Dynamic Jump Intensities and Risk Premiums in Crude Oil Futures and Options Markets*

Peter Christoffersen
University of Toronto
CBS, and CREATES

Kris Jacobs
Bauer College of Business
University of Houston

Bingxin Li
West Virginia University

Journal of Derivatives, 2016, 24(2), 8-30

Abstract

Options on crude oil futures are the most actively traded commodity options. We develop a class of computationally efficient discrete-time jump models that allow for closed-form option valuation, and we use crude oil futures and options data to investigate the economic importance of jumps and dynamic jump intensities in these markets. Allowing for jumps is crucial for modeling crude oil futures and futures options, and we find evidence in favor of time-varying jump intensities. During crisis periods, jumps occur more frequently. The properties of the jump processes implied by the option data differ from those implied by the futures data, which may be due to improved parameter identification.

JEL Classification: G12; G13

Keywords: Crude oil; Futures; Options; Discrete-time models; Jump intensities; Risk premiums.

*We would like to thank Hitesh Doshi, Garland Durham, Stephen Figlewski (the editor), Pascal Francois, Yoontae Jeon, Robert Kimmel, Xingguo Luo, Kasper Olesen, Chay Ornthanalai, Craig Pirrong, Ken Singleton, Anders Trolle, Stuart Turnbull, Chu Zhang, an anonymous referee, and seminar participants at the NBER Meeting on Economics of Commodities Markets, the University of Houston, the Singapore International Conference on Finance, the Montreal IFSID conference, the 23rd Annual Derivatives Securities and Risk Management Conference, the China International Conference on Finance, and the FMA Asian Conference for helpful comments. Christoffersen: Professor of Finance, Rotman School of Management, University of Toronto, Phone: 416-946-5511, Email: peter.christoffersen@rotman.utoronto.ca; Jacobs: Professor of Finance, Bauer College of Business, University of Houston, Phone: 713-743-2826, Email: kjacobs@bauer.uh.edu; Li: Assistant Professor of Finance, West Virginia University, Phone: 304-293-7832, Email: bingxin.li@mail.wvu.edu.

Crude oil is the most important commodity in international trade, and the crude oil derivatives market constitutes the most liquid commodity derivatives market. In December 2013, WTI and Brent crude oil futures accounted for almost half of the dollar weight in the S&P GSCI commodity index. Perhaps more importantly, Panels A and B in Exhibit 1 suggest that there is a low-frequency relationship the WTI spot price and the S&P500 equity index. Indeed, for our sample period (January 1990 through May 2014) the correlation between (overlapping) annual log returns on oil and equity is 25.9%. Comparing at-the-money implied volatility from 1-month oil futures options in Panel C with the VIX in Panel D reveals that the "fear gauges" in the two markets share common features as well. Volatility tends to spike at around the same time in the two markets although the magnitudes of the spikes of course vary across markets. The risk and return properties of crude oil are thus clearly related to those of the most commonly used proxy for aggregate wealth.

To price and hedge this increasingly important commodity, it is crucial to model crude oil futures and options and to better understand their dynamics. Surprisingly though, there are relatively few studies on pricing crude oil derivatives, especially when compared with the existing literature on equity derivatives. The extensive literature on modeling equity derivatives, which mainly focuses on index options, concludes that jumps are needed to capture the higher moments of index returns. It seems plausible that jumps are also present in crude oil futures and options, but the existing literature does not address this, perhaps due to model complexity and the size of the datasets. We use discrete-time models in which the conditional variance of the normal innovation and the conditional jump intensity of a compound Poisson process are governed by GARCH-type dynamics. In these models, closed-form option valuation formulas are available. We empirically investigate the importance of jumps and time-varying jump intensities using an extensive panel data set of crude oil futures and option prices.

For futures-based estimation of these models, we rely on the particle filter and maximum likelihood estimation. For option-based estimation, we use an option-based filter motivated by Mazzoni [2015]. Filtering the normal component and the jump component in this case is relatively simple and extremely fast, even when the jump intensity is time-varying. It takes about three seconds to filter volatility from 392,380 option contracts using Matlab on a standard PC.

We investigate the economic importance of jumps and dynamic jump intensities in the crude oil market and compare the fit of models with jumps with that of a benchmark model without jumps. We find strong evidence of the presence of dynamic jump intensities and relatively infrequent jumps in the crude oil market. Our preferred jump model contains a time-varying jump intensity and a time-varying conditional variance driven by the same dynamics. This model is related to the most complex SVJ dynamics studied in the literature (Eraker [2004], Santa-Clara

and Yan [2010]). This model significantly outperforms the benchmark GARCH model for futures data, and allowing for jumps also leads to important gains in option fit. During crisis periods, when market risk is high, jumps occur more frequently. The properties of the jump processes estimated from options differ from those estimated from futures. The option-implied properties are more plausible, which may be due to improved parameter identification driven by the richer option data.

To the best of our knowledge, no existing studies have implemented jump models using extensive cross-sections of crude oil futures and options, but the literature on commodity derivatives contains several related studies. Trolle and Schwartz [2009] estimate a continuous-time stochastic volatility model using NYMEX crude oil futures and options and find evidence for two, predominantly unspanned, volatility factors. They do not consider jump processes. Hamilton and Wu [2014] model crude oil futures with an affine term structure model and document significant changes in oil futures risk premia since 2005. Pan [2012] uses options on crude oil futures to study the impact of speculation on returns.

Models for Commodity Futures Markets

This section specifies the benchmark model as well as the dynamic jump intensity model.

The Benchmark Model

We formulate a new jump model for commodities markets. To provide a benchmark for the model that can capture important stylized facts in commodity markets, we first consider a GARCH model for futures returns. In commodity futures markets, we observe futures prices for different maturities. We proceed by directly modeling this futures price. Appendix A shows how the specification of the spot price and the cost of carry is embedded in this specification in a no-arbitrage setup.[1] Under the physical measure, the benchmark model is given by

$$\log \frac{F_{t+1,T}}{F_{t,T}} = (\lambda - \frac{1}{2})h_{t+1} + z_{t+1} = (\lambda - \frac{1}{2})h_{t+1} + \sqrt{h_{t+1}}\varepsilon_{t+1}, \tag{1}$$

where $F_{t,T}$ is the time t price of the futures contract maturing at time T, ε_{t+1} is an innovation of a standard normal distribution, $N(0,1)$, h_{t+1} is the conditional variance known at time t, and λh_{t+1} is the risk premium associated with the normal innovation.

The conditional variance of the normal innovation, h_{t+1}, is governed by a GARCH$(1,1)$

process as in Heston and Nandi [2000]

$$h_{t+1} = \omega + bh_t + a(\varepsilon_t - c\sqrt{h_t})^2. \tag{2}$$

GARCH models provide a convenient framework to capture stylized facts in financial markets such as conditional heteroskedasticity, volatility clustering, and mean reversion in volatility. These stylized facts are also very prominent in commodity futures markets. The GARCH dynamic in equation (2) is different from the more conventional GARCH specifications of Engle [1982] and Bollerslev [1986], and is explicitly designed to facilitate option valuation. We discuss the benefits of the specification in equation (2) in more detail below.

Consistent with other GARCH specifications, the conditional variance h_{t+1} in equation (2) is predictable conditional on information available at time t. The unconditional variance is given by

$$\sigma^2 = E[h_{t+1}] = (\omega + a)/(1 - b - ac^2), \tag{3}$$

where $b + ac^2$ is the variance persistence. Furthermore, given a positive estimate for a, the sign of c determines the correlation between the futures returns and the conditional variance. Equivalently, c can be thought of as controlling the skewness or asymmetry of the distribution of log returns, with a positive c resulting in a negatively skewed multi-day distribution.

Commodity Futures Returns with Dynamic Jump Intensities

The futures return process in equations (1)-(2) provides a benchmark model that can capture several important stylized facts using a simple setup with a single normal innovation. Building on the models for stock returns in Maheu and McCurdy [2004] and index returns in Ornthanalai [2014] and Christoffersen, Jacobs, and Ornthanalai [2012] we now formulate a much richer model with jumps in returns and volatilities, and with potentially time-varying jump intensities. Omitting the maturity subscripts, futures returns are given by

$$\log \frac{F_{t+1,T}}{F_{t,T}} = (\lambda_z - \frac{1}{2})h_{z,t+1} + (\lambda_y - \xi)h_{y,t+1} + z_{t+1} + y_{t+1}, \tag{4}$$

where the z subscript refers to the normal component specified as in the GARCH model; and the y subscript refers to the jump component.

The jump component y_{t+1} is specified as a Compound Poisson process denoted as $J(h_{y,t+1}, \theta, \delta^2)$. The Compound Poisson structure assumes that the jump size is independently drawn from a normal distribution with mean θ and variance δ^2. The number of jumps n_{t+1} arriving between times

t and $t+1$ is a Poisson counting process with intensity $h_{y,t+1}$. The jump component at $t+1$ is therefore given by

$$y_{t+1} = \sum_{j=0}^{n_{t+1}} x_{t+1}^{j}, \tag{5}$$

where x_{t+1}^{j}, $j = 0, 1, 2, \ldots$ is an $i.i.d.$ sequence of normally distributed random variables, $x_{t+1}^{j} \sim N(\theta, \delta^2)$. The conditional expectation of the number of jumps arriving over time interval $(t, t+1)$ equals the jump intensity, $E_t[n_{t+1}] = h_{y,t+1}$. The conditional mean and variance of the jump component y_{t+1} are given by $\theta h_{y,t+1}$ and $(\theta^2 + \delta^2)h_{y,t+1}$ respectively.

In equation (4) the convexity adjustment terms $\frac{1}{2}h_{z,t+1}$ and

$$\xi h_{y,t+1} \equiv E_t\left[\exp\left(y_{t+1}\right)\right] = (\exp(\theta + \frac{\delta^2}{2}) - 1)h_{y,t+1} \tag{6}$$

act as compensators to the normal and jump components respectively, and they ensure that the conditional risk premium is given by

$$E_t\left[\exp\left(\log\frac{F_{t+1,T}}{F_{t,T}}\right)\right] = \lambda_z h_{z,t+1} + \lambda_y h_{y,t+1}, \tag{7}$$

with λ_z and λ_y denoting the market prices of risks for the normal and jump components.

Variance and Jump Intensity Dynamics

Christoffersen, Jacobs, and Ornthanalai [2012] investigate a richly parameterized model with separate dynamics for the diffusive volatility $h_{z,t+1}$ and jump intensity $h_{y,t+1}$. Our initial empirical investigation indicated that for our sample several of this model's parameters were poorly identified. We therefore investigate a more tightly parameterized specification, which impose restrictions on $h_{z,t+1}$ and $h_{y,t+1}$ and thereby reduce the dimension of the parameter space.

The general form for the dynamics of $h_{z,t+1}$ and $h_{y,t+1}$ are given by the extended GARCH (1, 1) processes:

$$h_{z,t+1} = \omega_z + b_z h_{z,t} + \frac{a_z}{h_{z,t}}(z_t - c_z h_{z,t})^2 + d_z y_t, \tag{8}$$

$$h_{y,t+1} = \omega_y + b_y h_{y,t} + \frac{a_y}{h_{z,t}}(z_t - c_y h_{z,t})^2 + d_y y_t, \tag{9}$$

where ω_z, b_z, a_z, c_z, d_z, ω_y, b_y, a_y, c_y, and d_y are parameters to be determined.

In our preferred specification we assume that $h_{z,t+1}$ follows the dynamics in (8) and $h_{y,t+1}$ is time-varying but driven by the same dynamic. The conditional jump intensity is a linear function

of the conditional variance of the normal component

$$h_{y,t+1} = \omega_y + kh_{z,t+1}, \tag{10}$$

where k is an additional parameter to be estimated.

Below, we will refer to the dynamic-intensity jump model defined by (4), (8), and (10) as the DI-Jump model. We have investigated a number of other parameter specifications within the general framework defined by (8) and (9) but found that the DI-Jump model offers the best trade-off between fit and parsimony for the data in this study.

Our model is quite tightly parameterized, yet it allows for jumps in volatility as well as jumps in returns. It has been shown in the index option literature that jumps in volatility are useful to explain option volatility smiles and smirks (see for example Eraker, Johannes and Polson [2003], and Eraker [2004]). Note however that following Ornthanalai [2014], our model is designed to yield a closed-form solution for option prices, and in order to do so we have adopted a rather simple specification for jumps in volatility.

It is important to note that the normal and jump innovations, z_t and y_t, enter separately into the GARCH updating dynamic in (8). The model therefore allows the two innovations to impact the variance and thus the jump intensity separately.

