Driven by Fear? The Tail Risk Premium in the Crude Oil Futures Market*†

Reinhard Ellwanger‡

This version: July, 2015

Abstract

Oil prices are notoriously difficult to forecast and exhibit wild swings or "excess volatility" that are difficult to rationalize by changes in fundamentals alone. This paper offers an explanation for these phenomena based on time varying disaster probabilities and disaster fears. Using information from crude oil options and futures we document economically large jump tail premia in the crude oil derivative market. These premia vary substantially over time and significantly forecast crude oil futures and spot returns. Our results suggest that oil futures prices overshoot (undershoot) in the presence of upside (downside) tail fears in order to allow for smaller (larger) risk premia thereafter. We show that this overshooting (undershooting) is amplified for the spot price because of time varying benefits from holding inventory that work in the same direction. The novel oil price uncertainty measures yield additional insights into the relationship between the oil market and macroeconomic outcomes.

JEL classification: C53, C58, D84, E44, G12, G13, Q43

*The views expressed in this paper are those of the author and no responsibility for them should be attributed to the Bank of Canada.

†I would like to thank Peter Reinhard Hansen, Christiane Baumeister, Lutz Kilian, Marek Raczko, Lai Xu, Victor Todorov, Antoni Targa Barrera, Gabriela Galassi, Robert Goodhead as well as the participants at the Econometrics working group and the "Measuring and Modeling Financial Risk with High Frequency Data" conference at the European University Institute in Florence, the participants at the 10th BMRC-DEMS Conference on Macro and Financial Economics/Econometrics in London and at the One day workshop at CREATES, Aarhus, for their comments and helpful discussions.

‡Bank of Canada, International Economic Analysis Department; email: rellwanger@bankofcanada.ca

1 Introduction

Volatile oil spot and futures prices have drawn a lot of attention from academics, policy makers and investors over the last years. The origins of this volatility are hotly debated because the changes in oil demand and supply appear too smooth to explain the large swings in oil prices (see e.g. Tang and Xiong 2010; Alquist and Kilian 2010; Baumeister and Peersman 2013). One explanation proposed by Kilian (2009) is that precautionary demand shocks are an important driver of short run fluctuations in oil prices. He suggests that a key source of precautionary demand movements is uncertainty about shortfalls of expected supply relative to expected demand. A natural question is therefore how uncertainty shocks can be accurately measured and whether they help to predict the future evolution of oil prices. Yet another is if and how they relate to the risk premia embedded in oil derivatives. The explanations building on competitive storage models such as the one proposed by Alquist and Kilian (2010) are typically based on the assumption of risk neutrality and ignore potential effects arising from temporal variations in these premia. This stands in contrast to the convincing evidence of significant and time-varying risk premia in oil markets from the commodity finance literature (see e.g. Gorton and Rouwenhorst 2004; Erb and Harvey 2006; Singleton 2013; Hamilton and Wu 2014; Szymanowska et al. 2014). We bridge this gap by presenting a novel uncertainty measure that is explicitly related to the risk premium embedded in oil market derivatives.

The contribution of this paper is threefold. First, we present novel estimates of oil market uncertainty based on the left jump and right jump varation premia ($LJVP$ and $RJVP$, respectively) embedded in crude oil futures and options. Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and Xu (2014) show that these premia - defined as the part of the variance risk premium that is due to large sized upward and downward jumps - can be estimated in an essentially model free manner and contain important information about market participants' sentiments and expected stock market returns. We demonstrate that for the oil market, $LJVP$ and $RJVP$ are economically large and significantly vary over episodes of documented supply and demand uncertainty. The variation measures have strong predictive power for both crude oil futures and spot prices that is not contained in traditional oil price predictors. Our results suggest that oil futures and spot prices overshoot (undershoot) in the presence of upside (downside) tail fears in order to allow for smaller (larger) risk premia thereafter.

Second, we use a stylized no-arbitrage model of storage in order to show that relative to the futures price movement, this overshooting (undershooting) is amplified for the spot price due to the time varying value of holding physical oil in inventory. A relative increase in $RJVP$ is also associated with a rise in this value that pushes the spot price in the same direction as the futures risk premium. Consistently, $RJVP$ and $LJVP$ exhibit larger in-sample and out-of-sample predictability for spot price returns than for futures returns. This finding complements the model of Alquist and Kilian (2010), in which a similar no-arbitrage condition in the storage market is used to show that an increase in oil production volatility leads to an overshooting of spot prices. Their model is derived under the assumption of risk neutrality and the overshooting is due to convex adjustment costs of inventories. Since in our model the fear shock is associated with an additional movement of the risk premium this overshooting is magnified. Under risk aversion, the spot price not only increases with respect to the spot price in some future period, but also with respect to the rise implied by storage models that assume risk neutrality. Taken together, these results represent new evidence for the importance of time varying risk premia in explaining the perceived excess volatility in oil spot and futures prices.

Finally, to put our jump risk measures in perspective, we also investigate the link between $LJVP$, $RJVP$ and real economic activity as well as macroeconomic uncertainty. The oil fear measures appear not to be spanned by aggregate uncertainty. Also, there is little evidence of a stable linear relationship between stock market returns, $LJVP$ and $RJVP$. The results are consistent with the idea that the oil tail risk measures aggregate different types of uncertainty that are relevant for oil prices, e.g. oil supply and oil demand uncertainty, that individually have a very distinct relationship with aggregate uncertainty.

There are several reasons for focusing on the tails of the oil price distribution to quantify uncertainty. First, recent theoretical works show that models with tail risks can account for the high equity risk premium and excess market return volatility (Barro 2006; Wachter 2013). There is also increasing evidence that the index option implied compensation for aggregate market volatility and tail risks are closely connected to economic uncertainty and temporal variation in risk aversion (Bollerslev, Tauchen, and Zhou 2009; Bollerslev, Todorov, and Xu 2014). In contrast, most of the commodity finance literature has focused on futures, ignoring potential information from the related option prices. Among the exceptions are Trolle and Schwartz (2010), who document a significant and time-varying variance risk premium in the crude oil option market and Pan and Kang (2013) who show that this premium forecasts short term futures returns. We find that most, if not all of the variance risk premium and its forecasting power is due to time varying compensation for tail variations. One of the advantage of considering tail risk premia is that they are naturally separated into upside and downside uncertainty, thus providing additional information beyond that contained in the variance risk premium. We show that in particular the time varying asymmetry of the jump premia across the two tails improves the prediction spot and futures returns. Based on the variance risk decomposition proposed by Bollerslev and Todorov (2011b), we present evidence that this asymmetry reflects changes in effective risk aversion of oil market participants. On a more general level, these findings suggest that time varying disaster fears embedded in option prices of individual assets, not only on market indexes, convey important information on the their return dynamics.

Formally, $RJVP$ and $LJVP$ are defined as the difference of the conditional expected variation of jump tails under the statistical, objective probability measure and the risk neutral measure. In other words, they represent the difference of the expected actual jump variations and the option implied market price for the insurance against these jump variations. Our implemented model for the option implied jump-measures follows Bollerslev, Todorov, and Xu (2014). It is semi-parametric and flexible, allowing the tail distributions to differ across the left and the right tail and for independent time variation in the shape and level of the tails. The empirical estimation is based on panel of out-of-money (OTM) call and put options. Intuitively, short-maturity OTM options are most sensitive to large jumps, which allows us to separate the jump risk from the diffusive risk. The statistical jump variation is based on intraday futures data and non-parametric methods developed by Bollerslev and Todorov (2011a). Empirically, we find that the statistical jump variations in oil futures are significantly smaller than their risk neutral counterparts. Moreover, the actual jumps are approximately symmetric, while the option implied prices for these jumps display time-varying asymmetries. Thus the observed time variation in the relative size of $RJVP$ and $LJVP$ are largely void of influences from the actual jumps, and can be interpreted as a direct measure of investor fears (Bollerslev and Todorov 2011b).

On average, the $LJVP$ tends to be larger than $RJVP$, implying that oil investors are on average requesting a higher premium when downside-risks predominate. This is largely consistent with the theory of normal backwardation (Keynes 1976). Accordingly, when producers of the physical commodity want to hedge their price risk using derivatives, then arbitrageurs who take the other side of the contract have to be compensated in form of a risk premium order to take on the risk. Our empirical estimates imply that on average about half futures risk premium is due to large jump risks.

Last, we contribute to a strand of literature going back to the idea of Bernanke (1983) that oil market uncertainty - rather than oil price changes alone - is a key variable to understand the relationship between the oil market and macroeconomic outcomes. Our measures provide precise definitions and estimates for oil price uncertainty, and we discuss some interactions with aggregate variables towards the end of the paper.

The rest of this paper is structured as follows. The next section provides a formal definition of the tail risk variation measures, and explains the spot price overshooting. In section 3 we discuss the empirical implementation and data as well as the properties of our estimates. Section 4 presents the forecasting results and section 5 the interaction of the jump tail premia with aggregate outcomes. Section 6 concludes.

2 Theoretical Setup

2.1 Setup and Definitions

In this section, we present the general setup and provide formal definitions for our tail risk measures. This setup is inspired by Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and Xu (2014). Instead of considering the (aggregate) stock market, we will focus on the dynamics of an individual asset, namely oil futures.

