

Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 1122

October 2014

Generating Options-Implied Probability Densities to
Understand Oil Market Events

Deepa Dhume Datta

Juan M. Londono

Landon J. Ross

Board of Governors of the Federal Reserve System

NOTE: International Finance Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References to International Finance Discussion Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors. Recent IFDPs are available on the Web at www.federalreserve.gov/pubs/ifdp/. This paper can be downloaded without charge from Social Science Research Network electronic library at www.ssrn.com.

Generating Options-Implied Probability Densities to Understand Oil Market Events*

Deepa Dhume Datta, Juan M. Londono, and Landon J. Ross

Federal Reserve Board of Governors

October 29, 2014

Abstract

We investigate the informational content of options-implied probability density functions (PDFs) for the future price of oil. Using a semiparametric variant of the methodology in Breeden and Litzenberger (1978), we investigate the fit and smoothness of distributions derived from alternative PDF estimation methods, and develop a set of robust summary statistics. Using PDFs estimated around episodes of high geopolitical tensions, oil supply disruptions, and macroeconomic data releases, we explore the extent to which oil price movements are expected or unexpected, and whether agents believe these movements to be persistent or temporary.

JEL Classifications: G13, G14, C13

Keywords: options-implied PDFs, futures, options, oil.

*The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. The authors thank John Ammer, David Bowman, Andy Fiedler, Jeffrey Frankel, Will Melick, Charlie Thomas, and especially Carl J. Liebersohn for helpful discussions. They also thank seminar participants at the Federal Reserve Board and the NBER Summer Institute. Emails: deepa.d.datta@frb.gov, juan-miguel.londono-yarce@frb.gov, landon.j.ross@frb.gov.

1 Introduction

The ever-expanding set of financial derivatives allows agents to hedge against price fluctuations in a variety of markets. Therefore, investors, researchers, and policymakers use measures based on derivatives extensively to analyze market participants' beliefs about future movements in the prices of these derivatives' underlying assets and the probability these participants assign to unlikely events. The recent expansion of derivatives markets has prompted resurgent interest in estimating options-implied probability density functions (PDFs). Although a wide range of procedures to estimate options-implied PDFs has been developed in the literature, the effect of methodological choices on PDF estimates and pricing errors is rarely explored, and the research for oil derivatives is no exception.¹

Oil derivatives are particularly interesting for two main reasons. First, their underlying asset, crude oil, is the most widely traded commodity in the world. Relative to derivatives markets for stocks and currencies, oil derivatives are traded for a much larger number of strikes. While helpful, the volume of data poses additional challenges for PDF-estimation methods that have previously proven successful for stock and currency derivatives. By improving the application of these methods to oil derivatives, we can learn more about the methods' robustness and future applicability to other financial markets. Second, the price of crude oil both reflects and has implications for the global economic outlook. A better understanding of oil market dynamics can shed light on oil prices, which are a key source of tail risk for the global economy. In addition, the informational content of oil derivatives can be indicative of shifts in global economic expectations.

¹A few notable exceptions are Melick and Thomas (1997), Flamouris and Giamouridis (2002), and Askari and Krichene (2008).

In this paper, we investigate the estimation choices made when generating PDFs from options data for West Texas Intermediate (WTI) crude oil and analyze the stability of the information extracted from these options-implied PDFs. We also investigate the dynamics of the information extracted from PDFs around particular market events to assess the informational content of options-implied PDFs for changes in oil prices.

Our contribution to the literature is twofold. First, we propose a semiparametric method that generates accurate, interpretable, and stable options-implied distributions for changes in the price of oil. We compare alternative specifications for the two degrees of methodological freedom commonly discussed in the literature: fitting a curve to the observed implied volatilities and estimating the unobserved tails of the PDF. We choose a method that generates low pricing errors while still generating a sensible PDF to ensure accuracy and interpretability. We also investigate the robustness of key moments of the distribution. Our second contribution is using the information from estimated options-implied PDFs to analyze the changes in market participants' beliefs about oil price movements around important market events. We provide evidence on the extent to which price movements are expected by market participants and whether agents believe these movements in oil prices are persistent or temporary.

We confine our investigation to the class of semiparametric methods because semiparametric methods require fewer assumptions about the shape of the distribution function than purely parametric methods, while providing more reasonable and interpretable PDFs than purely nonparametric methods.² In particular, we follow the finite-difference method in Breeden and Litzenberger (1978), who show that the second difference of the price of a European call with respect to the strike price is equivalent to the risk-free rate discounted PDF.

²Jackwerth (1999), Jondeau and Rockinger (2000), and Bliss and Panigirtzoglou (2002) provide a comprehensive review of methods to calculate options-implied PDFs available in the literature.

The empirical implementation of the method in Breeden and Litzenberger (1978) creates the two methodological decisions we look at in detail in this paper. First, because call prices are observed for a discrete number of strike prices, we must decide on a curve-fitting method to interpolate the observed prices. Instead of interpolating the call price function directly, we follow Shimko (1993) and Dumas, Fleming, and Whaley (1998), and interpolate the implied volatility curve to reduce the sensitivity of the results to out-of-the-money options.³ We implement four options-implied volatility curve-fitting methods: cubic spline, quartic (fourth-order) spline, smoothing spline, and one-knot spline. We evaluate the alternative curve-fitting methods by comparing the smoothness of the estimated PDFs and the pricing errors of the fitted implied volatility curves. We find that, for the case of the highly traded crude oil options, the smoothing spline method allows for convenient visualization of the PDFs while keeping pricing errors moderately low.

The second methodological choice is choosing the method for generating the PDF outside the range of observed strikes. To estimate the tails of the distribution, we consider either extrapolating the implied volatility curve (in volatility-strike space) beyond the range of observed call prices, or attaching tails in probability-strike space to a PDF that is estimated from an implied volatility curve only within the range of observed strikes. Since the tails of the PDF cover the range of strike prices over which we have no observed call prices, we compare put pricing errors for the near- and far-future maturities to evaluate the alternative tail specifications. We find that extrapolating the implied volatility curve using a first-order polynomial is preferred, as it generates low put pricing errors in most contexts.

We also investigate the sensitivity of the informational content of the PDFs to choices in estimation methodology and to small changes in the observed options prices.

³Bliss and Panigirtzoglou (2002) fit the delta surface to put even more emphasis on at-the-money options.

In line with the literature, we find that higher-order moments, such as skewness and kurtosis, are more sensitive to the estimation method (see Flamouris and Giamouridis (2002), Campa, Chang, and Refalo (2002), and papers cited therein). We also find that quantile moments are preferable to standard moments, as the former are more robust to alternative methodological choices. Finally, we find that, at least for the smoothing spline method, confidence intervals, which are calculated using the Monte Carlo method in Bliss and Panigirtzoglou (2002), are quite narrow, especially for quantile volatility and cost of insurance measures.

In the final part of the paper, we investigate the informational content of the estimated PDFs around important events that have implications for global oil supply and demand. In particular, we calculate PDFs and their moments using the smoothing spline method with first-order polynomial tails for the days around the February 2011 Libyan oil supply disruption, the tensions with Iran in February 2012, and the change in market sentiment driven by surprising macroeconomic data releases in May 2012. We examine the changes in market beliefs by estimating daily PDFs around these events by focusing on our robust set of summary statistics, including quantile moments and cost of insurance measures. We find that examining the PDFs allows us to gain a more nuanced view of market beliefs before and after events. and helps us better understand market reactions.

This paper contributes to the general literature on the informational content of options prices (Weinberg (2001), Jiang and Tian (2005), and Kang, Kim, and Yoon (2010)). Although most of this literature has focused on S&P 500 options, options-implied PDFs have been used to study a wide range of markets, including inflation options (Kitsul and Wright (2012)), government bonds (Neuhaus (1995), Shiratsuka (2001), and Cheng (2010)), international stock indexes (Kang and Kim (2006), Birru and Figlewski (2011), Shiratsuka (2001), Glatzer and Scheicher (2005), Äijö (2008),

Kim and Kim (2003), and Syrdal (2002)), and exchange rates (Weinberg (2001), Bahra (1997), Campa, Chang, and Refalo (2002), and Londono and Zhou (2012)). In comparison, only a few papers have used options-implied PDFs to study commodities, including oil (Melick and Thomas (1997), Flamouris and Giamouridis (2002), and Askari and Krichene (2008)) and agricultural commodities (Facker and King (1990)).

Although options-implied measures have become widely used, the literature on the robustness and reliability of the methods used to calculate options-implied PDFs is surprisingly thin. Jondeau and Rockinger (2000) and Campa, Chang, and Refalo (2002) compare the moments obtained from PDFs calculated using several alternative methods. However, these papers do not take into account the variations in the estimated PDFs due to small changes in the parameters or in the prices of options used as inputs. Melick and Thomas (1999) investigate the stability of the double lognormal method by calculating confidence intervals for the moments. These confidence intervals are obtained by perturbing the parameters in the estimated distribution function. Bliss and Panigirtzoglou (2002) propose a stability method based on small perturbations to the price of observed options. They show that the smoothed implied volatility method yields more stable PDFs than the double lognormal method for sterling futures and stock options (FTSE).⁴ Using a stability test that perturbs the cross-section of fitted theoretical option data, instead of observed prices, Bu and Hadri (2007) find that currency options-implied PDFs generated using confluent hypergeometric functions (see Abadir and Rockinger (2003)) are more stable than those generated using the smoothed implied volatility method. We contribute to this literature by analyzing jointly the performance of alternative methods and the stability of the key options-implied moments extracted from the estimated PDFs.

⁴Using a Monte Carlo method, Cooper (1999) also finds that the smoothed implied volatility method is more stable than the mixture of normals.

The rest of this paper is organized as follows. In section 2, we describe our method to calculate options-implied PDFs and the methodological decisions we study. In section 3, we explore the robustness of the informational content of our PDFs, by examining the impact of alternative specifications and price perturbations on key moments extracted from the estimated PDFs. In section 4, we present a few case studies in which the oil PDFs moved considerably around the time of significant events. We discuss how the PDF movements indicate shifting beliefs held by market participants. Finally, section 5 concludes.

2 Estimating options-implied PDFs

In this section, we explore several alternative methods to estimate options-implied PDFs for crude oil futures. In the first part of the section, we explain the general semiparametric method used to calculate the PDFs. In the second part, we introduce the data used to calculate the PDFs. Next, we compare different functional forms for the fitted options-implied volatility curve within the range of the observed strikes and for the PDF's unobserved tails. We explore the trade-off required between the smoothness of the obtained PDFs and pricing errors, which are defined as the difference between the fitted and observed options prices.

2.1 General description of the method

We apply a standard semiparametric methodology for extracting the options-implied PDF based on Breeden and Litzenberger (1978). Their paper demonstrates that the second derivative of an asset's European call prices (C) with respect to strike prices

(K) yields a discounted options-implied probability distribution ($f(\cdot)$):

$$\frac{\partial^2 C(K, T)}{\partial K^2} = e^{-rT} f(K), \quad (1)$$

where r is the one-period risk-free rate and T is the options' maturity.⁵

To apply Breeden and Litzenberger (1978) in an empirical setting, we require two types of interpolation. First, since call prices are available for only a discrete number of strike prices, we must interpolate across strike prices before calculating the derivative in equation (1). In addition, because options have fixed expiration dates, in order to generate fixed-horizon options-implied PDFs, we must also interpolate across time. Throughout the rest of the paper, we use a fixed 90-day maturity.

To calculate PDFs, we collect prices for options contracts with different strikes for the contracts expiring just before and just after 90 days (near-future and far-future maturities). We then clean the data as described in the next subsection. Then, rather than fitting a curve to the options prices, we use Black's model to convert observed call prices to implied volatilities and fit a curve to the implied volatilities at each maturity.⁶ Interpolating in the volatility space reduces the sensitivity of the options-implied PDF to deeply in and out of the money options that are potentially thinly traded (Shimko (1993) and Ait-Sahalia and Lo (1998)). The choice of functional form for the fitted

⁵In our benchmark method, we assume that this relationship also holds for American options. In other words, we assume that, for a given strike and maturity, the price of a European call is exactly the same as that of an American call. This assumption implies that exercising the American call is never attractive—a reasonable assumption for short horizons and low interest rates (see Melick and Thomas (1992) and Chaudhury and Wei (1994)). Melick and Thomas (1997) derive bounds for American options that depend on the terminal PDF. Another implication of using American options is that the put-call parity does not hold exactly. Therefore, we rely exclusively on call options to obtain options-implied volatility curves and use put options only to measure the goodness of fit for the portion of the distribution for which we do not observe call prices.

⁶Note that Black's model is only used to provide a one-to-one mapping between options prices and options-implied volatility. Therefore, the methodology used in this paper is model-free in the sense of Jiang and Tian (2005), as it does not require the assumption that Black's model is the true model for options prices on futures.

volatility curve, which will be further discussed in section 2.3, is one of the two key choices in the estimation of options-implied PDFs.

After obtaining the near-future and far-future fitted volatility curves, we obtain the implied volatility curve for the fixed 90-day maturity by linearly interpolating across maturities. Finally, we convert the fitted fixed-horizon volatility curve back to fitted call prices using Black's model, and generate the PDF by taking the second derivative of the fitted call price function with respect to strike prices (equation (1)).

At times, we must also extrapolate beyond the observed strike prices to make sure that the obtained PDF is interpretable; that is, that the probabilities add up to one. Thus, the second key choice in the estimation of options-implied PDFs is the method used to extrapolate the unobserved tails. We either extrapolate the volatility curve in volatility-strike space (for the near- and far-future maturities) before interpolating to generate the fixed 90-day volatility curve, or we extrapolate the fixed 90-day maturity PDF in probability-strike space. Further details on these extrapolation methods are provided in section 2.4.

2.2 Data description

In this section, we give a brief description of the data before describing the various methods used for fitting the implied volatility curve and extrapolating the unobserved tails of the distribution.

To produce options-implied PDFs we use daily data on WTI crude oil futures and options on WTI futures traded on the New York Mercantile Exchange between January 2000 and December 2013. We use call prices to construct PDFs and use put prices to calculate out-of-sample errors that are helpful for comparing alternative curve-fitting and extrapolation methods. To reduce data and pricing errors, we clean the data by removing options that are priced at less than 10 cents, as very cheap options might

add too much noise to the estimation of the implied volatility curve. We also perform checks to remove observations representing two types of arbitrage opportunities in each day's data. First, we check that call options with a higher strike price are cheaper; in other words, that the empirical first derivative of call prices with respect to strikes is negative. Second, we check that call prices are falling at a decreasing rate with respect to strike prices. Since the second derivative of call prices with respect to strike prices is the options-implied probability distribution, this check for a positive empirical second derivative helps ensure that the PDF remains positive for all strikes.

Even after the data are cleaned, the number of available options contracts remains quite large. To give a sense of the breadth of the market for options on WTI futures contracts, figure 1 plots the monthly average of the number of options available, including put and call options for near- and far-future maturities. Table 1 reports additional summary statistics for these options. After cleaning the data, for the near-future maturity, there are on average 51 calls and 58 puts with different strikes available to trade on a given day. For the far-future maturity, there are on average 49 calls and 56 puts. The number of options available displays considerable variation, as one standard deviation in the number of traded strikes is 23 for calls and 39 for puts on the near-future. The number of puts and calls available has increased substantially in recent years, reaching a maximum number of strikes available on the near-future of 185 and 221 for calls and puts, respectively.

