

Rare events and investor risk aversion: Evidence from crude oil options

Marie-Hélène Gagnon[§]

Gabriel J. Power[†]

Abstract: The effect of four distinct events on investor risk aversion is evaluated using options data on the WTI crude oil futures contract during the 2007-2011 period. Absolute risk aversion functions and pricing kernels are estimated for eight fifteen-trading day windows--before and after each event. Risk-neutral densities are parameterised under two cases: normal and Generalised Beta 2 distributional assumptions for returns. Results show that investor risk aversion functions, but not pricing kernels, are significantly and differentially affected by the events. The effect on risk aversion depends on the nature of the shock. Anticipated events affect the real-world density more than the risk-neutral density, leading to flatter risk aversion functions (i.e., less decreasing in wealth). In contrast, unanticipated events affect the risk-neutral density more than the real-world density, leading to steeper risk aversion functions (i.e., decreasing more rapidly in wealth).

EFM classification: 350, 410

Presenting author: Gabriel J. Power ; research areas 410, 420, 430, 440, and 450

Keywords: risk aversion, pricing kernel, event study, generalised beta, option-implied density.

* This work was supported by the Institut de Finance Mathématique de Montréal (IFM2), the Centre Inter-universitaire sur le risque, les politiques économiques et l'emploi (CIRPÉE), and the Social Sciences and Humanities Research Council of Canada (SSHRC). We thank Patrick Chabot and Dominic Toupin for excellent research assistance.

[§] Assistant Professor, Département de finance, assurance et immobilier, Université Laval, Pavillon Palasis-Prince, Université Laval, Québec, Canada G1K 7P4. TEL 1 (418) 656 2131 ext. 4742, FAX 1 (418) 656 2624; Email; mhgagnon@fsa.ulaval.ca

[†] Corresponding author, Assistant Professor, Département de finance, assurance et immobilier, Université Laval, Centre interuniversitaire sur le risque, les politiques économiques et l'emploi (CIRPÉE), mailing address : Pavillon Palasis-Prince, Université Laval, Québec, Canada G1K 7P4. TEL 1 (418) 656 2131 ext. 4619, FAX 1 (418) 656 2624; Email; gabriel.power@fsa.ulaval.ca

1 Introduction

Risk aversion and the pricing kernel (stochastic discount factor) are among the most studied concepts in finance and financial economics, particularly in the areas of asset pricing and derivatives. For instance, a persistent problem in the literature, known as the pricing kernel puzzle or risk aversion puzzle², concerns the empirical findings of non-monotonically decreasing pricing kernels and non-monotonic and/or negative-valued risk aversion functions (in terms of future endowment or wealth levels). Related work on the equity premium puzzle and risk-free rate puzzle, which suggest risk aversion functions and pricing kernels that are incompatible with some of the more restrictive but frequently used expected utility consumption/investment models, has in recent years focused on allowing for “peso problem”-type rare disasters in financial asset markets (e.g. Barro, 2006). Lastly, a smaller but active literature has looked at the time-varying character of the risk aversion coefficient (e.g. assuming power utility and thus constant relative risk aversion) and the rate of time preference (e.g. Chabi-Yo, Garcia and Renault, 2008).

However, comparatively little research has been conducted on the effect of rare market shocks on the general shape of pricing kernels and risk aversion functions. The recent turmoil in capital asset markets resulting from the credit freeze crisis has provided researchers with rich opportunities to examine the impact of large, unforeseen shocks on financial markets and investors (e.g. Birru and Figlewski, 2011). In this paper, we investigate using options data whether and to what extent rare economic and political events can change investor risk attitudes

² E.g. Ait-Sahalia and Lo, 2000; Bakshi, Madan and Panayotov, 2010; Christoffersen, Heston and Jacobs, 2012; Dittmar, 2002; Engle and Rosenberg, 2002; Jackwerth 2000, 2004; Liu et al. 2009; Hens and Reichlin, 2010; Ziegler, 2007.

(i.e., risk aversion) and inter-temporal substitution of wealth (i.e., the pricing kernel or stochastic discount factor). Our motivation is that investors acquire new information revealed through financial or political events, which may influence their sentiment toward risk or inter-temporal wealth substitution. Indeed, Barone-Adesi, Mancini and Shefrin (2012) develop a theory in which investor sentiment affects risk aversion and the rate of time preference, then show empirically that under- or overconfidence helps explain the pricing kernel puzzle.

Thus, the importance of rare market events should be reflected in forward-looking measures such as the risk-neutral density, which incorporates investor perceptions of the distribution of the underlying asset, as well as risk aversion functions and the pricing kernel that are estimated using both historical (real-world) and risk-neutral densities implied from option prices.

While the financial literature on option-implied risk-neutral densities and state price densities is immense,³ the literature on the robustness of the pricing kernel and risk aversion functions to rare market events is scarce. Melick and Thomas (1997) look at the option-implied distribution of crude oil futures prices during the Persian Gulf crisis and find that the Black-Scholes model option-implied density overestimates the market's assessment of the likelihood of a major disruption. Jackwerth (2000) examines the long-term stability of the absolute risk aversion function implied by S&P 500 index options and finds that it changes dramatically after the 1987 crash, possibly due to mispricing in the options market. These two papers essentially consider events to be structural breaks, such that the effect is assumed to be long-term. In

³ E.g., Ait-Sahalia and Lo, 2000; Bahra, 1997; Bliss and Panigirtzoglou, 2002; Figlewski, 2008; Jackwerth, 1999, 2000, 2004; Jackwerth and Rubinstein, 1996; Jondeau, Poon and Rockinger, 2007; Liu et al. 2007; Pérignon and Villa, 2002; Rubinstein, 1994; Ziegler, 2007.

contrast, this paper investigates short-term reactions in risk aversion functions and pricing kernels to significant exogenous events by developing an event study approach allowing us to present graphical and statistical evidence. Our paper further differentiates itself from the literature in its focus on crude oil options, as the literature has focused mainly on equity and Forex markets (Coutant, Jondeau and Rockinger, 2001). Empirical studies of crude oil option-implied densities include, however, Flamouris and Giamouridis (2002), Giamouridis (2005), and Melick and Thomas (1997).

The main contribution of this paper is to provide evidence of significant and differential effects of infrequent financial and political events on risk aversion functions and pricing kernels, through the risk-neutral and historical densities. These are estimated from available options data on the West Texas Intermediate (WTI) Future contract for the period of 2007-2011. Based on financial news, we identify four events that have affected the crude oil market, namely the end of the commodity bull cycle, the trough of the credit freeze crisis, the BP Deepwater Horizon platform explosion, and the Libyan uprising in Tripoli. The first two events may be considered at least partly anticipated as they define the end of cycles, while the latter two events are idiosyncratic and wholly unanticipated shocks. Estimates are obtained in windows of fifteen trading days before and after each event (omitting the event date itself). A parametric non-structural approach is used to recover the option-implied risk-neutral and historical (real-world) densities for crude oil futures price returns. Two cases are considered, under the assumptions of a Gaussian normal distribution for returns, and a Generalised Beta 2 (GB2) distribution. The normal case provides a baseline case, while the GB2 case provides greater flexibility to capture in particular skewness and kurtosis. Once both densities have been estimated, we compute

pricing kernels and risk aversion functions following Jackwerth (2000) and present both graphical and statistical evidence as well as several robustness checks to support our findings.

Our results, briefly, are as follows. We find that exogenous events have a short-term impact on risk aversion (the investor's tolerance for risk), but not on the pricing kernel (inter-temporal substitution of wealth). The effect on risk aversion depends on the nature of the shock. End-of-cycle (anticipated) shocks affect the historical density but leave the risk-neutral density relatively unchanged, its information appearing to be anticipated by and incorporated into the option market. This implies flatter absolute risk aversion functions, such that the representative investor will be less risk-averse if his future wealth level is low, but more risk-averse if his future wealth level is high. On the other hand, idiosyncratic or unanticipated shocks affect the risk-neutral density more than they do the historical (real-world) density. As a result, unpredictable shocks cause the risk aversion function to become steeper, such that an investor becomes more risk-averse for lower future endowments and less risk-averse for higher future wealth. Several robustness checks are presented to provide support for the paper's claim that the events have an economically and statistically significant effect on risk aversion.

The remainder of this paper is organised as follows: section 2 presents the methodology. Section 3 presents the data. Section 4 presents the results and section 5 concludes.

2. METHODOLOGY

2.1 Practical considerations : event windows

Our paper aims to build on the work of Melick and Thomas (1997) and Jackwerth (2000). Our goal is to perform an event study type analysis identifying changes in the investor's attitude toward risk associated with infrequent political and financial events. In particular, we aim to identify changes in risk-neutral densities, pricing kernels and risk aversion functions associated to such events.

The exact window dates for each event are presented in Table 1. The four events considered are: i) the end of the commodity bull cycle on July 11th 2008 associated with uncertainty of the housing market in the USA, ii) the oil price trough that occurred on January 19th 2009 related to the global recession, iii) the explosion of the BP Deepwater Horizon oil platform on April 20th 2010 and iv) the beginning of the “Libyan uprising” (Jasmine revolution) in the spring of 2011 with expected repercussions on crude oil production and shipping. The day of the events was chosen according to the magnitude of the WTI price movement and the importance of news affecting commodity prices as reported on Bloomberg. In particular, there has been several notable events related to the “Jasmin revolution” prior to the chosen date (February 22nd, 2011), but the day the insurrection started in Tripoli is chosen because it cast a great deal of uncertainty on Future oil supplies.⁴ The first two events, the end of the bull commodity cycle and the credit freeze bottom, are characterised as end-of-cycle and were possibly anticipated by investors. The last two events, the BP platform explosion and the Libyan uprising, should be considered as idiosyncratic shocks that are very difficult to forecast.

⁴ News regarding oil price on Bloomberg show that this day was where the uncertainty was at its highest and where the markets were volatile in response to it.

Figure 1 presents the evolution of the WTI future contract price during the 2006-2012 period. It shows that the WTI futures contract price was relatively volatile during the period considered. Note that the event windows are identified on the figure using rectangles. The first two windows were chosen because they correspond to the end of a cycle for the future price of oil. The third window shows the oil movements around the explosion of the BP platform. Finally, the last window looks at the movement in oil prices around the beginning of the Libyan uprising. Table 2 reports that the average gross return for a two-month holding period for the whole period considered was 1.03 with a standard deviation of 16.69%.

2.2 *Densities, pricing kernel and risk aversion function*

The pre and post event densities, pricing kernel and risk aversion function need to be estimated for each event. We estimate the pre- and post-event densities using fifteen trading days before and after the event, respectively. For both the pre- and post-event densities, an historical and risk neutral density is estimated in each fifteen days.

2.2.1 *The historical density*

Two general approaches have been presented in the literature to estimate the historical or real-world density. First, one strand of literature aims at estimating the “conditional” density of historical returns using a short window of observations that are available to the representative investor at a given time. For example, Jackwerth (2000) uses one month of daily returns. The objective is to highlight the time-varying component of risk aversion. This method has the advantage of allowing for a time-varying structure in the historical density. Another approach presented in Ait-Sahalia and Lo (1998) is to use longer time series to relax the distributional assumption and obtain the most accurate estimate of the risk aversion function. Given the event

study nature of our objective, we adopt a methodology similar to Jackwerth (2000). However, we use the longest non-overlapping historical window available for each event to improve the precision of our estimates.

The historical densities for returns on a two-month investment horizon are estimated for each day in the pre- and post-event window using data from the WTI third nearby futures contract price. We use a moving window approach, such that the number of observations varies with each window, but is kept constant within a given window. The number of observations is kept constant to avoid spurious improvements in the goodness-of-fit associated with a growing number of observations in later measurements. Only the information available to investors on a given day (i.e., past prices) is used to estimate the historical densities. Moreover, an important issue is to isolate the impact of one event from the other events in the sample. The historical densities are non-overlapping across the events in order to isolate the event's effect. Table 2 reports the maximum number of non-overlapping observations available to fit the historical densities. Given that the historical returns must be non-overlapping across events, there is not a large enough number of observations to implement for each event a fully non-parametric approach as in Ait-Sahalia and Lo (1998). Therefore, two parametric approaches are considered instead.

Firstly, a baseline case using a Gaussian normal parameterisation is estimated. Following the recommendation of Jackwerth (2000) and Merton (1980) the distribution is demeaned and the mean is fixed as the risk-free rate plus the equity premium, set at 6% for our sample based on Ivo Welch's 2009 survey of financial economists (see e.g. Welch, 2000, 2009). This approach avoids the problem of irregularities in the estimated pricing kernel and risk aversion function that are due to the difficulty of estimating the historical mean of returns. Moreover, the paper aims to

document and explain changes in the shape and slope of the risk aversion function, rather than its actual level. Variance of returns is calibrated as the empirical variance of historical returns in each sub-sample.

Secondly, the distributional assumption of normality in returns is relaxed and replaced by the more general class of Beta densities. Given our objective of documenting the effect of several significant market events on the representative investor risk aversion, we emphasize the importance of making the distributional restrictions as similar as possible for the two densities (i.e., historical and option-implied). For example, it would be misleading, in terms of our pricing kernel and risk aversion function results, to allow for non-zero skewness in one density but not the other. For this reason, we opt for a (scaled) Beta distribution for the historical density and a Generalized Beta 2 distribution for the risk-neutral density. The Beta distribution is a special case of the Generalised Beta 2 distribution that will be used in the risk-neutral density. This somewhat more restrictive parameterisation than the Generalised Beta distribution is chosen because the number of non-overlapping observations related to each event limits our capacity to adequately parameterize the latter.