To fix ideas, let $(\Omega, \mathcal{F}, \mathbb{P})$ be a filtered probability space with the filtration $(\mathcal{F}_t)_{t\geq 0}$ and let F_t denote the price of a crude oil futures contract. The dynamics of the futures price are described by the following jump diffusion process

$$\frac{dF_t}{F_{t-}} = \alpha_t dt + \sigma_t dW_t + \int_{\mathbb{R}} (e^x - 1)\tilde{\mu}(dt, dx), \tag{1}$$

where the drift α_t and the stochastic volatility σ_t are assumed to be locally bounded càdlàg processes, and W_t is a standard Brownian Motion. Here $\tilde{\mu}(dt, dx) = \mu(dt, dx) - v_t^{\mathbb{P}}(dx)dt$ denotes a compensated jump measure, with $\mu(dt, dx)$ the counting measure and $v_t^{\mathbb{P}}(dx)dt$ the compensator of jumps, where $\mathbb{P}$ denotes the statistical, objective measure .[1]

Under standard non-arbitrage assumptions, there exists a risk-neutral measure denoted $\mathbb{Q}$, under which the futures price follows a martingale of the form

$$\frac{dF_t}{F_{t-}} = a_t dt + \sigma_t dW_t^{\mathbb{Q}} + \int_{\mathbb{R}} (e^x - 1)\tilde{\mu}^{\mathbb{Q}}(dt, dx), \tag{2}$$

where a_t denotes the drift, $dW_t^{\mathbb{Q}}$ is a Brownian motion with respect to the risk neutral measure and $\tilde{\mu}^{\mathbb{Q}} = \mu(dt, dx) - v_t^{\mathbb{Q}}(dx)dt$ denotes the jump measure under $\mathbb{Q}$ following the previous decomposition. In general, the change of measure alters both the drift and the jump intensity describing the dynamics of the futures price while the volatility associated with the Brownian motions remains the same under both measures. This reflects the special pricing of jumps in comparison with continuous movements.

Our interest will be in both the futures risks premium (FRP) - a premium reflecting risk associated with holding a (long) futures contract - and the variance risk premium (VRP) - a premium reflecting risks associated with holding a (long) variance swap - that are associated with the jump part of the futures price. Following Bollerslev and Todorov (2011b), the FRP at time t and for some $T > t$ is defined as

$$FRP_{t,T} \equiv \frac{1}{T-t}\left(E_t^{\mathbb{P}}\left(\frac{F_T - F_t}{F_t}\right) - E_t^{\mathbb{Q}}\left(\frac{F_T - F_t}{F_t}\right)\right). \tag{3}$$

[1] The compensator $v_t^{\mathbb{P}}(dx)dt$ ensures that the jump measure $\tilde{\mu}(dt, dx)$ is a martingale.

Since the futures price F_t is a martingale under the Q-measure, $FRP_{t,T}$ is effectively determined by the difference of the objective expectation of the futures price at some future date T and the current futures price.

Given our jump diffusion model in equation (1) we can, without loss of generality, define the FRP due to large jumps above some threshold $k_t > 0$,

$$FRP_{t,T}(k_t) \equiv \frac{1}{T-t} E_t^{\mathbb{P}} \left(\int_t^T \int_{|x|>k_t} (e^x - 1) v_s^{\mathbb{P}}(dx) ds \right) - \frac{1}{T-t} E_t^{\mathbb{Q}} \left(\int_t^T \int_{|x|>k_t} (e^x - 1) v_s^{\mathbb{Q}}(dx) ds \right). \tag{4}$$

Going one step further, we can decompose $FRP_t(k_t)$ into the contributions from large positive and large negative jumps

$$FRP_{t,T}(k_t) = FRP^+_{t,T}(k_t) + FRP^-_{t,T}(k_t), \tag{5}$$

where $FRP^+_{t,T}(k_t)$ captures the futures risk premia due to $x > k_t$ and $FRP^-_{t,T}(k_t)$ captures the premia due to $x < -k_t$.

The variability of the futures price is measured by the quadratic variation QV of its log-price process of the interval $[t, T]$

$$QV_{[t,T]} = \int_t^T \sigma_s^2 ds + \int_t^T \int_{\mathbb{R}} x^2 \mu(ds, dx). \tag{6}$$

Similar to the futures risk premium, $VRP_{t,T}$ is formally defined as as the difference of the expected quadratic variation over the $T-t$ period under the respective probability measure.

$$VRP_t = \frac{1}{T-t} (E_t^P(QV_{[t,T]}) - E_t^Q(QV_{[t,T]})). \tag{7}$$

Under this definition of the variance risk premium, VRP_t equals the expected payoff from a long variance swap contract (Carr and Wu 2009). The variance risk premium is also naturally decomposed into a part associated with the continuous-time stochastic volatility process σ_s and a part that is due to jumps. We denote $RJV^{\mathbb{P}}_{t,T}(k_t)$ and $LJV^{\mathbb{P}}_{t,T}(k_t)$ the predictable component of the quadratic variation arising through large positive and large negative jumps under the $\mathbb{P}$ measure

$$RJV^{\mathbb{P}}_{t,T}(k_t) = \int_t^T \int_{x>k_t} x^2 v_s^{\mathbb{P}}(dx) ds, \quad LJV^{\mathbb{P}}_{t,T}(k_t) = \int_t^T \int_{x<-k_t} x^2 v_s^{\mathbb{P}}(dx) ds \tag{8}$$

and their counterparts under the risk neutral measure $\mathbb{Q}$

$$RJV^{\mathbb{Q}}_{t,T}(k_t) = \int_t^T \int_{x>k_t} x^2 v^{\mathbb{Q}}_s(dx)ds, \quad LJV^{\mathbb{Q}}_{t,T}(k_t) = \int_t^T \int_{x<-k_t} x^2 v^{\mathbb{Q}}_s(dx)ds. \tag{9}$$

The part of the variance risk premium due to large positive jumps is then

$$RJVP(k_t) \equiv \frac{1}{T-t}\left(E^{\mathbb{P}}_t\left(RJV^{\mathbb{P}}_{t,T}(k_t)\right) - E^{\mathbb{Q}}_t\left(RJV^{\mathbb{Q}}_{t,T}(k_t)\right)\right), \tag{10}$$

while the part due to large negative jumps is

$$LJVP(k_t) \equiv \frac{1}{T-t}\left(E^{\mathbb{P}}_t\left(LJV^{\mathbb{P}}_{t,T}(k_t)\right) - E^{\mathbb{Q}}_t\left(LJV^{\mathbb{Q}}_{t,T}(k_t)\right)\right). \tag{11}$$

As suggested by Bollerslev and Todorov (2011b), for equity index option the difference between $LJVP(k_t)$ and $RJVP(k_t)$ is naturally associated with investors' fear. In this paper we investigate this hypothesis for the oil market and define

$$FI_t(k_t) \equiv LJVP(k_t) - RJVP(k_t). \tag{12}$$

The index FI_t measures the asymmetry between the premium requested for the downside variance risk and the premium charged for upside variance risk that is due to large jumps. Under the above definition, a relatively large left jump variation premia or "downside fear" is associated with a low value of $FI_t(k)$.

The respective premia can then be estimated in an essentially model free manner using high-frequency returns and options data. The next section discusses this procedure in more detail. Before, we turn to the relationship between risk premia and spot prices.

2.2 Risk Premia and Spot Prices

In this section, we present a theoretical framework based on no-arbitrage conditions that relates the risk premia to oil spot prices. We show that in general, an increase in the futures risk premium associated with a rise in market participants fears - as proxied by the tail risk premia - will drive a temporary wedge between the current and expected prices for both futures and spot prices. Since this wedge subsequently reverts to zero, our model predicts that futures and spot prices overshoot with respect to future prices in the wake of disaster fears.

For our analysis we draw on non-arbitrage conditions derived from two different approaches to commodity derivative pricing that allow to relate the current futures price and the current spot price to the expected spot price some period ahead.[2] The first approach is based on the basic definition of the futures risk premium as described in equation (3). By non-arbitrage, the value of the futures price at the time of maturity must be equal to the spot price of the commodity. Hence $F_{T,T} = S_T$, where S_T stands for the spot price of oil at time T, and where in slight abuse of notation we let T denote the contract's terminal date for the remainder of this subsection. Moreover, since $F_{t,T}$ is a martingale under the $\mathbb{Q}$-measure, it follows that

$$1 + (T - t)FRP_{t,T} = \frac{E_t(S_T)}{F_{t,T}} = \frac{E_t(F_{t,T})}{F_{t,T}}. \tag{13}$$

Equation (13) reflects the non-arbitrage condition that the price of a futures has to be equal to the expected spot price discounted by the premium associated with holding the futures contract. This decomposition of the futures price into the expected spot price and a risk premium has been used frequently in the analysis of the futures risk premium. The empirical and theoretical evidence overwhelmingly suggests that in the oil market, and in commodity markets in general, the (net) premium is on average positive and fluctuating over time (e.g. Keynes 1976; Bessembinder 1992; Hamilton and Wu 2014; Baumeister and Kilian 2014).

The second non-arbitrage condition is based on the theory of storage, going back to the works of Kaldor (1939) and Working (1949). A distinguishing feature of commodities as an asset class is the significance of the convenience yield, defined as the benefit of immediate availability of a physical commodity rather than a time T contingent claim on the commodity (see e.g. Fama and French 1987). This benefit links the futures to the spot price through the following relationship

$$F_{t,T} = S_t(1 - (T - t)CY_{t,T}), \tag{14}$$

where $CY_{t,T}$ is the (net of storage costs and interest rate outlays) equilibrium convenience yield, in annualized terms. This equation has to hold under non-arbitrage, since the price of a futures contract has to be equal to the cost of buying the commodity now minus the net benefits of carrying the commodity to maturity. Such benefits can arise due to a variety reasons such as temporary stock-outs and associated price spikes or convex adjustment costs. Importantly, the convenience yield can also be interpreted as a call option on a futures contract (Milonas and Thomadakis 1997). Given that our fear measured is derived from actual option prices on the futures, we expect a close link between oil market fears and the convenience yield.