These options cover a wide range of degrees of moneyness, which is the ratio of an option's strike price to the underlying asset's current price. For the near-future, call moneyness ranges from an average daily minimum of 0.71 to an average daily maximum of 1.37. The average daily maximum and minimum moneyness of near-future puts, far-future calls, and far-future puts is similar. This indicates that market participants can,

on average, hedge against price changes between a roughly 30 percent decrease and 40 percent increase in the price of oil during the next 90 days.

Lastly, the minimum and maximum moneyness in the sample indicates market participants can sometimes hedge against much greater changes in the price of oil. For instance, the minimum degree of moneyness for puts is 0.36, which implies that on at least one day in our sample agents pay to hedge the risk of a 64 percent decrease in oil prices over the next 90 days. The maximum degree of moneyness for calls is also quite high at 2.90.

2.3 The trade-off between pricing errors and PDF smoothness

As explained in section 2.1, there are two key choices in the estimation of options-implied PDFs. In this section, we investigate the curve-fitting method used to obtain the implied volatility curve within the range of observed options. In the following section, we explore alternative methods used to extrapolate the unobserved tails.

We compare four curve-fitting methods: the cubic spline, the quartic spline, the smoothing spline, and the one-knot spline.⁷ Each spline interpolates the implied volatility curve using piecewise polynomials and is required to have continuous first and second derivatives everywhere, including at the knot points where the polynomial pieces connect. Figure 2 shows the PDFs for June 1, 2012, generated using each of the four curve-fitting methods.

Unlike options on many assets, WTI futures options are actively traded at many strikes (see table 1). On the one hand, a simple cubic or quartic spline can minimize pricing errors by fitting a spline that passes through every observed data point. However, the translation of such a fitted implied volatility curve into a call price function

⁷Our cubic and smoothing splines both implement the boundary conditions of the "natural" cubic spline, in which the second derivative is set to zero at the boundaries.

and a PDF can result in an extremely jagged distribution. On the other hand, we can relax the requirement of passing through every data point to obtain a smoother distribution. For example, our smoothing spline minimizes the sum of the squared distances between the curve and the data points, and maximizes the smoothness of the function, which is measured by the inverse of the sum of the squared numerical second derivative.⁸ Our one-knot spline generates an even smoother distribution, by fitting a cubic spline that minimizes the sum of the squared distances between the curve and the data points and limiting the piecewise polynomial to just one knot. These functions result in smoother distributions that are easier to interpret, but generate greater pricing errors, since the fitted implied volatility curves are not forced to go through the observed data points.

Table 2 summarizes the results for the trade-off between pricing errors and smoothness for the four curve-fitting methods. We first discuss the pricing errors. Each day between January 2000 and December 2013, we calculate the root mean squared error (RMSE) over all the observed options prices at the near and far maturities. The simple cubic and quartic splines yield null pricing errors as the implied volatility curve goes through all observed data points. The median of the daily RMSE values, in column A, is 0.75 cents for the smoothing spline and 2.63 cents for the one-knot spline. Moreover, the one-knot spline does a particularly poor job on the days with the worst-fitting curve, since the maximum of the daily RMSE values for this method is 76.18 cents.

To get an idea of the magnitude of these pricing errors, and because the RMSE is likely to overweight the pricing error of at-the-money options, we also calculate the mean of the absolute percentage errors (APE) for each day. The median value over all the days, reported in column C, is 0.34 percent for the smoothing spline and 1.27 percent for the one-knot spline. Finally, to reflect the worst pricing error on each

⁸In appendix A, we discuss the specification of this dual objective function in more detail.

day, column D reports the median of daily maximum absolute percentage errors. The one-knot spline's value is 8.82 percent and the smoothing spline's value is 2.33.

While cubic and quartic splines result in lower pricing errors, they generate less smooth PDFs. The average smoothness, measured as the inverse of the sum of the squared second derivative of the 90-day PDF, in column E and reported in logs, is greater for the smoothing spline (8.77) than for simple cubic and quartic splines (2.84 and 3.07, respectively). One-knot spline PDFs are, on average, even smoother than smoothing spline PDFs (13.62). The smoothness of PDFs obtained using these methods is illustrated in figure 2. Fitting implied volatility curves using cubic and quartic splines generates extremely jagged distributions (smoothness of 3.46 and 3.84 for the cubic and quartic splines, respectively) not only on this day, but on almost every day in the sample. In contrast, PDFs calculated using smoothing or one-knot splines are smoother (17.16 and 18.40, respectively) and, therefore, easier to interpret. Nevertheless, the one-knot spline PDFs come at the cost of much greater pricing errors, as can be seen in the bottom right panel of figure 2.

In sum, our results suggest that, for the case of highly traded WTI options, the smoothing spline method allows for convenient visualization of the PDFs while keeping pricing errors low.

2.4 Functional form of unobserved tails

The final step in generating a meaningful probability distribution is to make some assumptions on the functional form of the unobserved tails of the distribution. This step is required because, as long as options prices for the minimum and maximum strike price are not zero, the probabilities over the range of observed strike prices will not add

up to one.⁹ In this section, we consider the following five functional forms to estimate the unobserved tails of the distribution: an exponential function, a power function, a flat or zero-order polynomial function, a linear or first-order polynomial function, and second-order polynomial function. The exponential and power functions are attached directly to the estimated PDF in probability-strike space, and the polynomial functions are attached to the fitted implied volatility curve in volatility-strike space.

Since the tail region, by definition, is the region in which we have no observed call options prices, we cannot generate a measure of fit for call options prices. However, as noted in section 2.2, we often observe put options prices outside of the range of observed calls, especially put options that hedge against severe drops in oil prices. We generate put pricing errors by comparing the full set of observed put prices to those implied by the options-implied PDF.

Table 3 reports a summary of these errors for all possible combinations of curve-fitting and tail-fitting methods. For almost all curve-fitting methods, the median of the daily root mean square pricing errors, reported in column A, is minimized by the first-order polynomial tail-fitting method (3.19, 2.87, and 6.37 cents for the cubic, smoothing and one-knot spline methods, respectively).¹⁰ Relative pricing errors in column C are also generally smallest for the first-order polynomial method (2.98, 3.01, and 5.38 percent for the cubic, smoothing, and one-knot spline, respectively). The first-order polynomial method performs best for all splines, except the quartic spline, based on maximum relative pricing errors in column D. The average maximum pricing errors for the first-order polynomial tail-fitting method are 30.85, 29.99, and 46.76 percent for the cubic, smoothing, and one-knot splines, respectively.

⁹An important related choice is how far to extrapolate beyond the last observed strike. We explore alternative maximum extrapolation ranges in appendix B.

¹⁰Note that the first- and second-order extrapolations are the same for the cubic and smoothing splines, because we implement the "natural" spline boundary conditions in these methods. Specifically, because the second derivative is set to zero at the boundaries, extrapolation using first- and second-order polynomials will both result in first-order polynomial extrapolation.

In sum, while the best choice of functional form for the unobserved tail can depend on the choice of the curve-fitting method, the first-order polynomial generally provides the lowest put pricing errors.

3 Robust options-implied moments

Once we have estimated the shape of the options-implied PDF, we can obtain summary measures including options-implied moments and measures of how much investors are willing to pay to hedge against large changes in oil prices. In this section, we investigate the stability of these types of statistics. In particular, we analyze which methods yield moments that are more likely to reflect meaningful underlying daily market movements rather than data anomalies. We also investigate which moments are less sensitive to the curve- and tail-fitting methods used. Finally, we analyze the sensitivity of options-implied moments to small price perturbations.

We first examine the central moments of the distribution of returns on oil futures. To calculate these, we convert our options-implied PDF from strike-probability space into return-probability space by using the futures price and annualizing returns. In addition to using the standard measures of volatility, skewness, and kurtosis, we also calculate the comparable quantile moments.¹¹ Figure 3 shows the moments implied by the daily PDFs for the four curve-fitting methods described in section 2.3 using first-order polynomial tails. Although we calculate these moments over the entire sample, we illustrate them here for only the final two years of our sample.

¹¹Quantile volatility is calculated as the difference between the returns at the 75th and 25th percentiles of the options-implied PDF. Quantile skewness is calculated as the difference between the returns at the 75th and 50th percentiles, less the difference between the returns at the 50th and 25th percentiles, with this difference divided by the difference between the returns at the 75th and 25th percentiles. Finally, quantile kurtosis is calculated as the difference between the returns at the 95th and 5th percentiles, divided by the difference between the returns at the 75th and 25th percentiles, less a coefficient adjustment.

One noticeable feature of the plots in figure 3 is the large amount of daily variation. The high variability and large number of outliers may not be representative of the underlying state of the market; instead, they likely represent sensitivity of these measures to outliers in the underlying data. In particular, we find that moments calculated from cubic and quartic spline volatility curves are more prone to outliers, precisely because of the jaggedness of these methods' distributions.

Especially when abstracting from these outliers, the options-implied moments for all methods are highly correlated. This is particularly true for lower-order moments, such as volatility (panels A1 and A2). After trimming the series for outliers, the average pairwise correlation across all four curve-fitting methods is 0.85 for the standard volatilities and 0.96 for the quantile volatilities over the entire sample period. While still highly correlated, the levels of the higher-order moments tend to decouple, as can be seen in the charts for skewness (panels B1 and B2) and kurtosis (panels C1 and C2). The average correlation is 0.48 across the skewness series, but is a bit higher (0.67) across the quantile skewness series. For the kurtosis series, the average correlation is 0.43, while the average correlation across all methods' quantile skewness series is 0.66. The evidence also suggests that, at least for the smoothing and one-knot splines, the quantile moments are less dependent on the method than the standard moments. In other words, the level and dynamics of quantile moments are more stable across methods, likely because quantile moments rely on only a few data points of the PDF instead of the whole distribution.

Turning to the study of extreme events, figure 4 shows the cost of insurance implied by each curve-fitting method. The cost of insurance is calculated as the price of a binary option that pays \$1 if the futures price drops or rises a given percent or more in the next 90 days, and zero otherwise. Panels A, B, and C illustrate the cost of insurance against 15, 25, and 30 percent changes in the futures price, respectively. The cost of insurance is

only moderately affected by the method considered, and the difference among the cost of insurance measures across methods increases with the magnitude of the percentage change considered. However, if we exclude the obvious outliers, it seems that the cost of insurance is more similar across methods than the quantile moments, and provides summary information that is more robust to the use of different curve-fitting methods. Across all four methods, the average pairwise correlation for each of the six cost of insurance series plotted is above 0.90.

Finally, for the oil market, we are often interested in the likelihood of the oil price crossing a particular threshold level. To study this type of extreme event, we define the cost of insurance against reaching a benchmark level as the price of a binary option that pays \$1 if the futures price rises to or above a given price in the next 90 days, and zero otherwise. In figure 5, we report the cost of insurance against the \$120 and \$135 benchmark levels. The figures suggest that the cost of insurance against reaching a certain price threshold is little affected by the method considered, except for a few days of outliers. After removing the outliers, the correlation between the cost of insurance across all methods is above 0.97.

While the options-implied moments and cost of insurance measures are sensitive to the curve-fitting method, they are less sensitive to the choice of tail-fitting method. This is particularly true when using the smoothing spline, as shown in the next set of figures. In figure 6, we can see very little difference in the quantile moments when using the smoothing spline with various tail-fitting methods.¹² In figures 7 and 8, the tail-fitting method only slightly affects the cost of insurance against extreme events. As expected, the cost of insurance differs more across tail-fitting methods for more extreme events, as moments calculated using information from the range of observed options prices will generally be more robust to the curve- and tail-fitting methods used.

¹²Recall that for the smoothing spline, the first- and second-order extrapolation methods are equivalent, as mentioned in footnote 10.

We now explore the sensitivity of options-implied moments to price perturbations. To do so, we follow the Monte Carlo procedure proposed in Bliss and Panigirtzoglou (2002) to examine the robustness of options-implied PDFs. At each point in time, we randomly perturb observed prices using a uniformly distributed perturbation between plus and minus one half of the contract's tick size, or minimum price fluctuation. For WTI futures options, the tick size is \$0.01. We repeat the perturbation procedure 500 times and, for each simulation, calculate the options-implied moments. Thus, we are able to obtain confidence intervals for each moment.

We show for the smoothing spline method the options-implied quantile moments and their confidence intervals in figure 9. For the smoothing spline method, quantile volatility is hardly sensitive to price perturbations. The sensitivity to price perturbations increases with the order of the moment, although the confidence intervals are still relatively narrow for skewness and kurtosis. A similar pattern holds for the confidence intervals for the cost of insurance. In unreported results, we find that confidence intervals are extremely narrow for the cost of insurance against 15 and 20 percent price changes. Although they widen slightly for more extreme events (30 percent price changes), they remain quite narrow, as can be seen in figure 10. The cost of insurance against reaching high price levels (panels C and D) is not very sensitive to price perturbations, although, as expected, the confidence intervals widen slightly for more extreme prices.

4 Options-implied PDFs around important market events

In this section, we use our estimated PDFs to analyze the changes in market participants' beliefs about oil price movements around important market events. Specifically,

we study the interruption of Libyan production in February 2011, the increasing geopolitical tensions with Iran in February 2012, and the negative macroeconomic surprises in May 2012. At each point in time, the estimated 90-day PDF reflects market participants' beliefs about the price of the oil futures contract on its expiration date. In general, the the futures price should be equal to the spot price on this expiration date, save any transaction or transportation costs. Consequently, we interpret the PDF as reflecting market participants' beliefs about the future spot price.

By examining the information extracted from options-implied PDFs, we can investigate whether an observed price movement reflects the realization of market expectations or an unexpected market shift. For example, an unexpected event that prompts a sharp price increase would not only prompt an increase in the mean, but would also likely shift the entire distribution. Since the shape of the distribution is unchanged, the cost of insurance against a percent increase in prices would not change. However, the cost of insurance against reaching a \$120 or \$135 price level would increase. In contrast, consider a similarly sharp price increase prompted by an event that was somewhat expected by market participants. In this case, while the futures price would increase, the cost of insurance against reaching a benchmark price level could remain constant. Furthermore, the cost of insurance against an increase in prices could actually fall as the realized price increase reflects the resolution of a previously uncertain possibility. Similar analysis can help us assess whether agents believe the observed movements in oil prices are persistent or temporary. Contrasting the changes in our full set of options-implied information can therefore contribute to a more nuanced understanding of oil price beliefs around major events.

Case 1: February 2011

Oil price uncertainty rises with interruption of Libyan production

From mid-February through late April of 2011, oil prices trended upward as conflict broke out and escalated in Libya. Local peaks in prices (visible in figure 11) coincided with news related to increasing probability of lasting interruption to Libyan oil supply, including the first reports of interruptions in production (February 23), the halting of exports (March 7), and the sabotaging of oil fields (April 8). These price increases in February and March were accompanied by increasing volatility, skewness, and kurtosis, as well as increases in the cost of insurance against further price movements (in both directions) and against reaching the benchmark price levels of \$120 and \$135.