Note also that our specification in equation (8) allows for negative jumps to volatility. In recent work, Amengual and Xiu [2013] find negative volatility jumps to be an important feature of equity market volatility, often due to policy news. Negative volatility jumps appear to be an important feature of oil volatility as well. For example, on November 28, 2011, CBOE's oil volatility index (OVX) closed at 44.5% down from 69.1% at the end of the previous trading day. When allowing for negative jumps to volatility, an important concern is that volatility itself should not be negative. This can easily be handled in model implementation by restricting the parameters to effectively impose a small positive lower bound on volatility.

Crude Oil Futures and Options Data

We now discuss the crude oil futures and options data used in the empirical analysis, and present summary statistics.

We use a data set of Chicago Mercantile Exchange (CME group, formerly NYMEX) crude oil futures and options data. We use a sample of daily data from January 2nd, 1990 to May 30, 2014. The CME crude oil derivatives market is the world's largest and most liquid commodity derivatives market. The range of maturities covered by futures and options and the range of

option strike prices are also wider than for other commodities (for a discussion see Trolle and Schwartz [2009]).

We screen futures contracts based on patterns in trading activity. Open interest for futures contracts tends to peak approximately two weeks before expiration. Among futures and options with more than two weeks to expiration, the first six monthly contracts tend to be very liquid. For contracts with maturities of more than six months, trading activity is concentrated in the contracts expiring in March, June, September, and December.

Following Trolle and Schwartz [2009], we therefore screen the available futures and options data according to the following procedure: we discard all futures contracts with 14 or fewer days to expiration. Among the remaining, we retain the six shortest maturities. Furthermore, we choose the two shortest maturity contracts from those with expiration either in March, June, September or December. This procedure leaves us with eight futures contract series which we label M1, M2, M3, M4, M5, M6, Q1, and Q2.[2]

We include the following options on these eight futures contracts. For each option maturity, we consider eleven moneyness intervals: 0.78-0.82, 0.82-0.86, 0.86-0.90, 0.90-0.94, 0.94-0.98, 0.98-1.02, 1.02-1.06, 1.06-1.10, 1.10-1.14, 1.14-1.18, and 1.18-1.22. Moneyness is defined as option strike divided by the price of the underlying futures contract. Among the options within a given moneyness interval, we select the one that is closest to the mean of the interval.

Our data consists of American options on crude oil futures contracts. The CME has also introduced European-style crude oil options, which are easier to analyze. However, the trading history is much shorter and liquidity is much lower than for the American options. Since the pricing formulae are designed for European options, we have to convert the American option prices to European option prices. Assuming that the price of the underlying futures contract follows a geometric Brownian motion, we can accurately price American options using the Barone-Adesi and Whaley [1987] formula. Inverting this formula yields an implied volatility, from which we can subsequently obtain the European option price using Black's [1976] formula. To minimize the effect of errors in the early exercise approximation, we use only OTM and ATM options, i.e., puts with moneyness less than one and calls with moneyness greater than one. In addition, we only consider options that have open interest in excess of 100 contracts and prices higher than ten cents.

This data filtering procedure yields $49,008$ futures contracts and $392,380$ option contracts observed over $6,126$ business days. The number of futures contracts is eight on every day in the sample, while the number of option contracts is between 23 and 87. All futures and options prices used in the empirical work are settlement prices.[3]

The risk-free discount function for option valuation is obtained by fitting a Nelson and Siegel

[1987] curve each trading day to a LIBOR curve consisting of the 1, 2, 3, 6, 9 and 12 month LIBOR rates.

Panel A in Exhibit 1 shows dramatic swings in WTI crude oil prices over the last decade as documented in the literature (see, for example, Hamilton and Wu [2012]). Oil prices increase dramatically between 2003 and 2007, subsequently decline sharply, followed by another run-up toward the end of the sample. Some of the key geopolitical events in the sample are indicated in Panel A.

The M1 futures price is very highly correlated with the WTI spot price in Panel A, so we do not plot the M1 futures price. Panel E of Exhibit 1 displays the futures term structure captured by the Q2 less the M1 futures price. For the first part of the sample, the price of the long maturity futures contract (Q2) is typically lower than that of the short maturity contract (M1): the crude oil market tends to be in backwardation. This finding is consistent with existing studies such as Trolle and Schwartz [2009] and Litzenberger and Rabinowitz [1995]. For the second half of the sample, the price of the long maturity futures contract is typically higher than that of the short maturity futures contract: the crude oil market tends to be in contango. Panel A of Exhibit 2 reports the time series average of the futures prices for each maturity. Over the entire sample, futures prices are slightly hump-shaped as a function of maturity. The average price of the M3 futures contract is the highest among all maturities.

The two top panels of Exhibit 5 plots the daily returns, $log(F_{t+1,T}/F_{t,T})$, for the M1 and Q2 futures contracts, and Panel B of Exhibit 2 provides summary statistics for futures returns for all contracts. Futures returns on longer maturity futures contracts, e.g., Q2 futures contracts, are higher and less volatile than futures returns on shorter maturity contracts. On average across maturities, daily skewness is -0.70 and daily excess kurtosis is 10.33. The daily crude oil futures return series is thus skewed towards the left, indicating that there are more negative than positive outlying returns in the crude oil market. The skewness and kurtosis are smaller in magnitude for longer maturities. Panel C shows that futures trading volume is much larger for short than long maturities.

Panel D of Exhibit 2 lists the average number of option contracts in the sample across maturity and moneyness. Recall that each day, we select the option that is closest to the mean of the moneyness interval for each maturity. Among the 11 moneyness intervals, the number of option contracts is highest in the ATM interval. Panel E reports on average daily option trading volume. Volume is highest for the shortest maturities and for at-the-money contracts as well.

Panel C in Exhibit 1 displays the M1 ATM implied volatilities over time. Large spikes in the M1 option-implied volatility appear around the end of 1990 and beginning of 1991, at the time of the first Gulf War, around the September 2001 terrorist attack, the second Gulf

War in March 2003, and during the financial crisis in 2008. Exhibit 3 reports the IV averages across moneyness for each maturity. Options with short maturities have relatively high implied volatilities, consistent with the patterns in physical volatility documented in Panel B of Exhibit 2. Panel F in Exhibit 1 shows that throughout the sample, the M6-M1 ATM implied volatility term structure is most often either downward-sloping or flat.

Exhibit 3 shows that implied volatilities for the shortest maturities display a "smile" shape on average similar to that found in currency options, whereas for the longer maturities the shape is closer to the "smirk" commonly found in equity index options. These patterns suggest the relative importance of return skewness and kurtosis at different maturities for characterizing the option data, because excess kurtosis is sufficient to generate a smile but not a smirk. Panels G and H in Exhibit 1 plot the time-series of option-based measures of skewness and kurtosis for M1, computed using implied volatilities at different moneyness. They provide model-free evidence for the importance of allowing for dynamics in risk-neutral higher moments in the option-valuation model.

In summary, the data set displays substantial richness and variation, which we attempt to capture with the proposed class of models.

Evidence from Futures Prices

We first discuss maximum likelihood estimation of the model using futures returns. We then present parameter estimates for the DI-Jump model as well as for the benchmark GARCH model. Subsequently we use the parameter estimates to investigate the models' most important implications for futures valuation.

Maximum Likelihood Estimation using Futures Data

We estimate the model parameters using an approximate maximum likelihood procedure constructed from conditional variances and jump intensities that are filtered using the particle filter. The likelihood function for returns depends on the normal and Compound Poisson distributions. The conditional density of the returns process in equation (4) with time-to-maturity T_i, given that there are $n_{t+1} = j$ jumps occurring between period t and $t+1$, is given by

$$f_t\left(R_{t+1,i}|n_{t+1}=j\right) = \frac{1}{\sqrt{2\pi(\widetilde{h}_{z,t+1}+j\delta^2)}} \exp\left(-\frac{\left(R_{t+1,i}-\widetilde{\mu}_{t+1}-j\theta\right)^2}{2(\widetilde{h}_{z,t+1}+j\delta^2)}\right), \tag{11}$$

where $R_{t+1,i} = \log \frac{F_{t+1,T_i}}{F_{t,T_i}}$, and $\widetilde{\mu}_{t+1} = (\lambda_z - \frac{1}{2})\widetilde{h}_{z,t+1} + (\lambda_y - \xi)\widetilde{h}_{y,t+1}$ denotes the expected return using filtered values for the conditional variance and jump intensity.

The conditional probability density of returns can be derived by summing over the number of jumps

$$f_t\left(R_{t+1,i}\right) = \sum_{j=0}^{\infty} f_t\left(R_{t+1,i} | n_{t+1} = j\right) \Pr(n_{t+1} = j), \tag{12}$$

where $\Pr(n_{t+1} = j) = (\widetilde{h}_{y,t+1})^j \exp(-\widetilde{h}_{y,t+1})/j!$ is the probability of having j jumps, which is distributed as a Poisson counting process.

We can write the log likelihood function as the sum of the log likelihoods for all eight futures contracts

$$L_{Fut}\left(\Theta\right) = \frac{1}{8}\sum_{i=1}^{8}\sum_{t=1}^{T-1} \log(f_t\left(R_{t+1,i}\right)), \tag{13}$$

where Θ denotes the vector of parameters in the model.

When implementing maximum likelihood estimation, the summation in equation (12) must be truncated. We truncate the summation at 50 jumps per day, which is large enough to ensure that it will not be binding for the empirical results which turn out to display at most 3 jumps per day.

Equations (8) and (9) show that we need to separately identify the two unobserved shocks z_{t+1} and y_{t+1}. The nonlinear structure of the model renders the particle filter appropriate for this task. In the particle filter multiple price paths (particles) are sampled and resampled from the model and its likelihood, and the paths are averaged using probability weights to generate expected values for the objects of interest. Using the particle filter, calculating the expected (filtered) values $\widetilde{z}_{t+1}$ and $\widetilde{y}_{t+1}$ is thus relatively straightforward. As the number of particles goes to infinity the particle filtered values converge to their true values. We report results based on 10,000 particles. Experimentation with a higher number of particles indicated that our results are robust. We refer to Ornthanalai [2014] for a detailed discussion.

Finally, we impose variance targeting (Engle and Mezrich [1996], Francq, Horvath, and Zakoian [2011]). Existing studies demonstrate that variance targeting helps with parameter identification, which greatly facilitates the search for optimal parameter values and ensures that model properties closely match the data along critical dimensions. For example, for the GARCH model, instead of estimating a, we infer it from the historical sample variance of futures return and other parameter estimates, according to equation (3). Variance targeting for the jump model proceeds along the same lines, where we infer the value of the a_z parameter using the data and the model

moments. The unconditional variance for the DI-Jump model is given by

$$\sigma^2 = (\theta^2 + \delta^2)\omega_y + [1 + (\theta^2 + \delta^2)k]\frac{\omega_z + a_z + d_z\theta\omega_y}{1 - b_z - a_z c_z^2 - d_z\theta k}. \tag{14}$$

Estimation Results

Panel A of Exhibit 4 presents the return-based maximum likelihood parameter estimates for the GARCH benchmark model and the DI-Jump model. The results are obtained using all eight futures contracts jointly in estimation for the time period 1990-2014. For the DI-Jump model, we separate the parameters into two columns. The parameters with subscript y are reported in the column labeled "Jump". The parameters with subscript z are reported in the column labeled "Normal". We imposed a zero intercept in the variance and jump intensity dynamics, because this has a negligible impact on model fit and it improves parameter identification. We therefore have $\omega = 0$ in the GARCH model, and $\omega_z = 0$ in the DI-Jump model. Because it is difficult to estimate reliably using only futures returns, we also set $\lambda_y = 0$ in the DI-Jump model in Panel A. Below each parameter estimate, we report its standard error calculated using the Hessian matrix. Under "Properties", we report the expected number of jumps per year implied by the parameter estimates, the implied long-run risk premiums for the normal and jump components, the percent of total variance captured by the normal and the jump component, the persistence of the normal and jump components, the average annual volatility, and the log-likelihood.

The log-likelihood value for the DI-Jump model is larger than that of the GARCH model. To examine whether the jump model significantly improves over the GARCH model, we test the null hypothesis of no jumps. A likelihood ratio test of the null hypothesis of no jumps does not have the usual limiting chi-square distribution because the jump parameters are unidentified under the null. To implement this test, we use the standardized likelihood ratio test proposed by Hansen [1992, 1994]. Hansen's test is able to provide an upper bound to the asymptotic distribution of standardized likelihood ratio statistics, even when conventional regularity conditions (for example, due to unidentified parameters) are violated. Using Hansen's standardized LR test, we find that the jump model significantly improves on the GARCH model, suggesting that the null hypothesis of no jumps is rejected. Hansen's standardized LR test statistic for the jump model is 4.66, which is significantly larger than the simulated 1% critical value, 3.07.