From the decomposition of the futures risk premium into contributions from positive and

[2] The setup is similar to the one in Gospodinov and Ng (2013). In our model, however, we explicitly focus on the effect of the tail premia and their interaction with the futures risk premium and the convenience yield, whereas Gospodinov and Ng (2013) focus exclusively on the latter.

negative large jumps, equations (4) and (5), it follows that an exogenous, relative shift in the tail risk premium associated with the right tail above that of the left tail will decrease $FRP_{t,T}$. Equation (13) then implies that the futures price will rise above the expected spot price, and is expected to decline towards to spot price thereafter. Moreover, the size of the shift in the futures price will be determined by the change in the futures risk premium. In order to investigate the relationship between current and expected spot prices, we combine equation (13) and (14), obtaining the following relationship between the current and the expected spot price

$$\frac{E_t(S_T)}{S_t} = (1 + (T-t)FRP_{t,T})(1 - (T-t)CY_{t,T}). \tag{15}$$

Thus under no-arbitrage, the (relative to the expected spot price) price of current spot oil is determined by the net value of holding the physical oil in storage and the futures risk premium. Moreover, a tail fear induced the change in the spot price is unambiguously larger than the change in the futures price if the net convenience yield varies negatively with the relative size of the tail risk measures. In this case, the spot prices not only overshoot with respect to expected spot prices, but also with respect to the current futures price. The expected spot and futures returns, $r_{(T-t),S}$ and $r_{(T-t),F}$, respectively, that are associated with a relative increase in the right tail variation measure should therefore exhibit

$$E_t(r_{(T-t),S}|RJVP_t > LJVP_t) < E_t(r_{(T-t),F}|RJVP_t > LJVP_t) < 0, \tag{16}$$

with an inverse relationship for relative increases in the left tail measure. Considering a temporary, mean preserving change in the tail risk measures, i.e. a change such that $E_t(S_T|RJVP_t > LJVP_t) = E_t(S_T|RJVP_t = LJVP_t)$, equation (16) implies that both the spot and the futures price have to adjust immediately by an upward movement.[3] Moreover, the underlying mechanism does not rely on any shifts in inventories. This is important, since some of the large price fluctuations, e.g. around the 1991 Gulf War, are difficult to reconcile with the observed changes in inventory (Kilian and Murphy 2014).

Of course, equation (16) rests on the assumption that the net convenience yield varies negatively with FI_t. Subsequently we show that FI_t is inversely related to the market prices of out-of-the-money options. The interpretation of the convenience yield as an op-

[3]This comparative static thought experiment resembles the setting of Alquist and Kilian (2010), who describe the effects from a change in the conditional variance of oil supply shocks. In their model, an increase in the conditional variance of these shocks lead to a similar overshooting of the spot price. However, these results are derived under the assumption of risk neutrality so that the futures price is an unbiased predictor of the spot price. In contrast, our framework allows for an interaction of a time-varying risk premia with the convenience yield and suggest that conditional on $RJVP_t$, $LJVP_t$, futures prices are not necessarily an unbiased predictor.

tion suggests therefore that we should expect this inverse relationship to hold theoretically. Empirically this assumption is justified by our data, as we find a strong and statistically significant correlation of almost -40% between our measure of FI_t and the log of the net convenience yield measure as implied by equation (14).

Taken together, the non-arbitrage conditions linking the contemporaneous futures price and spot price with future prices, suggest that oil futures and spot prices overshoot relative to the expected spot price in the wake of upside (downside) fears. This implies that our tail variation measures should forecast futures and spot market returns. Moreover, this overshooting, reflected by the expected return, is larger for the spot price then for the futures futures price. We address this hypothesis by direct forecasts and a structural VAR analysis after presenting the estimation methodology and results for the tail variation measures.

3 Empirical Implementation

3.1 Estimation of the Variation Measures

As noted by Carr and Wu (2009), the implied total variation $QV_{t,T}^{\mathbb{Q}}$, also known as the variance swap rate, can be well approximated by a portfolio of out-of-the-money put and call options. For our calculation we follow the methodology for the CBOE Volatility Index (VIX).[4]

Our specification of the jump tails follows Bollerslev, Todorov, and Xu (2014). In particular, the jump distribution and intensity under the $\mathbb{Q}$-measure are based on the semi-parametric model

$$v_t^{\mathbb{Q}}(dx) = \left(\phi_t^+ \times e^{-\alpha_t^+ x} 1_{(x>0)} + \phi_t^- \times e^{-\alpha_t^- |x|} 1_{(x<0)}\right) dx. \tag{17}$$

Relative to other existing models, this specification imposes only minimal restrictions on the jump tail dynamics since (a) the left and right jump tails are allowed to differ and (b) the level shift parameters $\phi^{\pm}$ and the shape parameters $\alpha^{\pm}$ are allowed to vary independently over time.

The estimation of $\alpha^+(\alpha^-)$ and $\phi^+(\phi^-)$ is based on the observation that for $(T-t)\downarrow 0$ and $k \uparrow \infty (k \downarrow -\infty)$

$$\frac{e^r O_{t,T}(K)}{(T-t)F_{t,T}} \approx \frac{\phi_t^{\pm} e^{k(1 \pm \alpha_t^{\pm})}}{\alpha_t^{\pm}(\alpha_t^{\pm} \pm 1)}, \tag{18}$$

where $O_{t,T}(K)$ denotes the price of a call (put) option with strike K and $k = \log(K/F_{t,T})$. This reflects the intuition that for close to maturity, deep OTM options the risks associated

[4] See the white paper on the CBOE website for details regarding the VIX methodology.

with the diffusive part become negligible and their price therefore reflects jump risks.

From equation (18) it follows that the ratio of two OTM options does not depend on $\phi_t^{\pm}$, leading to the natural estimator suggested by Bollerslev and Todorov (2013):

$$\hat{\alpha}_t^{\pm} = \operatorname{argmin}_{\alpha^{\pm}} \frac{1}{N_t^{\pm}} \sum_{i=1}^{N_t^{\pm}} \left| \log\left(\frac{O_{t,\tau}(k_{t,i})}{O_{t,\tau}(k_{t,i-1})}\right)(k_{t,i} - k_{t,i-1})^{-1} - (1 \pm (-\alpha^{\pm}))\right|, \tag{19}$$

where $O_{t,\tau}$ is the time t price of an OTM option on the futures with log-moneyness k, $N_t^{\pm}$ denotes the total number of options used in the estimation and $0 < |k_{t,1}| < ... < |k_{t,N_t^{\pm}}|$. In practice we will pool options such that t refers to a given month which implicitly assumes that $\alpha^{\pm}$ is approximately constant during this period.

For a given $\alpha^{\pm}$, we then use equation (18) to estimate

$$\hat{\phi}_t^{\pm} = \operatorname{argmin}_{\phi^{\pm}} \frac{1}{N_t^{\pm}} \sum_{i=1}^{N_t^{\pm}} \left| \log\left(\frac{e^{r_\tau} O_{t,\tau}(k_{t,i})}{(T-t)F_{t-,\tau}}\right) + (\pm\hat{\alpha}_t^{\pm} - 1)k_{t,i} + \log(\hat{\alpha}_t^{\pm} \mp 1) + log(\hat{\alpha}_t^{\pm}) - \log(\phi^{\pm})\right|. \tag{20}$$

From the definition of the tail risk premia in equation (**??**) and our assumptions for the large jumps dynamics in (17), it follows that for time to maturity $T-t$ and threshold k_t

$$RJV_{t,T}^{\mathbb{Q}} = (T-t)\phi_t^{+} e^{-\alpha_t^{+} k_t}(\alpha_t^{+} k_t(\alpha_t^{+} k_t + 2) + 2)/(\alpha_t^{+})^3 \quad \text{and} \tag{21}$$

$$LJV_{t,T}^{\mathbb{Q}} = (T-t)\phi_t^{-} e^{-\alpha_t^{-} k_t}(\alpha_t^{-} k_t(\alpha_t^{-} k_t + 2) + 2)/(\alpha_t^{-})^3. \tag{22}$$

The $\mathbb{Q}$ tail measures are then computed by replacing the population quantities in (21) by their estimates.

The estimation of the corresponding quantities under the objective measure are based on high-frequency intraday data. We use the notation of Bollerslev and Todorov 2011b and divide the trading day t into the $[t, t+\pi_t]$ overnight period and the $[t+\pi_t, t+1]$ active trading period. Hence π_t denotes the length of the close to open interval.[5] Dividing the effective trading time in equally spaced intervals, we obtain n returns $\Delta_{t,i} f \equiv f_{t+\pi+i} - f_{t+\pi+i-1}$, where f denotes the logarithm of the futures price. We denote RV_t the realized variation on day t, which is consistently estimated by summing the squared intraday returns

$$RV_t \equiv \sum_{i=1}^{n} (\Delta_{t,i} f)^2 \xrightarrow{p} \int_{t+\pi_t}^{t+1} \sigma_s^2 ds + \int_{t+\pi_t}^{t+1} \int_{\mathbb{R}} x^2 \mu(ds, dx). \tag{23}$$

Realized jumps under the statistical measure are estimated using the threshold technique

[5] Although the trading hours are non-stochastic, it is convenient to treat π_t as stochastic. The theoretical derivations presented below are valid under mild conditions regarding the stochastic process for π_t. See Bollerslev and Todorov 2011b for details.

first proposed by Mancini (2001). Under the threshold estimation, we first compute an estimate for the continuous part of the volatility, σ_t, and then filter out jumps by identifying a threshold separating jumps that appear incompatible with the underlying normal distribution. The truncation threshold for large jumps is time-varying and captures the effects of well-described volatility clustering as well as intraday volatility. Out of the returns that are identified as jumps, we select the large and medium-sized ones for the tail estimation. The reader is referred to Appendix A regarding for further details regarding the estimation procedure.

Due to the lack of observations for sufficiently large jumps under the statistical measure it is infeasible to the estimate the same flexible jump tail specification as under the $\mathbb{Q}$-measure. Instead we follow Bollerslev and Todorov (2011b) in assuming that[6]

$$v_t^{\mathbb{P}} = \left((\alpha_0^- 1_{x<0} + \alpha_0^+ 1_{x>0}) + (\alpha_1^- 1_{x<0} + \alpha_1^+ 1_{x>0})\sigma_t^2\right) v^{\mathbb{P}}(x)dx, \tag{24}$$

where $v^{\mathbb{P}}(x)dx$ is the time-invariant distribution of empirical jumps and the level parameters $\alpha_0^{\pm}$ and $\alpha_1^{\pm}$ relate linearly to the time varying continuous volatility σ_t^2.[7]

The estimation of the time invariant jump distribution draws on the insight by Bollerslev and Todorov (2011a) that the tails of an arbitrary distribution are approximately distributed according to a Generalized Pareto Distribution. Given our empirical jumps, the two parameters of the Generalized Pareto Distribution along with α_0 and α_1 are estimated seperately for each tail via GMM as outlined in Appendix A.