The Beta distribution is versatile, particularly in allowing for either positive or negative skewness (see e.g. Moitra, 1990). The density is:

$$f(x|a, b) = \frac{1}{B(a, b)} x^{a-1} (a - x)^{b-1} \mathbf{I}_{(0,1)}(x) \quad (1)$$

where a and b are strictly positive parameters, B is the Beta function, and $\mathbf{I}_{(0,1)}$ denotes the indicator function such that the only values having positive probability are those in the range of 0 to 1. The support is rescaled over the range of 0 to 2, consistent with gross returns being bounded

below at 0 (i.e., non-negative prices) and above at 2, consistent with the full sample of data that did not include any observations corresponding to gross returns above 1.6. Computationally, it is more straightforward to obtain the Beta function using Gamma functions, such that:

$$B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)} \quad (2)$$

where Γ is the Gamma function. Moments of the Beta distribution can be also obtained using the Gamma function:

$$E[X^n] = \frac{\Gamma(a+n)\Gamma(a+b)}{\Gamma(a+b+n)\Gamma(a)} \quad (3)$$

for the n -th moment. In particular, mean and variance are obtained as $E[X] = \frac{a}{a+b}$ and $Var(X) = \frac{ab}{(a+b)^2(a+b+1)}$. Parameters of the density are obtained by maximum likelihood estimation.

2.2.2 *The implied volatility surface*

Our methodology to recover risk-neutral densities for the baseline case of Gaussian normal futures price returns is based on Shimko (1993). A volatility surface is estimated as a quadratic form (see e.g. Dumas, Fleming and Whaley, 1998) and used to compute the risk-neutral density. Specifically, Black-Scholes (Black, 1976) implied volatilities are regressed over corresponding values of (log) moneyness, time-to-maturity, their squared values, and a constant. Time-to-maturity is calculated as the ratio of business days before contract expiry over 252. Log-moneyness is the natural logarithm of the ratio of strike price over futures price. Taking logs helps avoid the problem of non-spherical disturbances in the linear regression. The estimated regression is as follows:

$$\sigma_{BSi} = \beta_0 + \beta_1 TM + \beta_2 MN + \beta_3 TM^2 + \beta_4 MN^2 + \beta_5(TM \times MN) + e_i \quad (4)$$

From the results of this estimated equation, we can plot the volatility smile and the volatility surface in terms of time to maturity and moneyness. To estimate the risk-neutral density, we use $\hat{\sigma}_{IV}$, the fitted implied volatility:

$$\hat{\sigma}_{IV} = \hat{\beta}_0 + \hat{\beta}_1 t + \hat{\beta}_2 MN + \hat{\beta}_3 TM^2 + \hat{\beta}_4 MN^2 + \hat{\beta}_5(TM \times MN) \quad (5)$$

2.2.3 The risk-neutral density

There exists a rich literature relating options price data to state prices, and through state price probabilities, risk-neutral densities. Breeden and Litzenberger (1978) show that, assuming a continuum of strike prices, state prices (and thus, the risk-neutral density underlying asset prices for a terminal date T) can be recovered from the second partial derivative of the Black-Scholes option pricing solution with respect to the exercise price. Then, the cumulative distribution function is obtained by differentiating once with respect to the strike price (K). State prices are obtained by differentiating a second time with respect to the strike price. A numerical approximation for the risk-neutral density q using C as the call price and K the strike price is:⁵

$$q(S_T) = e^{r\tau} \frac{C_{n+1} - 2C_n + C_{n-1}}{(\Delta K)^2} . \quad (6)$$

However, it is well known that since in practice the data do not contain a continuum of strike prices, directly computing the risk-neutral density from option price data fails to yield

⁵ For more details regarding the methodology used to extract a risk-neutral density from options data, see the survey by Figlewski (2008) as well as Bahra (1997), Barthunek and Chowdury (1997), Cox and Ross (1976), Engle and Rosenberg (2002), Jackwerth (2004), and Jarrow and Rudd (1982).

sensible results (e.g., Ait-Sahalia and Lo, 2000; Bahra, 1997; Bliss et Panigirtzoglou, 2004; Jackwerth, 1999, 2000; Jackwerth and Rubinstein, 1996; Rubinstein, 1994). The literature has followed different approaches to overcome this significant challenge. Structural models aim to capture both the underlying price dynamics as well as the terminal period density of prices, while non-structural methods focus on the density without explicitly modeling the price dynamics. Given that this paper is concerned with the short-run effect of exogenous market events on changes in investor risk aversion rather than the estimation of a single underlying process for the entire sample, data constraints lead us to use a parametric, non-structural approach.

In the baseline case, a Gaussian normal density function is fitted to the return data, where μ is the mean and the variance σ :

$$f(x|\mu, \sigma(\beta)) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{(\ln x - \mu)^2}{2\sigma^2}} \quad (7)$$

where the parameters of equation (7) are the risk-neutral drift, measured by the risk-free rate over the horizon of investment, while the diffusion parameter $\sigma(\beta)$ is estimated using the implied volatility, $\hat{\sigma}_{IV}$, computed from the option volatility surface as presented in equation (5).

Subsequently, a more flexible distributional assumption is allowed by using the Beta distribution for the historical (real-world) density of returns and the Generalised Beta 2 (hereafter, GB2) distribution for the risk-neutral density. This distribution, developed for financial applications by Bookstaber and McDonald (1987) and McDonald and Xu (1995), has the advantage of flexibility and nests other distributions, including Beta and Normal. More recently, it has been used in the option-implied density literature by e.g., Aparicio and Hodges (1998), Anagnou-Basioudis et al. (2005), Liu et al. (2007, 2009) and Shackleton, Taylor and Liu

(2010). The GB2 can capture a wide range of distributional shapes through non-Gaussian values of skewness and kurtosis but, being parametric, does not require as many observations as do non-parametric approaches. The latter advantage is significant given the event study character of the analysis, whereby the number of observations can be fairly limited. The GB2's probability density function may be written as follows:

$$f(x|a, b, p, q) = \frac{ax^{ap-1}}{b^{ap}B(p,q)\left[1+\left(\frac{x}{b}\right)^a\right]^{p+q}} \quad (8)$$

where $B(p, q)$ is the Beta function and (a, b, p, q) are parameters to be estimated. The location parameter a is associated with the speed at which the tails of the distribution approach the abscissa. The scale parameter b tends toward the mean when a is large. Skewness is determined by the parameters p and q while kurtosis is determined by parameters a and q . The cumulative distribution function is:

$$F(y; a, b, p, q) = z^p {}_1F_2[p, 1 - q, 1 + p; z]/pB(p, q) \quad (9)$$

where $z = \left(\frac{y}{b}\right)^a / \left(1 + \left(\frac{y}{b}\right)^a\right)$ and ${}_1F_2[\cdot]$ is a hypergeometric function. Following McDonald and Xu (1995), moments of the GB2 distribution are given by:

$$E[Y^k] = \frac{b^k B\left(p + \frac{k}{a}, q - \frac{k}{a}\right)}{B(p, q)} \quad (10)$$

where $B(p, q)$ is the Beta function. Thus, the mean of the distribution is given by:

$$E[Y] = bB\left(p + \frac{1}{a}, q - \frac{1}{a}\right)/B(p, q) \quad (11)$$

while the variance can be computed as $E[Y^2] - E[Y]^2$. Note that the number of moments is limited by the value of aq such that only moments of order less than aq will exist. We therefore require $aq \geq 4$ ensuring existence of the first four moments.

For each date, a GB2 risk-neutral density is estimated from the daily cross-section of option data using all available strikes and a time-to-maturity matching the representative investor's time horizon of two months.⁶ A constrained optimisation program searches for the optimal GB2 distribution parameters (a, b, p, q) that minimizes average squared errors, defined as the difference between the observed option price $C_{t,T,M}$ and the option price predicted by the option pricing formula (given below) for a GB2 risk-neutral density $\hat{C}_{t,T,M}^{GB2}$:

$$(a, b, p, q) \in \operatorname{argmin} \frac{1}{N} \sum_{i=1}^N \left(C_{t,T,M}(X_i) - \hat{C}_{t,T,M}^{GB2}(X_i) \right)^2 \quad (12)$$

A grid of starting values is used to verify that a global optimum is obtained and that there are no active boundary conditions. The martingale condition for futures prices, on which options are written, is applied to ensure risk-neutrality of the GB2 density by constraining the mean of the distribution to equal the corresponding futures price P^f (i.e., for a distribution of returns, the mean is set to the risk-free rate of return):

$$P^f = \frac{bB\left(p + \frac{1}{a}, q - \frac{1}{a}\right)}{B(p, q)} \quad (13)$$

⁶ Note that as we estimate a RND for each of 15 business days before and after the event date, the contract time-to-maturity is precisely equal to two months on one date only, and for other dates is approximately equal to the two-month investment horizon as the closest contract is used. As a robustness check, we also estimate RNDs for a one-year investment horizon using all strikes associated with one-year contracts.

Using the GB2 distribution, the option pricing formula for European call options is given by:

$$\hat{C}_{t,T,M}^{GB}(a, b, p, q) = P^f e^{-rT} \left[1 - F^{GB} \left(z \left| p + \frac{1}{a}, q - \frac{1}{a} \right. \right) \right] - X e^{-rT} [1 - F^{GB}(z|p, q)] \quad (14)$$

where $z = \left(1 + \left(\frac{x}{b} \right)^{-a} \right)^{-1}$ under the martingale restriction defined above (Liu et al., 2007). Note further that the GB2 distribution has the advantage of allowing for a straightforward transformation from risk-neutral to real-world density by positing a constant relative risk aversion parameter γ , i.e., assuming a power utility function for the representative investor (Liu et al., 2007).

2.2.4 Pricing kernel and risk aversion function

Leland (1980) show how the risk-neutral density, historical or real-world density, and pricing kernel (or risk aversion function) are related such that knowledge of two quantities allows the researcher to compute the third. Letting $q(S)$, $p(S)$ and $m(S)$ respectively denote the risk-neutral density, historical density and pricing kernel associated with the payoff of an asset price or wealth (S), then we may write (see Jondeau, Poon and Rockinger, 2007, for a review):

$$q(S) = m(S)p(S) \quad (15)$$

Taking log-derivatives, the relationship may be expressed in a form that brings out the Arrow-Pratt coefficient of (absolute) risk aversion computed from a utility function increasing in future consumption ($C_T(S)$):

$$\frac{q'(S)}{q(S)} = \frac{U''(C_T(S))}{U'(C_T(S))} + \frac{p'(S)}{p(S)} \quad (16)$$

such that:

$$ARA(C_T(S)) = - \left(\frac{U'''(C_T(S))}{U''(C_T(S))} \right) = \frac{p'(S)}{p(S)} - \frac{q'(S)}{q(S)} \quad (17)$$

where $ARA(C_T(S))$ is the Arrow-Pratt coefficient of absolute risk aversion for the representative investor. The resulting risk-aversion function is plotted against a future period wealth index S for a given horizon of two months.

2.3 *Statistical robustness assessment*

In this section, a robustness analysis is presented. In line with the event study literature, the objective is to provide support for the graphical evidence of changes in the risk aversion function using statistical tests suggesting that the impact observed is also statistically and economically significant. This exercise is not as simple as it may seem. Indeed, the above theoretical results provide us with a risk aversion function rather than a risk aversion density, complicating statistical assessment because no “weight” or likelihood is assigned to each risk aversion/future wealth pair obtained from (17). Classical statistical tools such as the Kolmogorov-Smirnov test of equality of densities or multivariate regression analysis are not directly applicable in this case. We address this problem using several statistical approaches in parallel to build a case that changes in the risk aversion function are significant.

As a first step, we perform statistical tests on the equality of the pre- and post-event historical price densities. We use the Kolmogorov-Smirnov and the Cramer-Von-Mises two-sample tests to assess whether the historical density has changed after the event. For both tests, the null hypothesis is that the pre- and post-event observations are both drawn from the same distribution. Again, we are careful to use only non-overlapping observations. To build our subsamples, we divide the between-event observations in two halves. For each event date j , the first

half of the observations after the given event j is used to estimate the post-event j density while the second half is used to estimate the pre-event density associated with the next event ($j-1$). Although this step is only intermediary, in that it informs us only about one of the two components that determine the risk aversion function, it is useful to establish the degree to which historical returns were affected by the exogenous event.

The next step is to conduct a statistical test of differences in risk aversion before and after the event. To this end, pre- and post-event densities of risk aversion values are generated from the options and futures data. Therefore, a methodology similar to the one used to derive the volatility surface (section 2.2) is implemented. Specifically, we use option moneyness as a proxy for future wealth for each pre- and post-event sub-sample. Using equation (17) for the Gaussian normal baseline case, each future wealth level between 0.45 and 1.48, as a proportion of initial wealth and using increments of 0.005, is matched with the corresponding value of absolute risk aversion implied by the risk aversion functions estimated from risk-neutral and historical densities. The support [0.45, 1.48] is given by the range of two-month horizon gross returns in the sample data. Values of absolute risk aversion in terms of future wealth are determined assuming a quadratic form and using a linear regression to obtain coefficients of slope and curvature, such that for levels of future wealth S we have:

$$ARA_i = b_0 + b_1 S_i + b_2 S_i^2 + e_i \quad (18)$$

This approach can be compared with the estimation of a volatility surface, whereby coefficients on S and S^2 on the original “unweighted” function are similar to the coefficients in (4) fitting the implied volatility to moneyness and time-to-maturity. The dependence of risk aversion on wealth levels is well established in the literature since Arrow (1971). Moreover, the

inclusion of a squared term is motivated by the empirical literature, according to which risk aversion “smiles” and is not monotonically related to wealth (e.g. Ziegler, 2007, and references therein).

The estimated coefficients in (18) are used to map, for all option observations, moneyness to absolute risk aversion, under the assumption that each option’s moneyness is a proxy for a future wealth level. We obtain the fitted risk aversion associated with each option j using the following correspondence:

$$\text{fitted } ARA_j = \hat{b}_0 + \hat{b}_1 MN_j + \hat{b}_2 MN_j^2. \quad (19)$$

The last step generates a distribution of ARA values corresponding to observations of future wealth levels implied by traded options, before and after each event. This approach therefore serves as an approximation of the distribution of risk aversion values pre- and post-event, based on the observed levels of future wealth implied by options data through their moneyness.