In contrast, as prices reached almost \$115 in early April, kurtosis and cost of insurance against further price movements declined, even as the cost of insurance against reaching the benchmark price levels remained elevated due to the high mean of the distribution. The decline in the cost of insurance against a 15 percent increase in prices is consistent with the view that market prices had fully adjusted to the disruption of Libyan oil production, to the point that further price increases were seen as less likely. At the same time, the early April decline in the cost of insurance against a 15 percent decrease in prices indicated that market participants seemed to think that a rapid resumption of Libyan production and export was also not very likely.

Case 2: February 2012

Oil prices spike amid threats of conflict with Iran

Oil PDFs can also help us understand how increasing geopolitical tensions can affect the oil market even before any actual supply disruption occurs. For example, in February 2012, the likelihood of conflict with Iran seemed to have risen. According to news reports around this time, there were fears that Israel would bomb Iran, and that Iran would block oil shipments through the Strait of Hormuz. Oil prices increased over \$10 per barrel over three weeks, with spot prices peaking near \$110 by late February.

While the direct effects of these events on the oil spot and futures prices are easily measured, the PDFs and options-implied measures in figure 12 tell us more about the beliefs surrounding future prices. As tensions with Iran rose steadily through early February, the probability of an oil price increase decreased slightly. When the oil price climbed sharply on February 24, the cost of insurance against a further 15 percent increase dropped, indicating that the probability of increase was likely driven by the tensions with Iran. Once oil prices increased, these potential events had been realized, and the probability of further increases fell.

Case 3: May 2012

Oil prices fall amid negative data surprises

In some instances, the options-implied PDFs can clarify market expectations surrounding sharp declines in prices. One such example is the price decline in the first week of May 2012. On May 2, it was announced that after six weeks of rising, oil inventories were at their highest level in 20 years. The next day, OPEC announced that it wanted to increase oil production to bring prices down to \$100. Finally, on May 4, the employment data release in the United States was interpreted as providing evidence that the labor market was much weaker than previously thought. Cumulatively, the spot price of oil decreased by almost \$8 over these three days.

Figure 13 presents some analysis of the change in PDFs over this period. While volatility declined after the report of high inventories, the news of both potentially higher supply and lower demand increased volatility in the days following. The cost of insurance against a further 15 percent decline slightly increased, indicating that participants believed that some negative possibilities remained, even as all of these events were realized. At the same time, the higher volatility also contributed to an increase in the cost of insurance against a future increase in prices. Finally, driven

largely by the negative mean-shift in the distribution, the news about high inventory levels and potentially higher OPEC production sharply reduced the cost of insurance against a reaching oil prices of \$120.

Although we have presented a few case studies that focus on rapid and unexpected changes in market sentiments, the time series of options-implied moments can also be used to identify slower evolutions in market beliefs, including market reactions to events that unfold over time. Another application of these methods is to compare across similar events. For example, we can compare PDFs from different instances of unexpected cuts in production. Even more generally, the PDFs and moments developed in this paper can be studied in relation to other oil market signals, including data on crude oil inventories, production, and demand. Finally, when applied to data in real time, these methods are potentially useful for predicting future price movements as well.

5 Conclusion

Motivated by the thin literature on the robustness and stability of options-implied PDF methods, in this paper, we explore alternative methods to calculate oil PDFs using options on WTI crude oil futures contracts. Using these PDFs, we analyze the changes in market participants' beliefs about oil price movements around a few important events.

We find that using the smoothing spline method to fit the options-implied volatility curve allows for convenient visualization of the estimated PDFs, while keeping pricing errors low. Other traditional spline methods yield, by construction, null pricing errors at the cost of extremely jagged distributions. In contrast, extremely smoothed methods, such as the one-knot spline, yield extremely large pricing errors. We also find that using a first-order polynomial to extrapolate the unobserved tails of the distribution yields, on

average, the lowest pricing errors across all curve-fitting methods. Combining all curve- and tail-fitting methods, we find that higher-order options-implied moments, such as skewness and kurtosis, are more sensitive to the method selected and to small price perturbations. However, measures of quantile moments and cost of insurance against extreme events are more robust to price perturbations and across methods.

Having established confidence in the smoothing spline method with first-order polynomial tails, we apply this method to options data from days around important market events. Further, using a set of moments shown to be relatively robust, we examine market participants' beliefs about oil price movements around important market events. We find that when the oil price increases, the cost of insurance against a further price increase sometimes increases (when the realized price increase is either unexpected or seen as foreshadowing further increases), and sometimes decreases (when the realized price increase is seen as temporary). Examining the PDFs therefore allows us to gain a more nuanced view of market beliefs before and after events, and helps us better understand realized market reactions.

References

- ABADIR, K. M., AND M. ROCKINGER (2003): “Density functionals, with an option-pricing application,” *Econometric Theory*, 19, 778–811.
- ÄIJÖ, J. (2008): “Impact of US and UK macroeconomic news announcements on the return distribution implied by FTSE-100 index options,” *International Review of Financial Analysis*, 17(2), 242–258.
- AÏT-SAHALIA, Y., AND J. DUARTE (2003): “Nonparametric option pricing under shape restrictions,” *Journal of Econometrics*, 116(1), 9–47.
- AÏT-SAHALIA, Y., AND A. LO (1998): “Nonparametric estimation of state-price densities implicit in financial asset prices,” *The Journal of Finance*, 53(2), 499–547.
- ALQUIST, R., L. KILIAN, AND R. VIGFUSSON (2012): “Forecasting the Price of Oil,” *ECB Working Paper Series No. 198*.
- ANDERSEN, A., AND T. WAGENER (2002): “Extracting risk neutral probability densities by fitting implied volatility smiles: Some methodological points and an application to the 3M Euribor futures options prices,” *ECB Working Paper Series No. 198*.
- ANDERSSON, M., AND M. LOMAKKA (2005): “Evaluating implied RNDs by some new confidence interval estimation techniques,” *Journal of Banking & Finance*, 29(6), 1535–1557.
- ASKARI, H., AND N. KRICHENE (2008): “Oil price dynamics (2002-2006),” *Energy Economics*, 30(5), 2134–2153.
- BAHRA, B. (1996): “Probability distributions of future asset prices implied by option prices,” *Bank of England Quarterly Bulletin*, 36(3), 299–311.
- (1997): “Implied risk-neutral probability density functions from option prices: theory and application,” *Bank of England Working Paper*.
- BAKSHI, G., AND N. KAPADIA (2003): “Delta-hedged gains and the negative market volatility risk premium,” *Review of Financial Studies*, 16(2), 527–566.
- BARONE-ADESI, G., AND R. WHALEY (1987): “Efficient analytic approximation of American option values,” *Journal of Finance*, 42, 301–201.
- BEKAERT, G., AND E. ENGSTROM (2010): “Asset Return Dynamics Under Bad Environment-Good Environment Fundamentals,” *CEPR Discussion Paper No. DP8150*.
- BERKOWITZ, J. (2001): “Testing density forecasts, with applications to risk management,” *Journal of Business & Economic Statistics*, 19(4), 465–474.

- BIRRU, J., AND S. FIGLEWSKI (2011): “Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008,” *Journal of Financial Markets*, 15(2), 151–180.
- BLACK, F. (1978): “The pricing of commodity contracts,” *Journal of Financial Economics*, 3, 167–179.
- BLISS, R., AND N. PANIGIRTZOGLU (2002): “Testing the Stability of Implied Probability Density Functions,” *Journal of Banking and Finance*, 26, 381–422.
- BOLLERSLEV, T., G. TAUCHEN, AND H. ZHOU (2009): “Expected stock returns and variance risk premia,” *Review of Financial Studies*, 22(11), 4463–4492.
- BREEDEN, D., AND R. LITZENBERGER (1978): “Prices of state-contingent claims implicit in option prices,” *Journal of Business*, pp. 621–651.
- BU, R., AND K. HADRI (2007): “Estimating option implied risk-neutral densities using spline and hypergeometric functions,” *The Econometrics Journal*, 10(2), 216–244.
- CAMPA, J., P. CHANG, AND J. REFALO (2002): “An options-based analysis of emerging market exchange rate expectations: Brazil’s *Real Plan*, 1994–1999,” *Journal of Development Economics*, 69(1), 227–253.
- CHAUDHURY, M., AND J. WEI (1994): “Upper bounds for American Futures Options: A Note,” *Journal of Futures Markets*, 14, 111–116.
- CHENG, K. (2010): “A New Framework to Estimate the Risk-Neutral Probability Density Functions Embedded in Options Prices,” *IMF Working Papers*, pp. 1–31.
- CINCIBUCH, M. (2003): “Distribution implied by american currency futures options: a ghost’s smile,” *Journal of Futures Markets*, 24(2), 147–178.
- COCHRANE, J. H. (2005): *Asset Pricing (Revised Edition)*. Princeton University Press, Princeton, N.J.
- COOPER, N. (1999): “Testing techniques for estimating implied RNDs from the prices of european options,” *Working Paper, Bank of England*.
- DUMAS, B., J. FLEMING, AND R. WHALEY (1998): “Implied volatility functions: empirical tests,” *Journal of Finance*, 53, 1269–1300.
- FACKER, P. L., AND R. KING (1990): “Calibration of option-based probability assessments in agricultural commodity markets,” *Journal of Agricultural Economics*, 72(1), 73–83.
- FLAMOURIS, D., AND D. GIAMOURIDIS (2002): “Estimating implied PDFs from American options on futures: a new semiparametric approach,” *Journal of Futures Markets*, 22(1), 1–30.

- GAI, P., AND N. VAUSE (2006): “Measuring Investors’ Risk Appetite,” *International Journal of Central Banking*, 2, 167–188.
- GEREBEN, Á. (2002): “Extracting market expectations from option prices: an application to over-the-counter New Zealand dollar options,” *Reserve Bank of New Zealand Discussion Paper No. DP2002/04*.
- GLATZER, E., AND M. SCHEICHER (2005): “What moves the tail? The determinants of the option-implied probability density function of the DAX index,” *Journal of Futures Markets*, 25(6), 515–536.
- JACKWERTH, J. (1999): “Option implied risk-neutral distributions and implied binomial trees: a literature review,” *Journal of Derivatives*, 7, 66–82.
- JARROW, R., AND A. RUDD (1982): “Approximate option valuation for arbitrary stochastic processes,” *Journal of financial Economics*, 10(3), 347–369.
- JIANG, G., AND Y. TIAN (2005): “The model-free implied volatility and its information content,” *Review of Financial Studies*, 18(4), 1305–1342.
- JONDEAU, E., AND M. ROCKINGER (2000): “Reading the smile: the message conveyed by methods which infer risk neutral densities,” *Journal of International Money and Finance*, 19(6), 885–915.
- KANG, B., AND T. KIM (2006): “Option-implied risk preferences: An extension to wider classes of utility functions,” *Journal of Financial Markets*, 9(2), 180–198.
- KANG, B., T. KIM, AND S. YOON (2010): “Information content of volatility spreads,” *Journal of Futures Markets*, 30(6), 533–558.
- KIM, I., AND S. KIM (2003): “On the usefulness of implied risk-neutral distributions—evidence from the Korean KOSPI 200 Index options market,” *Journal of Risk*, 6, 93–110.
- KITSUL, Y., AND J. H. WRIGHT (2012): “The Economics of Options-Implied Inflation Probability Density Functions,” *NBER Working Paper 18195*.
- LONDONO, J., AND H. ZHOU (2012): “Variance Risk Premiums and the Forward Premium Puzzle,” *Available at SSRN 2133569*.
- MELICK, W., AND C. THOMAS (1992): “War and peace: recovering the market’s probability distribution of crude oil futures prices during the gulf crisis,” *Federal Reserve Board International Finance Discussion Paper 437*.
- (1997): “Recovering an asset’s implied PDF from option prices: an application to crude oil during the gulf crisis,” *Journal of Financial and Quantitative Analysis*, 32(1), 91–115.

- (1999): “Confidence intervals and constant-maturity series for probability measures extracted from options prices,” in *Information in Financial Asset Prices: Proceedings of a conference held by the Bank of Canada*, pp. 293–320. Bank of Canada, Ottawa, Canada.
- NEUHAUS, H. (1995): “The information content of derivatives for monetary policy,” *Bundesbank Monthly Bulletin*, 47(11), 17–32.
- POLKOVNICHENKO, V., AND F. ZHAO (forthcoming): “Empirical Pricing Kernels: Implications for Probability Weights,” *Journal of Financial Economics*.
- ROMPOLIS, L., AND E. TZAVALIS (2008): “Recovering risk neutral densities from option prices: A new approach,” *Journal of Financial and Quantitative Analysis*, 43(4), 1037.
- SHIMKO, D. (1993): “Bounds of Probability,” *RISK*, 6, 33–37.
- SHIRATSUKA, S. (2001): “Information Content of Implied Probability Distributions: Empirical Studies of Japanese Stock Price Index Options,” *Monetary and Economic Studies*, 19(3), 143–170.
- SODERLIND, P., AND L. SVENSSON (1997): “New techniques to extract market expectations from financial instruments,” *Journal of Monetary Economics*, 40, 383–429.
- SYRDAL, S. (2002): “A study of implied risk-neutral density functions in the Norwegian option market,” *Norges Bank Working Paper ANO 2002/13*.
- TODOROV, V. (2010): “Variance risk-premium dynamics: The role of jumps,” *Review of Financial Studies*, 23(1), 345–383.
- WEINBERG, S. (2001): “Interpreting the volatility smile: an examination of the information content of option prices,” *FRB International Finance Discussion Paper*, (706).

Table 1: Summary statistics for options on WTI futures contracts

This table shows a set of summary statistics for the availability of put and call options for the maturity just before (near future) and just after (far future) our benchmark fixed maturity of 90 days. In columns (A) and (B), we report the summary statistics for the total number of strikes available for calls and puts, respectively. In columns (C) to (F), we report the summary statistics for the minimum and maximum degree of moneyness (calculated as a ratio between the option's strike price and the underlying futures price) for calls and puts. Our sample period runs from January 2000 to December 2013.

	Calls (A)	Puts (B)	Min. Call (C)	Max. Call (D)	Min. Put (E)	Max. Put (F)
<i>Near Future</i>						
Mean	50.85	57.90	0.71	1.37	0.71	1.34
Standard Deviation	23.41	38.86	0.16	0.19	0.07	0.36
Minimum	13.00	9.00	0.11	1.04	0.36	0.97
Maximum	185.00	221.00	1.10	2.57	0.89	3.49
<i>Far Future</i>						
Mean	48.59	55.70	0.74	1.43	0.68	1.30
Standard Deviation	23.34	36.87	0.16	0.23	0.07	0.35
Minimum	13.00	9.00	0.10	1.01	0.40	0.94
Maximum	182.00	213.00	1.11	2.90	0.89	3.54

Table 2: Pricing errors and PDF smoothness

This table reports measures of fit in columns (A) to (D) and smoothness in columns (E) to (G) for the curve-fitting methods described in section 2.3. Each day between January 2000 and December 2013, we calculate the root mean squared error (RMSE) over all the observed options prices for the near- and far-future maturities. The median of these daily RMSE values in column (A), reported in cents, reflects the typical goodness-of-fit for each method, while the maximum of the daily RMSE values in column (B) reflects how poorly the data are fit on the worst days. We also calculate the absolute percentage error (APE) for each observed option price on each day between January 2000 and December 2013. Column (C) reports the median over all the daily mean absolute percentage errors reported in percentage points. For the final measure of fit, we calculate for each day the maximum absolute percentage error, which reflects the worst pricing error on each day. We report the median of these daily maximum absolute percentage errors in column (D). The final three columns present measures of smoothness of the 90-day PDF. For each PDF, smoothness is calculated as the inverse of the sum of the square of the PDF's numerical second derivative. Columns (E) to (G) report the mean, minimum, and maximum of smoothness over all daily smoothness measures, respectively. These measures are reported in logs for readability.