We also report the persistence of the conditional variance and jump intensities, which are identical in the DI-Jump specification. The persistence of the processes is rather high and similar in magnitude in both models.

The average number of jumps in the DI-Jump model is 1.37 jumps per year. Most existing

estimates in equity index markets find less than four jumps per year (see for example Bates [2006], and Andersen, Benzoni, and Lund [2002]). The estimate of the average jump size θ for the DI-Jump model is -0.0485. The estimate of k is statistically significant, suggesting the dependence of the jump arrival rate on the level of volatility. Overall, the results for the DI-Jump model suggest that allowing for time-varying jump intensities can greatly improve model fit beyond the standard GARCH model.

In the DI-Jump model, the total unconditional return variance, σ^2, is the sum of normal variance and jump variance

$$\sigma^2 = \sigma_z^2 + (\theta^2 + \delta^2)\sigma_y^2,$$

where σ_z^2 and σ_y^2 are computed as the time series averages of $h_{z,t+1}$ and $h_{y,t+1}$. The fraction of the total return variance due to jumps is substantial at 48.77%. While the likelihood ratio tests indicate that allowing for time-varying jump intensities is supported statistically, these findings suggest that this model feature is also economically important.

The variation in jump intensities affects the risk premiums, $\lambda_z h_{z,t+1} + \lambda_y h_{y,t+1}$. Under "Properties" in Exhibit 4 we report the averages of the normal and jump components of risk premiums over the sample. The GARCH and DI-Jump risk premia are similar at around 8%, perhaps due to the fact that we have set $\lambda_y = 0$ in Panel A.

To further evaluate model specification and performance, we construct QQ plot for futures returns standardized by estimates of the unconditional variance, the GARCH conditional variance, and the conditional variance from the DI-Jump model. Exhibit 5 reports the results for M1 (left column) and Q2 (right column) futures returns. QQ plots for the other maturities yield very similar results.

When standardizing returns by their unconditional variance, QQ plots for M1 and Q2 significantly deviate from the 45-degree line, especially for quantiles far from zero, indicating that the empirical distribution significantly deviates from the standard normal distribution. When standardizing instead by the GARCH variance the QQ plot is somewhat closer to the 45-degree line. Finally, when standardizing by $h_{z,t}$ in the DI-Jump model, the distribution is closer yet to the standard normal. Exhibit 5 suggests that including jumps helps capture the crude oil futures return distribution.

There is some concern in the existing literature regarding the role of roll-over days of futures contracts. We have investigated this issue and found that roll-over days appear to be no different from other days in the QQ plots, indicating that this issue does not impact the empirical results for our sample.

In summary, we find that the new DI-Jump model with dynamic variance and jump-intensity

does a good job of fitting observed oil futures volatility and tail return behavior. We therefore now proceed to deriving the option pricing implications of the model.

Option Valuation Theory for Crude Oil Futures

We first characterize the risk-neutral dynamics. Subsequently we derive the closed-form option valuation formula.

The Equivalent Martingale Measure and the Risk-Neutral Dynamics

The estimates obtained from futures prices in above are physical parameters. To value crude oil options, we need return dynamics under the equivalent martingale or risk-neutral measure. In a framework with compound Poisson processes, the futures price can jump to an infinite set of values in a single period, and the equivalent martingale measure is therefore not unique. We assume a pricing kernel that is commonly used in the affine literature and proceed by specifying the conditional Radon-Nikodym derivative:

$$\frac{\frac{dQ_{t+1}}{dP_{t+1}}}{\frac{dQ_t}{dP_t}} = \frac{\exp(\Lambda_z z_{t+1} + \Lambda_y y_{t+1})}{E_t[\exp(\Lambda_z z_{t+1} + \Lambda_y y_{t+1})]}, \tag{15}$$

where Λ_z and Λ_y are the equivalent martingale measure (EMM) coefficients that capture the wedge between the physical and the risk-neutral measure. As in Ornthanalai [2014], this Radon-Nikodym derivative specifies a risk-neutral probability measure if and only if Λ_z and Λ_y are determined by

$$\Lambda_z + \lambda_z = 0, \tag{16}$$

so that $\Lambda_z = -\lambda_z$, and

$$\lambda_y - (\exp(\theta + \frac{\delta^2}{2}) - 1) - \exp(\Lambda_y \theta + \frac{\Lambda_y^2 \delta^2}{2})(1 - \exp((\Lambda_y + 0.5)\delta^2 + \theta)) = 0. \tag{17}$$

The solution for Λ_y is not analytical but it is well behaved and can reliably and efficiently be computed using a numerical approach.

The risk-neutral dynamics take the same functional form as the physical dynamics. Denote the normal and jump components under the risk neutral measure by $z^*_{t+1} \sim N(-\lambda_z h_{z,t+1}, h_{z,t+1})$ and $y^*_{t+1} \sim J(h^*_{y,t+1}, \theta^*, \delta^2)$. The change of measure shifts the mean of the normal component, $z^*_{t+1} = z_{t+1} - \lambda_z h_{z,t+1}$. The futures return process under the risk-neutral dynamic can then be

written as

$$\log \frac{F_{t+1,T}}{F_{t,T}} = -\frac{1}{2}h_{z,t+1} - \xi^* h^*_{y,t+1} + z_{t+1} + y^*_{t+1}. \tag{18}$$

The general specifications of the variance dynamics and jump intensity dynamics under the risk-neutral measure are given by

$$h_{z,t+1} = \omega_z + b_z h_{z,t} + \frac{a_z}{h_{z,t}}(z_t - c^*_z h_{z,t})^2 + d_z y^*_t, \tag{19}$$

$$h^*_{y,t+1} = \omega^*_y + b_y h^*_{y,t} + \frac{a^*_y}{h_{z,t}}(z_t - c^*_y h_{z,t})^2 + d^*_y y^*_t, \tag{20}$$

where $h^*_{y,t+1} = h_{y,t+1}\Pi$, $\Pi = \exp(\Lambda_y \theta + \frac{\Lambda_y^2 \delta^2}{2})$, $\theta^* = \theta + \Lambda_y \delta^2$, $\xi^* = \exp(\theta^* + \frac{\delta^2}{2}) - 1$, $\omega^*_y = \omega_y \Pi$, $a^*_y = a_y \Pi$, $c^*_z = c_z - \Lambda_z$, $c^*_y = c_y - \Lambda_z$, and $d^*_y = d_y \Pi$.

The DI-Jump model uses the variance dynamic in (19). For the jump intensity it is a special case of the general specification, with the following restrictions

$$\omega_z = 0,\ b_y = b_z,\ a^*_y = a_z k^*,\ c^*_y = \ c^*_z,\ d^*_y = k^* d_z. \tag{21}$$

where $k^* = k\Pi$.

The risk neutral dynamic for the GARCH benchmark model is a special case of (19)-(20) with $h^*_{y,t+1} = 0$.

Closed-Form Option Valuation

Under the risk-neutral measure, the generating function for the asset process in (18)-(19)-(20) takes the following form

$$f^*(\varphi; t, T) \equiv E^Q_t[F^\varphi_T] = F^\varphi_{t,T} \exp(A(\varphi; t, T) + B(\varphi; t, T) h_{z,t+1} + C(\varphi; t, T) h^*_{y,t+1}). \tag{22}$$

The derivation of the affine coefficients $A(\varphi; t, T)$, $B(\varphi; t, T)$, and $C(\varphi; t, T)$ is provided in Appendix B.

Using the risk neutral generating function, we can value European options using the Fourier inversion method as in Heston [1993], Heston and Nandi [2000], and Duffie, Pan and Singleton

[2000]. The price of a European call option on a futures contract is given by

$$CO(t,T,K) = E_t^Q[\exp(-\int_t^T r(s)ds)(F(T,T)-K)^+] = \tag{23}$$

$$F(t,T)(\frac{1}{2}+\frac{1}{\pi}\int_0^\infty \mathrm{Re}[\frac{K^{-i\varphi}f^*(i\varphi+1)}{i\varphi f^*(1)}]d\varphi) - \exp(-\int_t^T r(s)ds)K(\frac{1}{2}+\frac{1}{\pi}\int_0^\infty \mathrm{Re}[\frac{K^{-i\varphi}f^*(i\varphi)}{i\varphi}]d\varphi),$$

where $CO(t,T,K)$ is the time t price of a European call option on a futures contract expiring at time T,[4] and where K is the strike price of the option.

Joint Estimation Using Futures and Options Data

It is possible to use the parameter estimates in Panel A of Exhibit 4, obtained using MLE on futures data, to compute option prices using the option valuation formulae. However, this procedure exclusively uses historical information and ignores expectations about the future evolution of the futures prices that are embedded in option prices. It is therefore also interesting to study the models' option valuation performance by specifying an objective function based on option contracts, and matching model option values as closely as possible to observed market prices.

While such an exercise imposes considerable discipline upon the models, it has nevertheless two important drawbacks. First, if a model is richly parameterized, only fitting the option data may result in over-fitting. Second, the price of risk parameters, which are some of the most economically important model parameters, cannot be reliably identified using option data only.

We therefore follow Bates [1996], who suggests that the most stringent test of an option pricing model lies in its ability to jointly fit the option data and the underlying returns. In our case, this means that we have to construct an objective function that contains a crude oil option component as well as a futures return component. We first discuss the option-based filtering and the likelihood function used to estimate the model parameters from option data. We then explain how to combine the option data with the underlying futures data and construct a joint likelihood. Subsequently we discuss the parameter estimates, and compare the most important model properties with the properties implied by the physical parameters from Panel A of Exhibit 4.

Option-Based Filtering

To obtain the fitted option prices, we first need to filter the conditional variance and jump intensity. In order to give the model the best possible chance of fitting the data and at the same time improving the estimation speed, we use an option-based filter. This allows for the information in the option market at the end of date t to get incorporated into the conditional variance, $h_{z,t+1}$, and the conditional jump intensity, $h_{y,t+1}$, in a way that is consistent with the underlying model dynamics.

Recent studies such as Broadie, Chernov, and Johannes [2007] implement option-based filtering by extracting conditional variance from the implied volatility of option data. Mazzoni [2015] applies a Brownian bridge process to parameterize the local volatility surface based on expressing local variance as the conditional expectation of instantaneous variance. We apply this approach to the models in this paper.

First consider the GARCH model. Denote the current date as $t = 0$. We are interested in obtaining the next day conditional variance h_1, given all available information at time t. On each trading day for $0 \leqslant t \leqslant T-1$, the Brownian bridge process imposes structures by forcing the process $x_{t+1} \equiv \log \frac{F_{t+1,T}}{F_{0,T}}$ to end up at a particular value $k' = \log \frac{F_{T,T}}{F_{0,T}} = \log \frac{K}{F_{0,T}}$ at time T. Using the relationship between local volatility and the conditional expectation of the variance, as well as the relationship between local variance and implied volatility, the implied volatility can be approximated as

$$\widehat{\sigma}_{imp}(k', T) \approx \sqrt{\frac{\alpha(k', T)}{1-\beta} + \frac{1}{T}(h_1 - \frac{\alpha(k', T)}{1-\beta})\frac{1-\beta^T}{1-\beta}}, \tag{24}$$

where $\alpha(k', T) = 2a\frac{k'}{T}(\frac{1}{2} - c^*) + \frac{a}{2}$, $\beta = b + a(\frac{1}{2} - c^*)^2$. Appendix C provides the details. From equation (24), we can calibrate the next day's conditional variance, h_1, by minimizing the implied volatility root mean squared error (IV RMSE) each day.

For the DI-Jump model, we proceed in a similar way. Since futures returns have both normal and jump components in the futures return innovation, the bridge process is a more general one than the Brownian bridge for the GARCH model. The implied volatility as a function of $h_{z,1}$ for the DI-Jump model is given by

$$\widehat{\sigma}_{imp|Jump}(k', T) \approx \sqrt{\frac{\alpha_J}{1-\beta_J} + \frac{1}{T}(h_{z,1} - \frac{\alpha_J}{1-\beta_J})\frac{1-\beta_J^T}{1-\beta_J}}, \tag{25}$$

where α_J and β_J depend on the specification of the DI-Jump model and are derived in Appendix

C. Similar to the case of the GARCH model, using (25) one can obtain the next day conditional variance, $h_{z,1}$, by minimizing the IV RMSE.

This option based filtering technique not only improves model fit, it also expedites the estimation process by providing an analytical approximation for filtering the conditional variance and conditional jump intensity. The improvement in implementing the model using the option based filter rather than the particle filter is dramatic, especially for large data sets. In our application, it takes about 220 seconds to filter volatility on $49,008$ futures contracts using a particle filter with $10,000$ particles, while it takes only about 3 seconds to filter volatility on $392,380$ option contracts using the option based filter.