Although this method is only a proxy for the true distribution of risk aversion values in the economy and need not represent precisely the aggregate risk aversion distribution, it is nonetheless based on an accepted approach to recover implied volatility from options data. What is more, it is informative because it results in a distribution of risk aversion values associated with crude oil futures markets from which empirical moments can be recovered. Intuitively, it provides a means to assign “weight” to various points in the risk aversion function and therefore conduct statistical testing to support the graphical evidence of changes in the slope and curvature of ARA. Indeed, to determine whether ARA functions before and after each event are similar at any or all points in the support (i.e., future wealth levels), the likelihood of each future wealth level occurring in the sample must be taken into account. Mapping to ARA fitted values provides

additional evidence to investigate possible changes in investor risk sentiment resulting from exogenous events.

A natural starting point for statistical testing is to perform Kolmogorov-Smirnov and Cramer-Von-Mises two-sample distributional tests on the pre- and post-event sub-samples of fitted risk aversion values.

Moreover, we test statistical significance in the change in risk aversion using regression analysis. The following multivariate regression is performed to enable a statistical comparison in the estimated coefficients on future wealth (proxied by MN) and on the curvature coefficient of risk aversion (i.e., the smile, proxied by MN^2) :

$$\begin{aligned}
 fittedARA_{jprior} &= \beta_0^{prior} + \beta_S^{prior} MN_j^{prior} + \beta_{SS}^{prior} MN_j^{2post} + \varepsilon_j^{prior} \\
 fittedARA_{jpost} &= \beta_0^{post} + \beta_{SS}^{post} MN_j^{post} + \beta_{SS}^{post} MN_j^{2post} + \varepsilon_j^{post}
 \end{aligned} \tag{20}$$

This approach is analogous to the estimation of volatility surface coefficients using a second-order polynomial in TM and MN as in (5). It has the advantage of documenting more precisely the nature of changes in risk aversion resulting from exogenous events. In this framework, a change in risk aversion can be associated either with a change in the slope associated with wealth or the curvature of risk aversion in wealth. For example, if the representative investor has decreasing absolute risk aversion (e.g., preferences described by a power utility function), a change in the slope coefficients implies a flatter or steeper risk aversion function. A better understanding of the impact of market events on the coefficients of risk aversion curvature could help explain the pricing kernel puzzle (i.e., why empirically the pricing

kernel is not monotonically decreasing and strictly positive). Therefore, the following hypotheses are tested using a standard F-test procedure:

$$H_0^S: \beta_S^{prior} = \beta_S^{post}, \quad (21)$$

$$H_0^{SS}: \beta_{SS}^{prior} = \beta_{SS}^{post}, \quad (22)$$

$$H_0^{sum}: \beta_S^{prior} + \beta_{SS}^{prior} = \beta_S^{post} + \beta_{SS}^{post}, \quad (23)$$

$$H_0^{joint}: \beta_S^{prior} = \beta_S^{post} \text{ and } \beta_{SS}^{prior} = \beta_{SS}^{post}. \quad (24)$$

where (21) and (22), respectively, test the equality of the slope and smile coefficients before and after the event. Hypotheses (23) and (24) are introduced to conduct a statistical test of whether the relationship between wealth and risk aversion has changed as a result of the event, without relying on the general two-sample density approach. (23) tests whether the global effect of the slope and curvature coefficient is the same after the event while (24) is more restrictive and tests the joint equality of the coefficients before and after the event. Thus, both hypotheses serve to verify whether the general shape of the risk aversion function has changed.

A final robustness check is in order to better understand the nature of changes in the distribution of implied absolute risk aversion values derived above. Note that the number of traded options observed in the pooled samples (15 days before and after the event) is very large, nearing 18,000 observations. It could therefore be argued that statistical significance found in our samples may not be economically significant, but rather an artifact of having a very large sample (e.g. Berger and Sellke, 1987). This potential issue is not solved by the use of a single day for statistical testing, instead of pooling all data for the 15 days before or after an event, because the

number of observations is over 1000 for most trading days. Instead, the analysis is performed on smaller random samples obtained from draws without replacement from the pooled samples. Using this sampling approach, we report results, for a range of possible numbers of draws and numbers of observations in each draw, for the Kolmogorov-Smirnov and Cramer-von-Mises tests as well as multivariate regression hypothesis tests.

3. Data

We use business daily data on futures and options on futures for the West Texas Intermediate (WTI) crude oil on the NYMEX market over the period of June 1st, 2007 to October 25, 2011. The data were extracted from the Bloomberg database and from the Commodity Research Bureau. The WTI futures contract was chosen for several reasons. First, it is well established in the literature that oil prices tend to be pro-cyclical and endogenous to the real business cycle, both having a deep influence on the economy and also being affected by the state of the economy (e.g. Ewing and Thompson, 2007; Kilian, 2009; Sadorsky, 2001).

Second, crude oil by far is the most liquid commodity futures contract, with a very large transaction volume and a wide range of strike prices and maturities having nontrivial trading volume. Although more liquid financial futures do exist, a commodity futures contract is more promising for an event study-type analysis as the effect of exogenous events on supply or demand can be identified more clearly. Based on an analysis of financial market news, it is more intuitive to identify relevant market events for crude oil than it would be for stock market indices. Indeed, options on futures reflect market expectations relating to realised and potential

international economic and geopolitical turmoil. Thus, it is likely to reflect changes in the representative investor's risk aversion following events affecting world oil markets.

To be sure, it should be acknowledged that oil is an imperfect proxy for wealth endowment, and our analysis may be subject the same critique as presented in Ziegler (2007). Note however that this critique also applies to stock market indices, which are only proxies for the true market portfolio. Nonetheless, using crude oil appears a reasonable choice for two reasons. First, the paper's contribution is to provide a richer and more robust framework for an event study analysis of risk aversion. Our interest lies in documenting and testing changes in risk aversion functions and their distribution, and less in their precise level. It is therefore essential to identify significant market shocks on the asset, whereas an aggregate index in contrast is diversified against idiosyncratic shocks. Second, the pro-cyclical character of crude oil price and its importance for the economy suggests it is a reasonable proxy for investor wealth endowment.

To estimate historical densities, gross returns of the WTI futures contract over a two-month horizon are used. Table 2 presents the empirical moments for the gross returns used in each estimated historical density. The average gross return for the whole period was slightly positive and above the risk free rate at 1.02. In general, the sample shows empirical evidence of negative skewness in gross returns, but surprisingly no evidence of excess kurtosis. This is likely because two month horizon gross returns are computed rather than daily net returns, which are typically discussed in the oil futures literature. Table 3 reports on changes in the empirical moments for the historical returns in each of the four event windows. The sign and magnitude of the changes depend on the nature of the shock. In particular, unexpected shocks entailing uncertainty about oil supply, such as the BP explosion or the Libyan uprising, induce a more positive skew in the density of historical returns. This also suggests an increased possibility of

shortage. Moreover, shocks related to systemic uncertainty in the stock market and broader economic conditions, such as the end of the bull cycle and the trough of credit freeze crisis, increase variance.

To recover the risk-neutral density, data on NYMEX American options written on the WTI futures contract are extracted from the Commodity Research Bureau database from June 2007 until April 2012. As our options data are American-style, not European-style, it is necessary to consider the effect of early-exercise value on option prices. It is well understood that because no dividends are yielded by options on commodity futures, early exercise of a call cannot ever be optimal. For put options, however, we appeal to the sufficient condition proposed independently by Carr, Jarrow and Myneni (1992), Jacka (1991) and Kim (1990) that ensures a negligible early-exercise value when the risk-free rate is particularly low. This condition is empirically verified for the sample period studied in the paper. Moreover, Pan (2011) computes early-exercise values for options on crude oil futures contracts to adjust option prices but finds that his results are unchanged as early-exercise values over the same period are negligible.

Prior to analysis, several steps must be taken to clean and filter the options dataset. First, the data need to be refined because the procedure does not allow for redundant assets. Following established practice (e.g., Rebonato, 2004, p.254), we exclude in-the-money options, which tend to be thinly traded. Out-of-the-money call options are used together with out-of-the-money put options. Using Put-Call parity, out-of-the-money puts are converted into in-the-money calls. Furthermore, options for which the implied volatility is negative or above 100% are excluded. Moreover, data for the last three days before maturity are excluded from the analysis for each contract. To avoid spurious effects from stale prices, we also exclude options that have a price of \$0.01. As a result, we obtain a set of 155,071 options for the full sample, out of which 57% are

call options and are 43% put options. Table 4 presents descriptive statistics related to options data for each event window. On average, our sample contains nearly 1000 option observations for each trading day. The number of active strikes ranges between 157 and 220 daily. As the WTI futures price decreases, the number of options and the range of strikes also falls.

4. Results

4.1 *Baseline Case*

Panels 1 through 4 present figures for the pre- and post-event historical and risk-neutral densities, pricing kernel and risk aversion function for each event. Following Jackwerth (2000), we present the mean distribution for both the fifteen trading days before and the fifteen days after each event, with empirical confidence intervals representing ± 0.5 standard deviations from the mean. As Jackwerth (2000) explains, the empirical confidence intervals are not meant to be used for purposes of statistical testing. Rather, they document the dispersion of the daily densities and further allow aggregating information from two sub-samples of fifteen days into one meaningful figure. We note that our estimated densities, pricing kernel and risk aversion functions are not always as smooth as those presented in related, previous studies (see e.g. Christoffersen, Heston and Jacobs, 2012; Jackwerth, 2000; Kang and Kim, 2008; Pérignon and Villa, 2002). The explanation has to do with the scope of the study and data limitations. We look at risk aversion functions and pricing kernels estimated over a short sample period before and after significant event dates. In contrast, the focus of the literature has typically been on using all available data to obtain the smoothest and most accurate estimates of the pricing kernel or risk aversion. To identify short-run changes due to market shocks, our samples are necessarily much smaller.

Thus, some accuracy is sacrificed in the interest of highlighting and testing for the existence and nature of changes.

The end of the commodity bull cycle and the trough of the credit freeze crisis are characterised by heavier tails in the historical densities after the event,⁷ reflecting higher volatility in the oil Futures price as reported in Table 3, as well as a greater likelihood of extreme returns. However, this increased volatility has only a marginal impact on the risk-neutral densities. Indeed, both events seem to “crystallise” anticipations in the economy and leave the distribution of the options market expectations relatively intact. The evidence suggests that although the market experienced an important change, the event was to some extent anticipated as the forward-looking densities (i.e., the investors’ expectations) were not substantially affected.

The estimated pricing kernel is U-shaped both before and after the event, consistent with much of the empirical literature (the “pricing kernel puzzle”, see e.g., Christoffersen, Heston and Jacobs, 2012; Jackwerth, 2004). Although the U-shaped pricing kernel is flatter after such “anticipated” events, the pricing kernel puzzle persists, as the U-shape and a small interval of negative values persist after the events.

For both anticipated events, we document an important change in the absolute risk aversion function. For the first event, the absolute risk aversion function is decreasing prior the event but nearly flat (i.e. constant) after the event. For the credit freeze bottom, the absolute risk aversion function also becomes flatter after the event, going from strongly decreasing to only lightly

⁷ The upper left quadrant of Panels 1 to 8 reports the historical densities using the empirical mean. However, as mentioned above, the historical density used for the calculation of ARA is demeaned and then fixed at the risk-free rate plus a 6% equity risk premium following Welch (2000, 2009).

decreasing. In both cases, the event seems to increase the propensity of a representative investor whose future wealth endowment is low to undertake more risk, a finding consistent with a “double or nothing” strategy and loss aversion theory (e.g. Benartzi and Thaler, 1995). On the other hand, a representative investor endowed with higher positive future wealth becomes more risk averse after the event and is less willing to expose his portfolio to additional risk.

The findings suggest that during end-of-cycle events, successful, wealthy investors will hold more liquid and riskless assets while investors with lower return portfolios will try to use the end of a cycle to recover from their loss. The end of the bull commodity cycle can be interpreted as a structural break not unlike the 1987 stock market crash, though less dramatic. In that sense, our results are consistent with those of Jackwerth (2000) who reported increasing risk aversion in wealth and negative values of risk aversion after the 1987 crash. Likewise, our results document a significant though less severe change in the shape of the risk aversion function, which becomes less decreasing (flatter) in wealth following the end of the bull market in commodity prices.

In contrast, the third and fourth events, namely the BP Deepwater Horizon explosion and the Libyan uprising, were unpredictable and unanticipated by the markets. We document that their effect on densities, pricing kernels and risk aversion functions was different than the effect of the first two, partly anticipated events. First, the effect of the two unanticipated events on the historical densities is limited. In the case of the Libyan crisis, the distribution shifts slightly to the right, reflecting a temporary increase in oil prices not entirely buffered by oil inventories. Given that both events were supply shocks affecting only a small part of worldwide production, this is not surprising. Indeed, Libyan production of crude oil accounts for about 0.6% of the world’s production (Energy Information Administration, 2012). Even if production was temporarily

halted, worldwide inventories of crude oil could lessen the effect of the shock. Indeed, the American Petroleum Institute (API) reported a decrease of more than a million of barrels of crude in the American inventories in that week.

While both events had limited and temporary effects on oil prices, they also added noise and uncertainty to option markets, due to a fear of possible disruptions in other crude oil-producing countries (for Libya) or other platforms (for BP). This is reflected in the important changes observed in the risk-neutral densities. Both densities exhibit fatter tails after the event and underscore changes in the nature of uncertainty in future oil supplies as perceived by investors. The pricing kernels for the third and fourth events do not change dramatically after each event, although the U-shape becomes steeper after the event. However, risk aversion functions do change after the events. Risk aversion is steeper (i.e., more decreasing in wealth) after the unanticipated events.