	Pricing Errors				Smoothness		
	RMSE (cents)		Mean APE	Max. APE	Mean	Min.	Max.
	Median	Max.	Median	Median			
Spline	(A)	(B)	(C)	(D)	(E)	(F)	(G)
Cubic	0.00	0.00	0.00	0.00	2.84	0.84	5.15
Quartic	0.00	0.00	0.00	0.00	3.07	0.15	6.25
Smoothing	0.75	1.85	0.34	2.33	8.77	2.82	22.19
One-knot	2.63	76.18	1.27	8.82	13.62	6.50	22.89

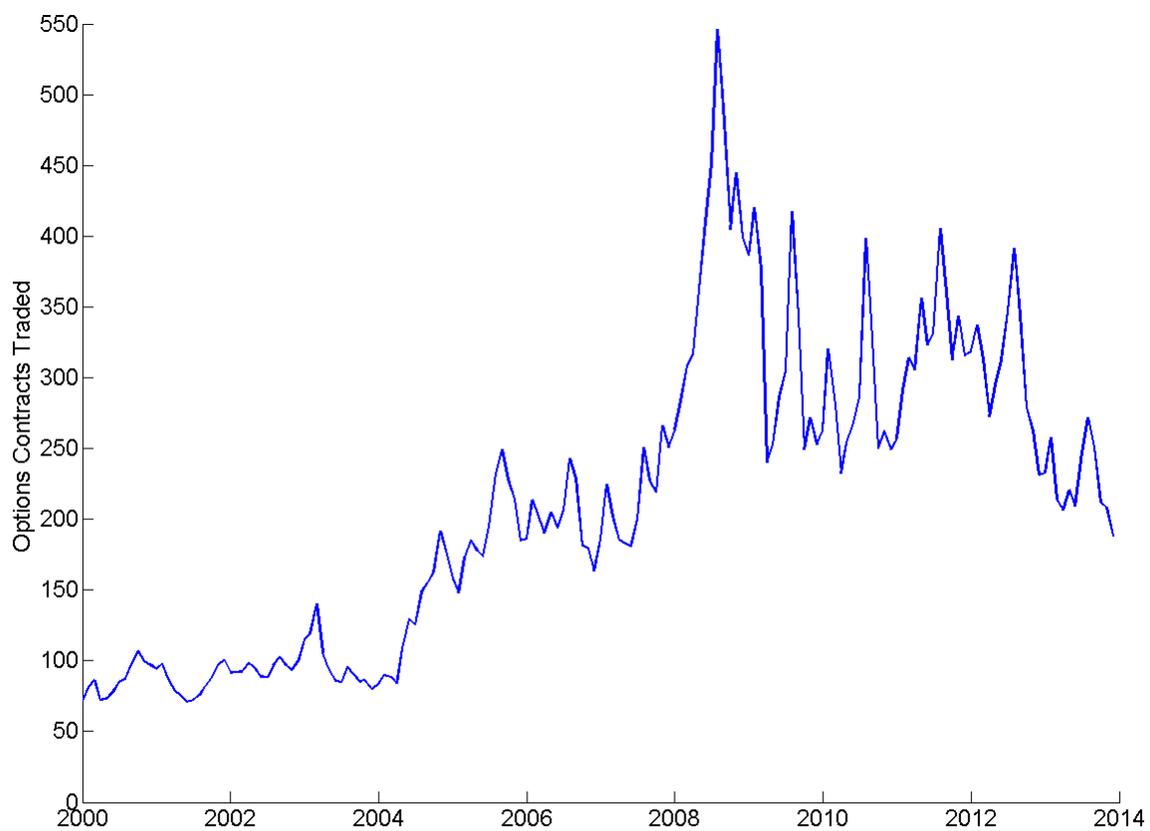
Table 3: Put pricing errors

This table reports put pricing errors for several curve- and tail-fitting methods as described in sections 2.3 and 2.4. For each method, we estimate daily PDFs between January 2000 and December 2013. Each day, we calculate the various pricing error statistics by comparing observed put prices to those mapped from the fitted options-implied volatility curve using the model in Black (1978). See table 2 for a description of how the pricing error summary statistics are calculated.

Method	RMSE (cents)		Mean APE	Max. APE
	Median (A)	Max. (B)	Median (C)	Mean (D)
<i>Cubic Spline</i>				
Exponential	3.76	393.58	4.10	56.73
Power	3.63	393.58	3.96	52.28
0-D polynomial	3.74	393.58	4.42	50.56
1-D polynomial	3.19	393.58	2.98	30.85
2-D polynomial	3.19	393.58	2.98	30.85
<i>Quartic Spline</i>				
Exponential	6.40	2027.16	7.05	94.73
Power	5.76	2027.16	6.39	89.65
0-D polynomial	4.57	2027.16	5.08	55.07
1-D polynomial	5.25	2027.16	5.63	68.30
2-D polynomial	16.22	3419.59	9.53	100.00
<i>Smoothing Spline</i>				
Exponential	2.90	393.57	3.03	32.98
Power	3.11	393.57	3.45	42.12
0-D polynomial	3.83	393.57	4.58	51.01
1-D polynomial	2.87	393.57	3.01	29.99
2-D polynomial	2.87	393.57	3.01	29.99
<i>One-knot Spline</i>				
Exponential	6.55	392.59	5.90	55.49
Power	6.72	392.59	6.10	59.69
0-D polynomial	6.43	392.59	6.02	53.05
1-D polynomial	6.37	392.59	5.38	46.76
2-D polynomial	6.75	392.59	5.83	53.35

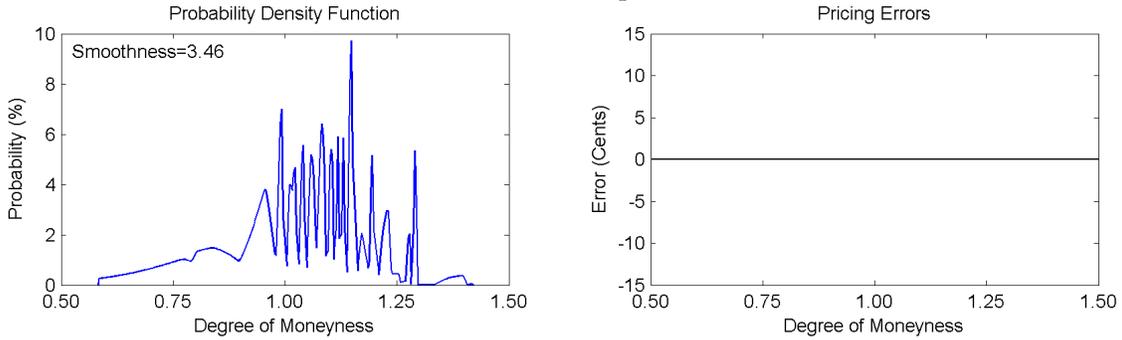
Figure 1: Availability of options on WTI futures contracts

This figure shows the monthly average of the number of American put and call options traded on the WTI futures maturing just before and just after 90 days over our sample period from January 2000 to December 2013.

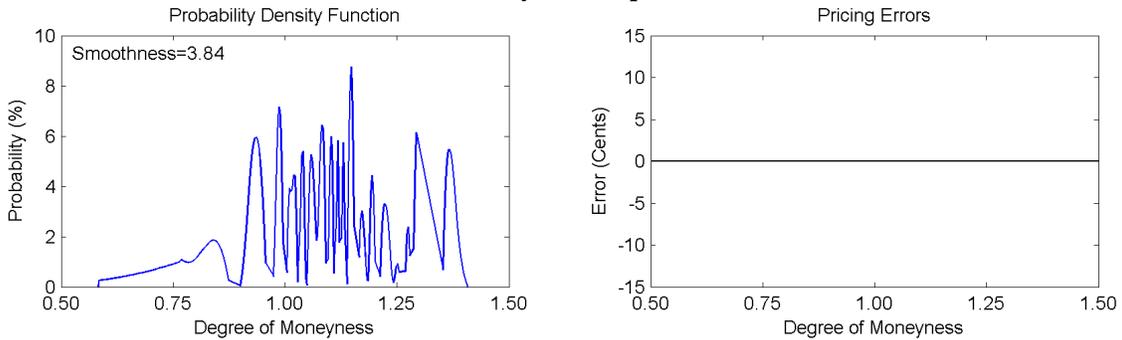


The figures on the left show the near-future PDFs for June 1, 2012 generated using the four curve-fitting methods explained in section 2.3. The smoothness is calculated as the inverse of the sum of the square of the PDF's numerical second derivative and is reported in logs for readability. The figures on the right show the pricing errors implied by the corresponding PDFs, calculated as the difference between the observed call prices and those mapped from the fitted implied volatility curve.

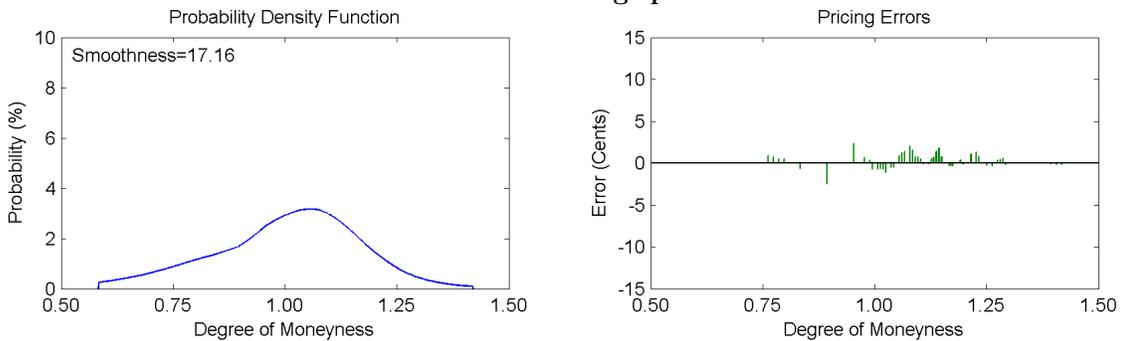
A. Cubic spline



B. Quartic spline



C. Smoothing spline



D. One-knot spline

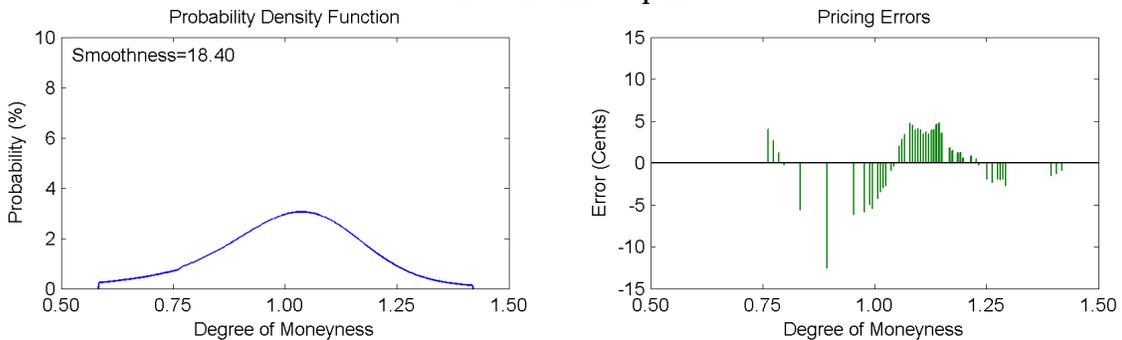
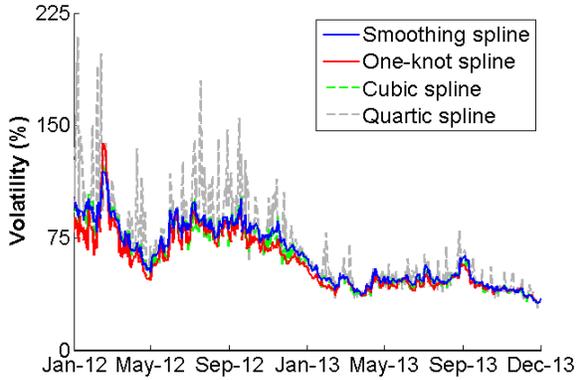


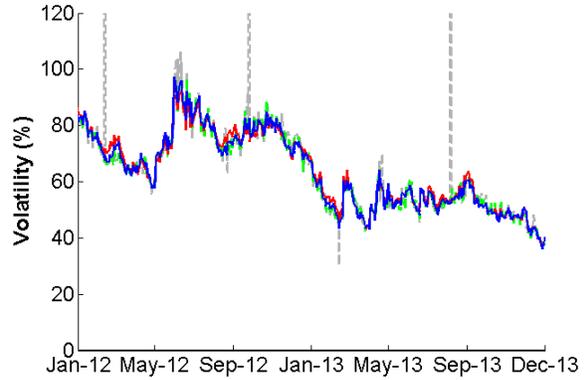
Figure 3: Quantile and standard moments

This figure shows the daily moments of the distribution for the four curve-fitting methods described in section 2.3 for the sample period January 2012 to December 2013. For all curve-fitting methods, we use the first-order polynomial tail-fitting method. The left panels illustrate the standard measures of the volatility, skewness, and kurtosis, while the right panels show the quantile moments calculated as described in section 3.