3.2 Data Description

Our empirical analysis is based on light sweet crude oil (West Texas Intermediate) futures and options.[8] Crude oil derivatives are traded in extremely liquid markets and available historical data goes back to the 1980s. Trading of oil futures started in April 1983 for contracts with maturities up to three months, and for options on futures in November 1986. For the estimation of the $\mathbb{Q}$ jump tails we use an option data set obtained from the Chicago Mercantile Exchange (CME group, formerly NYMEX) that contains all historical end-of-the-day settlement prices.[9] Conveniently these crude oil options are quoted for a variety of

[6]Empirical evidence that the jump distribution is in the oil market is approximately proportional to the continuous volatility is presented in Doran and Ronn (2008).

[7]The specification for the jump tails under the $\mathbb{P}$-measure is significantly more restrictive than the corresponding $\mathbb{Q}$-measure specification. Empirically the $\mathbb{P}$ jump tails are dwarfed by the risk neutral analogues and we will therefore only work with the more general $\mathbb{Q}$-measures in the later part of this paper. The empirical evidence for a significant difference in the expected jump tail variations under the different measures are presented in the following section.

[8]The CME tickers for the futures and options is CL and LO, respectively.

[9]NYMEX crude oil options are American style. For short maturity, deep OTM options the difference between European and American options is negligible, so we use the original options for the jump tail

strike prices and expiration dates - one for each calendar month of the year - thus ensuring a sufficient number of short maturity deep OTM options for the empirical implementation of our estimator. The derivation of the tail parameters formally relies on a decreasing time to maturity $(T-t) \downarrow 0$. We therefore only take the contract with the shortest time to maturity, whenever the maturity is larger than 9 days. The last trading days of a given option contract can be characterized by prizing abnormalities due to the lack of trading volume, which makes it necessary to discard this data and resort to the first back contract for those days.[10]

In order to mitigate potential influences from the diffusive risk, we retain only OTM call (put) options with at log-moneyness more than plus (minus) twice the maturity-normalized Black-Scholes at-the-money implied volatility. We clean the data by discarding all options with a settlement price of less than 3 cents and those violating the monotonicity condition in the strike dimension. The dataset comprises daily data from December 1987 to December 2013. For the estimation of the jump tails we pool all clean deep OTM options for a given calendar month. This leaves us with 313 monthly sets of pooled option data with maturities 9-40 days.[11]

Our futures data comprises 5-minute intraday price quotes from Tickdatamarket. We retain the future contracts corresponding to the option contracts used for the jump tail estimation under the risk neutral measures in order to ensure an exact matching. After standard data cleaning procedures our sample comprises 6,327 trading days.[12] For the computation of the realized measures we use the part of the day were trading was actively carried out throughout the sample. Thus the first price observation is taken to be 10:00 (CST), and the last price observation 14:30. This leaves us with 54 price observations for each trading day.

3.3 Empirical Tail Risk Measures

The estimates for the tail variation measures under $\mathbb{Q}$ implied by equation (21) require a choice of the threshold k_t that separates large from small jumps. Similar to Bollerslev, Todorov, and Xu (2014) we allow k_t to vary as a linear function of the implied volatility. This form of time variation mimics the estimation procedure for the statistical jumps and accounts for the idea that what is classified as an "extreme" event can differs with economic conditions and corresponding market volatility. The following results are presented for the

estimation. For the computation of the expected quadratic variation under the $\mathbb{Q}$-measure we convert the option prices into corresponding European style values following Barone-Adesi and Whaley (1987).

[10] This "cleaning procedure" is standard, see e.g. Trolle and Schwartz (2010) or Bollerslev and Todorov (2011b).

[11] The monthly pooling ensures a sufficient number of options for estimation throughout the sample. It also has the advantage that potential monthly seasonalities are averaged out.

[12] The original sample constitutes 6519 trading days. Some days around Christmas, Thanksgiving and July 4th feature irregular trading hours and were discarded.

threshold k_t equal to three times the at-the-money Black-Scholes (BS) implied volatility.[13] This specification of k_t corresponds to a median threshold of 25 %.

The estimates for FI_t are presented in figure (5). Most of the noticeable movements in the series correspond to well known periods of oil price or aggregate uncertainty, such as the 1st Gulf War in 1990 and 1991, the financial crisis in 2008 and NBER recessions. Similar intuitive results are presented in figure (5) for the individual tail measures, alongside the risk neutral and statistical expected quadratic variations.

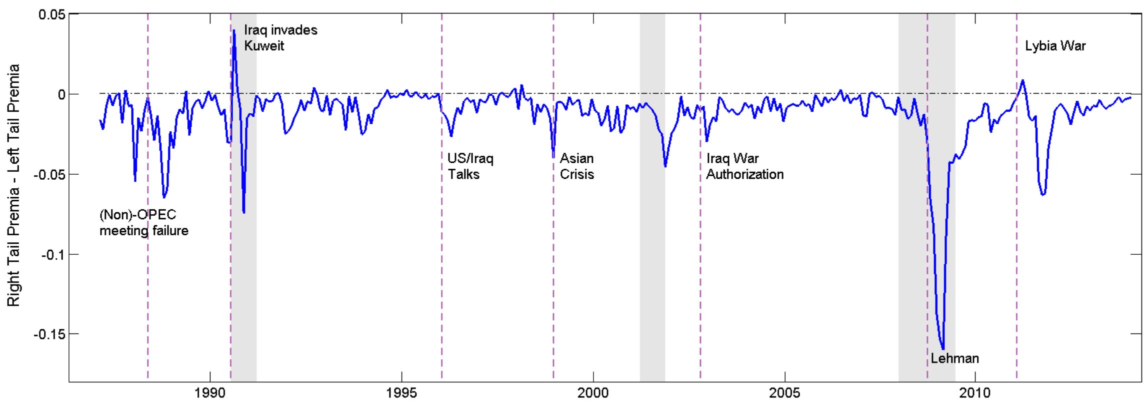

Figure 1: This figure displays the difference between the tail measures, $RJV_{t,T}^{\mathbb{Q}}$-$LJV_{t,T}^{\mathbb{Q}}$, in annualized form. The the tail measures are computed for the threshold $k_t = 3\times$ BS implied volatility. Shaded areas represent NBER recessions.

Similar to results from the equity index market presented by Bollerslev and Todorov (2011b), we find that in the crude oil futures markets the $\mathbb{P}$-tail distribution implied by the futures data is dwarfed by the corresponding $\mathbb{Q}$ measures. The estimates presented in table (1) indicate that on average, the statistical variation measure for the left tail is about 200 times smaller than the $\mathbb{Q}$ measure counterpart. The corresponding ratio for the right tail is around 50 and also sufficiently small in order to conclude that changes in the tail premia are almost entirely driven by movements in the tail variations under $\mathbb{Q}$ measure. Thus the tail variation premia appear well approximated by the risk neutral variation measures only:[14]

$$RJVP(k_t) \approx -RJV_{t,T}^{\mathbb{Q}} \qquad \text{and} \tag{25}$$

$$LJVP(k_t) \approx -LJV_{t,T}^{\mathbb{Q}}. \tag{26}$$

[13]We also experimented with other thresholds, obtaining qualitatively similar results. Table (6) in Appendix C displays the result for k_t four times the at-the-money Black-Scholes implied volatility.

[14]A similar approximation empirically holds for the equity market. As Bollerslev, Todorov, and Xu (2014) point out, this conveniently avoids peso-type estimation problems.

	Mean	SD	AR(1)
$LJV_{t,T}^{\mathbb{Q}}$	0.0233	0.0253	0.8444
$RJV_{t,T}^{\mathbb{Q}}$	0.0081	0.0101	0.7730
$LJV_{t,T}^{\mathbb{P}}$	0.0001	0.0001	0.6698
$RJV_{t,T}^{\mathbb{P}}$	0.0002	0.0002	0.7103
OILVIX2	0.1353	0.1293	0.8374
VRP	-0.0244	0.0741	0.1897

Table 1: Summary statistics for the monthly estimates of the tail variation measures and the traditional variation measures. SD stands for the standard deviation, AR(1) for first order autocorrelation. The sample period is 1987 to 2013, comprising 324 observations. All measures are presented in annualized form. The tail variation measures are evaluated at $k_t = 3 \times$ Black-Scholes implied volatility.

Interestingly, this indicates that changes in the objective jump distributions play a minor role in explaining the time variation in the size of the tail premia.[15]

A second important finding, reported in Appendix A, is that the statistical left and right tail variation measures are approximately symmetric. The symmetry implies that in each point of time, the conditional probabilities of a large upward jump is roughly equal to the conditional probability of a large downward jump. Together with strong time variation of the difference of $RJV_{t,T}^{\mathbb{Q}}$ and $LJV_{t,T}^{\mathbb{Q}}$ documented in figure (5), this provides additional evidence that the tail risk premia are only loosely connected to the statistical tail variation measures.

On average, the variation risk premia for the left tail is much larger than the premia for the right tail (table 1). This is consistent with commodity futures markets being in normal "backwardation", an idea first put forth by Keynes (1976). He postulated that producers of the physical commodity that want to hedge their output will have to pay a risk premium for speculators that take on the matching long positions in futures markets. Accordingly, speculators will demand a larger premium for their exposure to downside tail risk.

The implied total variation $QV_{t,T}^{\mathbb{Q}}$ is computed using all OTM options on a given day. For comparability with the tail measures, the monthly series we present in figure (5) is calculated by taking the average over the respective calendar month. In the computation of the realized total variation we also account for the overnight returns. The contribution of the squared overnight returns to the entire daily observation, π_t, is about 50% on average.[16] We compute the daily series by an appropriate scaling of the intraday realized variance RV_t and

[15]There is a similar finding for the equity index market. See e.g. Bollerslev and Todorov (2011a), Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and Xu (2014).