The Likelihood Function from Option Data

We use an objective function based on implied volatilities. This approach uses market data that is of similar magnitude along the moneyness, maturity, and time-series dimensions, which is attractive from a statistical perspective. Option prices differ significantly along these dimensions. Define the model error

$$u_{k,t} = \sigma_{k,t} - \widehat{\sigma}_{k,t}(O_{k,t}(\Theta)), \tag{26}$$

where $\sigma_{k,t}$ is the Black [1976] implied volatility of the kth observed option price at time t, and $\widehat{\sigma}_{k,t}$ is the implied volatility converted from each model option price, $O_{k,t}(\Theta)$, using Black's formula.

Assuming normality of the implied volatility errors, $u_{k,t} \sim N(0, \sigma_u^2)$, the log-likelihood function based on options is

$$L_{Opt}(\Theta) = -\frac{N}{2}\ln(2\pi\sigma_u^2) - \frac{1}{2}\sum_{t,k}^{N}\frac{u_{k,t}^2}{\sigma_u^2}, \tag{27}$$

where $N = 392,380$ is the total number of option contracts. Risk-neutral parameters can be estimated from options data only by maximizing L_{Opt}. We do not proceed this way, instead we estimate the model using the joint log likelihood.

The Joint Log Likelihood Function

The log-likelihood functions for futures and options are defined in equations (13) and (27) respectively. We maximize the weighted average of the log-likelihoods of futures and options to obtain parameter estimates for the GARCH and DI-Jump models that are determined by the option data as well as the underlying futures data.

The number of option contracts in the data set is much larger than the number of futures contracts. To ensure that joint parameter estimates are not dominated by options, we use a

contract-weighted log-likelihood. The resulting joint log-likelihood is

$$L_{Joint}(\Theta) = \frac{M+N}{2}\frac{L_{Fut}(\Theta)}{M} + \frac{M+N}{2}\frac{L_{Opt}(\Theta)}{N}, \tag{28}$$

where $M = 49,008$ is the total number of futures contracts and $N = 392,380$ is the total number of option contracts.

We report the optimized joint likelihoods as well as the corresponding IV RMSE

$$IV\ RMSE = \frac{1}{T}\sqrt{\frac{1}{N_t}\sum_{t,k}^{N_t} u_{k,t}^2}. \tag{29}$$

Ideally one would fit the model directly to implied Black volatilities. However, since the optimization routine requires computing implied volatility from model prices at every function evaluation, this approach is quite slow. We follow Trolle and Schwartz [2009] and fit the model to option prices scaled by their Black [1976] vega, that is, the sensitivities of the option prices to variations in implied volatilities. This approach is motivated by the approximation $\sigma_{k,t} \approx O_{k,t}/\nu_{k,t}$, where $\nu_{k,t}$ is the Black [1976] vega associated with the kth observed option price at time t. This approximation has been shown to work well in existing work. Due to the use of the option-based filter, the quasi-closed form option valuation formula, and the use of the vega-scaled prices, the optimization problem is feasible on our large data set. Finally, we again use variance targeting in this estimation exercise.

Estimates and Model Implications

Panel B of Exhibit 4 reports the parameter estimates obtained by maximizing the joint log-likelihood in (28) for the GARCH model and the DI-Jump model. At the bottom of the table, we report the log-likelihood, the IV bias and RMSE, and several other model properties implied by the parameters such as the expected number of jumps per year, the long-run risk premium, the percent of total annual variance explained by the normal and the jump component, the persistence, and the average annual volatility. The estimates reported in Exhibit 4 are physical parameters. For option valuation, we risk-neutralize these parameters using the transformation provided above. We impose positivity of the state variables in implementation due to the possibility of negative jumps in our specification.

The DI-Jump model in Panel B of Exhibit 4 outperforms the benchmark GARCH model. The DI-Jump model has an IV RMSE of 3.03 compared with that of 3.79 for GARCH which is an improvement of around 20%. This confirms that incorporating jumps in addition to dynamic

volatilities improves model fit.

The persistence of the conditional variance and conditional jump intensity estimated using both futures and options are much higher than that estimated using futures only. Using option information in model estimation reveals that oil futures volatility is predictable at much longer horizons than suggested by the evidence from return-based estimation.

Panels A and B in Exhibit 6 report the conditional variance and jump intensity paths for the DI-Jump model based on the estimates from futures and futures options in Panel B of Exhibit 4. In the crisis period around the first Gulf war, the conditional variance paths contain pronounced spikes. Note that when using options in estimation, the average number of jumps per year is significantly higher. It is 4.32 in Panel B compared to 1.37 in Panel A of Exhibit 4.

Since option prices contain important information about the pricing kernel that cannot necessarily be inferred from the underlying futures returns dynamics, the market prices of risk for the normal and jump components and the implied risk premiums in Panel B of Exhibit 4 are of particular interest. The jump risk premium is relatively modest in Panel B and the share of total variation accounted for by jumps turns out to be quite small as well.

Panels C-E in Exhibit 6 plot the resulting time variation in the conditional normal risk premium, the conditional jump risk premium, and the total risk premium. Overall, Exhibit 6 indicates that the normal component is economically most important.

Panels F-H in Exhibit 6 plot the sample paths of the number of jumps, as well as that of the filtered jump component and normal component. We find evidence of multiple jumps per day, especially during the first Gulf War in late 1991. Although normal components dominate return innovations most of the time, the jump component explains a substantial amount of the variation in returns in crisis periods.

Skewness and kurtosis are determined by the jump parameters as follows

$$Skew_t(R_{t+1}) = \frac{\theta(\theta^2 + 3\delta^2)h_{y,t+1}}{(h_{z,t+1} + (\theta^2 + \delta^2)h_{y,t+1})^{\frac{3}{2}}}, \tag{30}$$

$$Kurt_t(R_{t+1}) = 3 + \frac{(\theta^4 + 6\theta^2\delta^2 + 3\delta^4)h_{y,t+1}}{(h_{z,t+1} + (\theta^2 + \delta^2)h_{y,t+1})^2}. \tag{31}$$

Panels I and J in Exhibit 6 plots the time series of skewness and excess kurtosis for the DI-Jump model. Note that these measures of skewness and kurtosis are annualized, consistent with the variances in Exhibit 6.

Exhibit 7 investigates the differences in fit between the models. We report IV RMSEs (left column) and IV bias (right column) by moneyness for four of our eight maturity categories. For almost all moneyness and maturity levels in Exhibit 7, the average IV RMSE is considerably

lower for the DI-Jump model than the GARCH model. Almost everywhere, the magnitude of the IV bias is substantially lower in the DI-Jump model than in the GARCH model as well. The outperformance of the DI-Jump model is particularly evident at longer maturities. The IV bias differences between the two models are typically larger than the IV RMSE differences. Since RMSEs reflect model bias and variance, we conclude that the data may be rather noisy, and that some contracts are poorly fit by both models. The much improved bias for the DI-Jump model is therefore very important.

Conclusion

We estimate a new discrete-time jump model for CME crude oil futures and options on futures. The proposed model allows for a normal innovation with dynamic volatility and a jump component with dynamic intensity. The new model is tractable and allow for quasi-analytical option valuation.

We find evidence for the presence of jumps in the crude oil derivatives market, using futures data as well as options data. Both the analysis of futures data and the analysis based on the joint estimation of futures and options suggests that jumps are relatively rare events in the crude oil market, but the option data point to relatively more–but slightly smaller–jumps than the futures data.

Our jump model with dynamic jump intensity dramatically improves model performance compared with the benchmark GARCH model. This is the case whether or not options are used in estimation. During crisis periods, when market risk is high, jumps occur more frequently.

We also find that the joint estimation using both crude oil futures and options provides more plausible results than the estimation using futures data only. This suggests that option data are needed to reliably identify the models, and dynamic jump model in particular. The primary benefit of modeling jumps in crude oil markets seems to be that the excess kurtosis of the distribution is modeled better, whereas the modeling of skewness is a second-order effect.

Several extensions of our results are possible. The main benefit of combining futures and options data in estimation is the identification of jump intensity parameters and risk premiums associated with both the jump and diffusive components, but of course the results depend on model specification. Recently several studies have argued that risk premiums in oil markets have experienced a structural break around 2005 due to the increased presence of speculators (see for instance Acharya, Lochstoer, and Ramadorai [2013], Hamilton and Wu [2014], and Singleton [2013]). It may prove interesting to characterize the evolution in risk premiums with option data using variations on the simple risk-premium specification used in this study.

Notes

[1]See also Trolle and Schwartz [2009] for the specification of the cost of carry and the spot price. See Casassus and Collin-Dufresne [2005] for an analysis of the most general specification for convenience yields allowed within an affine structure.

[2]Futures contracts expire on the third business day prior to the 25th calendar day (or the previous business day if the 25th is not a business day) of the month that precedes the delivery month.

[3]The CME light sweet crude oil futures contract trades in units of 1000 barrels. Prices are quoted in US dollars per barrel.

[4]Crude oil futures options expire three business days before the termination of trading in the underlying futures contract. For notional simplicity, we use the same T for crude oil futures and options on the corresponding futures.

Appendix

A. Futures Specification

In our empirical work, we use a discrete-time implementation of the following continuous-time economy. Let $S(t)$ be the time t spot price. The spot price process under the risk-neutral measure is given by

$$\frac{dS(t)}{S(t)} = \delta(t)dt + \sqrt{v(t)}d\omega_t^Q, \tag{A.1}$$

where $v(t)$ is the variance and $\delta(t)$ is the instantaneous spot cost of carry.

Let $y(t,T)$ be the time t forward cost of carry to time T, with $y(t,t) = \delta(t)$. The literature documents a large negative correlation between shocks to spot returns and the cost of carry. For example, Trolle and Schwartz [2009] estimate a correlation of approximately -0.89. Schwartz [1997] models the convenience yield instead of the cost of carry, which switches the sign of the correlation, and obtains a correlation of approximately 0.8. In order to simplify the model and facilitate identification, we assume the forward cost of carry and the spot price have a correlation of -1. The stochastic process of $y(t,T)$ is then given by

$$dy(t,T) = \mu_y(t,T)dt - \sigma_y\sqrt{v(t)}d\omega_t^Q. \tag{A.2}$$

Trolle and Schwartz [2009] observe that long-term forward cost of carry rates are usually less volatile than short-term forward cost of carry rates. We assume the volatility of the forward cost

of carry, σ_y, satisfies:

$$\sigma_y = \exp(-\sigma\tau), \tag{A.3}$$

where $\tau = T - t$ measures time to maturity and σ is a parameter to be estimated. Let $F(t,T)$ denote the time t price of a futures contract maturing at time T. By definition,

$$F(t,T) \equiv S(t)\exp\{\int_t^T y(t,u)du\}. \tag{A.4}$$

Introduce $Y(t,T) = \int_t^T y(t,u)du$. With the cost of carry dynamics (A.2), we have

$$dY(t,T) = \{-\delta(t) + \int_t^T \mu_y(t,u)du\}dt + \frac{1-\exp(-\sigma\tau)}{\sigma}\sqrt{v(t)}d\omega_t^Q. \tag{A.5}$$

Apply Ito's lemma to (A.4),

$$\begin{aligned}
\frac{dF(t,T)}{F(t,T)} &= \frac{dS(t)}{S(t)} + dY(t,T) + \frac{1}{2}(dY(t,T))^2 + \frac{dS(t)}{S(t)}dY(t,T) \\
&= \{\int_t^T \mu_y(t,u)du + v(t)(\frac{1}{2}(\frac{\exp(-\sigma\tau)-1}{\sigma})^2 \\
&\quad + \exp(-\sigma\tau)\tau)\}dt + \sqrt{v(t)}(1 + \frac{\exp(-\sigma\tau)-1}{\sigma})d\omega_t^Q.
\end{aligned} \tag{A.6}$$

In the absence of arbitrage opportunities, the futures price process must be a martingale under the risk-neutral measure, see, e.g. Duffie [2001]. Setting the drift of (A.6) to zero, we have

$$\frac{dF(t,T)}{F(t,T)} = \sqrt{v(t)}(1 + \frac{\exp(-\sigma\tau)-1}{\sigma})d\omega_t^Q. \tag{A.7}$$

Apply Ito's lemma to $\log F(t,T)$,

$$\begin{aligned}
d\log F(t,T) &= \frac{dF(t,T)}{F(t,T)} - \frac{1}{2}(\frac{dF(t,T)}{F(t,T)})^2 \\
&= -\frac{1}{2}v(t)(1 + \frac{\exp(-\sigma\tau)-1}{\sigma})^2 dt + \sqrt{v(t)}(1 + \frac{\exp(-\sigma\tau)-1}{\sigma})d\omega_t^Q.
\end{aligned} \tag{A.8}$$

The futures dynamics under the physical measure, P, is achieved by specifying the market price of risk, $\Lambda(t)$, that links the Wiener process under P and Q measures:

$$d\omega_t^Q = \Lambda(t)dt + d\omega_t^P. \tag{A.9}$$

Duffee [2002] specifies a flexible essentially affine specification for the price of risk; Casassus and Collin-Dufresne [2005] allows risk premium to be an affine function of the state variables. We use an affine specification for $\Lambda(t)$:

$$\Lambda(t) = \lambda\sqrt{v(t)}(1 + \frac{\exp(-\sigma\tau) - 1}{\sigma}). \tag{A.10}$$

Substituting (A.9) and (A.10) into (A.8), we get the futures dynamics under the physical measure

$$d\log F(t,T) = [(\lambda - \frac{1}{2})v(t)(1 + \frac{\exp(-\sigma\tau) - 1}{\sigma})^2]dt + \sqrt{v(t)}(1 + \frac{\exp(-\sigma\tau) - 1}{\sigma})d\omega_t^P. \tag{A.11}$$

In this specification the parameter σ captures the term structure of futures volatility. In our empirical implementation, we found that within a reasonable range of σ, this parameter did not improve model fit substantially. Note that when σ is large, $\frac{\exp(-\sigma\tau)-1}{\sigma}$ approaches zero. The resulting simplification of equation (A.11) yields the discrete-time model in equation (1).