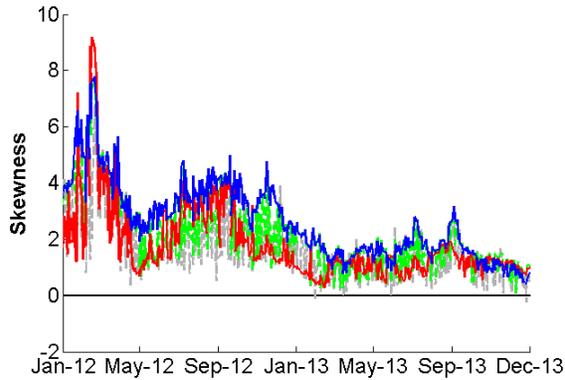
A1. Volatility



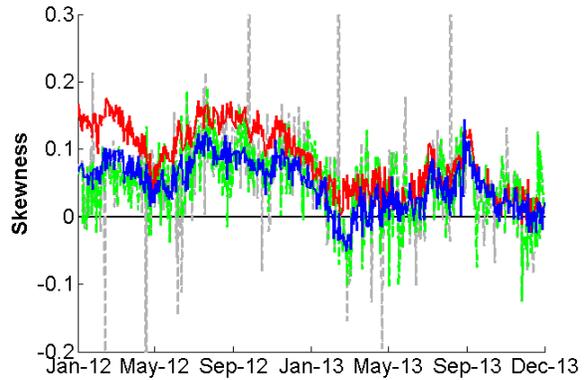
A2. Quantile volatility



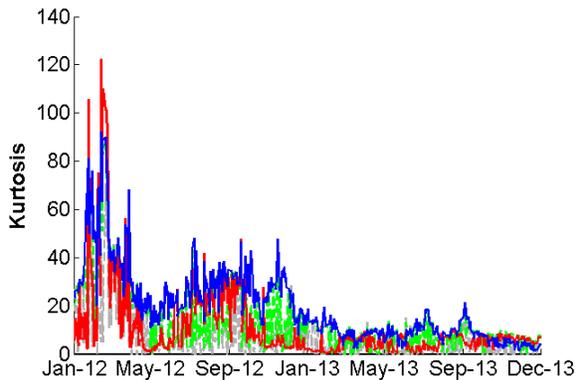
B1. Skewness



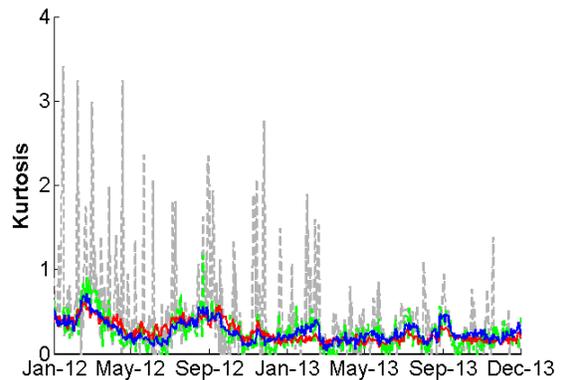
B2. Quantile skewness



C1. Kurtosis



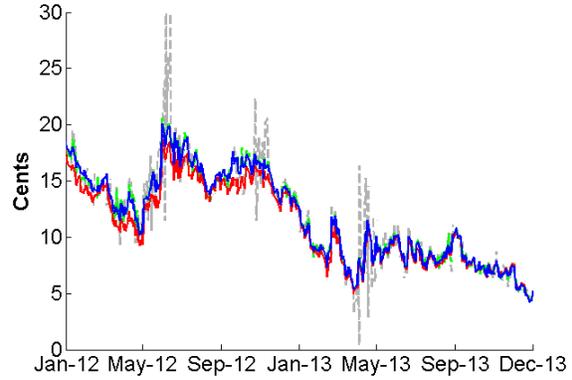
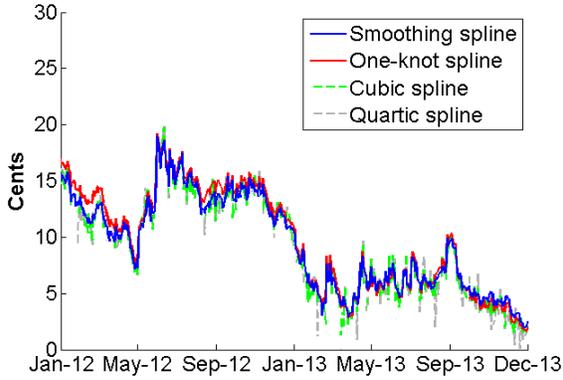
C2. Quantile kurtosis



This figure shows the cost of insurance against extreme price changes for the four curve-fitting methods described in section 2.3 for the sample period January 2012 to December 2013. The cost of insurance is calculated as the price of a binary option that pays \$1 if the futures price rises (left panels) or drops (right panels) a given percent or more in the next 90 days and zero otherwise.

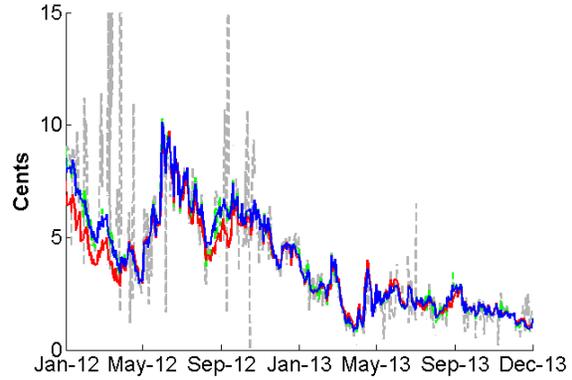
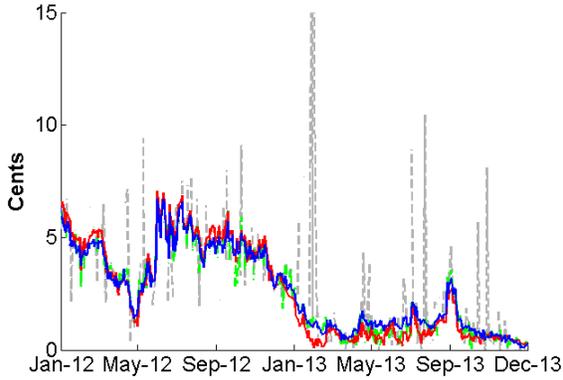
A1. 15 percent increase

A2. 15 percent decrease



B1. 25 percent increase

B2. 25 percent decrease



C1. 30 percent increase

C2. 30 percent decrease

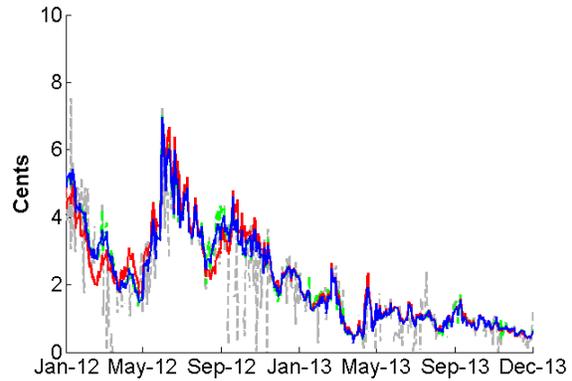
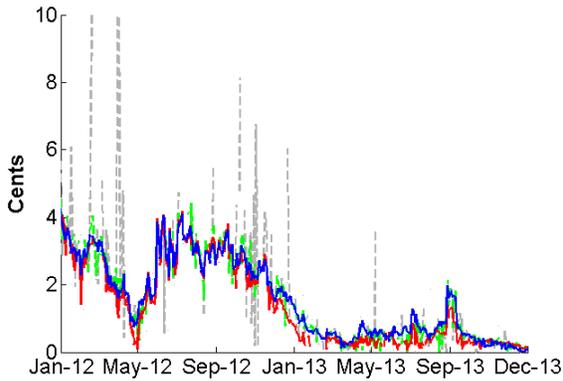
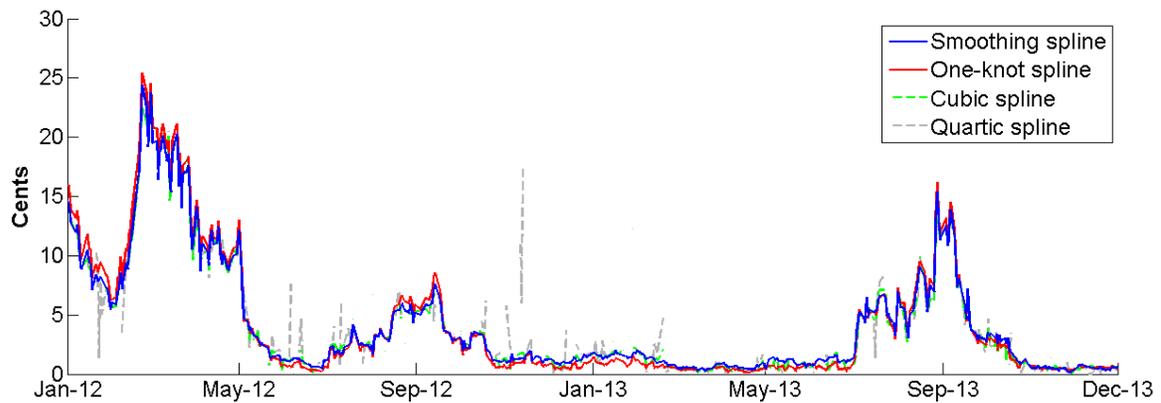


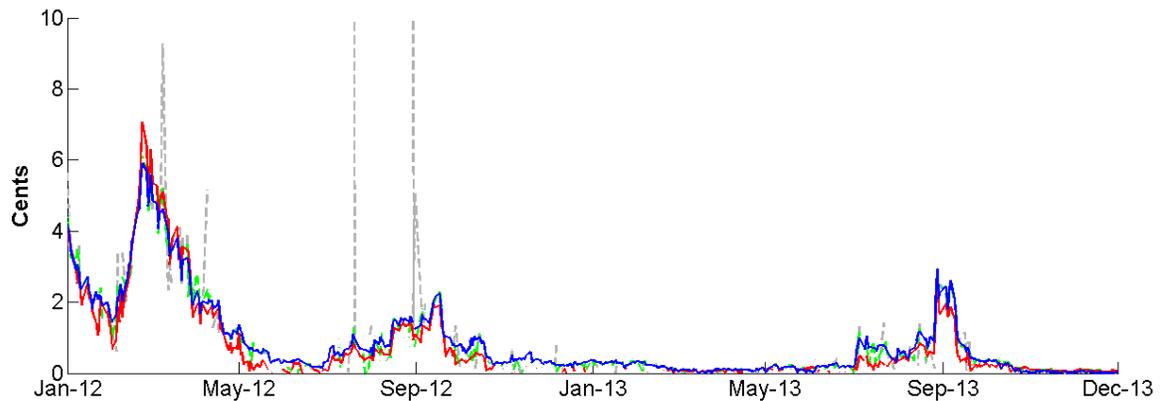
Figure 5: Cost of insurance against extreme price levels

This figure shows the cost of insurance against reaching extreme price levels for the four curve-fitting methods described in section 2.3 for the sample period January 2012 to December 2013. The cost of insurance is calculated as the price of a binary option that pays \$1 if the futures price rises to a given level in the next 90 days and zero otherwise.

A. Crude oil price of \$120



B. Crude oil price of \$135

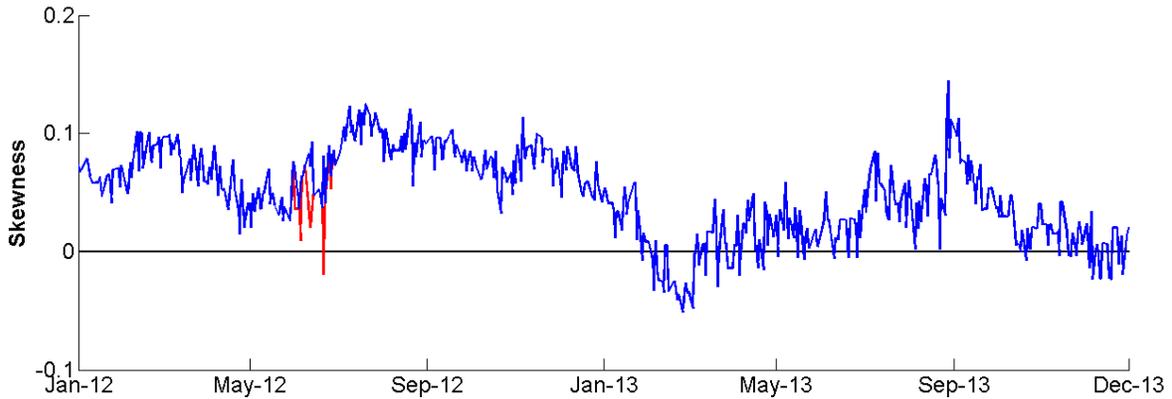


This figure shows the daily quantile moments of the PDF estimated using the smoothing spline for the five tail-fitting methods described in section 2.4 for the sample period January 2012 to December 2013. The quantile moments are calculated as described in section 3.

A. Quantile volatility



B. Quantile skewness



C. Quantile kurtosis

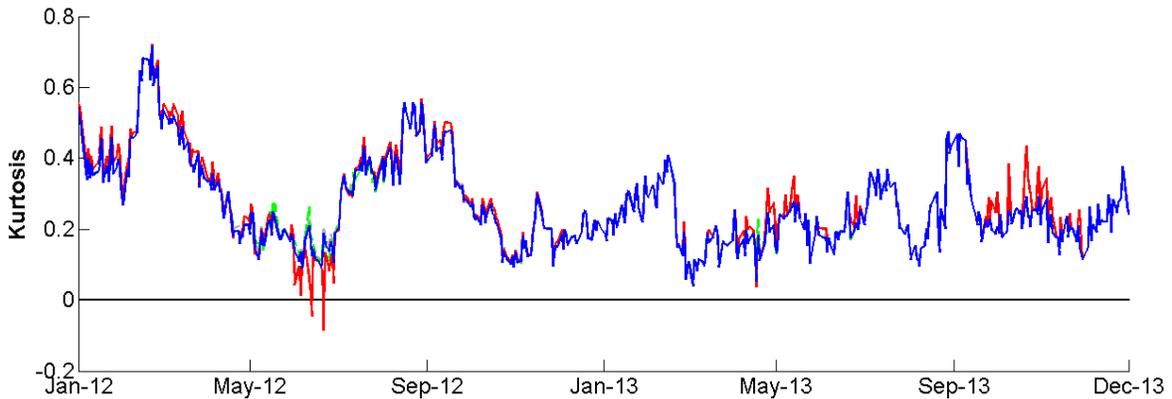


Figure 7: Cost of insurance against extreme price changes: Comparing tail-fitting methods

This figure shows the cost of insurance against extreme price changes for the five tail-fitting methods described in section 2.4 for the sample period January 2012 to December 2013. The cost of insurance is calculated as the price of a binary option that pays \$1 if the futures price rises (left panels) or drops (right panels) a given percent or more in the next 90 days and zero otherwise.