[16]This number is larger than then the average contribution of the squared overnight to the daily volatility, see e.g. Ahoniemi and Lanne (2013). Part of this is due to the relatively small active trading window we are considering for the realized measures in the previous section. For additional details the reader is referred to Appendix A.

	FI_t	FI_{t-1}	FI_{t-2}	FI_{t-3}
$r_{S,t}$	0.1426**	-0.0319	-0.119**	-0.148***
	(0.0552)	(0.0558)	(0.0555)	(0.0594)
$r_{F,t}$	0.2117***	0.0491	-0.0462	-0.1002*
	(0.0545)	(0.0557)	(0.0558)	(0.0562)

Table 2: Contemporaneous and 1-3 months spot return ($r_{S,t}$) and futures return ($r_{F,t}$) correlation with $FI_t \approx RJV^{\mathbb{Q}} - LJV^{\mathbb{Q}}$. The estimates are based on monthly observations from 1988 to 2013.

obtain the monthly series by averaging over days.[17] Our estimates of VRP_t are then based on the difference between the expected quadratic variation under the $\mathbb{Q}$-measure and the realized variation for the respective month.[18] Over our sample period, VRP_t is about -2%. This figure right between the numbers of Trolle and Schwartz (2010), who report an average premium of almost -3% over the 1996 - 2006 sample period and Pan and Kang (2013), who report a premium of -1.65% for a sample period slightly shorter than ours.

For a threshold of $k_t = 3 \times$ BS-implied volatility used here, the sum of the two tail premia is on average about 3% and slightly larger than the absolute value of the average VRP_t. Of the total VRP_t, about $E(RJV^{\mathbb{Q}})/E(VRP_t) = 33\%$ come from the right tail and about $E(LJV^{\mathbb{Q}})/E(VRP_t) = 94\%$ from the left tail. This suggests that the entire variance risk premium is due to compensation for tail variations, while variations due to continuous price movements and small jumps earn no or even a small positive premium.

In order to investigate how the tail premia interact with the spot prices, we also document the correlations of the relative movements in the premia - captured by the fear index FI_t - with current and future spot returns. Table (2) shows that the contemporaneous spot and future returns and FI_t are positively correlated, while this correlation turns negative for future spot and futures returns. Thus, these prices seem to exhibit an "overshooting" in the wake of large upside risks, allowing the price to decrease over the subsequent periods. We will address this issue more carefully in the next section where we discuss the forecasting properties of our indicators.

It is also interesting to see how the estimated jump intensities relate to the futures risk premium due to large jumps. Using our separate estimates for the left and the right tails under the respective measure, we obtain this premium through equation (5). Similar to

[17]Since there is no consensus in the literature, we experimented also with different forms of the scaling the contributions of the overnight returns. The level of the average variation presented in this section is of course not affected by the scaling, while the results for the monthly series used for forecasting in the next sections are qualitatively similar.

[18]This obviously differs from the definition of VRP_t which is based on the expectation of the realized variation under the $\mathbb{P}$-measure rather than the ex-post realized variation. For the unconditional estimates tabulated here this distinction does not matter if we assume that the expectation error is zero.

our results from the tail variation premia, the futures premia for the left tail tends to be larger than that for the right tail. The average of the (continuously compounded) futures premium due to large jumps period is 3.7% over our sample period. In comparison, the average (continuously compounded) total futures risk premium computed from the 1st back contract is 7.5%, suggesting that about half of the short maturity futures risk premium is due to tail risk.

4 Predictability of Futures and Spot Returns

4.1 Predictive Regression Framework and Control Variables

This section presents new predictability evidence of the tail premia for crude oil spot and futures returns. The baseline forecasts are performed in-sample, while cross-validation techniques are employed for the out-of-sample robustness check. If the premia capture oil market agents' attitudes toward tail risks, we would expect a large upside (downside) tail variation risk premium to be associated with relatively small (large) returns. We test this hypothesis in a regression framework for predictability of futures and spot returns of the form

$$r_{j,t+i} = \beta_{0,i} + \beta_{1,i} \cdot LJV^{\mathbb{Q}}_{t,T} + \beta_{2,i} \cdot RJV^{\mathbb{Q}}_{t,T} + Controls'_t \cdot \beta_{3,i} + \epsilon_{t,i}, \quad j = \{S, F\}, \tag{27}$$

where $r_{S,t+i}$ is the ith-month ahead spot return and $r_{F,t+i}$ is the ith-month ahead futures return, while $LJV^{\mathbb{Q}}_{t,T}$ and $RJV^{\mathbb{Q}}_{t,T}$ represent the left and the right tail premia, respectively, approximated by $\mathbb{Q}$-measure variations. The vector $Controls_t$ include both macroeconomic-financial variables and crude oil market specific variables that are potential predictors of commodity spot and futures returns (see e.g. Bessembinder and Chan 1992; Hong and Yogo 2012). For $i > 1$, we employ overlapping regressions in order to enhance the efficiency of our estimates, using robust Newey-West standard errors so as to account for the autocorrelation in the residuals induced by the overlap. The lag length for the computation of the standard errors is chosen twice the length of the overlap.

We first describe the set of oil market specific control variables. As suggested by the recent literature, we include the estimated VRP_t. We compute VRP_t as the difference between the lagged expected variation under the risk neutral and the actual variation in period t as suggested in Bollerslev, Tauchen, and Zhou (2009) and as described in the previous section. For additional robustness analysis, we also include the contemporaneous expectation $QV^{\mathbb{Q}}_t$ and the contemporaneous realized variation separately. The oil market specific variables further include changes in oil inventories, obtained as the monthly storage level from the web site of the U.S. Energy Information Administration (EIA) and open interest growth.

The computation of the open interest variable is based on the open interest of futures and options combined obtained from the CFTC website and computed as the 12-month growth rates taking geometric averages as suggested by Hong and Yogo (2012). Finally, we include the slope of the term structure as measured by the net ratio of the current spot price over the 1st back futures contract.[19]

We control for macroeconomic conditions by including the short term interest rate, computed as the yield of a 3-month T-Bill, and the Aruoba-Diebold-Scotti Business Conditions Index (ADS) published by the Federal Bank of Philadelphia. The index is based on may economic indicators in the U.S. and a higher value is associated with better economic conditions. Since crude oil prices might be also driven by global rather than US-specific factors, we include the Real Activity Index developed in Kilian (2009). The other variables include the yield spread, computed as the difference between Moody's Aaa and Baa corporate bond yields and the CBOE VIX, a measure of the implied volatility of S&P 500 index options.

4.2 Forecasting Results

We first discuss the forecasting results for crude oil futures returns, presented in columns (5) - (8) of table (3). For all regressions, the coefficients have the expected sign. A relatively high right tail premium is associated with negative futures returns, and the left tail premium with positive returns. For the model without controls, all coefficients are statistically significant at the 5% significance level and most at the 1% significance level. This confirms our intuition that the tail risk premia are associated with a substantial change in the oil futures premia. The results are robust to the inclusion of standard predictors of crude oil prices. The only exception is the noticeable rise in the standard error for the left tail premium in the three month horizon regression, which renders the coefficient statistically insignificant at conventional significance levels. This indicates certain degree of correlation with some of the predictor variables, which is also noted through the increasing coefficient for the coefficient associated with the right tail premium when controlling for the other predictors. Jointly, the tail premia are always significant at the 1% significance level as measured through F-test. The adjusted R^2 is 3.7% for the three month horizon and and 6.8% for the six month horizon regression, which amounts to almost one fourth of the predictability associated with all regressors.

We now describe the forecasting results for crude oil spot returns, presented in columns (1) - (4) of table (3). For all regressions, the coefficients associated with the tail premia have similar sizes to those of the futures regressions. Again, the significance of the coefficient

[19] We also experimented with other definitions of the slope of the term structure, e.g. the net ratio of the spot price with the 3rd back futures contract and the net ratio of the first two futures contracts, yielding almost identical results.

Variables	(1) $r_{S,t+3}$	(2) $r_{S,t+3}$	(3) $r_{S,t+6}$	(4) $r_{S,t+6}$	(5) $r_{F,t+3}$	(6) $r_{F,t+3}$	(7) $r_{F,t+6}$	(8) $r_{F,t+6}$
L3.$RJV_{t,T}^{\mathbb{Q}}$	-7.70***	-10.13***			-5.27***	-9.54***		
	(2.044)	(2.254)			(1.872)	(2.364)		
L3.$LJV_{t,T}^{\mathbb{Q}}$	2.56***	1.85*			1.59***	1.49		
	(0.510)	(1.059)			(0.546)	(1.025)		
L6.$RJV_{t,T}^{\mathbb{Q}}$			-11.03***	-16.83***			-9.76**	-18.32***
			(2.845)	(2.901)			(3.875)	(3.539)
L6.$LJV_{t,T}^{\mathbb{Q}}$			5.27***	3.80***			3.47***	3.99***
			(0.698)	(1.413)			(0.860)	(1.460)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Adj. R^2	0.0877	0.2371	0.1419	0.3131	0.0374	0.1634	0.0685	0.2827
Obs.	321	283	318	282	321	283	318	282
Wald test (p-value)	<0.001	<0.001	<0.001	<0.001	0.005	<0.001	<0.001	<0.001

Newey-West standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3: Forecasting results for $i = 3$ and $i = 6$ months. The dependent variable $r_{S,t+i}$ stands for spot returns, $r_{F,t+i}$ for futures returns. Wald test stands for an F-test of the joint significance of $RJV_{t,T}^{\mathbb{Q}}$ and $LJV_{t,T}^{\mathbb{Q}}$. The estimation period is 1987 - 2013. The differences in the number of observations is due to data availability for the control variables.

for the left tail premium in the three month regression when the other predictor variables are is slightly decreased due to the rise in the standard error relative to the model without predictors. For all other regressions, the individual coefficients are significant at the 1% significance level. Again, a joint F-test for the premia allows us to reject the null of no significance at all conventional significance levels for all regressions. The adjusted R^2s associated with the spot return regressions are much larger than those for futures returns. At the three month horizon, the tail measures account for about 9% of the variability in the spot return, while this number increases to 14% at the six month horizon, constituting almost half of the in-sample variability explained by all regressors.[20] Thus the tail risk measures have strong predictability for futures and spot price returns in the expected direction, and, judging from the fit of the regression, this predictability is larger for the spot than for the futures price. This is consistent with the idea that fears about large oil price fluctuations affect both the futures and the spot price, and that the convenience yield moves in the same direction as the futures risk premium.