B. The Generating Function and the Option Valuation Formula

We solve for the coefficients $A(\varphi; t, T)$, $B(\varphi; t, T)$, and $C(\varphi; t, T)$ in equation (22) as in Ingersoll [1987] and Heston and Nandi [2000], using the fact that the conditional moment generating function is exponential affine in the state variables $h_{z,t+1}$ and $h^*_{y,t+1}$.

Since F_T is known at time T, equation (22) requires the terminal condition

$$A(\varphi; T, T) = B(\varphi; T, T) = C(\varphi; T, T) = 0. \tag{B.1}$$

Applying the law of iterated expectations to $f^*(\varphi; t, T)$, we get

$$\begin{aligned} f^*(\varphi; t, T) &= E_t^Q[f^*(\varphi; t+1, T)] \qquad (B.2) \\ &= F_t^\varphi E_t^Q[\exp(\varphi F_{t+1} + A(\varphi; t+1, T) + B(\varphi; t+1, T)h_{z,t+2} + C(\varphi; t+1, T)h^*_{y,t+2})]. \end{aligned}$$

We can rewrite the futures return process in (18) as

$$R_{t+1} = \mu_z h_{z,t+1} + \mu_y h^*_{y,t+1} + z_{t+1} + y^*_{t+1}, \tag{B.3}$$

where $\mu_z = -\frac{1}{2}, \mu_y = -\xi^*$.

Substituting the futures return process (B.3), the conditional normal variance dynamic (19),

and the conditional jump intensity dynamic (20) into (B.2), we get

$$\begin{aligned} f^*(\varphi;t,T) &= F_t^{\varphi} E_t^Q[\exp(\varphi(\mu_z h_{z,t+1} + \mu_y h^*_{y,t+1} + z_{t+1} + y^*_{t+1}) + A_2(\varphi;t+1,T) \\ &\quad + B_2(\varphi;t+1,T)(\omega_z + b_z h_{z,t+1} + \frac{a_z}{h_{z,t+1}}(z_{t+1} - c^*_z h_{z,t+1})^2 + d_z y^*_{t+1}) \\ &\quad + C_2(\varphi;t+1,T)k^*(\omega_z + b_z h_{z,t+1} + \frac{a_z}{h_{z,t+1}}(z_{t+1} - c^*_z h_{z,t+1})^2 + d_z y^*_{t+1}))]. \end{aligned} \tag{B.4}$$

After rearranging terms and completing squares we get

$$\begin{aligned} f^*(\varphi;t,T) &= F_t^{\varphi} E_t^Q[\exp(A(\varphi;t+1,T) + B(\varphi;t+1,T)\omega_z + C(\varphi;t+1,T)\omega^*_y \\ &\quad + (\varphi\mu_z + (b_z + a_z c_z^{*2})B(\varphi;t+1,T))h_{z,t+1} + (\varphi\mu_y + (b_z + a_z c_z^{*2})C(\varphi;t+1,T))h^*_{y,t+1} \\ &\quad + (a_z B(\varphi;t+1,T) + a^*_y C(\varphi;t+1,T))\frac{z^2_{t+1}}{h_{z,t+1}} \\ &\quad + (\varphi - 2a_z c^*_z B(\varphi;t+1,T) - 2a^*_y c^*_y C(\varphi;t+1,T))z_{t+1} \\ &\quad + (\varphi + d_z B(\varphi;t+1,T) + d^*_y C(\varphi;t+1,T))y^*_{t+1})], \end{aligned} \tag{B.5}$$

where we use the following results for a normal variable z and a Poisson variable y^*:

$$E_t^Q[\exp(\alpha z_{t+1} + \beta z^2_{t+1})] = \exp(\frac{\alpha^2 h_{z,t+1}}{2(1-2\beta h_{z,t+1})} - \frac{1}{2}\log(1 - 2\beta h_{z,t+1}), \tag{B.6}$$

$$E_t^Q[\exp(\Xi y^*_{t+1})] = \exp(\xi^*(\Xi)h^*_{y,t+1}). \tag{B.7}$$

Applying (B.6) and (B.7) to (B.5) and subsequently equating terms on the right hand side of (B.5) and (22) gives the following analytical solutions for the affine coefficients $A(\varphi;t,T)$, $B(\varphi;t,T)$, and $C(\varphi;t,T)$:

$$\begin{aligned} A(\varphi;t,T) &= A(\varphi;t+1,T) + B(\varphi;t+1,T)\omega_z + C(\varphi;t+1,T)\omega^*_y \\ &\quad - \frac{1}{2}\log(1 - 2B(\varphi;t+1,T)a_z - 2C(\varphi;t+1,T)a^*_y), \end{aligned} \tag{B.8}$$

$$\begin{aligned} B(\varphi;t,T) &= -\frac{1}{2}\varphi + B(\varphi;t+1,T)(b_z + a_z c_z^{*2}) + C(\varphi;t+1,T)a^*_y c_z^{*2} \\ &\quad + \frac{(\varphi - 2B(\varphi;t+1,T)a_z c^*_z - 2C(\varphi;t+1,T)a^*_y c^*_z)^2}{2(1 - 2B(\varphi;t+1,T)a_z - 2C(\varphi;t+1,T)a^*_y)}, \end{aligned} \tag{B.9}$$

$$C(\varphi;t,T) = b_y C(\varphi;t+1,T) - \varphi\xi^*(1) + \xi^*(\Xi), \tag{B.10}$$

where $\xi^*(\Xi) = \exp(\theta^*\Xi + \frac{\Xi^2\delta^2}{2}) - 1$, with $\Xi = \varphi + B(\varphi;t+1,T)d_z + C(\varphi;t+1,T)d^*_y$.

C. Option Based Filtering: Retrieving State Variables from Option Data

First consider the GARCH model. Let the current date be $t = 0$, for $0 \leqslant t \leqslant T-1$, and rewrite the return process under the risk neutral measure as

$$x_{t+1} \equiv \log \frac{F_{t+1,T}}{F_{0,T}} = \log \frac{F_{t,T}}{F_{0,T}} - \frac{1}{2} h_{t+1} + \sqrt{h_{t+1}} \varepsilon_{t+1}. \tag{C.1}$$

The Brownian bridge process in Mazzoni [2015], B_t, imposes structure at the end point by forcing the process (C.1) to end up at a particular value $k' = \log \frac{F_{T,T}}{F_{0,T}} = \log \frac{K}{F_{0,T}}$ at time T. The conditional expectation of the log return process (C.1) is obtained by iterated expectations

$$E^Q[x_{t+1} | x_T = k', \mathcal{F}_0] = \frac{t+1}{T} k'. \tag{C.2}$$

The conditional expectation and variance of the Brownian bridge increment are given by

$$E^Q[\triangle\ B_{t+1} | \mathcal{F}_t] = \frac{1}{\sqrt{h_{t+1}}} \frac{k' - x_t}{T-t} + \frac{1}{2} \sqrt{h_{t+1}}, \tag{C.3}$$

$$Var[\triangle\ B_{t+1} | \mathcal{F}_t] = 1 - \frac{1}{T-t}. \tag{C.4}$$

See Mazzoni [2015] for the derivation of (C.2)-(C.4). Since all expectations are taken with respect to the risk neutral probability measure Q and are conditional on the filtration $x_T = k'$ and $\mathcal{F}_0$, for notational convenience we omit both arguments, $E^Q[...|\mathcal{F}_0] = E[...]$.

Using the fact that $E[\varepsilon_t^2] = Var(\varepsilon_t) + E[\varepsilon_t]^2$ and $\varepsilon_t = \triangle\ B_t$. Applying equations (C.2)-(C.4), the conditional expectation of the variance is given by

$$\begin{aligned} E[h_{t+1}] &= bE[h_t] + a\frac{T-t}{T-(t-1)} + a(E[\frac{1}{\sqrt{h_t}} \frac{k' - x_{t-1}}{T-t+1} + \frac{1}{2}\sqrt{h_t}])^2 \\ &\quad -2ac^* \frac{k'}{T} - ac^* E[h_t] + ac^{*2} E[h_t]. \end{aligned} \tag{C.5}$$

To compute the local variance $E[h_T]$ we set $t = T-1$. Since the variance of x_T vanishes, we assume that $Cov(x_{T-2}, \sqrt{h_{T-1}}) \to 0$ and $Var(\sqrt{h_{T-1}}) \to 0$. This gives

$$E[h_T] \approx \alpha + \beta E[h_{T-1}] + \frac{a\eta^2}{E[h_{T-1}]}, \tag{C.6}$$

Introduce the substitution $a\eta^2 \to \epsilon$. If the GARCH model is stationary and $|\epsilon| < 1$, we have the

convergent perturbation series

$$E[h_T] = \sum_{n=0}^{\infty} \epsilon^n E_n[h_T]. \tag{C.7}$$

Substituting (C.7) into (C.6), taking a Taylor expansion at 0 for the third term in (C.6), and equating terms according to their orders in ϵ, one obtains a sequence of equations, starting with the unperturbed problem

$$E_0[h_T] = \alpha + \beta E_0[h_{T-1}]. \tag{C.8}$$

We tested the accuracy of the unperturbed solution for the GARCH model. The relative errors of the unperturbed solution, $E_0[h_T]$, and the true one, $E[h_T]$, range from $1e^{-3}$ to $1e^{-12}$. This suggests that the unperturbed solution is sufficiently accurate, and that it is unnecessary to include higher order terms in the perturbation series.

Iterating (C.8) backwards from time T to 1 and using the geometric series representation, we have:

$$\sigma^2_{loc}(k', T) = E[h_T] \approx \alpha(k', T)\frac{1-\beta^{T-1}}{1-\beta} + \beta^{T-1}h_1, \tag{C.9}$$

where $\alpha = 2a\eta(\frac{1}{2} - c^*) + \frac{a}{2}$, $\beta = b + a(\frac{1}{2} - c^*)^2$, and $\eta = \frac{k'}{T}$.

Mazzoni [2015] uses a straight line assumption for x_t from $(0, 0)$ to (k', T). The implied variance can be thought of as a summation over all local variances

$$\sigma^2_{imp}(k', T) \approx \frac{1}{T}\sum_{t=0}^{T-1}(\alpha\frac{1-\beta^t}{1-\beta} + \beta^t h_1), \tag{C.10}$$

which gives us the approximation for implied volatility in (24).