This finding suggests that a representative investor with poorer future wealth prospects will be more risk-averse after the event, but that an investor with good future wealth prospects has a greater appetite for risk. The latter investor is more willing to take risks to profit from the possibility of oil shortages.

4.2 Alternative Case under the Generalised Beta 2 Distributional Assumption

We present results under the distributional assumption of a GB2 density to show that the changes observed in risk aversion functions in the baseline case are not an artifact of its restrictive parametric assumptions. Panels 5 to 8 report figures for the historical and risk-neutral distributions, pricing kernel, and absolute risk aversion functions, before and after each event. A Beta density is used to parameterise the historical densities and a GB2 density is used for the

risk-neutral distribution. Although we find greater variation in the shape of the functions and the densities estimated, as well as larger empirical confidence intervals (especially for the pricing kernel), the results corroborate our findings for the baseline case.

First, the GB2 results confirm, under more flexible distributional assumptions, the findings for the historical densities, whose tails are heavier after the first two, partly anticipated events, but largely unchanged after the two unpredictable events. For the end of the commodity bull market cycle, the risk-neutral density is more concentrated around the mean after the event, suggesting once more that the market anticipated to some extent the end of the bull cycle and therefore “crystallised” investor perceptions. As for the credit freeze bottom, the risk-neutral density picks up more of a positive skew after the event, a finding that cannot be captured in the Gaussian normal baseline case and which can be interpreted as evidence that investors believed the trough had been reached.

Our GB2 findings for absolute risk aversion functions exhibit the same pattern as in the normal baseline case. For the end-of-cycle events, risk aversion functions are less decreasing in wealth after the events. For the BP platform explosion, the absolute risk aversion function before the event is highly non-monotonic when the GB2 density is used to estimate the risk-neutral density. However, this distortion is due to problematic observations in two of the fifteen trading days before the event. Risk-neutral densities estimated for these two dates were not robust, as parameter solutions obtained from the optimisation routine were sensitive to the starting values used.⁸ Nonetheless, both risk aversion figures for the unanticipated events suggest similar

⁸ Results are smoother when the two problematic days are taken out of the analysis. However, we report the results without manipulations given there was no clear indication that these days were measurement errors. It is a limitation

changes in risk aversion as in the normal baseline case. That is, after the event a representative agent with lower future wealth is more risk-averse, while a representative agent with higher future wealth has a greater appetite for risk.

The pricing kernels estimated for days around all four events have a U-shape as found in the baseline case and much of the empirical literature. However, estimation of GB2 risk-neutral densities leads to pricing kernels that are less smooth and which have larger empirical confidence intervals. This is because the GB2 distributional assumption allows for greater flexibility in the parameterisation of the densities which, combined with the relatively small number of observations in a given day, generates more dispersion among densities estimated for the 15 trading days before and after the events, and thus among pricing kernels, than in the baseline case. We acknowledge the lack of precision as a limitation of our short window approach, especially when the parameterisation process is more flexible.

The pricing kernel does not change dramatically after each of the events, especially if we focus on the interval of more plausible values of future wealth in the short term horizon between, e.g. future wealth between 0.85 and 1.20. Interestingly, the pricing kernel U-shape is different under the GB2 assumption than in the normal baseline case. Indeed, the function is highly non-monotonic in wealth for extreme values of future wealth, below 0.2 and above 1.8, but it is relatively flat and monotonically decreasing over the interval between these extremes. Given that our results are for a two-month horizon, it is fair to say gross returns below 0.2 and above 1.8 are highly unlikely. Therefore, although the pricing kernel puzzle is not resolved by the use of the

of our analysis using GB2 that for these two days, the optimisation process is not robust to change in the starting values.

flexible GB2 distributional approach, given its U-shape and interval of negative values, it is however monotonic over an interval of plausible values of future wealth for a short time horizon.

Finally, given that the GB2 findings relied on cross-sections of options data using only options data corresponding to a two-month maturity, we perform as a robustness check the same analysis using options data corresponding to a one-year time-to-maturity. The results are largely similar to those for the two month horizon.⁹

4.3 Robustness analysis

Table 5 reports results for non-parametric specification tests on the historical densities. The null hypothesis is that the prior and post-event historical densities are drawn from the same continuous distribution. The Kolmogorov-Smirnov and Cramer-Von-Mises tests reject the null in the first two events at a 5% confidence level, supporting the claim of a change in the historical densities as reported in panels 1, 2, 5 and 6. For the unexpected events, the null is also rejected at the 5% level for the BP explosion event, but not for the Libyan uprising event. These tests support our finding that most of the effect of events on the representative investor's risk attitudes after unexpected events is captured by the risk-neutral density.

A Kolmogorov-Smirnov test of normality of returns is presented in Table 5 for the historical densities as an indication of whether the baseline case assumption is appropriate. Although normality is rejected for the complete sample, it is not rejected at the 5% level for each of the sub-periods considered. This result has to be treated with caution, however, since the Kolmogorov-Smirnov test is known to possess relatively low power.

⁹ For brevity, only results for the two month horizon are reported, as results for the one-year horizon were largely similar. Results for the one year horizon are available upon request.

Table 6 presents the results for the Kolmogorov-Smirnov and Cramer-Von-Mises tests of equality of the prior and post-event densities using the fitted risk aversion values computed from equation (19). The equality of the distribution of fitted risk aversion values before and after the event is strongly rejected for the whole sample. However, as the rejection could be due to the large number of option observations in the pooled periods before and after each event, we also perform hypothesis tests on samplings drawn without replacement based on the full sample. Table 6 presents the proportion of rejections of the null hypothesis over the total number of samplings drawn. The samplings, using 100, 250 and 500 observations each, strongly reject the null for all events. This result does not vary with the number of samplings performed (M=1000, 2000 and 5000 samplings). Thus, we conclude that changes in risk aversion after the event as measured by a proxy distribution of risk aversion values are statistically significant.

To document changes observed in risk aversion values, fitted ARA observations are regressed on a proxy for endowment and the square of endowment. Table 7 presents the results for the whole sample of fitted ARAs and for smaller random samplings of the fitted ARAs. Two hypotheses (H_0^{sum} and H_0^{joint}) test the global change in the relationship between risk aversion and future wealth. Hypotheses H_0^S and H_0^{SS} consider whether the slope or curvature of the risk aversion function drive the adjustment process after the event. All tested hypotheses are strongly rejected for the full sample for every event.

To analyse results from the samplings, we adopt a conservative approach and set the rejection hurdle at 70% rejections while the non-rejection rate is chosen to be less than 30%. Test outcomes between 30% and 70% are considered inconclusive. First, all the hypotheses tested are rejected for the end of the commodity bull cycle event. Both the slope and curvature coefficients change significantly after the first event. Secondly, results for the credit freeze bottom show that

most of the adjustment in the risk aversion comes from the curvature coefficient. For instance, in the case of 200 observations and 10,000 samplings, we reject at the 1% level the null hypothesis of equality in the endowment coefficient (H_0^S) only 22% of the time, while we reject the null hypothesis of equality of the curvature coefficient (H_0^{SS}) in 71% of the samplings. In contrast, for unpredictable events (H_0^S) is rejected more than 70% of the time when $N=200$. For the BP event, the equality of the curvature coefficients cannot be rejected 28% of the time while the rejection rate is around 50% for the Libyan uprising. Taken collectively, these results suggest that the global effect (tested H_0^{joint} and H_0^{sum}) is statistically significant. Further results are needed to conclude on the adjustment process of the risk aversion function, but our results suggest there is a greater adjustment of the risk aversion smile than of the slope coefficient after a foreseeable event. For unexpected events, we find the reverse adjustment process.

4.4 Discussion of the results

Our analysis builds on Jackwerth's (2000) finding that major stock market events (the 1987 crash) affect investor risk aversion, by generalising the result to show that several market and political events have significant but differential effects on investor risk aversion. The results from the baseline case show that infrequent events affect risk aversion, but that the pricing kernel remains relatively unchanged. Thus, infrequent events have an impact on the representative investor's sentiment toward risk, but do not change how he values future over current wealth. Predictable events have a different impact on densities than do unforeseeable events. In predictable events, the historical density changes while the risk-neutral density remains the same after the event. Indeed, the information provided by the realization of such events is not "news" and the distribution of forward-looking expectations is unchanged. For unpredictable events,

however, the historical density is unchanged after the event, but the risk-neutral density shows fatter tails, reflecting greater uncertainty about extreme events and their likelihood.

Our results also confirm that the pricing kernel puzzle documented in equity and Forex markets is also found empirically in commodity futures markets. The nature of the shock influences how the risk aversion function adjusts to the event. A representative investor with a negative or smaller endowment has a greater appetite for risk (“risk-loving”) after predictable events, “doubling-down” on his strategy, but becomes more risk-averse after unexpected shocks. In contrast, a wealthier investor becomes more risk-averse after predictable shocks but more of a “risk-lover” after unforeseeable events. The theoretical concepts of loss aversion (e.g., Benartzi and Thaler, 1995) and ambiguity aversion (e.g., Mukerji and Tallon, 2001) may be useful to help reconcile the findings, but testing for their evidence in the data is beyond the scope of the paper.

Results obtained under the GB2 risk-neutral distributional assumption show a greater dispersion and less smoothness, likely a side effect from estimating more parameters and allowing for more flexibility in density estimation. However, the findings support the conclusions obtained in the baseline case regarding risk aversion. Moreover, the introduction, through the GB2 distribution, of greater flexibility in the parameterisation of the historical and risk-neutral densities may help explain the pricing kernel puzzle for a range of plausible values of future wealth for a short-term horizon.

Our robustness analysis supports the claim of statistically significant changes in the fitted risk aversion densities for all events as well as in the historical densities of the first three events. Hypothesis testing done using multivariate regression analysis also suggests that the events under consideration in this paper have a statistically significant impact on risk aversion. For anticipated

events, the adjustment of risk aversion functions appears to be mostly driven by the curvature coefficient of wealth. For unanticipated events, however, the slope coefficient of wealth seems to change more than does the curvature coefficient.

5. Conclusion

In this paper, we have investigated the impact of four major market events on the real-world and risk-neutral price densities, pricing kernel, and risk aversion functions derived from NYMEX WTI contract crude oil options on futures over the period 2007-2011. To provide a detailed event study-type analysis, we consider both a baseline Gaussian normal and an alternative Generalised Beta 2 case. Using estimated pricing kernels and risk aversion functions together with distributional and hypothesis tests and resampling robustness checks, we document significant and differential changes in the risk aversion function following each of the events.

First, we find that end-of-cycle events (i.e., the end of the bull commodity market and the bottom of the credit freeze) affect the historical density more than the risk-neutral density. Following these events, the risk aversion function is flatter (less decreasing) in wealth, such that a representative investor facing a poor future wealth endowment has a greater appetite for risk (tending toward risk-neutral and “risk-loving” preferences), while an investor with a good future wealth endowment becomes more risk-averse.

Second, unpredictable events affect the risk-neutral density more than the historical density. As a result, absolute risk aversion functions become steeper (more decreasing), such that risk aversion increases for investors facing poor future wealth outcomes, but decreases for investors with good financial perspectives, tending to risk neutrality or “risk-loving” preferences. In contrast, events do not significantly affect pricing kernels, as the U-shape, empirically

documented in options data for many equity and Forex markets, is present both before and after the event. It appears therefore that major shocks to the market affect risk sentiment but not preferences for inter-temporal substitution of wealth. Therefore, the pricing kernel puzzle is not resolved by looking at the short-term effect of major market events. However, a more flexible parameterisation generates pricing kernels that are more monotonic over a range of plausible future wealth levels for a short investment horizon. The statistical robustness checks provided in this paper support our main finding that rare events have a statistically significant and differential effect on risk aversion functions. In line with Ziegler (2007), we note that further work could be done in order to relax several of the assumptions made in this literature, and to further evaluate the robustness of the findings as well as their potential limitation.

REFERENCES

- Ait-Sahalia, Y., Lo, A.W., 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9-51.
- Aparicio, S. D., Hodges, S.D., 1998. Implied Risk-Neutral Distribution: A Comparison of Estimation Methods. FORC Preprint, University of Warwick.
- Anagnou-Basioudis, I., Bedendo, M., Hodges, S.D. Tompkins, R., 2005. Forecasting Accuracy of Implied and GARCH-Based Probability Density Functions. *Review of Futures Markets* 11, 41-66.
- Arrow, K.J., 1971, *Essays in the Theory of Risk bearing*. Markham Publishing Company, Chicago.
- Bahra, B., 1997, Implied risk-neutral probability density functions from options prices: Theory and application. Working paper, The Bank of England.
- Bakshi, G., Madan, D., Panayotov, G., 2010. Returns of claims on the upside and the viability of U-shaped pricing kernels. *Journal of Financial Economics* 97, 130-154.
- Barthunek, K.S., Chowdury, M., 1997. Implied Risk Aversion Parameter from Option Prices. *The Financial Review* 32, 1, 107-124.
- Barone-Adesi, G., Mancini, L., Shefrin, H. M., 2012. Sentiment, Asset Prices, and Systemic Risk. Swiss Finance Institute Research Paper No. 11-50.
- Barro, R., 2006. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics* 121, 3, 823-866.

Benartzi, S., Thaler, R.H., 1995. Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics* 110, 75-92.

Berger, J.O., Sellke, T., 1987. Testing a point null hypothesis: the irreconcilability of P values and evidence. *Journal of the American Statistical Association* 82, 112-122.

Birru, J., Figlewski, S., 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. 1976. The pricing of commodity contracts. *Journal of Financial Economics* 3, 167-179.

Bliss, R., Panigirtzoglou, N., 2004. Option-implied risk aversion estimates, *Journal of Finance* 109, 1, 407-446.

Bookstaber R.M., McDonald, J.B., 1987. A general distribution for describing security price returns. *Journal of Business* 60, 3, 401-424.

Breeden, D., Litzenberger, R., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621- 651.

