A theory of the term structure of interest rates. *Econometrica*, 1985, **53**, 385–407.

Curran, M., Valuing Asian and portfolio options by conditioning on the geometric mean price. *Management Science*, 1994, **40**, 1705–1711.

Dickey, D.A. and Fuller, W.A., Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 1979, **74**, 427–431.

Dickey, D.A. and Fuller, W.A., Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica*, 1981, **49**, 1057–1072.

Dingeç, K.D., Sak, H. and Hörmann, W., Variance reduction for Asian options under a general model framework. *Review of Finance*, 2015, **19**, 907–949.

Engle, R., Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 1982, **50**, 987–1007.

Ewald, C.O., Menkens, O. and Ting, S.H.M., Asian and Australian options: A common perspective. *Journal of Economic Dynamics and Control*, 2013, **37**, 1001–1018.

Fang, F. and Oosterlee, C.W., A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing*, 2008, **31**, 826–848.

Fang, F. and Oosterlee, C.W., A Fourier-based valuation method for Bermudan and barrier options under Heston's model. *SIAM Journal on Financial Mathematics*, 2011, **2**, 439–463.

Feller, W., Two singular diffusion problems. *Annals of Mathematics*, 1951, **54**, 173–182.

Fouque, J.P. and Han, C.H., Pricing Asian options with stochastic volatility. *Quantitative Finance*, 2003, **3**, 353–362.

Fouque, J.P. and Han, C.H., Variance reduction for Monte Carlo methods to evaluate option prices under multi-factor stochastic volatility models. *Quantitative Finance*, 2004, **4**, 597–606.

Fusai, G. and Kyriakou, I., General optimized lower and upper bounds for discrete and continuous arithmetic Asian options. *Mathematics of Operations Research*, 2016, **41**, 531–559.

Fusai, G., Marazzina, D. and Marena, M., Option pricing, maturity randomization and distributed computing. *Parallel Computing*, 2010, **36**, 403–414.

Geman, H., *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*, The Wiley Finance Series, 2005 (John Wiley & Sons, Ltd: Chichester).

Geman, H., Mean reversion versus random walk in oil and natural gas prices. In *Advances in Mathematical Finance*, Applied and Numerical Harmonic Analysis, edited by M.C. Fu, R.A. Jarrow, J.Y.J. Yen and R.J. Elliott, pp. 219–228, 2007 (Birkhäuser: Boston).

Geman, H., El-Karoui, N. and Rochet, J.C., Changes of numéraire, changes of probability measure and option pricing. *Journal of Applied Probability*, 1995, **32**, 443–458.

Geman, H. and Shih, Y.F., Modeling commodity prices under the CEV model. *Journal of Alternative Investments*, 2009, **11**, 65–84.

Glasserman, P., *Monte Carlo Methods in Financial Engineering*, Stochastic Modelling and Applied Probability, 2004 (Springer–Verlag: New York).

Heston, S.L., A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 1993, **6**, 327–343.

Hilliard, J.E. and Reis, J., Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot. *The Journal of Financial and Quantitative Analysis*, 1998, **33**, 61–86.

Hubalek, F., Keller-Ressel, M. and Sgarra, C., Geometric Asian option pricing in general affine stochastic volatility models with jumps. *Working Paper*, 2014.

Jarque, C.M. and Bera, A.K., Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, 1980, **6**, 255–259.

Kallsen, J. and Pauwels, A., Variance-optimal hedging for time-changed Lévy processes. *Applied Mathematical Finance*, 2011, **18**, 1–28.

Kim, B. and Wee, I.S., Pricing of geometric Asian options under Heston's stochastic volatility model. *Quantitative Finance*, 2014, **14**, 1795–1809.

Larsson, K. and Nossman, M., Jumps and stochastic volatility in oil prices: Time series evidence. *Energy Economics*, 2011, **33**, 504–514.

Li, L. and Linetsky, V., Time-changed Ornstein–Uhlenbeck processes and their applications in commodity derivative models. *Mathematical Finance*, 2014, **24**, 289–330.

Ljung, G.M. and Box, G.E.P., On a measure of lack of fit in time series models. *Biometrika*, 1978, **65**, 297–303.

Lord, R., Fang, F., Bervoets, F. and Oosterlee, C.W., A fast and accurate FFT-based method for pricing early-exercise options under Lévy processes. *SIAM Journal on Scientific Computing*, 2008, **30**, 1678–1705.

Nelson, C.R. and Siegel, A.F., Parsimonious modeling of yield curves. *The Journal of Business*, 1987, **60**, 473–489.

Nomikos, N.K., Kyriakou, I., Papapostolou, N.C. and Pouliasis, P.K., Freight options: Price modelling and empirical analysis. *Transportation Research Part E: Logistics and Transportation Review*, 2013, **51**, 82–94.

Phillips, P.C.B. and Perron, P., Testing for a unit root in time series regression. *Biometrika*, 1988, **75**, 335–346.

Pindyck, R.S., Volatility and commodity price dynamics. *Journal of Futures Markets*, 2004, **24**, 1029–1047.

Rogers, L.C.G. and Shi, Z., The value of an Asian option. *Journal of Applied Probability*, 1995, **32**, 1077–1088.

Samuelson, P.A., Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 1965, **6**, 41–49.

Shiraya, K. and Takahashi, A., Pricing average options on commodities. *Journal of Futures Markets*, 2011, **31**, 407–439.

Tahani, N., Exotic geometric average options pricing under stochastic volatility. *Applied Mathematical Finance*, 2013, **20**, 229–245.

Trolle, A.B. and Schwartz, E.S., Unspanned stochastic volatility and the pricing of commodity derivatives. *Review of Financial Studies*, 2009, **22**, 4423–4461.

Tse, S.T. and Wan, J.W.L., Low-bias simulation scheme for the Heston model by Inverse Gaussian approximation. *Quantitative Finance*, 2013, **13**, 919–937.

Večeř, J., Unified Asian pricing. *Risk*, 2002, **15**, 113–116.

Wong, H.Y. and Cheung, Y.L., Geometric Asian options: valuation and calibration with stochastic volatility. *Quantitative Finance*, 2004, **4**, 301–314.

Wystup, U., *FX Options and Structured Products*, The Wiley Finance Series, 2006 (John Wiley & Sons, Ltd: Chichester).

Yamazaki, A., Pricing average options under time-changed Lévy processes. *Review of Derivatives Research*, 2014, **17**, 79–111.