For the DI-Jump model, we proceed in a similar way and rewrite the risk neutral return process (18) by replacing $\log \frac{F_{t+1,T}}{F_{t,T}}$ with $x_{t+1} - x_t$.

$$x_{t+1} \equiv \log \frac{F_{t+1,T}}{F_{0,T}} = \log \frac{F_{t,T}}{F_{0,T}} - \frac{1}{2}h_{z,t+1} - \xi^* h^*_{y,t+1} + z_{t+1} + y^*_{t+1}. \tag{C.11}$$

Given a bridge process with $k' = \log \frac{K}{F_{0,T}}$ at time T, the conditional expectation of the log return process (C.11) is the same as that for the GARCH model $E[x_{t+1}] = \frac{t+1}{T}k'$. Since equation (C.11) has both normal and jump components in the futures return innovation, the bridge process is a more general one than the Brownian bridge used for the GARCH model. Denote the bridge increment for the total innovation as $\triangle\ BI_{t+1} = (z_{t+1} + y^*_{t+1}|x_T = k', \mathcal{F}_t)$. Its conditional expectation and variance are given by

$$E^Q[\triangle\ BI_{t+1}|\mathcal{F}_t] = \frac{k' - x_t}{T-t} + \frac{1}{2}h_{z,t+1} + \xi^* h^*_{y,t+1}, \tag{C.12}$$

$$Var[\triangle\ BI_{t+1}|\mathcal{F}_t] = (1 - \frac{1}{T-t})(h_{z,t+1} + (\theta^{*2} + \delta^2)h^*_{y,t+1}). \tag{C.13}$$

For the jump model, in order to obtain the expected value of $h_{z,t+1}$ or $h^*_{y,t+1}$, we need to find the expectation of z_t and y^*_t separately. We assume the mean shifting caused by the bridge assumption is allocated to the means of normal component and jump component proportional to their corresponding variance contribution. This assumption is motivated by the conditional expectations of the components in the absence of a bridge process. To verify if our results are robust with respect to this assumption, we computed option implied variance under the alternative, more extreme, assumption that the mean shifting caused by the bridge assumption is allocated completely to the normal component. The two assumptions yield similar results, suggesting that the errors induced by this assumption are negligible. Using this assumption, we get:

$$E[\widetilde{z}_{t+1}|\mathcal{F}_t] = \frac{h_{z,t+1}}{h_{z,t+1} + \delta^2 h^*_{y,t+1}} \left(\frac{k' - x_t}{T-t} + \frac{1}{2}h_{z,t+1} + (\xi^* - \theta^*)h^*_{y,t+1} \right), \tag{C.14}$$

$$E[\widetilde{y}_{t+1}|\mathcal{F}_t] = \theta^* h^*_{y,t+1} + \frac{\delta^2 h^*_{y,t+1}}{h_{z,t+1} + \delta^2 h^*_{y,t+1}} \left(\frac{k' - x_t}{T-t} + \frac{1}{2}h_{z,t+1} + (\xi^* - \theta^*)h^*_{y,t+1} \right). \tag{C.15}$$

Given this assumption, we calculate the conditional expectation of the risk neutral variance (19) using (C.2), (C.12), and (C.13)

$$\begin{aligned} E[h_{z,t+1}] \;=\; & b_z E[h_{z,t}] + a_z Var[\frac{\widetilde{z}_t}{\sqrt{h_{z,t}}}] + a_z(E[\frac{1}{\sqrt{h_{z,t}}}\frac{h_{z,t}}{h_{z,t} + \delta^2 h^*_{y,t}}(\frac{k' - x_{t-1}}{T-t+1} + \frac{1}{2}h_{z,t} \\ & +(\xi^* - \theta^*)h^*_{y,t})])^2 - 2a_z c^*_z E[\frac{h_{z,t}}{h_{z,t} + \delta^2 h^*_{y,t}}(\frac{k' - x_{t-1}}{T-t+1} + \frac{1}{2}h_{z,t} + (\xi^* - \theta^*)h^*_{y,t})] \\ & +a_z c^{*2}_z E[h_{z,t}] + d_z E[\theta^* h^*_{y,t} + \frac{\delta^2 \omega^*_y}{h_{z,t} + \delta^2 h^*_{y,t}} \left(\frac{k' - x_{t-1}}{T-t+1} + \frac{1}{2}h_{z,t} + (\xi^* - \theta^*)h^*_{y,t} \right)]. \end{aligned} \tag{C.16}$$

For the DI-Jump model, $h^*_{y,t+1} = \omega^*_y + k^* h_{z,t+1}$. The total conditional variance at T is given by

$$E[h_{z,T} + (\theta^{*2} + \delta^2)h^*_{y,T}] = \alpha_J + \beta_J E[h_{z,T-1}] + \Phi_J, \tag{C.17}$$

where $\alpha_J = 2a_z\eta_J c_J/\gamma + a_z/2 + d_z\theta^*\omega^*_y + \delta^2\omega^*_y\lambda_J(2a_z c^*_z/\gamma - 2a_z\lambda_J/\gamma + d_z)/\gamma + (\theta^{*2} + \delta^2)\omega^*_y$, $\beta_J = b_z + a_z c_J^2 + d_z\theta^* k^* + k^*(\theta^{*2} + \delta^2)$, $\eta_J = \frac{k'}{T} + (\xi^* - \theta^*)\omega^*_y$, $\lambda_J = \frac{1}{2} + (\xi^* - \theta^*)k^*$, and $c_J = \frac{1}{2} - c^*_z + (\xi^* - \theta^*)k^*$. Φ_J captures the terms which are neither constant nor the coefficient of $E[h_{z,T-1}]$.

We find that the unperturbed terms for (C.17) are also accurate enough to be good approximations for the true total variances, and we thus use them as approximations for the local

variances.

$$\sigma^2_{loc}(k', T) = E[h_{z,T} + (\theta^{*2} + \delta^2)h^*_{y,T}] \approx \alpha_J \frac{1 - \beta_J^{T-1}}{1 - \beta_J} + \beta_J^{T-1} h_{z,1}. \tag{C.18}$$

Again using the relationship between local volatility and the conditional expectation of the variance (C.9), as well as the relationship between the local variance and implied volatility (C.10), we obtain implied volatility as a function of $h_{z,1}$ for the DI-Jump model.

References

[1] Acharya, V., L. Lochstoer, T. Ramadorai. "Limits to Arbitrage and Hedging: Evidence from Commodity Markets." Journal of Financial Economics, 109 (2013), pp. 441-465.

[2] Amengual, D., D. Xiu. "Resolution of Policy Uncertainty and Sudden Declines in Volatility." Working paper, University of Chicago, 2013.

[3] Andersen, T., L. Benzoni, J. Lund. "An Empirical Investigation of Continuous-time Equity Return Models." Journal of Finance, 57 (2002), pp. 1239-1284.

[4] Barone-Adesi, G., R. Whaley. "Efficient Analytical Approximation of American Option Values." Journal of Finance, 42 (1987), pp. 301-320.

[5] Bates, D. "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options." Review of Financial Studies, 9 (1996), pp. 69-107.

[6] Bates, D. "Maximum Likelihood Estimation of Latent Affine Processes." Review of Financial Studies, 19 (2006), pp. 909-965.

[7] Black, F. "The Pricing of Commodity Contracts." Journal of Financial Economics, 3 (1976), pp. 167-179.

[8] Bollerslev, T. "Generalized Autoregressive Conditional Heteroskedasticity." Journal of Econometrics, 31 (1986), pp. 307-327.

[9] Broadie, M., M. Chernov, M. Johannes. "Model Specification and Risk Premia: Evidence from Futures Options." Journal of Finance, 62 (2007), pp. 1453-1490.

[10] Casassus, J., P. Collin-Dufresne. "Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates." Journal of Finance, 60 (2005), pp. 2283-2331.

[11] Christoffersen, P., K. Jacobs, C. Ornthanalai. "Dynamic Jump Intensities and Risk Premia: Evidence from S&P 500 Returns and Options." Journal of Financial Economics, 106 (2012), pp. 447-472.

[12] Duffie, D. Dynamic Asset Pricing Theory. Third Edition. Princeton, NJ: Princeton University Press, 2001.

[13] Duffie, D., J. Pan, K. Singleton. "Transform Analysis and Asset Pricing for Affine Jump-diffusions." Econometrica, 68 (2000), pp. 1343-1376.

[14] Engle, R. F. "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation." Econometrica, 50 (1982), pp. 987-1008.

[15] Engle, R.F., J. Mezrich. "GARCH for Groups." Risk, 9 (1996), pp. 36-40.

[16] Eraker, B. "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices." Journal of Finance, 59 (2004), pp. 1367-1403.

[17] Eraker, B., M. Johannes, N. Polson. "The Impact of Jumps in Volatility and Returns." Journal of Finance, 58 (2003), pp. 1269-1300.

[18] Francq, C., L. Horvath, J.-M. Zakoian. "Merits and Drawbacks of Variance Targeting in GARCH Models." Journal of Financial Econometrics, 9 (2011), pp. 619-656.

[19] Hamilton, J. D., J. C. Wu. "Risk Premia in Crude Oil Futures Prices." Journal of International Money and Finance, 42 (2014), pp. 9-37.

[20] Hansen, B. E. "The Likelihood Ratio Test under Non-standard Conditions: Testing the Markov Trend Model of GNP." Journal of Applied Econometrics, 7 (1992), pp. S61-S82.

[21] Hansen, B. E. "Autoregressive Conditional Density Estimation." International Economic Review, 35 (1994), pp. 705-730.

[22] Heston, S. "A Closed-form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." Review of Financial Studies, 6 (1993), pp. 327-343.

[23] Ingersoll, J. Theory of Financial Decision Making, Totowa, NJ: Rowman and Littlefield, 1987.

[24] Litzenberger, R., N. Rabinowitz. "Backwardation in Oil Futures Markets: Theory and Empirical Evidence." Journal of Finance, 50 (1995), pp. 1517-1545.

[25] Maheu, J., T. McCurdy. "News Arrival, Jump Dynamics and Volatility Components for Individual Stock Returns." Journal of Finance, 59 (2004), pp. 755-793.

[26] Mazzoni, T. "A GARCH Parametrization of the Volatility Surface." Journal of Derivatives, 23 (2015), pp. 9-24.

[27] Nelson, C., A. Siegel. "Parsimonious Modeling of Yield Curves." Journal of Business, 60 (1987), pp. 473-489.

[28] Ornthanalai, C. "Levy Jump Risk: Evidence from Options and Returns." Journal of Financial Economics, 112 (2014), pp. 69-90.

[29] Pan, X. N. "Investor Beliefs and State Price Densities in the Crude Oil Market." Working paper, Tulane University, 2012.

[30] Santa-Clara, P., S. Yan. "Crashes, Volatility, and the Equity Premium: Lessons from S&P500 Options." Review of Economics and Statistics, 92 (2010), pp. 435-451.

[31] Schwartz, E. S. "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging." Journal of Finance, 52 (1997), pp. 923-973.

[32] Singleton, K. "Investor Flows and the 2008 Boom/Bust in Oil Prices." Management Science, 60 (2013), pp. 300-318.

[33] Trolle, A. B., E. S. Schwartz. "Unspanned Stochastic Volatility and the Pricing of Commodity Derivatives." Review of Financial Studies, 22 (2009), pp. 4423-4461.

Exhibit 1

Crude Oil Futures Prices, Futures ATM Implied Volatilities, and Implied Skewness and Kurtosis