Chabi-Yo, F., Garcia, R., Renault, E., 2008. State dependence can explain the risk aversion puzzle. *Review of Financial Studies*, 21, 2, 973-1011.

Carr, P., Jarrow, R., Myneni, R., 1992. Alternative characterizations of American put options. *Mathematical Finance* 2, 87–106

Christoffersen, P., Heston, S., Jacobs, K., 2012, Capturing option anomalies with a variance-dependent pricing kernel, working paper.

- Cox, J.C., Ross, A., 1976. The valuation of options for alternative stochastic processes, *Journal of Financial Economics* 3, 1, 145-166.
- Coutant, S., Jondeau, E., Rockinger, M., 2001. Reading PIBOR futures options smiles: the 1997 snap election. *Journal of Banking and Finance* 25, 11, 1957-1987.
- Dittmar, R.F., 2002. Nonlinear pricing kernels, kurtosis preference, and evidence from the Cross section of equity returns. *Journal of Finance* 57, 1, 369-403.
- Dumas, B., Fleming, J., Whaley, R., 1998. Implied Volatility Function: Empirical Test. 53, 6, 2059-2106.
- Ewing, B., Thompson, M.A., 2007. Dynamic cyclical comovements of oil prices with industrial production, consumer prices, unemployment, and stock prices. *Energy Policy* 35, 11, 5535–5540.
- Engle, R.F., Rosenberg, J. V., 2002. Empirical pricing kernels. *Journal of Financial Economics* 64, 341–372.
- Figlewski, S., 2008. Estimating the implied risk neutral density for the U.S. market portfolio. In: *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, Bollerslev, J., R. Russell, M. Watson (eds.), Oxford University Press, Oxford, UK.
- Flamouris, D., Giamouridis, D., 2002. Estimating implied PDFs from American options on futures: A new semiparametric approach. *Journal of Futures Markets* 22,1, 1-30.
- Giamouridis, D., 2005. Inferring option-implied investors' risk preferences, *Applied Financial Economics* 15, 7, 479-488.
- Hens, T., Reichlinz, C., 2012. Three solutions to the pricing kernel puzzle, working paper.

- Jacka, S., 1991. Optimal stopping and the American put. *Mathematical Finance* 1, 1–14.
- Jackwerth, J., 1999. Option implied risk-neutral distributions and implied binomial trees: A literature review. *Journal of Derivatives* 7, 2, 66-82.
- Jackwerth, J., 2000. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13, 2, 433-451.
- Jackwerth, J., 2004. Option-Implied Risk-Neutral Distributions and Risk Aversion. The research foundation of the Association for Investment Management and Research, United States.
- Jackwerth, J., Rubinstein, M. 1996. Recovering probability distributions from option prices. *Journal of Finance* 51, 5, 1611-1631.
- Jarrow, R., Rudd, A., 1982. Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics* 10, 3, 347-369.
- Jondeau, E., Poon, S.H, Rockinger, M. (eds.), 2007. *Financial Modelling Under non-Gaussian Distributions*. Springer Finance.
- Kang, B.J., Kim, T. S., 2008. Empirical risk aversion functions-estimates and assessment of their reliability. *International Review of Financial Analysis* 17, 1123-1138.
- Kilian, L., 2009. “Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99, 3, 1053-1069.
- Kim, I. J., 1990. The analytic valuation of American options. *Review of Financial Studies* 3, 547-572.
- Leland, H. E., 1980. Who should buy portfolio insurance? *Journal of Finance* 35, 2, 581-594.

Liu, X., Shackleton, M. B., Taylor, S. J., Xu, X., 2007, Closed-form transformations from risk-neutral to real-world distributions. *Journal of Banking and Finance* 31, 1501- 1520.

Liu, X., Shackleton, M. B., Taylor, S. J., Xu, X., 2009, Empirical pricing kernels obtained from the UK index options market. *Applied Economics Letters* 16, 10, 989-993.

McDonald, J. B. , Xu,, Y.J., 1995. A generalization of the Beta distribution with applications. *Journal of Econometrics* 66, 133-152.

Melick, W., Thomas, C., 1997. Recovering an asset's implied SPD from option prices: An application to crude oil during the gulf crisis. *Journal of Financial and Quantitative Analysis* 32, 91-115.

Merton, R.C. , 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8, 323-361.

Moitra, S.D., 1990. Skewness and the Beta distribution. *Journal of the Operational Research Society* 10, 953–961.

Mukerji, S., Tallon, J.M., 2001. Ambiguity Aversion and Incompleteness of Financial Markets. *Review of Economic Studies* 68, 4, 883-904.

Pan, X., 2011. Investor beliefs and state price densities in the crude oil market. Working paper, McGill University Desautels Faculty of Management.

Pérignon, C., Villa, C., 2002, Extracting information from options markets: Smiles, state-price densities and risk aversion. *European Financial Management* 8, 4, 495–513.

Rebonato, R., 2004. *Volatility and Correlation*. John Wiley & Sons.

Rosenberg, J.V., Engle, R.F., 2002. Empirical pricing kernels. *Journal of Financial Economics* 64, 3, 341-372.

Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 49, 3, 771-818.

Sadorsky, P., 2001. Risk factors in stock returns of Canadian oil and gas companies. *Energy Economics* 23, 1, 17–28.

Shackleton, M.B., Taylor, S. J., Yu, P., 2010. A multi-horizon comparison of density forecasts for the S&P 500 using index returns and option prices. *Journal of Banking and Finance* 34, 2678–2693.

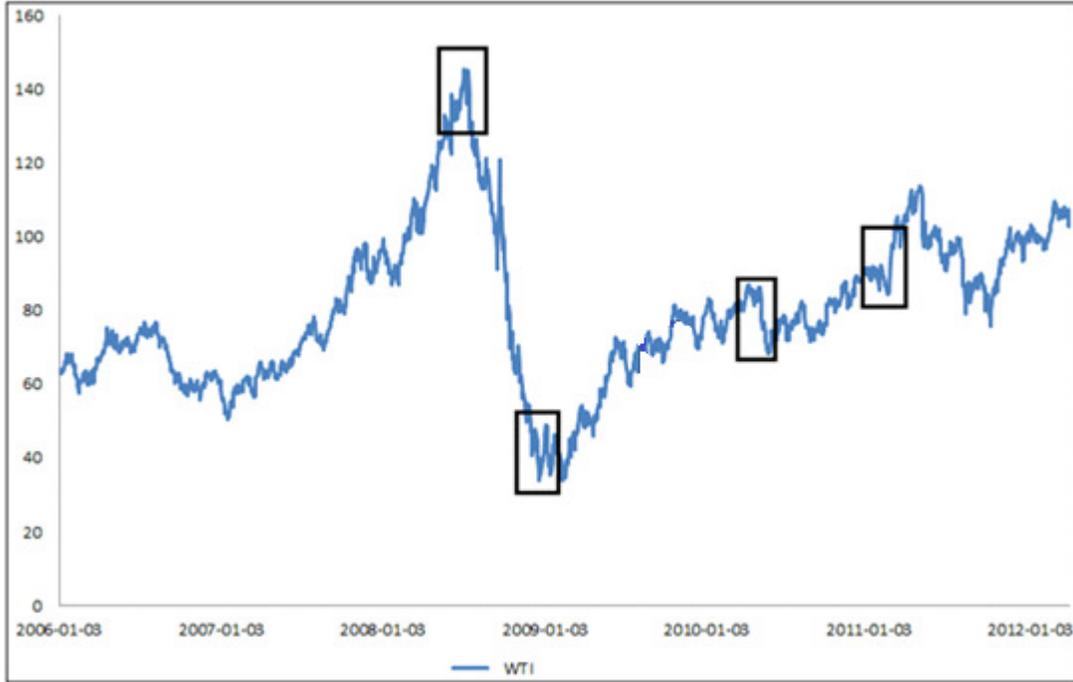
Shimko, D.C., 1993. Bounds of probability, *Risk Magazine* 6, 4, 33-37.

Welch, I., 2000. Views of financial economists on the equity premium and other issues. *Journal of Business* 73, 4, 501-537, with 2009 update.

Ziegler, A., 2007. Why does implied risk aversion smile?. *Review of Financial Studies* 20, 3, 859-904.

Figure 1

The historical price of the Future contract on the WTI during the 2006 -2012 period.



Note: Figure 1 presents the price of the WTI futures oil contract from January 2006 to January 2012. The four rectangle windows highlight the events in this paper. The first window is the end of the commodity bull cycle. The second window is the lowest point in the global 2008 credit freeze. The third window is the explosion of the Deepwater Horizon platform and finally, the fourth window is the beginning of the uprising in the Tripoli.

Table 1
The event windows

Window	End of the commodity bull cycle	Credit freeze bottom	Deepwater Horizon explosion	Lybian uprising
15 days prior the event	2008-06-12	2008-11-28	2010-03-29	2011-01-31
day of the event	2008-07-03	2008-12-19	2010-04-20	2011-02-22
15 days after the event	2008-07-25	2009-01-13	2010-05-11	2011-03-15

Note: Table 1 reports the time windows considered to estimate the historical and risk neutral densities before and after each of the four events considered.

Table 2**Descriptive statistics for the NYMEX WTI future price**

	Full period	End of the commodity bull cycle	Credit freeze bottom	Deepwater Horizon explosion	Libyan uprising
Number of observations	916	251	119	334	214
Minimum	0.46	0.80	0.46	0.84	0.84
Maximum	1.48	1.31	1.07	1.48	1.26
Mean	1.03	1.10	0.71	1.08	1.06
St. dev.	0.17	0.11	0.13	0.13	0.08
Skewness	-0.94	-0.61	0.35	0.67	-0.07
Kurtosis	4.25	2.85	2.50	3.39	2.79

Note: Table 2 presents the descriptive statistics for gross returns over a two month investment period of the WTI Future contract over the different sub-periods considered. Descriptive statistics for the full period and for the non-overlapping period prior the event used to obtain the historical density are presented. The first line reports the number of observations (trading days) used to obtain the historical densities for each event. The data on the WTI contract are extracted from Bloomberg.

Table 3**Empirical moments for the two month horizon gross return for each event window**

	End of the commodity bull cycle			Credit Freeze bottom			BP platform			Libyan crisis		
	prior	event	post	prior	event	post	prior	event	post	prior	event	post
mean	1.11	1.10	0.72	0.69	0.70	0.72	1.09	1.08	1.07	1.06	1.07	1.08
st. Dev.	0.11	0.12	0.13	0.13	0.14	0.17	0.12	0.13	0.13	0.07	0.07	0.07
skew.	-0.56	-0.71	-0.83	0.38	0.49	0.74	0.74	0.77	0.83	-0.12	0.20	0.15
kurt.	2.81	2.90	3.01	2.42	2.42	2.83	3.48	3.57	3.71	2.81	2.66	2.60

Note: Table 3 presents the empirical moments of the WTI future gross returns in the event windows. The empirical moments are computed using daily rolling windows. The prior column reports the mean of the daily empirical moments for a two month horizon of the empirical moments for the 15 days prior to the event. The event column reports the empirical moments based the two month horizon window the day of the event. The post column reports the mean of the daily empirical moments for a two month horizon of the empirical moments for the 15 days after the event.

Table 4
Descriptive statistics for the options data used to obtain the pre- and post- event risk-neutral densities.

		end of the commodity bull market	credit freeze bottom	BP deepwater horizon	Libyan uprising
number of options	mean	1195	1115	821	993
	min.	1122	1003	785	930
	max.	1247	1198	889	1083
strike price	mean	123.42	96.77	95.57	101.41
	min.	35.00	91.89	94.82	98.37
	max.	300.00	102.43	96.79	104.87
Implied volatility	mean	0.35	0.54	0.27	0.28
	min.	0.01	0.04	0.01	0.05
	max.	0.57	1.00	0.48	0.75
number of strike	mean	216.39	238.58	162.10	182.16
	min.	211.00	216.00	157.00	172.00
	max.	220.00	252.00	171.00	191.00
time to maturity	mean	0.93	0.96	0.87	0.75
	min.	0.12	0.13	0.13	0.12
	max.	4.56	4.98	3.76	2.92
number of maturities	mean	52.50	47.74	32.48	33.39
	min.	52	47	31	33
	max.	54	49	35	34

Note: Table 4 presents descriptive statistics for options on the WTI future contract used to obtain the risk neutral densities. The data are extracted from the Commodity Research Bureau database. Reported in Table 4 are the means of the daily means for the options inside the event windows.

Table 5
Statistical Tests on the historical densities of returns

	Full period	End of the commodity bull cycle	Credit freeze bottom	BP Deepwater Horizon explosion	Libyan uprising
Kolmogorov-Smirnov Normality test	.00	.05	.22	.85	.49
Kolmogorov-Smirnov Test of equality		.04	.01	.00	.31
Cramer-Von-Mises test of equality		.03	.02	.01	.03

Note: Table 5 presents p-values for several tests on the historical densities. The first row reports a Kolmogorov Smirnov test of normality performed on the historical returns leading to event. The second row presents the p-value for the two samples Kolmogorov-Smirnov. The null hypothesis is that the historical densities of the two months returns prior the event is drawn from the sample continuous distribution as the post event historical densities. The last row shows the p-values for the Cramer-Von-Mises statistic for a two-sample density test.

Table 6
Statistical tests on implied absolute risk aversion distributions

		P-values full sample	Simulations using the sampling without replacement									
			N=1000			N=2000			N=5000			
			100	250	500	100	250	500	100	250	500	
event 1	ks	.00	100	100	100	100	100	100	100	100	100	100
	cvm	.00	100	100	100	100	100	100	100	100	100	100
event 2	ks	.00	77	100	100	77	100	100	77	100	100	100
	cvm	.00	74	100	100	77	100	100	76	100	100	100
event 3	ks	.00	100	100	100	100	100	100	100	100	100	100
	cvm	.00	100	100	100	100	100	100	100	100	100	100
event 4	ks	.00	98	100	100	97	100	100	97	100	100	100
	cvm	.00	48	100	100	49	99	100	46	99	100	100

Table 6 presents the Kolmogorov-Smirnov (ks) and Cramer Von Mises (cvm) test of two sample densities. The tests are performed on the implied densities of the fitted ARA from (19). We obtain the implied absolute risk aversion function by mapping each option before and after the event to a level of wealth based on the moneyness of each option. The null hypothesis in both tests is that samples before and after event are drawn from same continuous distribution. First, we report the p-values for the test using the full distribution of ARA. We also report the results for a number of samplings without replacement with smaller samples. The number reported is the percentage of rejection of the null hypothesis for each case at a 1% level of significance over the total number of draws.