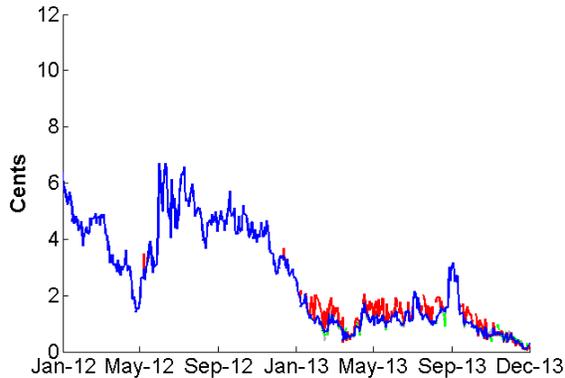
A1. 15 percent increase



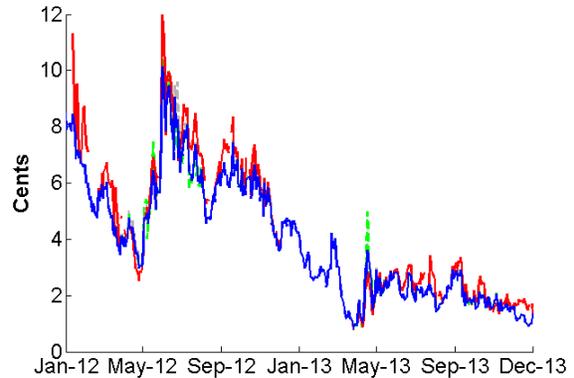
A2. 15 percent decrease



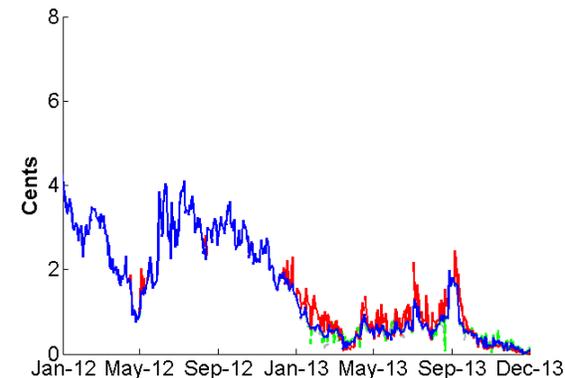
B1. 25 percent increase



B2. 25 percent decrease



C1. 30 percent increase



C2. 30 percent decrease

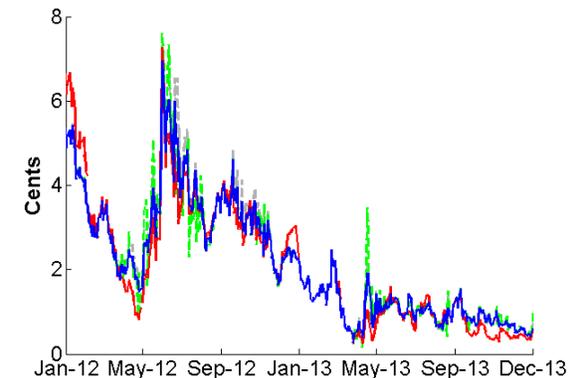
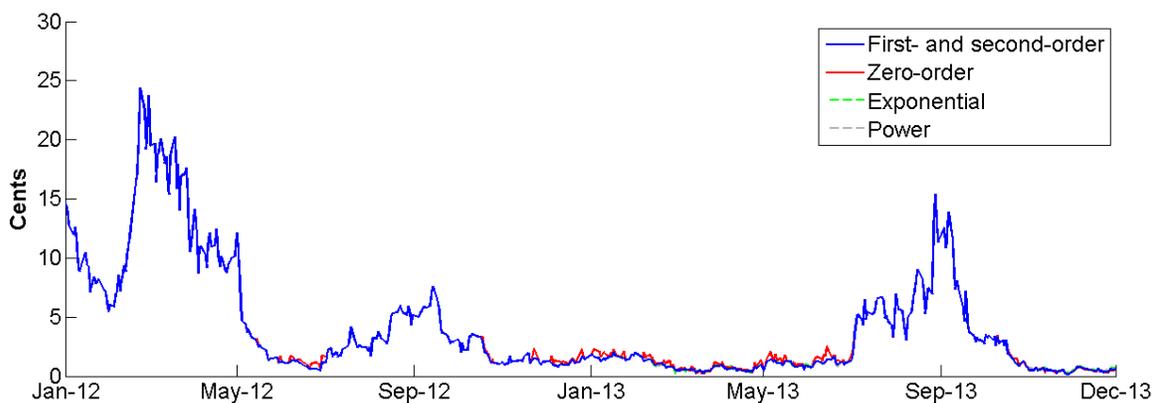


Figure 8: Cost of insurance against extreme price levels: Comparing tail-fitting methods

This figure shows the cost of insurance against reaching extreme price levels for the five tail-fitting methods described in section 2.4 for the sample period January 2012 to December 2013. The cost of insurance is calculated as the price of a binary option that pays \$1 if the futures price rises to a given level in the next 90 days and zero otherwise.

A. Crude oil price of \$120



B. Crude oil price of \$135

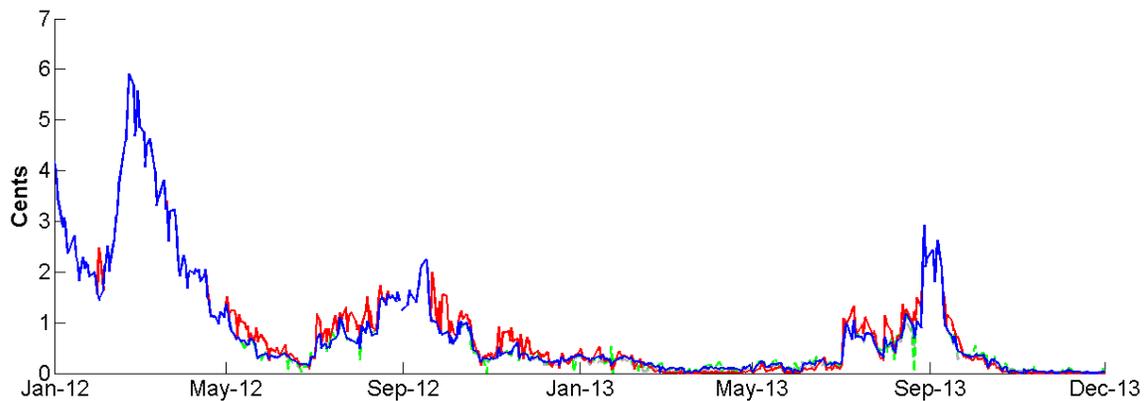
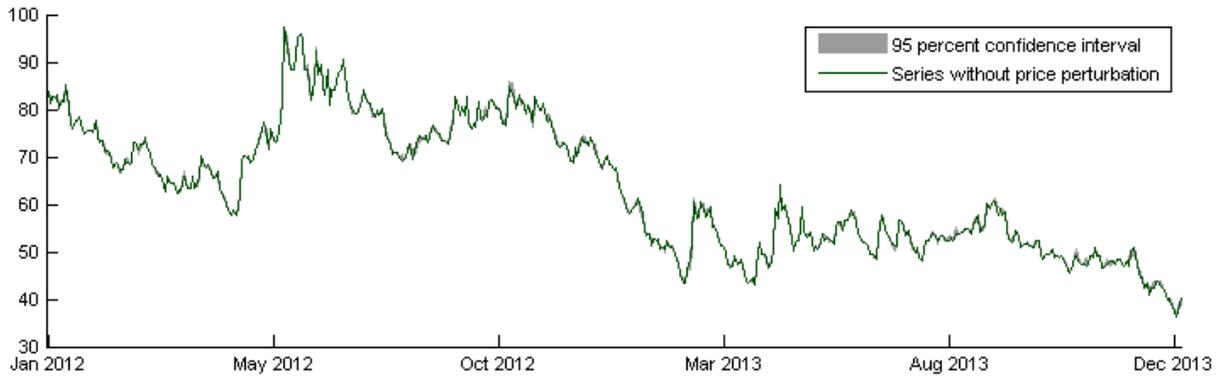


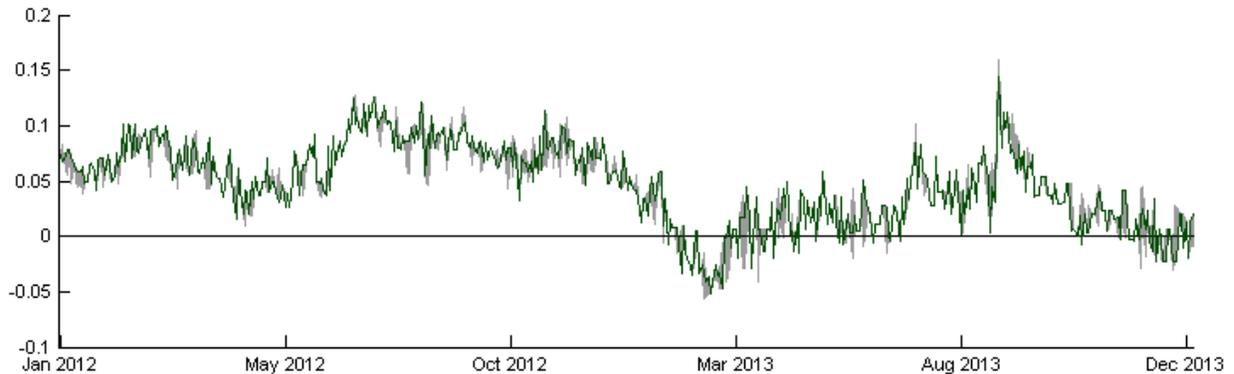
Figure 9: Confidence intervals for quantile moments

This figure shows the quantile moments calculated as described in section 3 and their corresponding 95 percent confidence intervals. The confidence intervals are calculated using the Monte Carlo options price perturbation method in Bliss and Panigirtzoglou (2002) for the sample period January 2012 to December 2013.

A. Quantile volatility



B. Quantile skewness



C. Quantile kurtosis

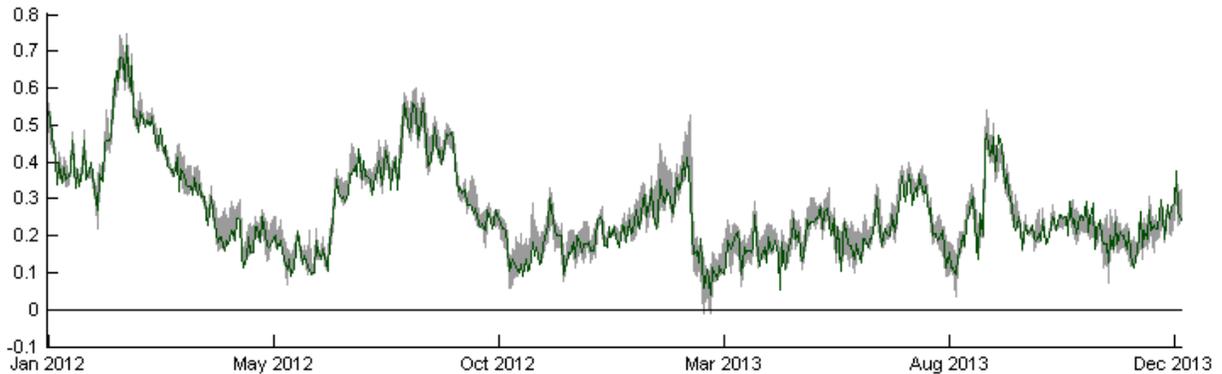
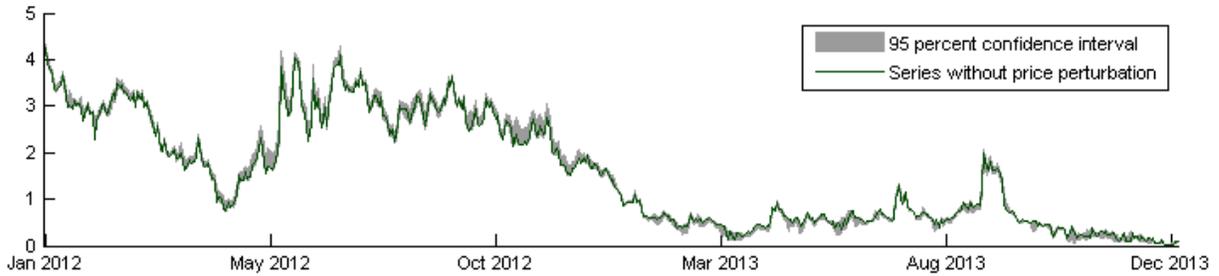


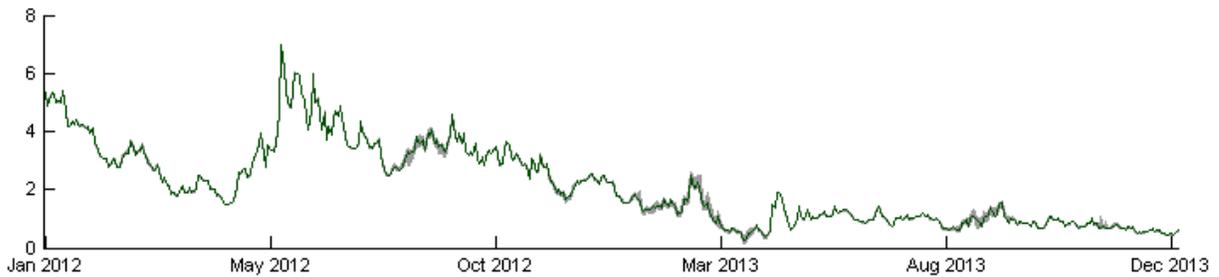
Figure 10: Confidence intervals for cost of insurance

This figure shows the cost of insurance against extreme prices, and the 95 percent confidence intervals for each cost of insurance series. The confidence intervals are calculated using the Monte Carlo options price perturbation method in Bliss and Panigirtzoglou (2002) for the sample period January 2012 to December 2013.

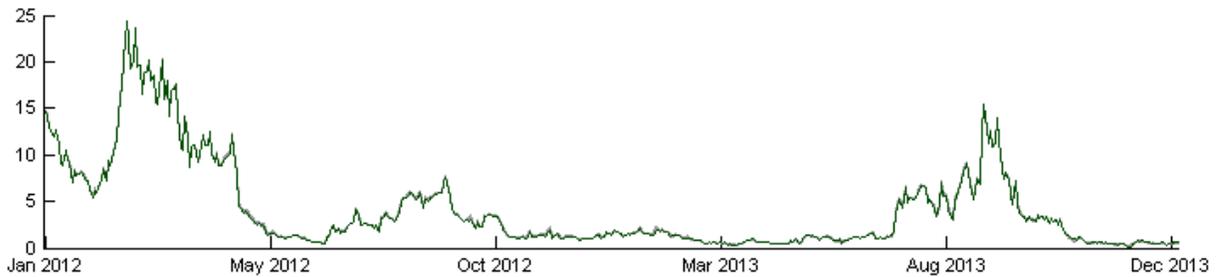
A. Cost of insurance against 30 percent price increase



B. Cost of insurance against 30 percent price decrease



C. Cost of insurance against prices above \$120



D. Cost of insurance against prices above \$135



Figure 11: Libyan civil war, February 2011

Figures 11 to 13 show a set of information extracted from the PDFs several days before and after significant market events. Each figure includes 6 panels. The top left panel shows the interpolated 90-day futures price. The bottom left panel shows the PDFs estimated using the smoothing spline with first-order polynomial tails, as described in section 2. The middle panels show quantile moments of volatility, skewness, and kurtosis for the distribution of returns. The top right panel shows the cost of insurance against a 15 percent increase or decrease in the future price. Finally, bottom right panel shows the cost of insurance against reaching extreme futures prices of \$120 and \$135.