For forecasting purposes, a high R^2 is not always indicative of a good model since it is always possible to increase the in-sample fit by adding additional regressors without improving out-of-sample forecasting power. In order to safeguard against a potential in-sample-overfit,

[20] We also assessed whether the tail risk measures posses incremental forecasting power with respect to the control variables. This is indeed the case: The incremental adjusted R^2 for the regressions including all variables in comparison with regressions using the control variables only is about 10% and 11% for futures and spot returns, respectively, for the six months horizon and about 5% and 6% for the three months horizon.

we perform an out-of-sample cross-validation. Our cross-validation procedure uses the entire series as both in-sample and out-of-sample data. Specifically, we run repeating regressions using a single return observation from the original sample as the validation data, and the remaining observations as the training data.[21] This is repeated until each observation in the sample is used once as the validation data, so that finally every return observation was treated "out-of-sample" once. The cross-validation statistics is based on the mean squared prediction error, hence lower values indicate a better fit. In order to evaluate the predictive out-of-sample performance of our predictors we estimate the forecasting model presented in equation (27) in the previous section first including the tail premia and then excluding it. The results are displayed in table (7) in Appendix C. Including the tail premia as predictive regressors reduces the cross validation statistics substantially for all models considered, indicating that their forecasting power is not only due to in-sample properties. Importantly, the out-of-sample MSPE for the model including $LJV_{t,T}^{\mathbb{Q}}$ and $RJV_{t,T}^{\mathbb{Q}}$ without additional control variables as regressors is always lower than the forecast assuming a constant return and the no-change forecast. In contrast, the model including only the control variables displays a higher out-of-sample MSPE against the no-change forecast for the spot price regressions, suggesting that the R^2 related to these variables is driven at least in part by in-sample overfitting.

4.3 Risk or Fear: Where does the predictability come from?

After having provided evidence for the forecasting power of our novel predictors, we now turn to the question whether the premia a compensation for potential risks or rather describe the oil market's attitudes towards risks. From the definition of FI_t it follows that for approximately symmetric jump tails under the $\mathbb{P}$ measure,

$$\begin{aligned} FI_t &= \frac{1}{T-t}\left[\left(E_t^{\mathbb{P}}(LJV_{t,T}^{\mathbb{P}}) - E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{Q}})\right) - \left(E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{P}}) - E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{Q}})\right)\right] \\ &\approx \frac{1}{T-t}\left[E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{Q}}) - E_t^{\mathbb{Q}}(RJV_{t,T}^{\mathbb{Q}})\right]. \end{aligned} \tag{28}$$

In this case, as shown by Bollerslev, Todorov, and Xu (2014), FI_t will be largely void of risk compensation associated with the temporal changes in the jump intensities and is therefore naturally interpreted as a proxy of oil market fears. The notion of a "fear index" is also warranted because the premia are not directly related to asymmetries in future actual jump probabilities. In table (8) presented in Appendix C we show that the $\mathbb{Q}$ tail variation

[21] We assure that these returns are completely non-overlapping with the remaining returns by leaving an out of sample window of ± 4 months around the respective return for $i = 3$ months and ± 7 months for $i = 6$ months.

measures do not display any forecasting power for the realized jump variation beyond that predicted by symmetric measures such as the realized variance. In contrast, the forecasting results shown in table (4) indicate that the return predictability through our novel predictors is mainly driven by its asymmetry. Here we use the 6 month ahead forecasting regression and decompose the tail measure in their difference, FI_t and a level component $RJV^{\mathbb{Q}}_{t,T} + LJV^{\mathbb{Q}}_{t,T}$. Using these variables as single regressors, displayed in columns (2) and (3), shows that the regression with the fear index as a predictor variable yields a statistically significant coefficient that indicates that upside fears are associated with decreasing spot price returns. The R^2 for this regression is almost 10%, while the R^2 for the regression with the tail risk level measure is less than 2% and yields an statistically insignificant coefficient.[22] These results are consistent with the idea that time variation in the effective risk aversion of oil market participants is an important driver of the short run fluctuation of the price of oil. Taken together, the results presented in this section are confirm evidence for the predictability of both futures and spot returns by our novel risk indicators.

Variables	(1) $r_{S,t+6}$	(2) $r_{S,t+6}$	(3) $r_{S,t+6}$	(4) $r_{S,t+6}$
L6.$RJV^{\mathbb{Q}}_{t,T}$	-11.03***			
	(2.626)			
L6.$LJV^{\mathbb{Q}}_{t,T}$	5.27***			
	(0.615)			
L6.FI_t		-3.84***		-8.15***
		(0.569)		(1.559)
L6.$(RJV^{\mathbb{Q}}_{t,T} + LJV^{\mathbb{Q}}_{t,T})$			1.01	-2.88***
			(0.881)	(1.099)
Adj. R^2	0.1419	0.0950	0.0180	0.1419
Obs	318	318	318	318

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 4: Forecasting results for 6 month oil spot returns, denoted by $r_{S,t+6}$. The estimation period is 1987 - 2013.

4.4 VAR Estimates on the Impact of the Oil Fear

In this section, we compare our forecasting results with estimates from VAR models, where we include FI_t as a directional measure of oil market uncertainty. The VAR framework has shown useful for structural analysis of oil price shocks and forecasting the real price of oil (Kilian 2009; Baumeister and Kilian 2012). Building on the reduced form version of

[22]Including both the fear index and the level component jointly yields, as in column (4), obviously yields the same R^2 as the baseline.

the model proposed by Kilian and Murphy (2014), the variables in the estimation order are oil production (in percentage changes), the global real activity proposed by Kilian (2009), changes in crude oil inventories, FI_t, the six month futures spread ($CY_{t,6}$) and the nominal price of oil (in percentage changes).[23] In a second exercise, we replace the price of oil by the percentage change of the first back futures contract.[24] The identification of shocks to FI_t is based on a Cholesky-decomposition of the reduced form errors. Including production, real activity and inventory as the first variables in the VAR ensures the impact of these variables is already controlled when looking at the impact of shocks to FI_t on the oil and futures price.[25] Given the tradeoff between overparameterization and allowing for sufficient lags to account for business cycle effects, we estimate the VAR with 12 lags.

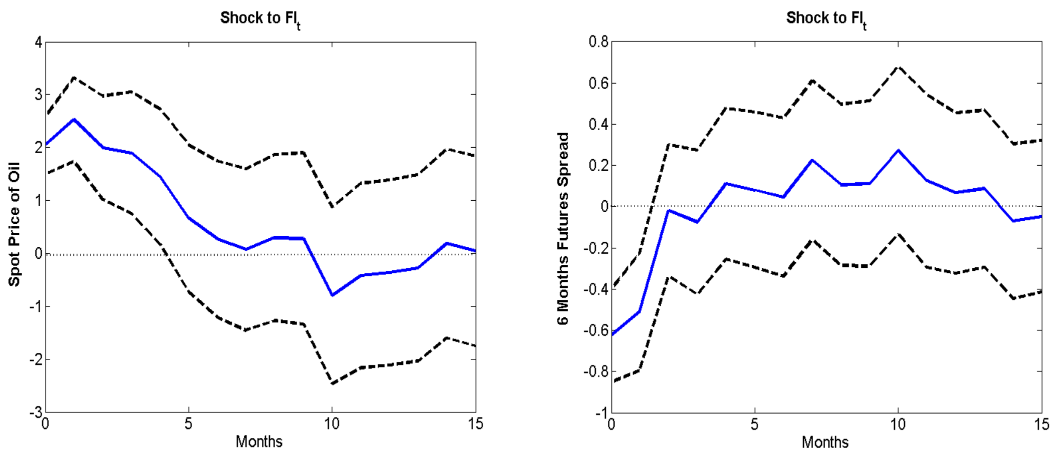

Figure 2: Responses of the spot price of oil and the six months futures spread to a one standard deviation shock. The estimates are based on monthly data, 1987:1 - 2012:6. Dashed lines represent 90% confidence intervals.

Figure (2) plots the impulse response function of oil spot prices and the six months futures spread to a shock in the oil fears index. Oil spot prices react with an instantaneous increase of about 2%, and a gradual reversion to previous levels over the following months. The 90 % confidence intervals are plotted around this, highlighting that this impact is statistically significant over the first months. The six months futures spread reacts with an immediate, yet much smaller decline and a faster reversion to previous levels. Instead, the response of inventories to fear shocks, presented in Appendix C, do not provide any evidence for a

[23]Expressing the price of in log percentage changes ensures that the cumulative impulse responses reflect the percentage change of oil over the entire period.

[24]The calculation of the percentage change of the first back futures contract is always based on the same contract for a given period. However across periods, contracts will differ.

[25]As pointed out by Kilian and Murphy (2014), inventory will in general also adjust quickly in response to uncertainty shocks. The results presented here are robust to a change in the ordering of the variables such that FI_t is included before inventories.

systematic reaction of inventories to FI_t.[26] These effects are consistent with the idea that shocks to upside (downside) fears are associated with an immediate increase (decrease) of the price of oil that is due to the combination of an increase (decrease) in the net convenience yield and an decrease (increase) in the risk premium. As such, changes in relative uncertainty do not require an immediate response of inventory in order to display discernible effects on prices.