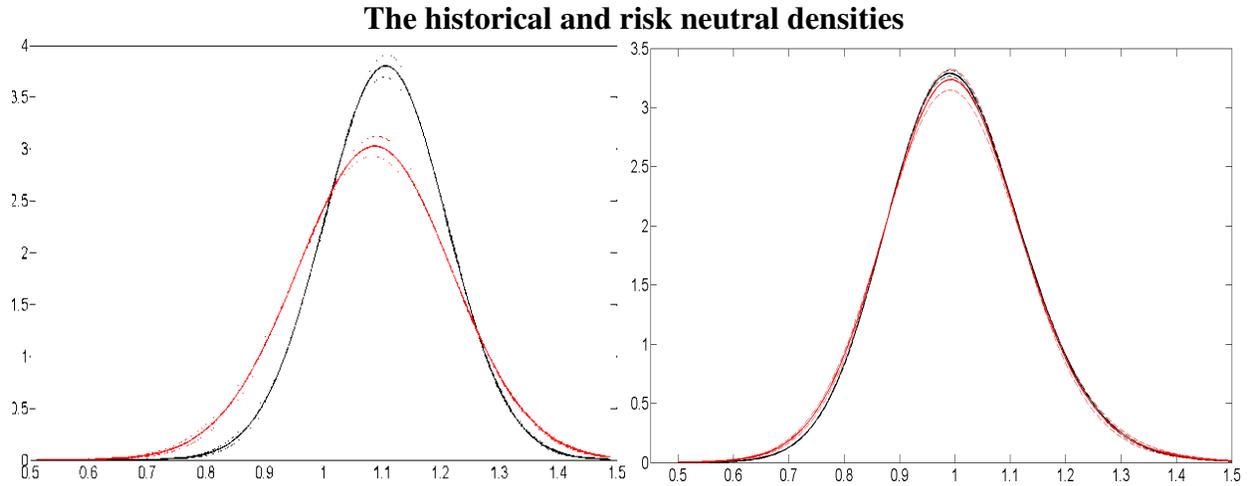
Table 7
Tests on the coefficients of the fitted absolute risk aversion function

		H_0^{joint}	H_0^S	H_0^{SS}	H_0^{sum}
event 1					
M=2000	N= 100	100	100	93	99
	N= 200	100	100	99	100
M=10000	N= 100	100	100	93	99
	N= 200	100	100	98	100
p-values full sample	Nbef= 17 547 Naft=18 273	.00	.00	.00	.00
event 2					
M=2000	N= 100	100	14	61	96
	N= 200	100	21	72	100
M=10000	N= 100	100	14	60	96
	N= 200	100	22	71	100
p-values full sample	Nbef= 16 240 Naft= 17 538	.00	.00	.00	.00
event 3					
M=2000	N= 100	100	83	27	38
	N= 200	100	97	28	41
M=10000	N= 100	100	84	28	39
	N= 200	100	98	27	44
p-values full sample	Nbef= 17 547 Naft=18 273	.00	.00	.00	.00
event 4					
M=2000	N= 100	63	62	46	57
	N= 200	78	77	52	64
M=10000	N= 100	62	61	45	56
	N= 200	78	78	52	65
p-values full sample	Nbef=13948 Naft= 15010	.00	.00	.00	.00

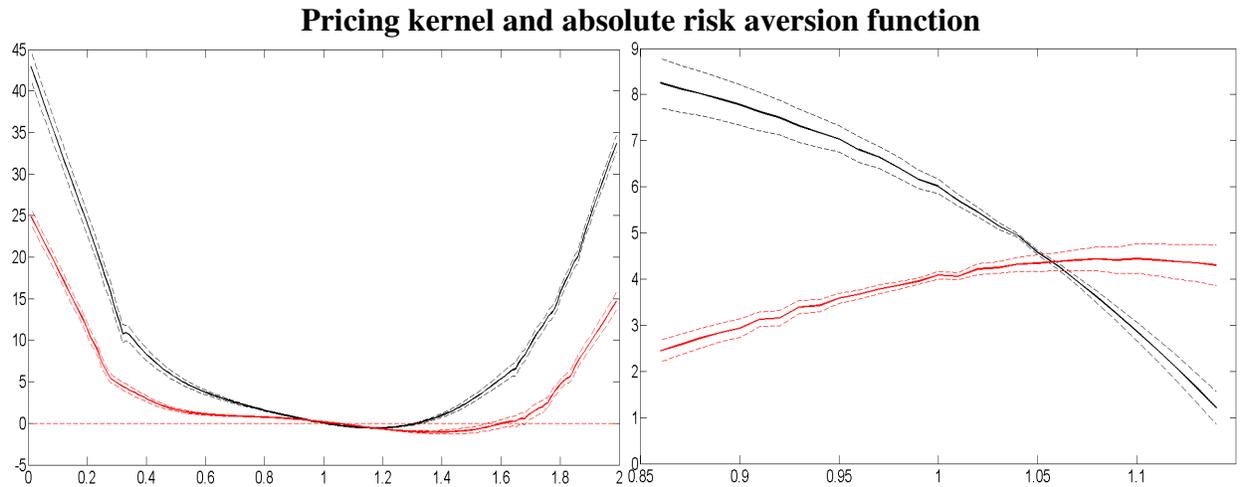
Table 7 presents the outcome for each event of the F test on hypotheses (21)-(24) performed following regression (20). Reported in Table 7 is the percentage of rejection of the null hypothesis for each case scenario at a 1% level of significance over the total number of draws. Also, we report the p-values of the test performed on the whole sample rather than on draws. H_0^{joint} tests the joint equality of all the coefficients in the two equation of (24). H_0^{sum} tests that the global effect of wealth (slope and curvature) is the same before and after the event. H_0^S tests the equality of the slope coefficients on future wealth in (20) and H_0^{SS} tests the equality of the curvature coefficients before and after the event.

The Gaussian baseline case

PANEL 1: Historical densities, risk neutral densities, pricing kernel and the mean absolute risk version functions for the end of the bull commodity cycle



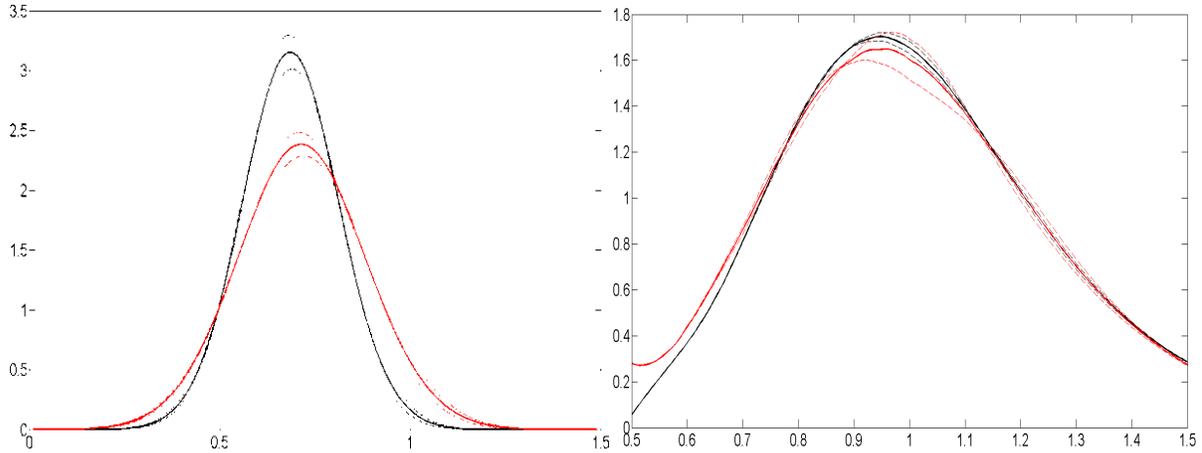
Note: the left figure represents the historical densities of the WTI futures and the right figure represents the risk neutral densities estimated with the options on the WTI future third nearby contract. The parametric assumption is the normality of the historical and risk neutral densities returns. The event occurs on July 3, 2008 and captures the end of the commodity bull cycle. The solid black curves represent the mean of the daily densities during the 15 trading days prior to the event and the solid red curves represent the mean of the daily densities during 15 trading days after the date of the event. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).



Note: the left figure represents the mean pricing kernel and the right figure represents the mean absolute risk aversion estimated around the event date using the historical WTI price and related options. The parametric assumption is the normality of the historical and risk neutral densities returns. The event occurs on July 3, 2008 and captures the end of the commodity bull cycle. The solid black curves represent the mean of the daily function during the 15 days prior to the event and the solid red curves represent the mean of the daily functions during the 15 trading days after the event date. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

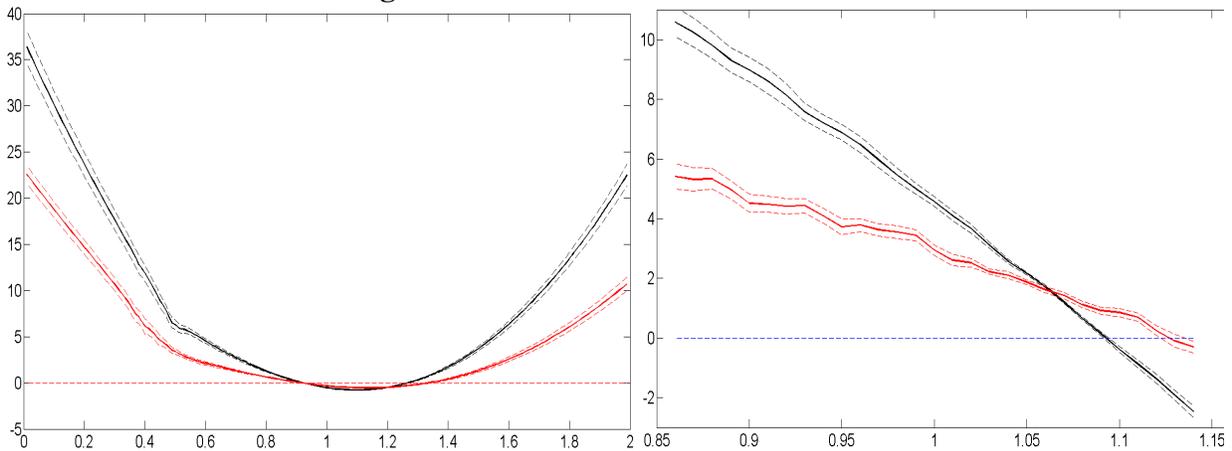
PANEL 2: Historical densities, risk neutral densities, pricing kernel and the mean absolute risk version functions for the credit freeze bottom.

The historical and risk neutral densities



Note: the left figure represents the historical densities of the WTI futures and the right figure represents the risk neutral densities estimated with options on the WTI future third nearby contract. The parametric assumption is the normality of the historical and risk neutral densities returns. The event occurs on December 19, 2008 and captures credit freeze bottom. The solid black curves represent the mean of the daily densities during the 15 trading days prior to the event and the solid red curves represent the mean of the daily densities during 15 trading days after the date of the event. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

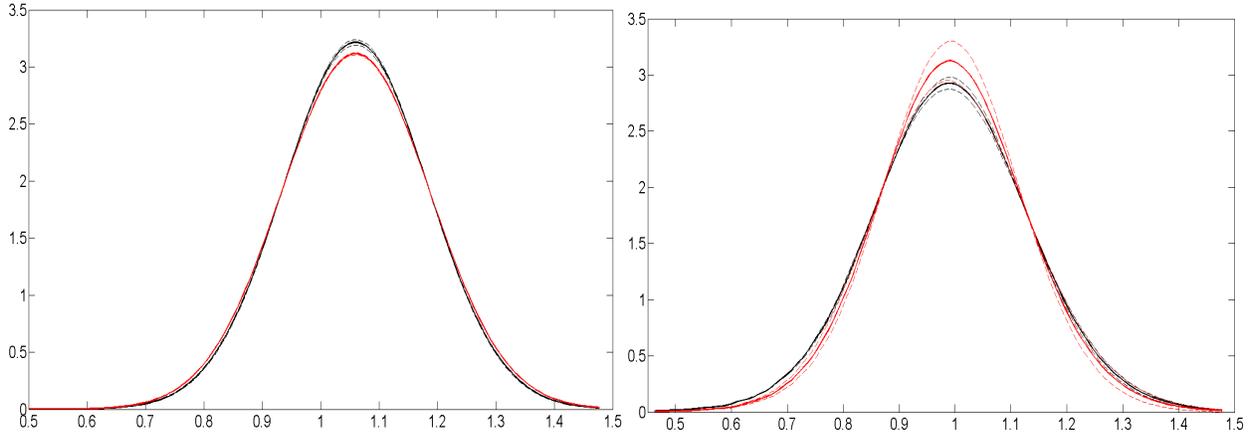
Pricing kernel and absolute risk aversion function



Note: the left figure represents the mean pricing kernel and the right figure represents the mean absolute risk aversion derived around the event date using the historical WTI price and related options. The parametric assumption is the normality of the historical and risk neutral densities returns. The event occurs at date on December 19, 2008 and captures credit freeze bottom. The solid black curves represent the mean of the daily functions during the 15 days prior to the event and the solid red curves represent the mean of the daily functions during the 15 trading days after the event date. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

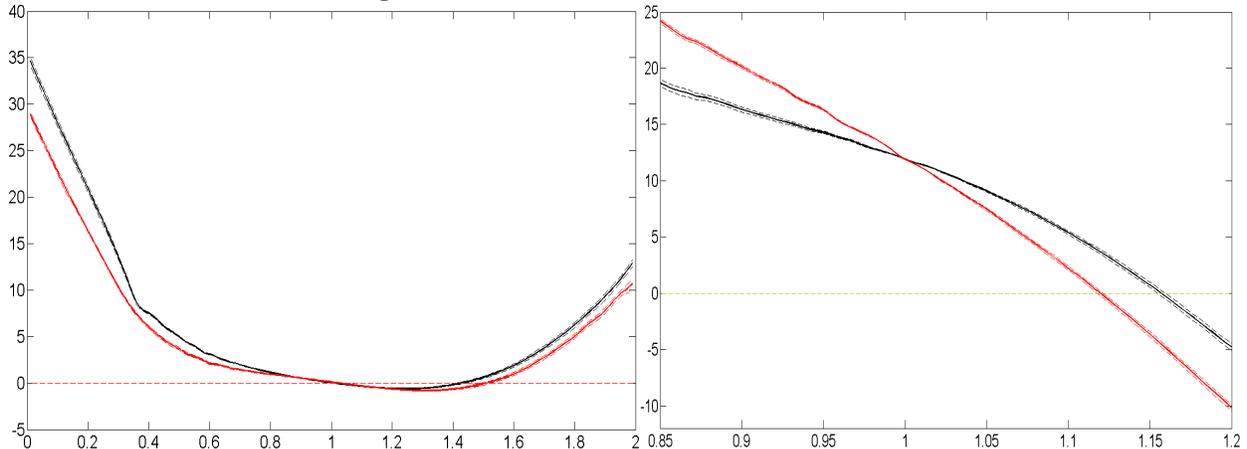
PANEL 3: Historical densities, risk neutral densities, pricing kernel and the mean absolute risk version functions for the explosion of the BP Deepwater Horizon platform.