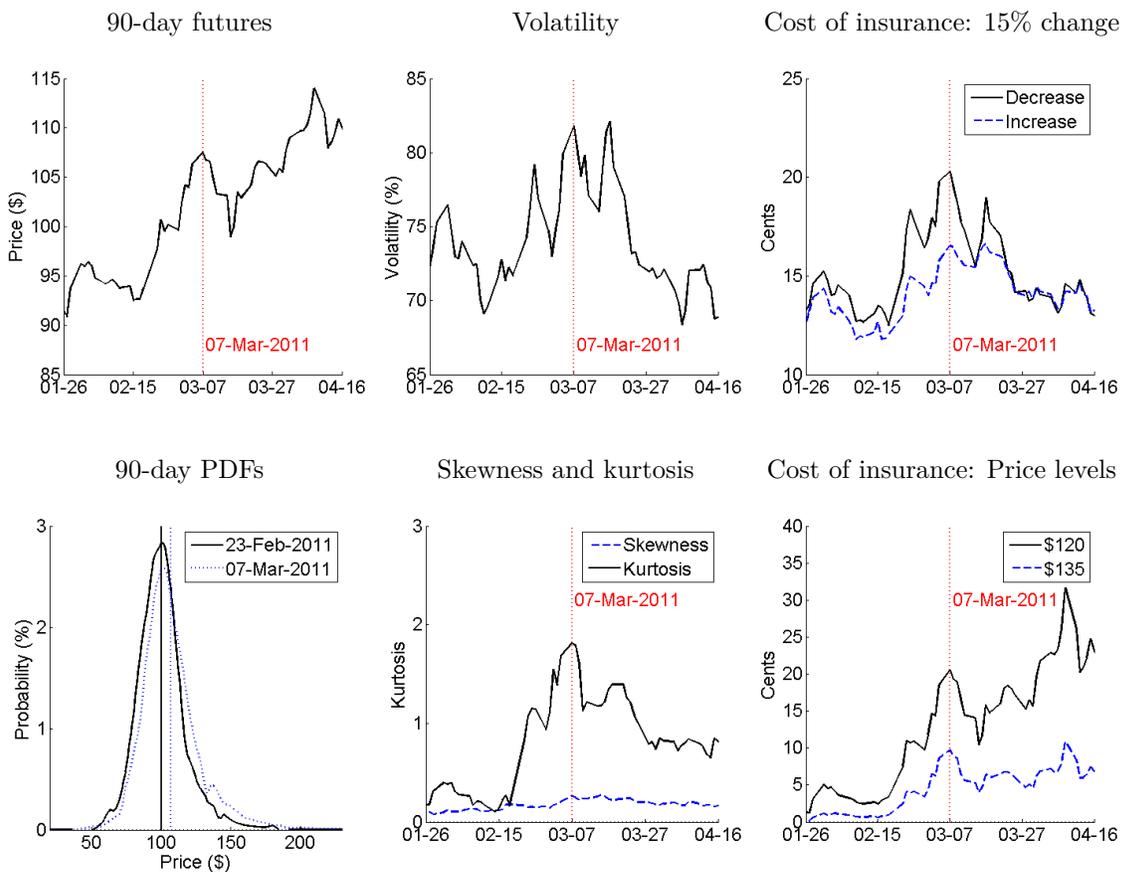


Figure 12: Conflict with Iran, February 2012

See notes to figure 11.

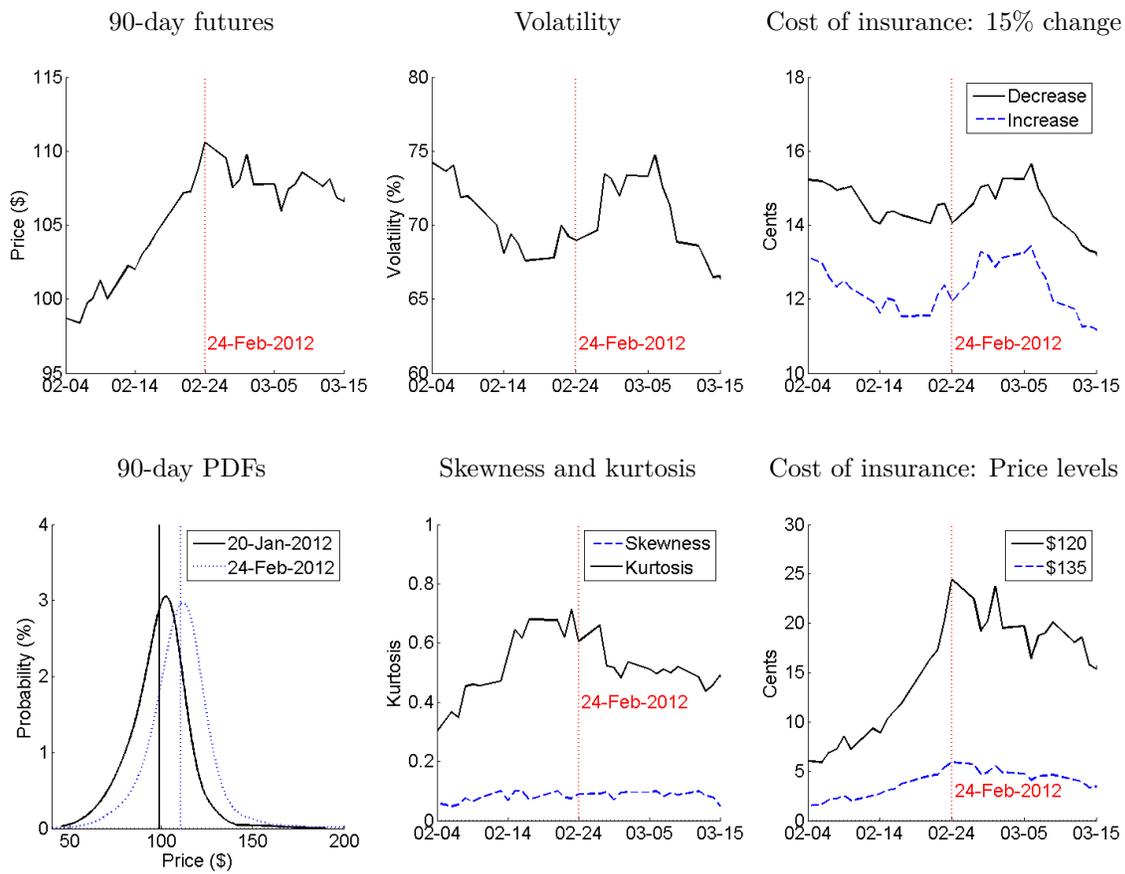
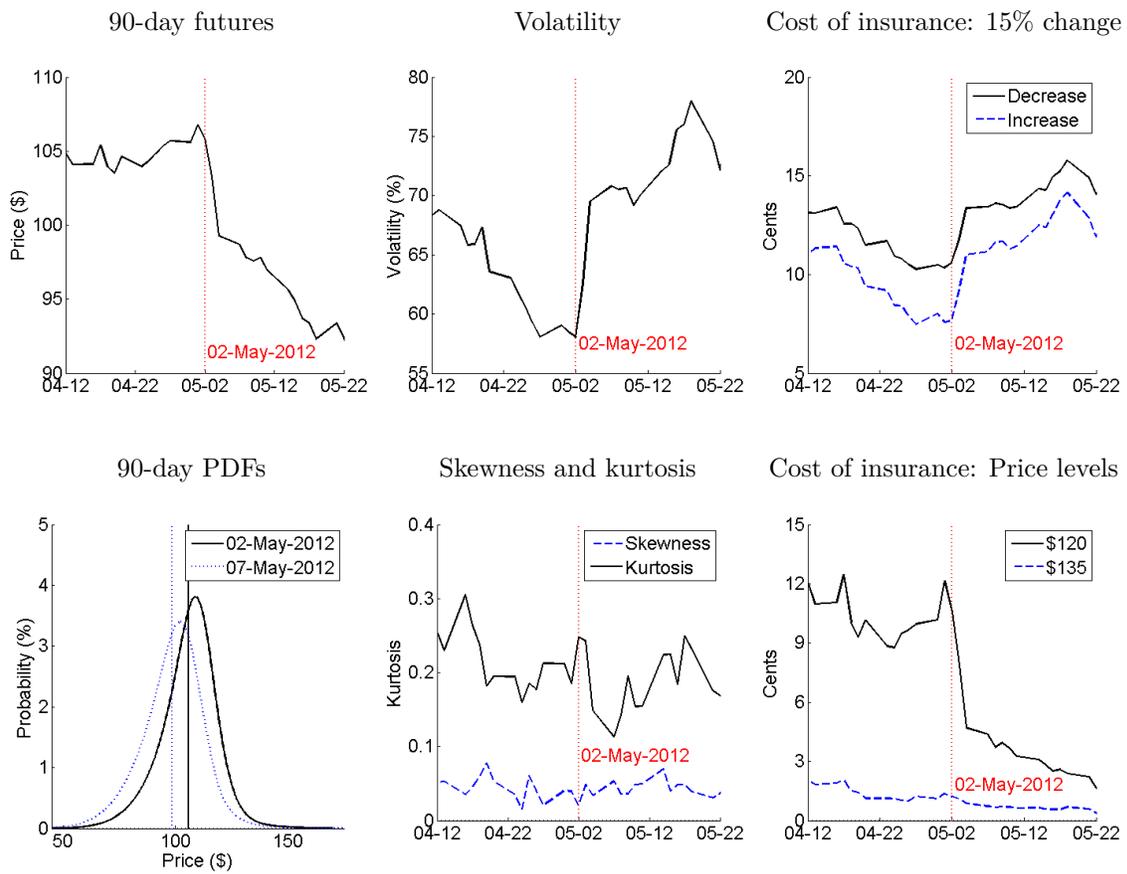


Figure 13: Negative data surprises, May 2012

See notes to figure 11.



A Smoothing spline error tolerance parameter

This appendix discusses the specification of the objective function used to fit the smoothing spline to the implied volatility data. Recall that the smoothing spline minimizes the sum of the squared distances between the curve and the data points, and maximizes the smoothness of the function, which is measured by the inverse of the sum of the squared numerical second derivative. In order to determine the tradeoff between the two objectives, we must select an error tolerance parameter. This parameter serves as an upper bound on the sum of squared errors, and thus determines how far the fitted spline's values can deviate from the observed data points to maximize the fitted spline's smoothness. The selection of the error tolerance parameter depends on the particular purpose and data to which smoothing spline is applied.

For this paper, we select an error tolerance parameter of 5×10^{-5} , which allows us to balance the feasibility of visual analysis of the PDFs with the accuracy of the fitted PDFs to the call price data. In figure A.1, we illustrate the influence of the error tolerance parameter on the shape of PDFs. In each panel, we fit the options price data for February 24, 2012 using the smoothing spline and an error tolerance parameter that ranges from a low value of 1×10^{-6} to a high value of 1×10^{-5} . The figure makes clear that as the error tolerance parameter increases, the smoothness of the PDF and the pricing errors both increase. The next two figures provide two additional examples. In figure A.2, we can see that while the higher error tolerance parameter of 1×10^{-5} is more likely to result in a smooth PDF, the errors are greater and, in some cases such as the one shown in figure A.2, a "smooth-enough" PDF could have been obtained with a lower error tolerance. Conversely, there are also instances in which the choice of 5×10^{-5} will result in an unsmooth PDF, such as the one shown in figure A.3.

Figures A.1 to A.3 show the near-future (before 90 days) fitted PDFs (left panels) and the pricing errors (right panels) for selected dates and alternative error tolerance parameters. In all figures, the smoothing spline is used to fit the options-implied volatility curve and the first-order polynomial is used to fit the unobserved tails.