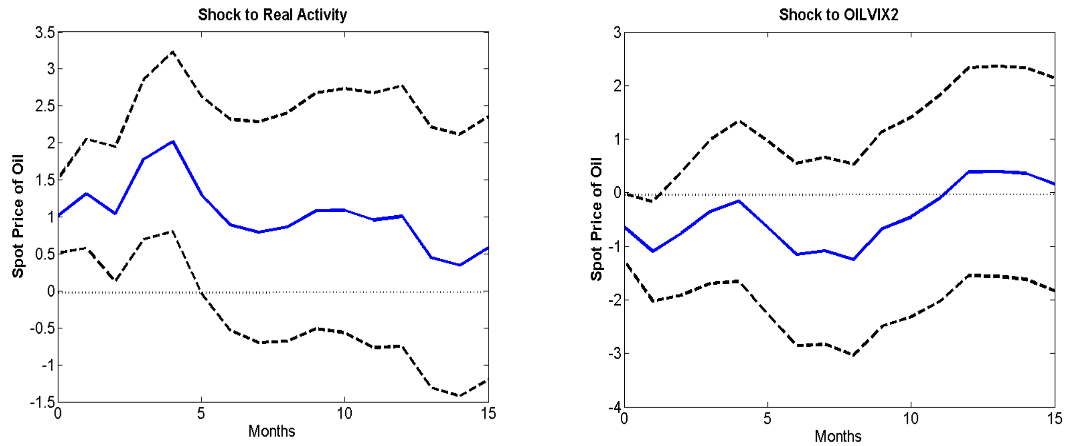

Figure 3: Responses of the spot price of oil to shocks in real activity and the option implied variance, OILVIX2. The estimates for the responses to the implied variance are based on a VAR model that contains the same ordering as described in the text and OILVIX2 instead of FI_t. The estimates are based on monthly data, 1987:1 - 2012:6. Dashed lines represent 90% confidence intervals.

For comparison, the responses of spot prices to a shock to real activity and and implied oil price volatility are presented in figure (3). The reaction of the price of oil to real activity appears more persistent, yet smaller on impact, whereas the estimated effect of the implied oil price volatility shocks is largely insignificant. The results are consistent with the idea that changes in the relative uncertainty about upside and downside fears are an important driver of short run fluctuations in oil prices and futures returns.

5 Interaction with the Macroeconomy

5.1 Oil and Aggregate Uncertainty

Thus far, we have treated the oil risk factors in isolation from the aggregate asset uncertainty. In this section we address their interaction. A recent strand of literature, building on

[26] This might also be an indication of the redundancy of inventories when the futures spread is included. I thank Christiane Baumeister for the comment.

Bollerslev, Tauchen, and Zhou (2009), has shown that the variance risk premium embedded in stock market index options and futures is a suitable measure of aggregate uncertainty. This work has been extended by Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and Xu (2014), suggesting that a substantial fraction of the variance risk premium and its forecasting power for stock market and portfolio returns is due to the aggregate "fears" as measured through the stock market fear index. This raises the question whether the aggregate risk measures proposed by these authors predicts oil spot and future prices beyond the oil risk measures proposed in this paper. Table (9) in Appendix D shows that the oil and stock market tail variation measures are indeed highly correlated.[27] The left tail oil variation measure $LJV_{t,T}^{\mathbb{Q}}$ and the stock market fear index $FI_{t,SPX}$ exhibit the highest correlation of 65%.

In table (10) in Appendix D we present the forecasting results for six months futures and spot price returns, respectively. The aggregate fear index and variance risk premium appear to contain some explanatory power for the 6 month spot price return regression, with an R^2 of about 5%.[28] However, this effect is completely dominated by the oil specific tail risk variation measures, implying that the oil specific measures already entail relevant information from the aggregate measures. On a more general level these results also suggest that time-varying disaster fears embedded in option prices on individual assets, not only on market indexes, convey important information on individual premia beyond that implied by the market.

In addition to these results, we also address the question whether $LJV_{t,T}^{\mathbb{Q}}$ and $RJV_{t,T}^{\mathbb{Q}}$ predict stock market returns. We find little evidence of a stable relationship, since these forecasting results depend crucially on the sample period, in particular on the inclusion of data during and after the financial crises. Table (11) in Appendix D shows that prior to the outbreak of the financial crisis, a relatively high left tail variation measure was associated with lower stock market returns, and a relatively high right tail variation with higher stock market returns. However, extending the sample beyond the outset of the financial crisis this relationship breaks down. These results - in terms of both the significance and signs of the coefficients - are consistent with the idea that $LJV_{t,T}^{\mathbb{Q}}$ and $RJV_{t,T}^{\mathbb{Q}}$ aggregate different types of uncertainty that are relevant for oil prices, e.g. oil supply and oil demand uncertainty, that individually have a very distinct relationship with the aggregate stock market. Prior to the financial crisis, the most notable event in terms of oil fears was the 1991 Gulf War

[27] The stock market variance risk premium $VRP_{t,SPX}$ and fear index $FI_{t,SPX}$ are based on S&P 500 index options and futures following the methodology of Bollerslev, Todorov, and Xu (2014) for the years 1996 to 2013. I would like to thank Lai Xu for providing the data on the S&P 500 fear index and Marek Raczko for providing the data on the S&P 500 variance risk premium.

[28] The important predictor seems to be $VRP_{t,SPX}$, a result that is consistent with the work of of Bollerslev, Todorov, and Xu (2014) who show that $VRP_{t,SPX} - FI_{t,SPX}$ is mainly associated with economic uncertainty, while $FI_{t,SPX}$ captures attitudes toward risk.

episode, which is clearly identified with supply risk (Alquist and Kilian 2010). On the other hand, the financial crisis and its aftermath were associated mainly with uncertainty about future demand for oil.[29].

5.2 Oil Uncertainty and Real Activity

Bernanke (1983) and more recently Elder and Serletis (2010) and Jo (2014) pointed out that uncertainty about oil prices, rather than the fluctuations per se, can have important effects on real economic activity. In this section we investigate this hypothesis by using our novel uncertainty measures to predict industrial production, which is available at a monthly frequency.

Our regressions, presented in table (12) in Appendix D, indicate that the right tail variation measure $RJV_{t,T}^{\mathbb{Q}}$ has a statistically significant impact on growth in industrial production, while the left tail variation measure does not contain any additional information. The incremental R^2 from the inclusion of the right tail measure is about 5%, is robust to the inclusion of oil market and aggregate control variables, and extends also to the 6 month horizon.[30] Intuitively, a large $RJV_{t,T}^{\mathbb{Q}}$ is always bad news for the oil importing economy because it can be due to (i) fears related to supply cuts (2) uncertainty about (global) economic growth, both of which should have a negative effect on economic activity. In contrast, a large premium for the left tail might be good news for the economy is the corresponding fears are related to oil market specific events.

[29]A different explanation for this result might be that oil derivative markets have only recently become more integrated in the broader financial system, and that this integration went alongside fundamental changes between the stock market and oil price relationship (Christoffersen and Pan 2014). We leave this hypothesis for further research

[30]These results are confirmed by VAR estimates in a framework proposed by Bloom (2009). The results are available upon request

6 Conclusion

Oil prices are difficult to forecast and exhibit wild swings or "excess volatility" that are difficult to rationalize by changes in fundamentals alone. We find that the jump risk premia embedded on crude oil future options contain important information on oil market fears and contribute to the explanation of oil price volatility. These premia are economically large, vary substantially over time and significantly forecast crude oil futures and spot returns. This result is robust after controlling for macro-finance and oil market specific variables, and importantly, for time-varying aggregate disaster fear as measured by $S\&P500$ option implied tail risk. Instead, our oil uncertainty measures appears to conveniently aggregate oil price uncertainty derived from different sources, e.g. oil supply and oil demand uncertainty, that individually have a very distinct relationship with aggregate uncertainty.

We show that oil futures prices overshoot (undershoot) in the presence of upside (downside) tail fears in order to allow for smaller (larger) risk premia thereafter. Consistent with the theory of storage, this overshooting (undershooting) is amplified for the spot price because of time varying benefits from holding inventory that work in the same direction. These results are complementary to storage models using risk neutrality, and stress the importance of time varying risk premia in explanations of large swings in oil prices. On a more general level it is shown that time-varying disaster fears embedded in option prices on individual assets, not only on market indexes, convey important information on the risk premia and price dynamics of these assets.

References

Ahoniemi, Katja and Markku Lanne (2013). "Overnight stock returns and realized volatility." *International Journal of Forecasting* 29.4, pp. 592–604.

Alquist, Ron and Lutz Kilian (2010). "What do we learn from the price of crude oil futures?" *Journal of Applied Econometrics* 25.4, pp. 539–573.

Barone-Adesi, Giovanni and Robert E Whaley (1987). "Efficient analytic approximation of American option values." *The Journal of Finance* 42.2, pp. 301–320.

Barro, Robert J (2006). "Rare disasters and asset markets in the twentieth century." *The Quarterly Journal of Economics*, pp. 823–866.

Baumeister, Christiane and Lutz Kilian (2012). "Real-time forecasts of the real price of oil." *Journal of Business & Economic Statistics* 30.2, pp. 326–336.

Baumeister, Christiane and Lutz Kilian (2014). "A general approach to recovering market expectations from futures prices with an application to crude oil." *CFS Working Paper*.

Baumeister, Christiane and Gert Peersman (2013). "The role of time-varying price elasticities in accounting for volatility changes in the crude oil market." *Journal of Applied Econometrics* 28.7, pp. 1087–1109.

Bernanke, Ben S (1983). "Irreversibility, uncertainty, and cyclical investment." *The Quarterly Journal of Economics* 98.1, pp. 85–106.

Bessembinder, Hendrik (1992). "Systematic risk, hedging pressure, and risk premiums in futures markets." *Review of Financial Studies* 5.4, pp. 637–667.

Bessembinder, Hendrik and Kalok Chan (1992). "Time-varying risk premia and forecastable returns in futures markets." *Journal of Financial Economics* 32.2, pp. 169–193.

Bloom, Nicholas (2009). "The impact of uncertainty shocks." *econometrica* 77.3, pp. 623–685.

Bollerslev, Tim, George Tauchen, and Hao Zhou (2009). "Expected stock returns and variance risk premia." *Review of Financial Studies* 22.11, pp. 4463–4492.

Bollerslev, Tim and Viktor Todorov (2011a). "Estimation of jump tails." *Econometrica* 79.6, pp. 1727–1783.