Historical and risk neutral densities



Note: the left figure represents the historical densities of the WTI futures and the right figure represents the risk neutral densities estimated using options on the WTI future third nearby contract. The parametric assumption is the normality of the historical and risk neutral densities returns. The event occurs on April 20, 2010 and captures the explosion of the Deepwater horizon platform. The solid black curves represent the mean of the daily densities during the 15 trading days prior to the event and the solid red curves represent the mean of the daily densities during 15 trading days after the date of the event. The dashed lines represent the associated empirical standard error (the mean ± 0.5 standard deviation).

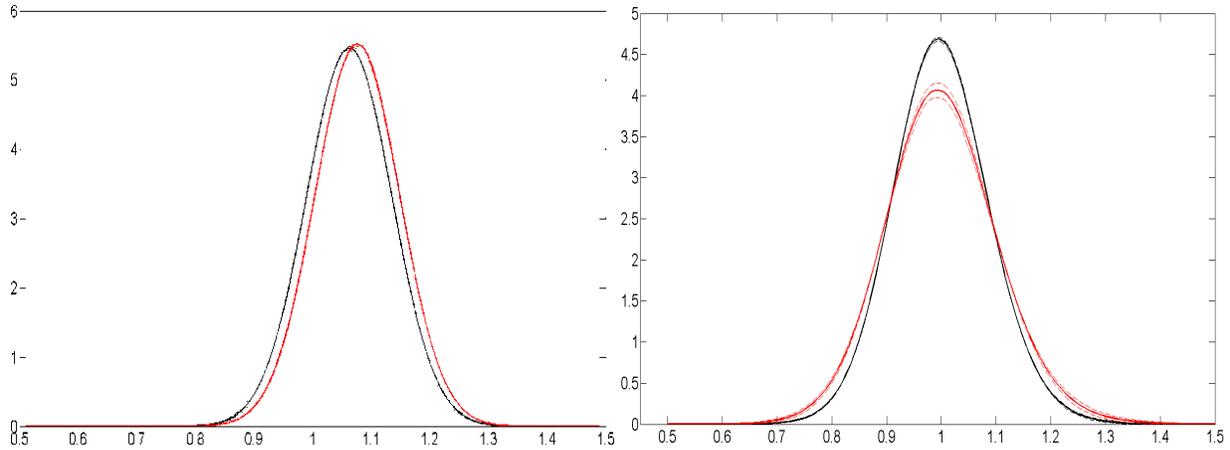
Pricing kernel and absolute risk aversion function



Note: the left figure represents the mean pricing kernel and the right figure represents the mean absolute risk aversion derived around the event date using the historical WTI price and related options. The parametric assumption is the normality of the historical and risk neutral densities returns. The event occurs on April 20, 2010 and captures the explosion of the Deepwater horizon platform. The solid black curves represent the mean of the daily functions during the 15 days prior to the event and the solid red curves represent the mean of the daily functions during the 15 trading days after the event date. The dashed lines represent the associated empirical standard error (the mean ± 0.5 standard deviation).

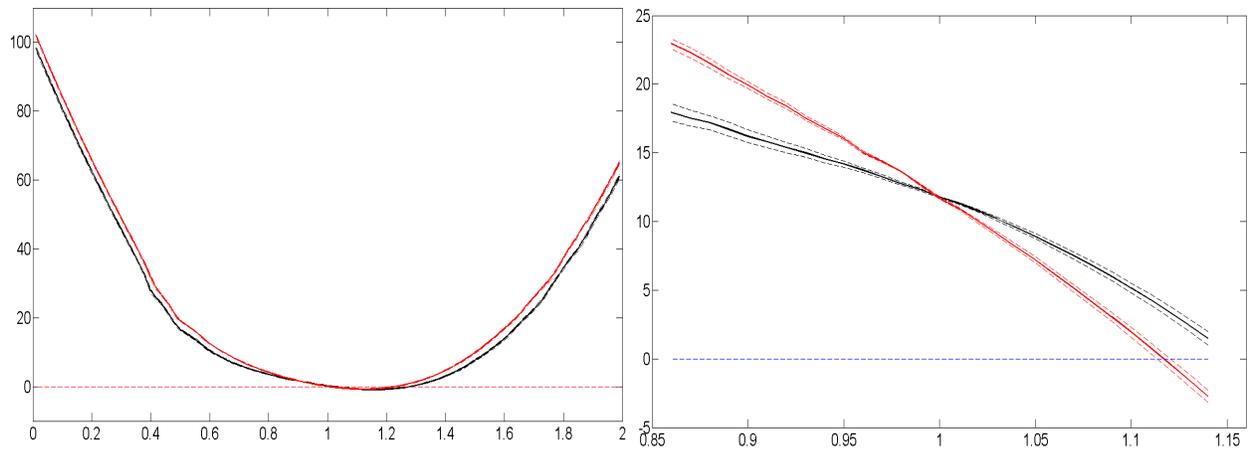
PANEL 4: Historical densities, risk neutral densities, pricing kernel and the mean absolute risk version functions for the Libyan uprising.

Historical and risk neutral densities



Note: the left figure represents the historical densities of the WTI futures and the right figure represents the risk neutral densities estimated using options on the WTI future third nearby contract. The parametric assumption is the normality of the historical and risk neutral densities returns. The event occurs on February 22, 2011 and captures the uprising in Libya. The solid black curves represent the mean of the daily densities during the 15 trading days prior to the event and the solid red curves represent the mean of the daily densities during 15 trading days after the date of the event. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

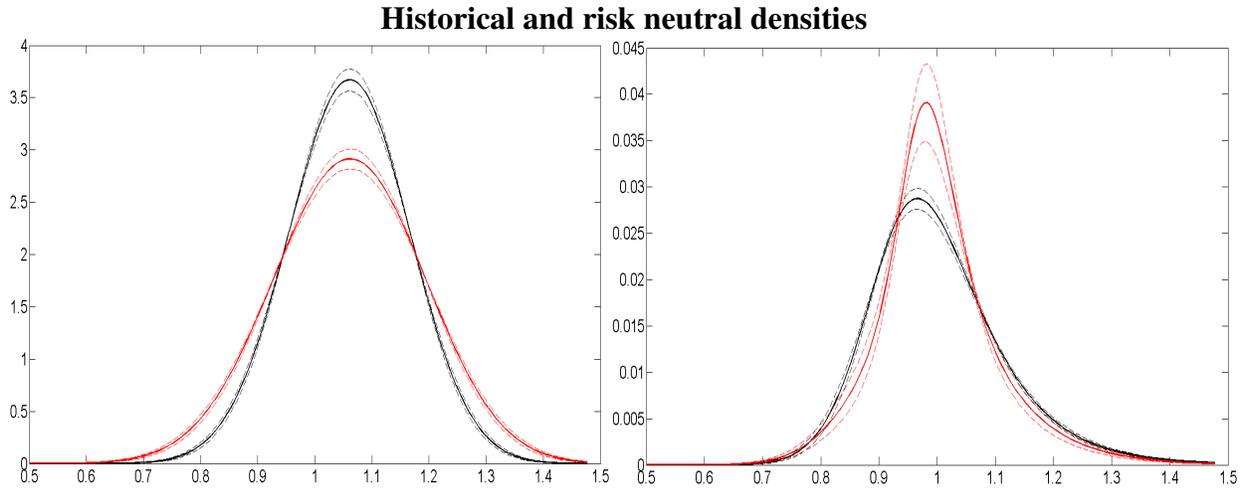
Pricing kernel and absolute risk aversion function



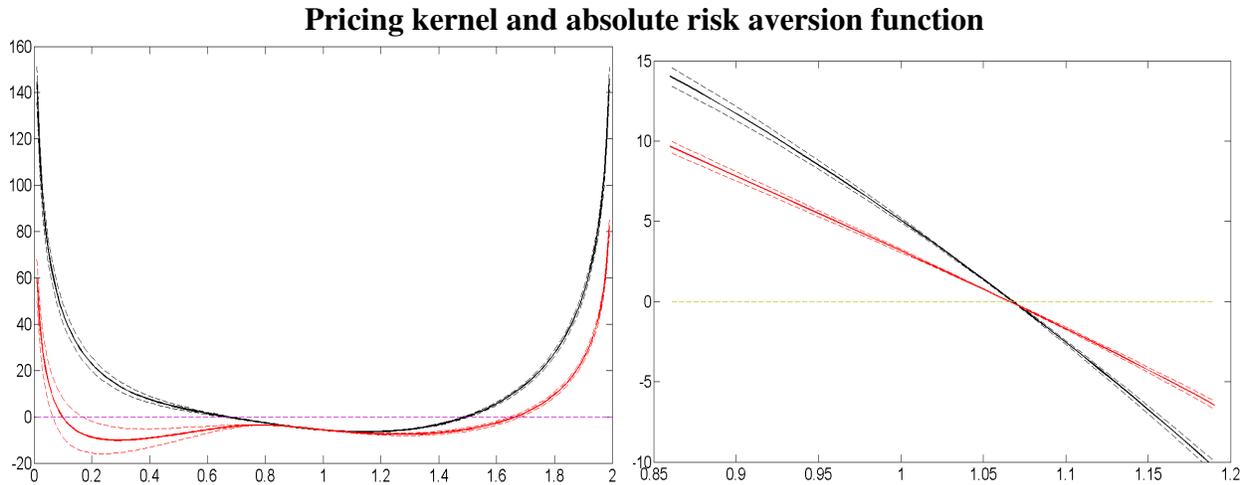
Note: the left figure represents the mean pricing kernel and the right figure represents the mean absolute risk aversion derived around the event date using the historical WTI price and related options. The parametric assumption is the normality of the historical and risk neutral densities returns. The event occurs on February 22, 2011 and captures the uprising in Libya. The solid black curves represent the mean of the daily functions during the 15 days prior to the event and the solid red curves represent the mean of the daily functions during the 15 trading days after the event date. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

Generalised Beta 2 case

PANEL 5: Historical densities, risk neutral densities, pricing kernel and the mean absolute risk version functions for the end of the bull commodity cycle



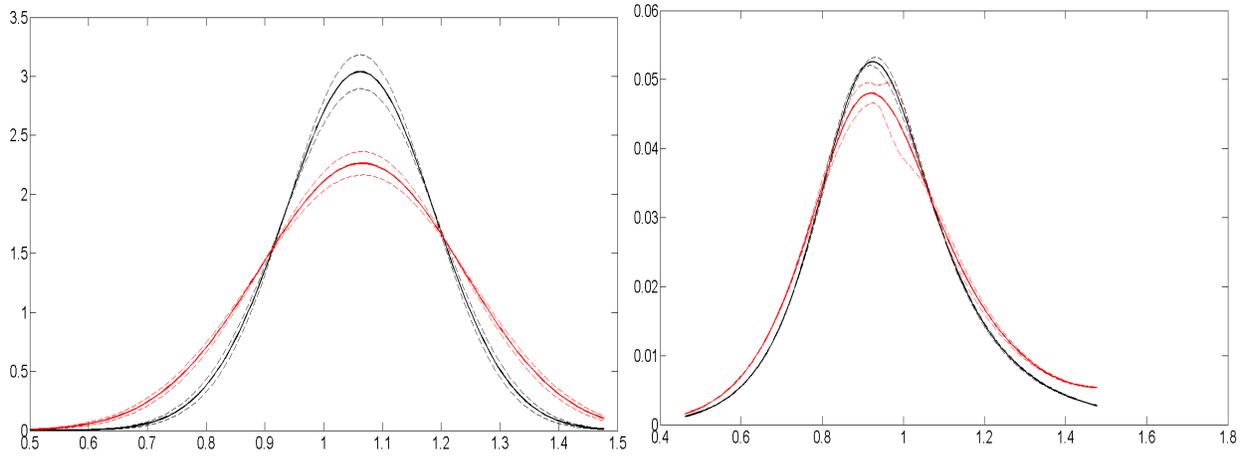
Note: the left figure represents the historical densities of the WTI futures and the right figure represents the risk neutral densities estimated with options on the WTI future third nearby contract. The parametric assumption is the Beta distribution for the historical densities and the generalised Beta 2 for the risk neutral densities. The event occurs on July 3, 2008 and captures the end of the commodity bull cycle. The solid black curves represent the mean of the daily densities during the 15 trading days prior to the event and the solid red curves represent the mean of the daily densities during 15 trading days after the date of the event. The dashed lines represent the associated empirical standard error (the mean ± 0.5 standard deviation).