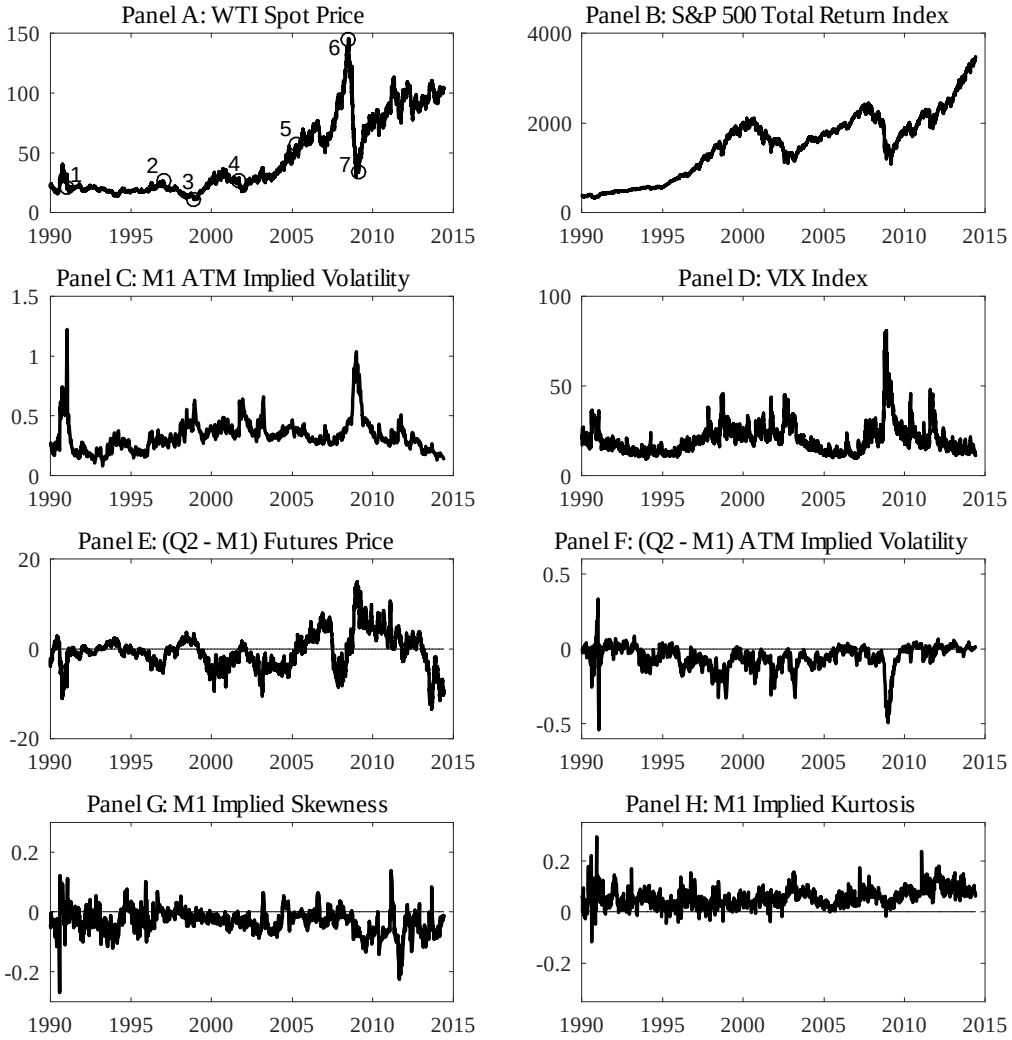

Notes: Annualized implied volatilities (IV) are computed from option prices by inverting the Barone-Adesi and Whaley [1987] formula. Implied skewness is computed as IV(X/F = 1.15) - IV(X/F = 0.85) and implied kurtosis is computed as IV(X/F = 1.15) + IV(X/F = 0.85) - 2 × IV(X/F = 1). Data spans 6,126 trading days from January 2, 1990 to May 30, 2014. The key geopolitical and economic events are: 1. Iraq invades Kuwait; 2, Asian financial crisis; 3. OPEC cuts production targets 1.7 million barrels of oil per day (mmbpd); 4. 9-11 attacks; 5. Low spare capacity; 6. Global financial collapse; 7. OPEC cuts production targets 4.2 mmbpd.

Exhibit 2

Summary Statistics

Panel A: Average Futures Prices

	Maturity								
	M1	**M2**	**M3**	**M4**	**M5**	**M6**	**Q1**	**Q2**	**Average**
01/1990 - 05/2014	45.96	46.05	46.06	46.03	45.98	45.92	45.78	45.55	45.92

Panel B: Daily Moments of Futures Returns

	M1	**M2**	**M3**	**M4**	**M5**	**M6**	**Q1**	**Q2**	**Average**
Mean	2.28E-04	2.30E-04	2.35E-04	2.38E-04	2.41E-04	2.44E-04	2.45E-04	2.47E-04	2.39E-04
Variance	4.97E-04	4.15E-04	3.66E-04	3.35E-04	3.11E-04	2.93E-04	2.70E-04	2.41E-04	3.41E-04
Skewness	-1.01	-0.92	-0.77	-0.72	-0.65	-0.60	-0.53	-0.39	-0.70
Excess Kurtosis	18.17	14.41	11.18	9.85	8.82	7.82	6.99	5.37	10.33

Panel C: Futures Average Daily Trading Volume

	M1	**M2**	**M3**	**M4**	**M5**	**M6**	**Q1**	**Q2**	**Average**
01/1990 - 05/2014	109,932	65,942	23,608	11,756	7,123	5,104	6,280	3,867	29,202

Panel D: Number of Options Contracts

Moneyness	**M1**	**M2**	**M3**	**M4**	**M5**	**M6**	**Q1**	**Q2**	**All**
0.78-0.82	1,265	3,539	4,127	4,172	4,009	3,644	3,672	2,736	27,164
0.82-0.86	2,357	4,357	4,736	4,755	4,528	4,021	4,111	3,070	31,935
0.86-0.90	3,646	5,102	5,362	5,417	5,118	4,537	4,491	3,382	37,055
0.90-0.94	4,814	5,761	5,799	5,608	5,347	4,724	4,600	3,495	40,148
0.94-0.98	5,611	5,988	5,818	5,543	5,346	4,756	4,712	3,606	41,380
0.98-1.02	5,909	5,902	5,728	5,466	5,255	4,671	4,676	3,521	41,128
1.02-1.06	5,617	5,979	5,800	5,558	5,243	4,496	4,538	3,371	40,602
1.06-1.10	4,793	5,821	5,727	5,455	5,081	4,505	4,400	3,087	38,869
1.10-1.14	3,740	5,195	5,569	5,333	4,818	4,094	4,314	2,939	36,002
1.14-1.18	2,649	4,485	4,803	4,771	4,404	3,810	3,982	2,725	31,629
1.18-1.22	1,741	3,773	4,225	3,951	3,668	3,135	3,464	2,511	26,468
All	42,142	55,902	57,694	56,029	52,817	46,393	46,960	34,443	392,380

Panel E: Option Average Daily Trading Volume

Moneyness	**M1**	**M2**	**M3**	**M4**	**M5**	**M6**	**Q1**	**Q2**	**Average**
0.78-0.82	539	1,225	713	464	298	233	318	149	492
0.82-0.86	1,020	1,876	932	570	403	326	346	169	705
0.86-0.90	2,076	2,765	1,243	707	410	294	341	145	998
0.90-0.94	4,041	3,710	1,429	750	437	295	353	129	1,393
0.94-0.98	6,910	4,234	1,424	719	448	306	368	159	1,821
0.98-1.02	10,461	5,810	2,373	1,429	997	758	1,091	604	2,940
1.02-1.06	5,730	3,594	1,254	588	379	219	265	138	1,521
1.06-1.10	2,931	2,809	1,099	526	277	184	228	101	1,019
1.10-1.14	1,501	2,008	855	492	271	173	245	106	706
1.14-1.18	820	1,431	651	388	239	151	195	88	495
1.18-1.22	482	997	512	308	168	123	189	85	358
Average	3,319	2,769	1,135	631	393	278	358	170	1,132

Notes: We report summary statistics for crude oil futures returns and options. M1 (M2, M3, M4, M5, M6) refers to futures contracts with expiration in 1 (2, 3, 4, 5, 6) months; Q1 and Q2 refer to the next two futures contracts with expiration in either March, June, September or December. Moneyness is defined as the option strike divided by the price of the underlying futures contract. Data spans 6,126 trading days from January 2, 1990 to May 30, 2014.

Exhibit 3

Implied Volatility Smiles and Smirks

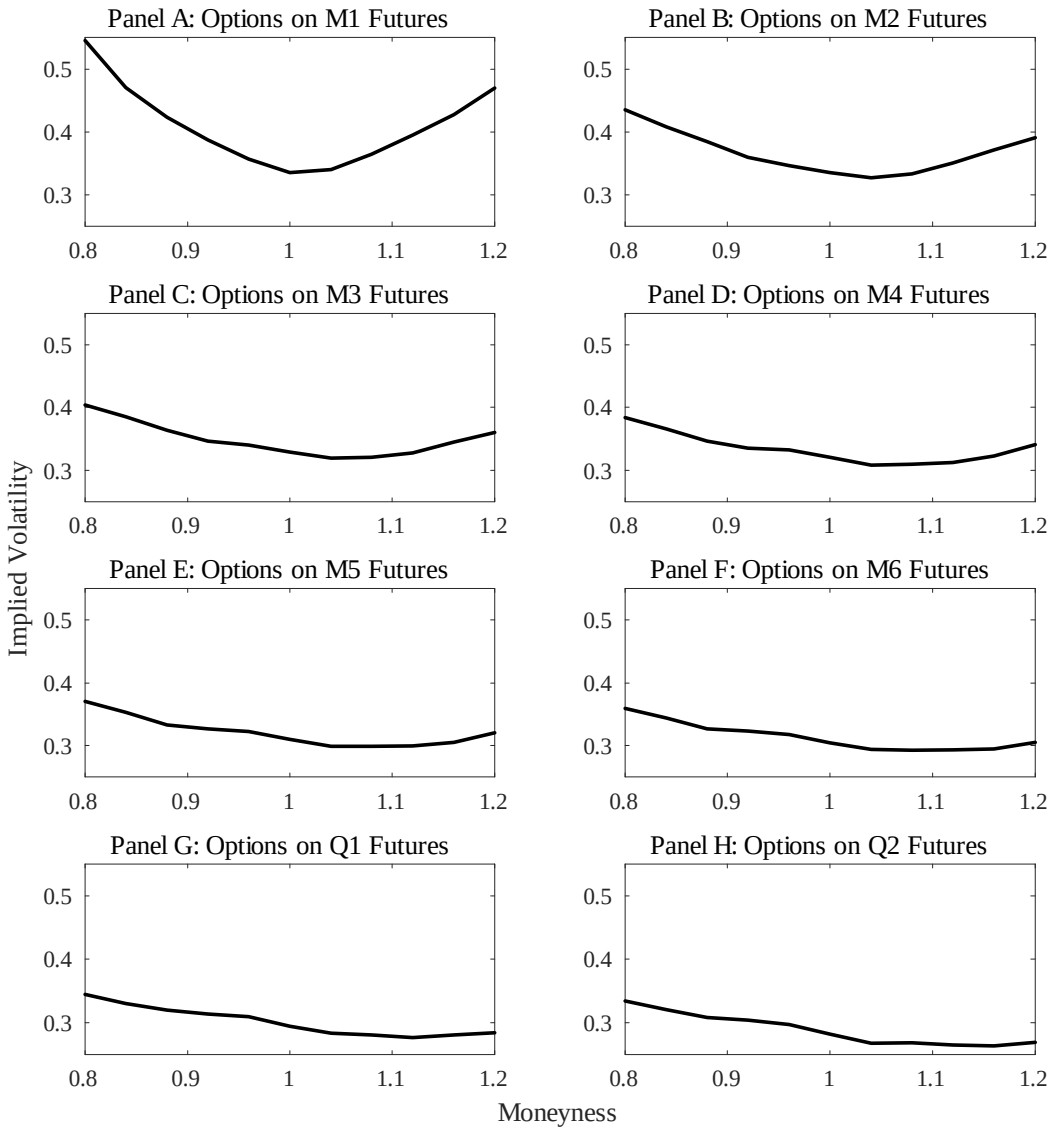

Notes: We plot average implied volatility across moneyness for M1-Q2 futures contracts. Averages are based on 6,126 daily observations from January 2, 1990 through May 30, 2014. Moneyness is defined as the option strike divided by the price of the underlying futures contract.

Exhibit 4

Maximum Likelihood Estimates

Parameters	Panel A: Using Futures only			Panel B: Using Futures and Options		
	GARCH	DI-Jump		GARCH	DI-Jump	
	Normal	Normal (z)	Jump (y)	Normal	Normal (z)	Jump (y)
λ	8.14E-01	1.16E+00		7.64E-01	7.30E-01	4.91E-04
	(2.53E-03)	(2.96E-03)		(8.90E-04)	(1.29E-04)	(1.70E-06)
w			8.28E-06			7.06E-04
			(9.94E-07)			(2.93E-06)
b	9.26E-01	9.53E-01		9.04E-01	1.00E+00	
	(2.82E-03)	(1.35E-03)		(5.00E-04)	(2.24E-06)	
c	2.66E+01	3.26E+01		4.91E+02	2.39E+02	
	(7.34E-02)	(4.79E-02)		(1.09E-01)	(4.35E-02)	
d		-1.67E-04			-1.40E-02	
		(1.58E-06)			(7.82E-06)	
θ			-4.85E-02			-3.84E-02
			(6.75E-05)			(2.32E-05)
δ			2.15E-01			1.17E-02
			(3.26E-04)			(6.96E-06)
k			1.95E+01			3.48E+01
			(4.13E-02)			(1.32E-02)
Properties						
Number of jumps / Yr		1.37			4.32	
Risk Premium (%)	7.99	8.09		8.78	8.59	0.21
% of Annual Variance	100.00	51.23	48.77	100.00	94.42	5.58
Persistence	0.9406	0.9600		0.9991	0.9999	
Avg. Annual Volatility	0.29	0.29		0.29	0.29	
IV Bias				0.91	-0.06	
IV RMSE				3.79	3.03	
Log-Likelihood	16,328.82	16,505.56		483,295.57	532,171.67	

Notes: We report the MLE estimation results on daily crude oil futures returns in the left three columns and the joint MLE estimation results using both daily crude oil futures and options in the right three columns. Columns labeled "Normal" contain estimates of the parameters governing the normal component; columns labeled "Jump" contain parameters governing the jump component. Standard errors computed using the Hessian matrix are reported in parentheses. Data spans 6,126 trading days from January 2, 1990 to May 30, 2014.