Figure A.1: February 24, 2012, near-future PDF

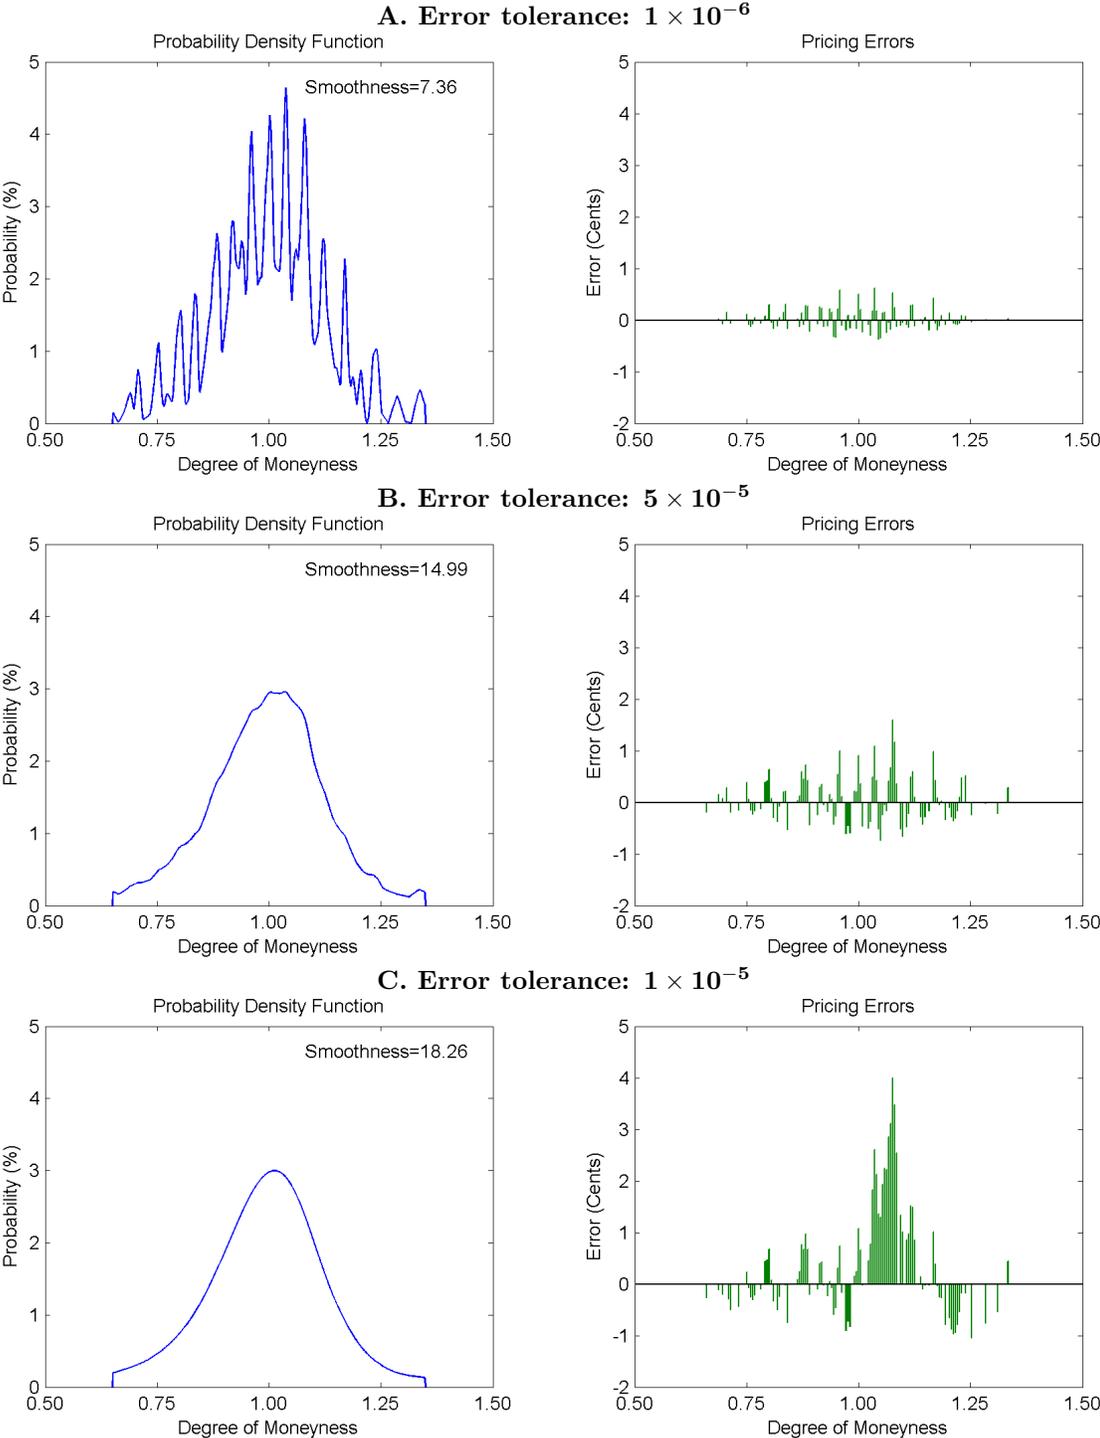


Figure A.2: May 2, 2012, near-future PDF

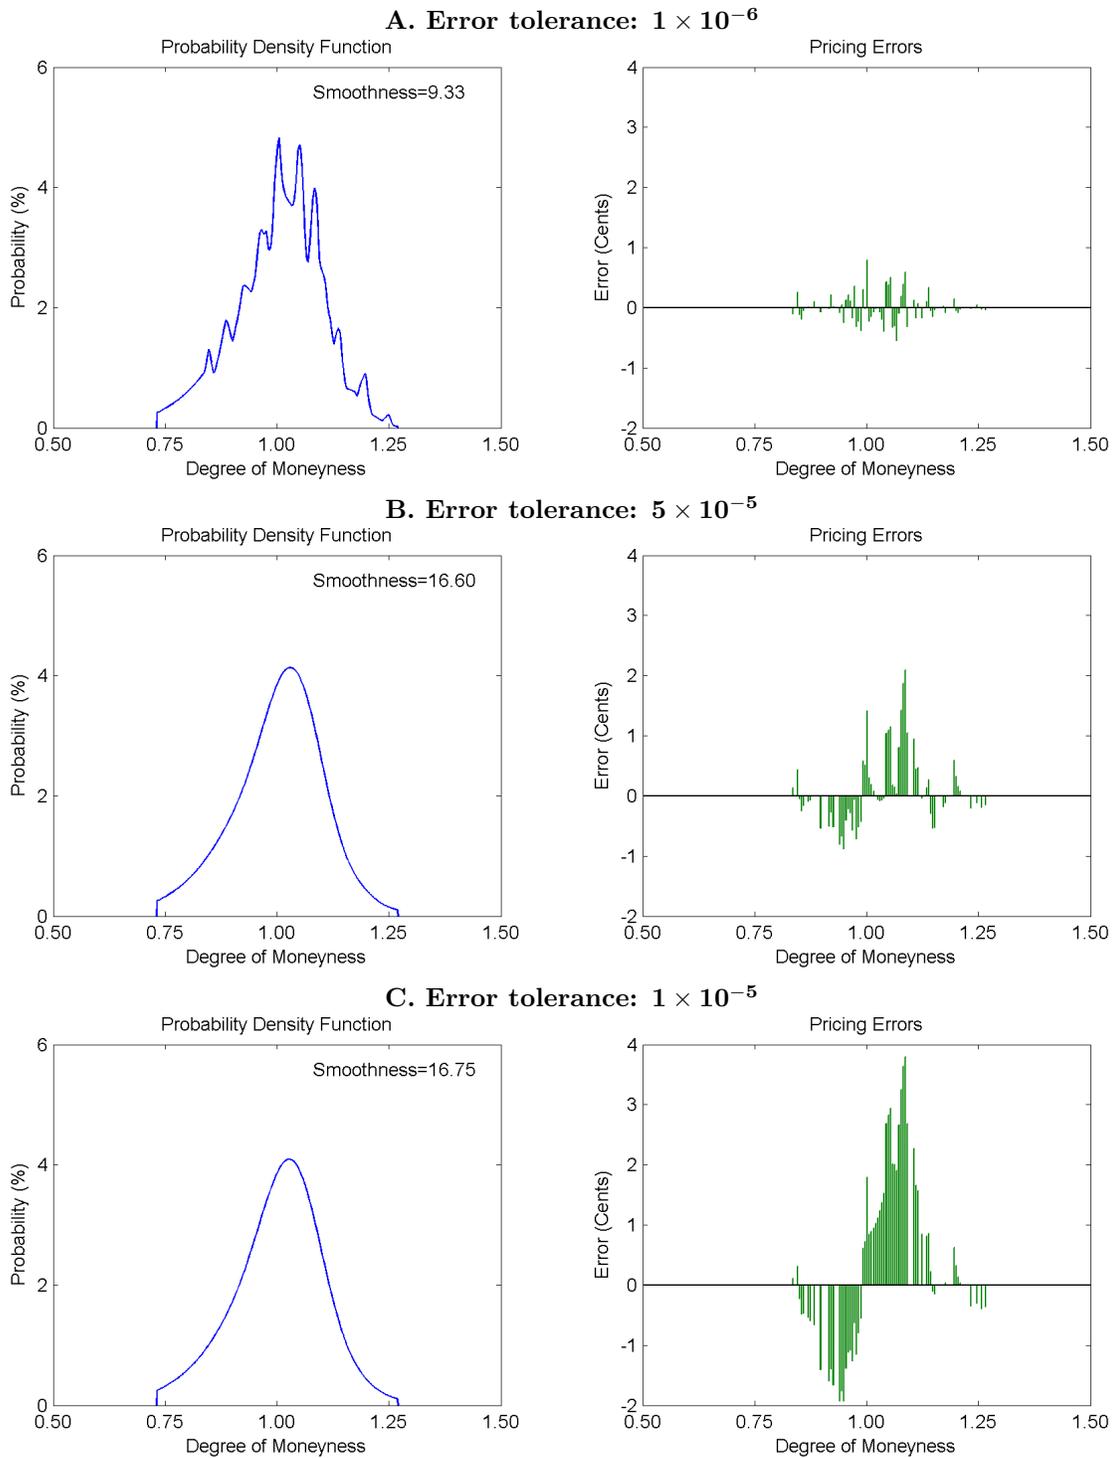
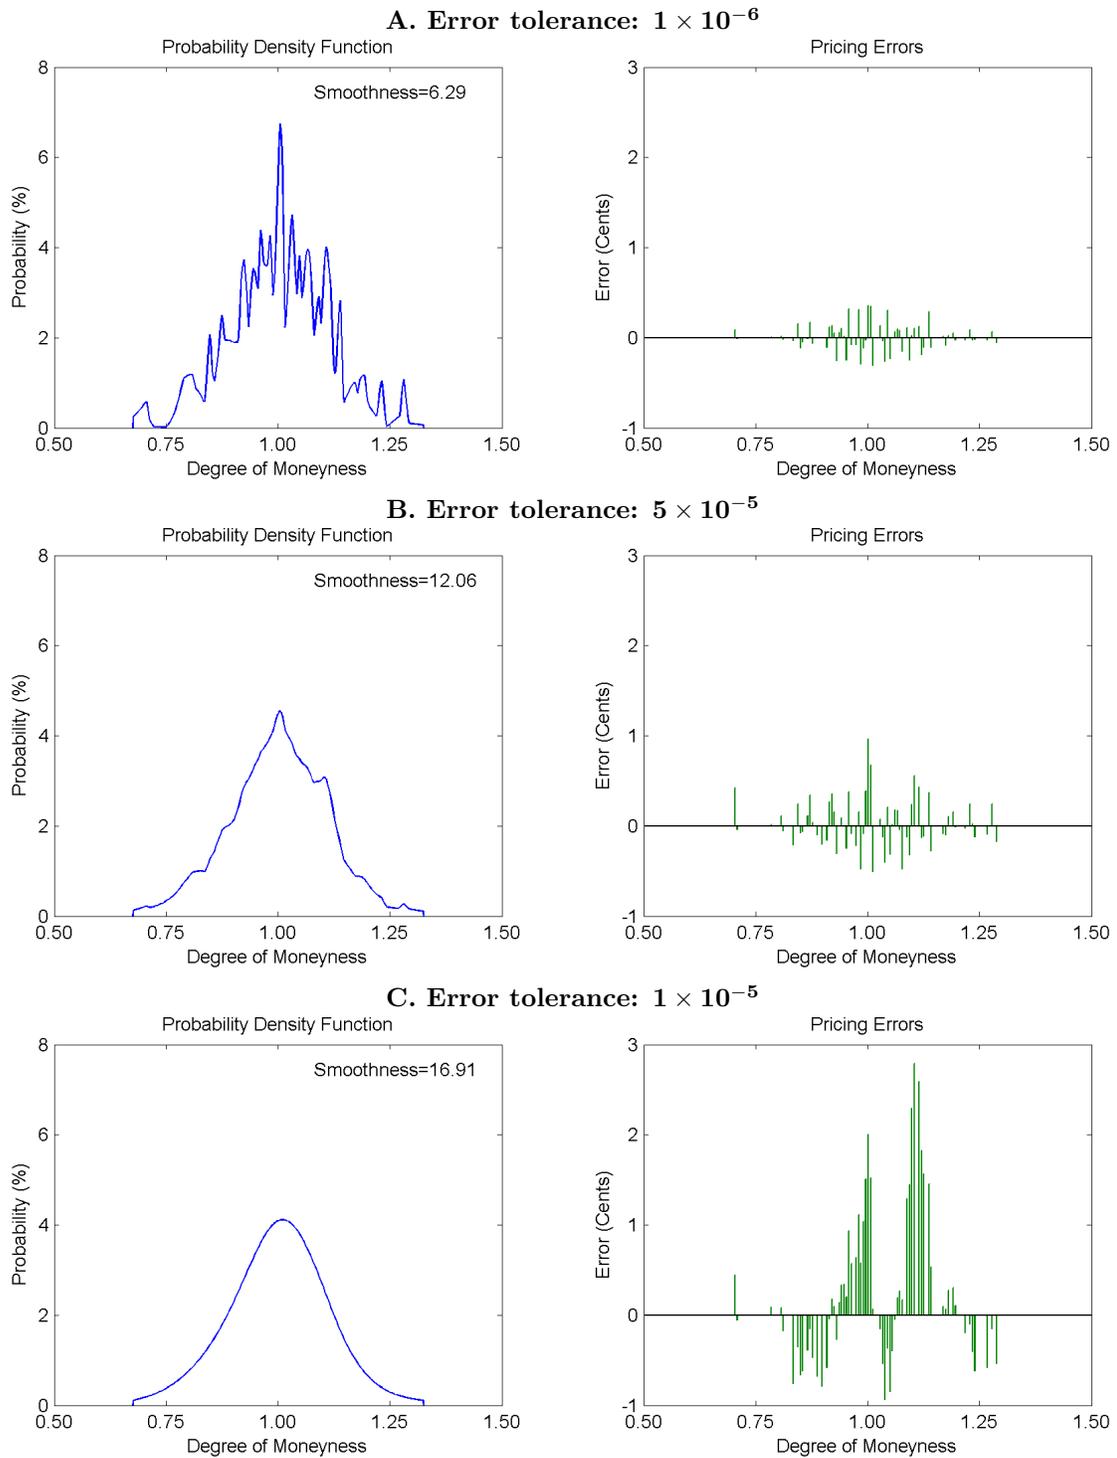


Figure A.3: February 11, 2011, near-future PDF



B Spline extrapolation parameter

To estimate the options-implied PDFs, we must extrapolate the implied volatility curve beyond the observed strikes to guarantee that the fitted probabilities add up to one. In this appendix, we investigate the optimal extrapolation range for the baseline smoothing spline with first-order polynomial tails specification. To do so, we fit the implied volatility curve extrapolating outside the observed range of strike prices between \$1 and \$40.

Table B.1 shows a set of summary statistics for the extrapolation range for alternative extrapolation values. On some days, an additional \$10 of extrapolation towards lower strike values is not possible, as the minimum strike price is bounded at \$0. Consequently, the minimum call strike (in column 1 of table B.1) declines by a bit less than \$10 in each row of the table. As expected, the maximum call strike (in column 4) increases by \$10 for each additional \$10 of extrapolation. The characteristics of the optimal extrapolation range should be zero probability density at both extremes of the distribution, which is equivalent to having null call prices in the extremes of the distribution, and a cumulative distribution function (CDF) that goes from zero to one. These characteristics are met for extrapolation values above \$10. However, as can be seen in table B.2, increasing the extrapolation value decreases the put pricing error at first, but, for very large extrapolation values, the put pricing error starts increasing again. Therefore, an extrapolation value of \$30 not only minimizes the pricing errors, but also satisfies the desired characteristics of the extrapolation range.

Table B.3 shows the impact of alternative extrapolation values on options-implied moments. We find that while standard moments are highly sensitive to the extrapolation values, quantile moments stabilize after an extrapolation value of \$20.

Table B.1: Extrapolation: CDF and PDF Values at the Extremes

This table reports the mean strike, CDF, and PDF at the minimum estimated call strike and the maximum estimated call strike, for various extrapolation ranges.

Dollars of Extrapolation	At Minimum Call Strike			At Maximum Call Strike		
	Mean Strike	Mean CDF	Mean PDF	Mean Strike	Mean CDF	Mean PDF
1	34.46	0.06	0.01	97.97	0.99	0.00
10	25.61	0.01	0.00	106.97	1.00	0.00
20	17.09	0.00	0.00	116.97	1.00	0.00
30	11.64	0.00	0.00	126.97	1.00	0.00
40	7.56	0.00	0.00	136.97	1.00	0.00

Table B.2: Extrapolation: Put Price Errors

This table reports put pricing errors for the smoothing spline curve-fitting method with first order polynomial extrapolation for various extrapolation ranges. We calculate daily PDFs between January 2009 and December 2013. Each day, we calculate the RMSE by comparing observed put prices to those mapped from the fitted implied volatility curve using the model in Black (1978). The reported mean of these daily RMSE values reflects the goodness-of-fit for each extrapolation range.

Dollars of Extrapolation	RMSE Mean
1	20.014
10	17.244
20	16.878
30	16.874
40	16.876

Table B.3: Effect of Extrapolation on PDF Moments

This table reports the average non-quantile and quantile moments for PDFs estimated from WTI future call options from 2009 to 2013 with between \$0 of extrapolation and \$40 of extrapolation using a first-order polynomial to extrapolate the implied volatility curve.

Dollars of Extrapolation	Volatility	Quantile Volatility	Skew	Quantile Skew	Kurtosis	Quantile Kurtosis
1	85.87	86.04	1.75	0.17	6.01	0.54
10	98.63	87.04	2.52	0.15	12.68	0.55
20	107.66	87.09	3.46	0.15	26.20	0.55
30	114.08	87.09	4.37	0.15	47.52	0.55
40	118.83	87.09	5.22	0.15	77.26	0.55