Note: the left figure represents the mean pricing kernel and the right figure represents the mean absolute risk aversion derived around the event date using the historical WTI returns and related options. The parametric assumption is the Beta distribution for the historical densities and the generalised Beta 2 for the risk neutral densities. The event occurs on July 3, 2008 and captures the end of the commodity bull cycle. The solid black curves represent the mean of the daily functions during the 15 days prior to the event and the solid red curves represent the mean of the daily functions during the 15 trading days after the event date. The dashed lines represent the associated empirical standard error (the mean ± 0.5 standard deviation).

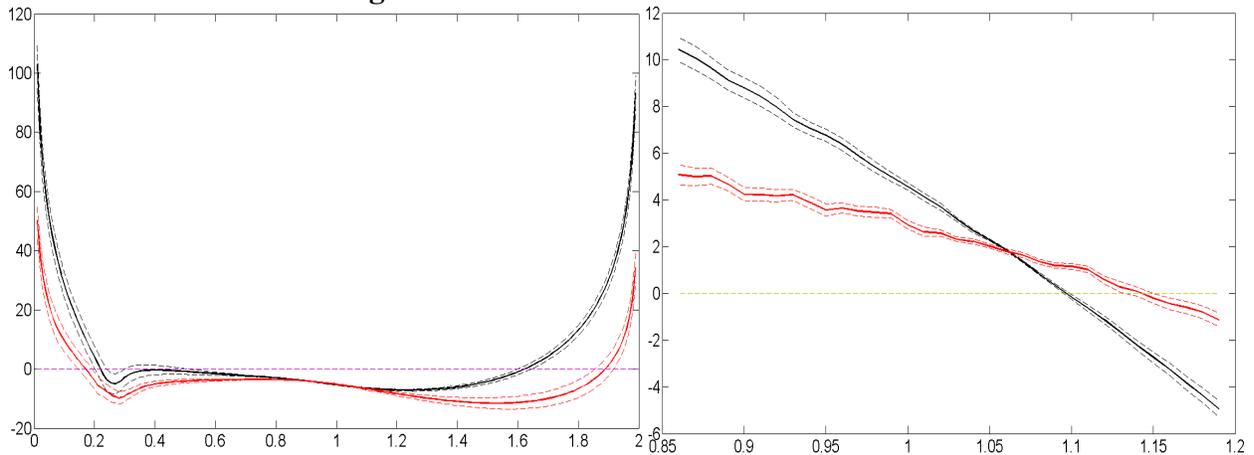
PANEL 6: Historical densities, risk neutral densities, pricing kernel and the mean absolute risk version functions for credit freeze bottom.

Historical and risk neutral densities



Note: the left figure represents the historical densities of the WTI futures and the right figure represents the risk neutral densities estimated with options on the WTI future third nearby contract. The parametric assumption is the Beta distribution for the historical densities and the generalised Beta 2 for the risk neutral densities. The event occurs at date on December 19, 2008 and captures credit freeze bottom. The solid black curves represent the mean of the daily densities during the 15 trading days prior to the event and the solid red curves represent the mean of the daily densities during 15 trading days after the date of the event. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

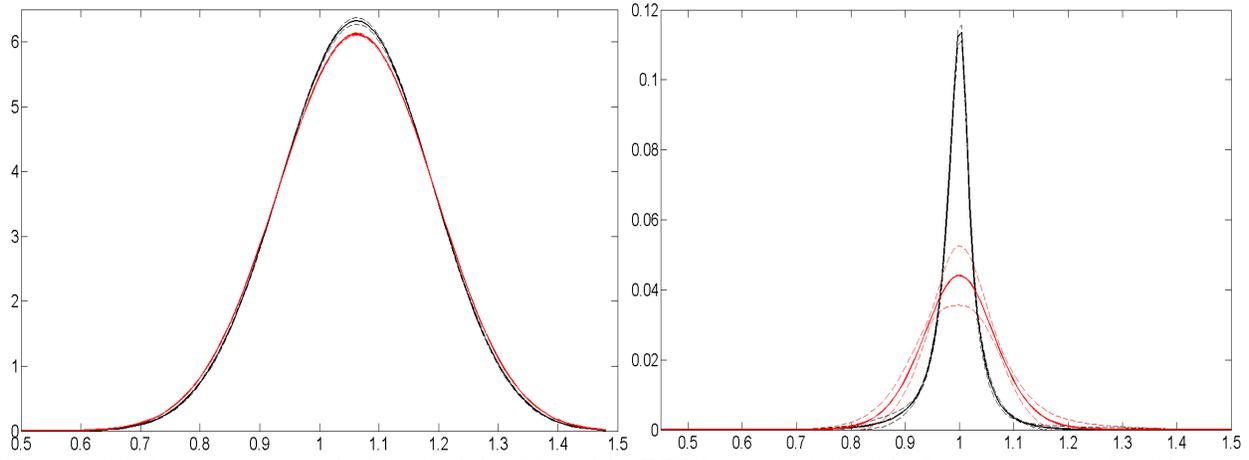
Pricing kernel and absolute risk aversion function



Note: the left figure represents the mean pricing kernel and the right figure represents the mean absolute risk aversion derived around the event date using the historical WTI returns and related options. The parametric assumption is the Beta distribution for the historical densities and the generalised Beta 2 for the risk neutral densities. The event occurs on December 19, 2008 and captures credit freeze bottom. The solid black curves represent the mean of the daily functions during the 15 days prior to the event and the solid red curves represent the mean of the daily functions during the 15 trading days after the event date. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

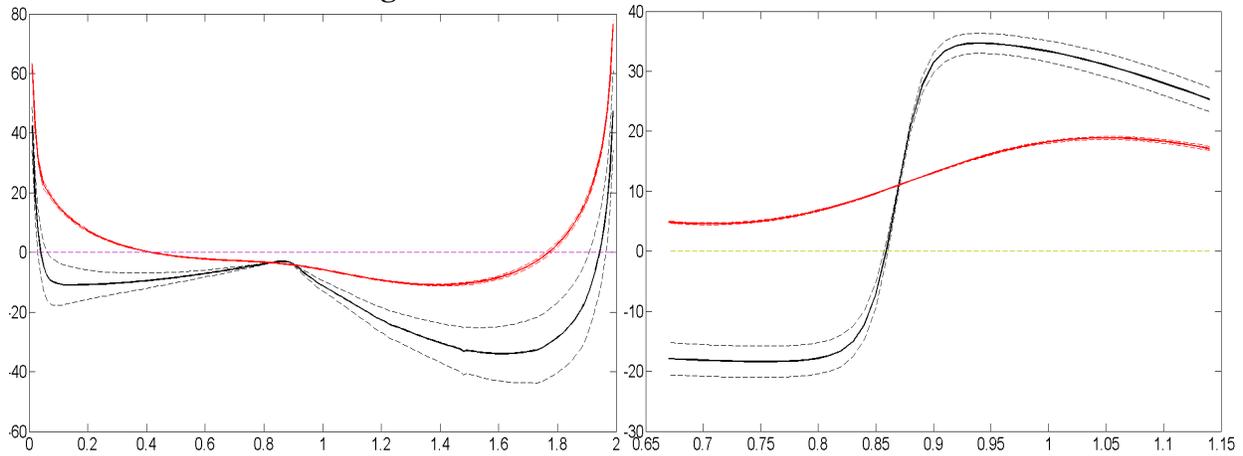
PANEL 7: Historical densities, risk neutral densities, pricing kernel and the mean absolute risk version functions for the explosion of the deepwater Horizon platform.

Historical and risk neutral densities



Note the left figure represents the historical densities of the WTI futures and the right figure represents the risk neutral densities estimated with options on the WTI future third nearby contract. The parametric assumption is the Beta distribution for the historical densities and the generalised Beta 2 for the risk neutral densities. The event occurs on April 20, 2010 and captures the explosion of the BP Deepwater Horizon platform. The solid black curves represent the mean of the daily densities during the 15 trading days prior to the event and the solid red curves represent the mean of the daily densities during 15 trading days after the date of the event. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

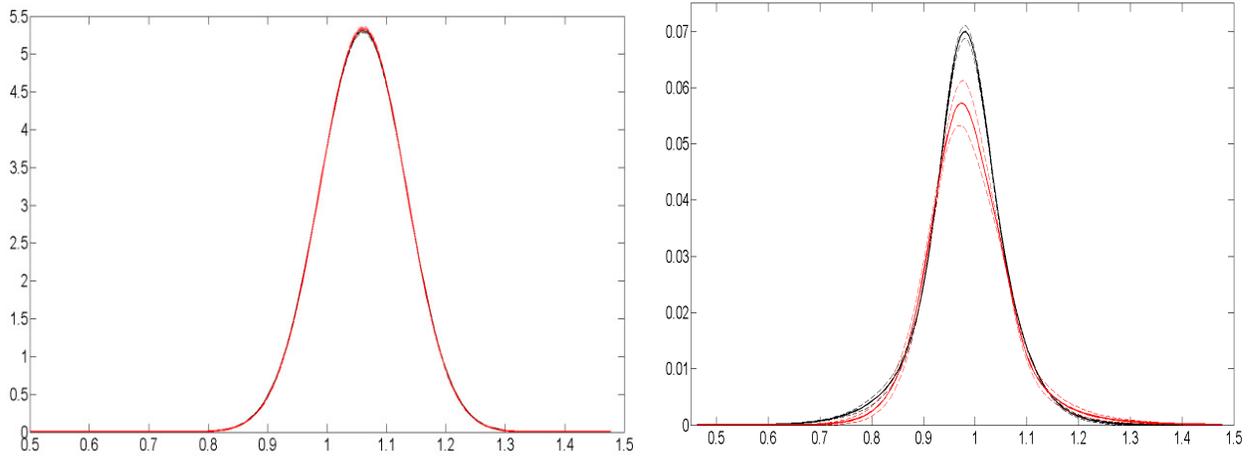
Pricing Kernel and absolute risk aversion function



Note: the left figure represents the mean pricing kernel and the right figure represents the mean absolute risk aversion derived around the event date using the historical WTI returns and related options. The parametric assumption is the Beta distribution for the historical densities and the generalised Beta 2 for the risk neutral densities. The event occurs on April 20, 2010 and captures the explosion of the BP Deepwater horizon platform. The solid black curves represent the mean of the daily functions during the 15 days prior to the event and the solid red curves represent the mean of the daily functions during the 15 trading days after the event date. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

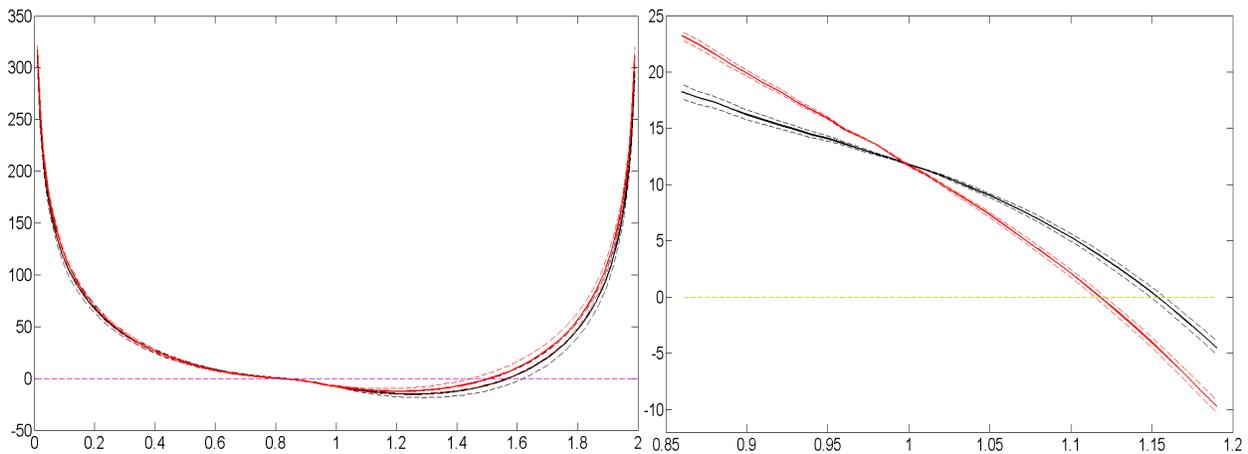
PANEL 8: The historical densities, risk neutral densities, pricing kernel and the mean absolute risk version functions for the Libyan uprising.

Historical and risk neutral densities



Note the left figure represents the historical densities of the WTI futures and the right figure represents the risk neutral densities estimated with options on the WTI future third nearby contract. The parametric assumption is the Beta distribution for the historical densities and the generalised Beta 2 for the risk neutral densities. The event occurs on February 22, 2011 and captures the uprising in Libya. The solid black curves represent the mean of the daily densities during the 15 trading days prior to the event and the solid red curves represent the mean of the daily densities during 15 trading days after the date of the event. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).

Pricing kernel and absolute risk aversion function



Note: the left figure represents the mean pricing kernel and the right figure represents the mean absolute risk aversion derived around the event date using the historical WTI price and related options. The parametric assumption is the Beta distribution for the historical densities and the generalised Beta 2 for the risk neutral densities. The event occurs on February 22, 2011 and captures the uprising in Libya. The solid black curves represent the mean of the daily functions during the 15 days prior to the event and the solid red curves represent the mean of the daily functions during the 15 trading days after the event date. The dashed lines represent the associated empirical standard error (the mean \pm 0.5 standard deviation).