Exhibit 5

Futures Returns and QQ plots for M1 and Q2 Contracts

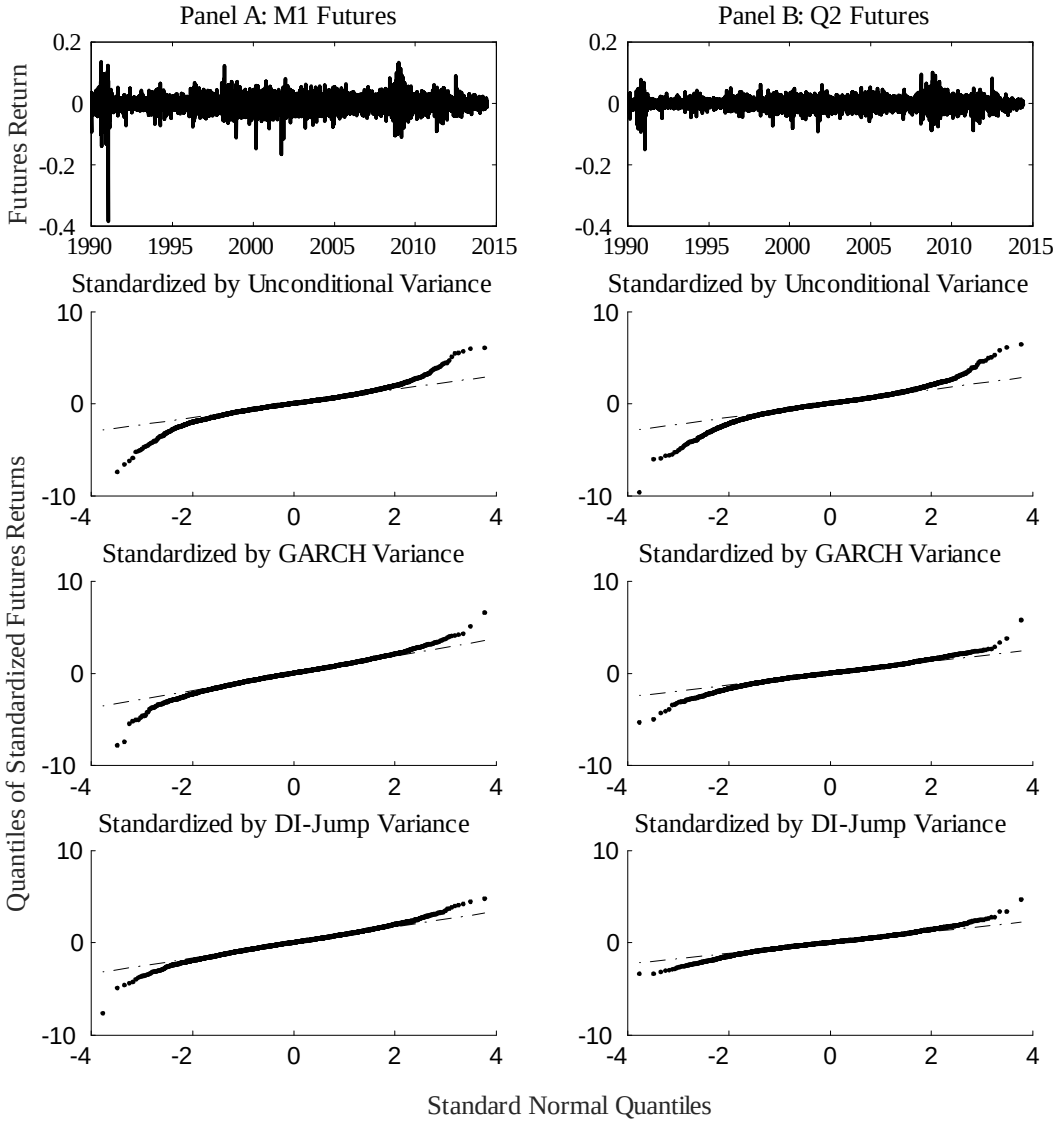

Notes: We construct QQ plots for the M1 futures contracts and Q2 futures contracts by standardizing futures returns using the unconditional variance, the GARCH conditional variance, and the conditional variance for the DI-Jump model estimated using results in Panel A of Exhibit 4, with respect to the standard normal distribution. Data spans 6,126 trading days from January 2, 1990 to May 30, 2014.

Exhibit 6

The DI-Jump Model

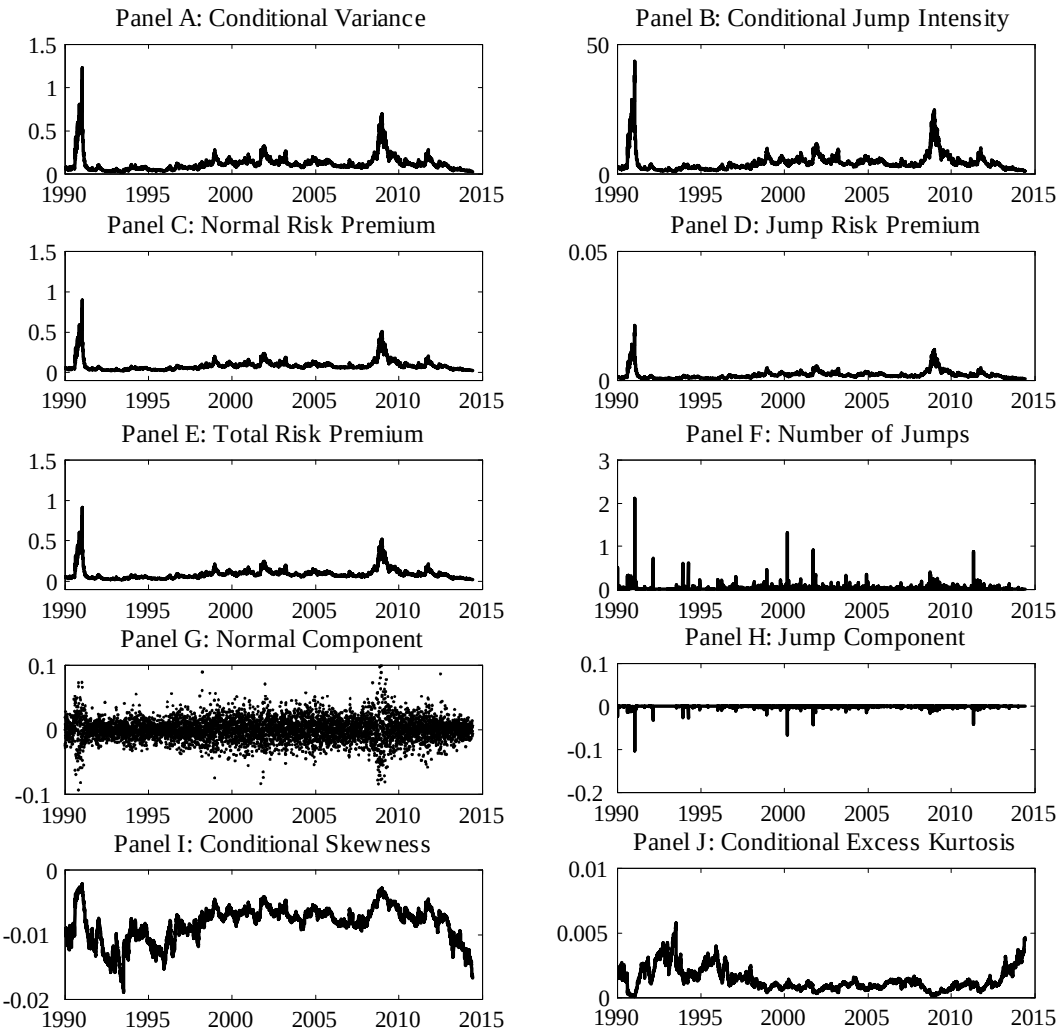

Notes: We plot the conditional variance, $h_{z,t+1}$, the conditional jump intensity, $h_{y,t+1}$, the normal risk premium, $\lambda_z h_{z,t+1}$, the jump risk premium, $\lambda_y h_{y,t+1}$, the total risk premium, $\lambda_z h_{z,t+1} + \lambda_y h_{y,t+1}$, the filtered number of jumps, n_t, the annualized conditional skewness, and the annualized conditional excess kurtosis for the DI-Jump model. Results are based on the estimates in Panel B of Exhibit 4 using joint MLE on futures and options data. All variables are reported in annualized terms. The standardized normal component, z_t, and the jump component, y_t, are obtained using option-based filtering.

Exhibit 7

IV RMSE and IV Bias by Moneyness and Maturity

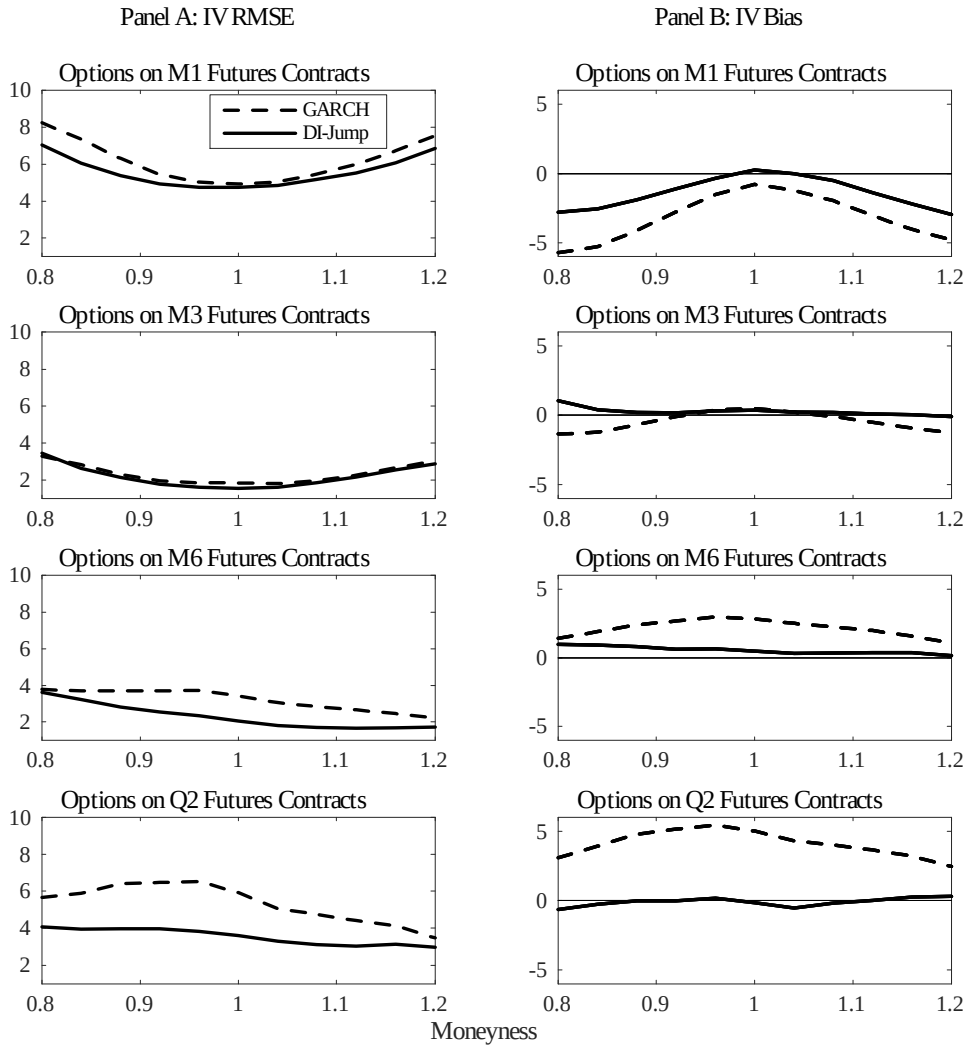

Notes: We plot the option implied volatility root mean squared errors (IV RMSE) and the implied volatility bias (IV Bias) within each moneyness-maturity category for the GARCH model and DI-Jump model. The models are estimated using daily crude oil futures returns and options jointly, using data for the period January 2, 1990 through May 30, 2014. The pricing errors are defined as the difference between fitted and actual implied volatilities and reported in percentages. M1, M3, and M6 refer to option contracts with expiration in 1, 3, 6 months, respectively; Q2 refers to the next second option contracts with expiration in either March, June, September, or December. Moneyness is defined as the option strike divided by the price of the underlying futures contract.