Bollerslev, Tim and Viktor Todorov (2011b). "Tails, fears, and risk premia." *The Journal of Finance* 66.6, pp. 2165–2211.

Bollerslev, Tim and Viktor Todorov (2013). "Time-Varying Jump Tails." *Journal of Econometrics, forthcoming.*

Bollerslev, Tim, Viktor Todorov, and Lai Xu (2014). "Tail Risk Premia and Return Predictability." *Working Paper.*

Carr, Peter and Liuren Wu (2009). "Variance risk premiums." *Review of Financial Studies* 22.3, pp. 1311–1341.

Christoffersen, Peter and Xuhui Nick Pan (2014). "Oil Risk Exposure and Expected Stock Returns." *SSRN Working Paper.*

Doran, James S and Ehud I Ronn (2008). "Computing the market price of volatility risk in the energy commodity markets." *Journal of Banking & Finance* 32.12, pp. 2541–2552.

Elder, John and Apostolos Serletis (2010). "Oil price uncertainty." *Journal of Money, Credit and Banking* 42.6, pp. 1137–1159.

Erb, Claude B and Campbell R Harvey (2006). "The strategic and tactical value of commodity futures." *Financial Analysts Journal*, pp. 69–97.

Fama, Eugene F and Kenneth R French (1987). "Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage." *Journal of Business*, pp. 55–73.

Gorton, G.B. and K.G. Rouwenhorst (2004). "Facts and fantasies about commodity futures." *NBER Working Paper.*

Gospodinov, Nikolay and Serena Ng (2013). "Commodity prices, convenience yields, and inflation." *The Review of Economics and Statistics* 95.1, pp. 206–219.

Hamilton, James D and Jing Cynthia Wu (2014). "Risk premia in crude oil futures prices." *Journal of International Money and Finance* 42, pp. 9–37.

Hong, Harrison and Motohiro Yogo (2012). "What does futures market interest tell us about the macroeconomy and asset prices?" *Journal of Financial Economics* 105.3, pp. 473–490.

Jo, Soojin (2014). "The Effects of Oil Price Uncertainty on Global Real Economic Activity." *Journal of Money, Credit and Banking* 46.6, pp. 1113–1135.

Kaldor, Nicholas (1939). "Speculation and economic stability." *The Review of Economic Studies* 7.1, pp. 1–27.

Keynes, J.M. (1976). *A treatise on money.* New York; AMS Press.

Kilian, Lutz (2009). "Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market." *The American Economic Review*, pp. 1053–1069.

Kilian, Lutz and Daniel P Murphy (2014). "The role of inventories and speculative trading in the global market for crude oil." *Journal of Applied Econometrics* 29.3, pp. 454–478.

Mancini, Cecilia (2001). "Disentangling the jumps of the diffusion in a geometric jumping Brownian motion." *Giornale dell' Istituto Italiano degli Attuari* 64, pp. 19–47.

Milonas, Nikolaos T and Stavros B Thomadakis (1997). "Convenience yields as call options: an empirical analysis." *Journal of Futures Markets* 17.1, pp. 1–15.

Pan, Xuhui Nick and Sang Baum Kang (2013). "Commodity Variance Risk Premia and Expected Futures Returns: Evidence from the Crude Oil Market." *Working Paper.*

Singleton, Kenneth J (2013). "Investor flows and the 2008 boom/bust in oil prices." *Management Science* 60.2, pp. 300–318.

Szymanowska, Marta et al. (2014). "An anatomy of commodity futures risk premia." *The Journal of Finance* 69.1, pp. 453–482.

Tang, Ke and Wei Xiong (2010). "Index investment and financialization of commodities." *NBER Working Paper*.

Trolle, Anders B and Eduardo S Schwartz (2010). "Variance risk premia in energy commodities." *The Journal of Derivatives* 17.3, pp. 15–32.

Wachter, Jessica A (2013). "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *The Journal of Finance* 68.3, pp. 987–1035.

Working, Holbrook (1949). "The theory of price of storage." *The American Economic Review*, pp. 1254–1262.

7 Appendix: Computation of the Jump Measures

Estimating realized and expected jumps under the statistical measure

This sections provides further details on the estimation of the jump properties under the statistical (objective) measure, which are computed from data on 5-minute intraday returns. We use the notation of Bollerslev and Todorov (2011b) and divide the trading day t into the $[t, t+\pi_t]$ overnight period and the $[t+\pi_t, t+1]$ active trading period, comprising $n+1 = 54$ price observations. Denoting $\Delta_{t,i} f \equiv f_{t+\pi+i} - f_{t+\pi+i-1}$, where f denotes the logarithm of the futures price, we have for a suitable threshold $\alpha_{t,i}$

$$\sum_{i=1}^{n} (\Delta_{t,i} f)^2 \overset{p}{\to} \int_{t+\pi_t}^{t+1} \sigma_s^2 ds + \int_{t+\pi_t}^{t+1} \int_{\mathbb{R}} x^2 \mu(ds, dx) \quad \text{and} \tag{29}$$

$$\sum_{i=1}^{n} (\Delta_{t,i} f)^2 \, 1_{|\Delta_{t,i} f \leq \alpha_{t,i}|} \overset{p}{\to} \int_{t+\pi_t}^{t+1} \sigma_s^2 ds \equiv CV_t. \tag{30}$$

We allow the truncation levels $\alpha_{t,i}$ to vary with both the daily and the intraday volatility following Bollerslev and Todorov (2011a). The time-of-day factor, TOD_j is then computed via

$$TOD_j = \frac{\sum_{m=0}^{N} (\Delta_{m(\pi+n),i} f)^2 1_{|\Delta_{m(\pi+n),i} f| < \bar{\alpha}}}{\sum_{m=0}^{N} |\Delta_{m(\pi+n),i} f| < \bar{\alpha}} \Big/ \frac{\sum_{m=0}^{N} \sum_{j=1}^{n} (\Delta_{m(\pi+n),i} f)^2}{\sum_{m=0}^{N} \sum_{j=1}^{n} 1_{|} \Delta_{m(\pi+n),i} f| < \bar{\alpha}}, \tag{31}$$

where $\bar{\alpha} = 3\sqrt{\Pi \cdot 0.5}\sqrt{\frac{1}{N}\sum_{m=0}^{N}\sum_{j=1}^{n-1} |\Delta_{m(\pi+n),j} f||\Delta_{m(\pi+n),j+1} f|}$. Given the time-of-day factor, the time-varying threshold $\alpha_{j,t}$ is then computed as

$$\alpha_{j,t} = 3(\frac{1}{n})^{0.49} \sqrt{CV_{t-n,t}} TOD_j. \tag{32}$$

The dynamics of our empirical jumps in equation (24), require an estimate of $v^{\mathbb{P}}$. This estimate is based on medium and large sized jump tails using the EVT proposed by Bollerslev and Todorov (2011a). In particular, defining $\psi^+(x) = e^x - 1$ and $\psi^-(x) = e^{-x}$, $v_\psi^+(y) = \frac{v(ln(y+1))}{y+1}$ and $v_\psi^-(y) = \frac{v(-ln\ y)}{y}$, $y > 0$, the jump tail measures are

$$\bar{v}_\psi^{\pm}(x) = \int_x^\infty v_\psi^{\pm}(u), \tag{33}$$

with $x > 0$ for $\bar{v}_\psi^+(x)$, and $x > 1$ for $\bar{v}_\psi^-(x)$. Under the assumption that $\bar{v}_\psi^{\pm}$ belong the the domain of attraction of an extreme value distribution (see Bollerslev and Todorov 2011b), it follows that

$$1 - \frac{\bar{v}_\psi^\pm(u+x)}{\bar{v}_\psi^\pm} \sim G(u;\sigma^\pm,\xi^\pm), \qquad u>0, x>0, \tag{34}$$

where $G(u;\sigma^\pm,\xi^\pm)$ is the CDF of a generalized Pareto distribution with

$$G(u;\sigma^\pm,\xi^\pm) = \begin{cases} 1-(1+\xi^\pm u/\sigma^\pm)^{-1/\xi^\pm}, & \xi^\pm \neq 0, \sigma^\pm > 0 \\ 1-e^{-u/\sigma^\pm}, & \xi^\pm = 0, \sigma^\pm > 0. \end{cases} \tag{35}$$

Now, for a large threshold $tr^\pm$, the integrals corresponding to the jump tail measures under $\mathbb{P}$ are a function of the parameter vector

$$\Theta \equiv [\sigma^\pm, \xi^\pm, \alpha_0^\pm \bar{v}_\psi^\pm(tr^\pm), \alpha_1^\pm \bar{v}_\psi^\pm(tr^\pm)], \tag{36}$$

which are estimated using the exact GMM framework suggested by Bollerslev and Todorov (2011a). The moment conditions used for estimation are

$$\begin{aligned} \frac{1}{N}\sum_{t=1}^{N}\sum_{i=1}^{n}\phi_j^\pm(\psi^\pm(\Delta_{t,i}f)-tr^\pm)1_{\psi^\pm(\Delta_{t,i}f)>tr^\pm} &= 0, \quad j=1,2 && (37)\\ \frac{1}{N}\sum_{t=1}^{N}\sum_{i=1}^{n}1_{\psi^\pm(\Delta_{t,i}f)>tr^\pm} - \alpha_0^\pm\bar{v}_\psi^\pm - \alpha_1^\pm\bar{\psi}^\pm(tr^\pm)CV_t &= 0 && (38)\\ \frac{1}{N}\sum_{t=2}^{N}\left(\sum_{i=1}^{n}1_{\psi^\pm(\Delta_{t,i}f)>tr^\pm} - \alpha_0^\pm\bar{v}_\psi^\pm(tr^\pm) - \alpha_1^\pm\bar{v}_\psi^\pm(tr^\pm)CV_t\right)CV_{t-1} &= 0, && (39) \end{aligned}$$

where

$$\begin{aligned} \phi_1^\pm(u) &= -\frac{1}{\sigma^\pm} + \frac{\xi^\pm}{(\sigma^\pm)^2}\left(1+\frac{1}{\xi^\pm}\right)\left(1+\frac{1+\xi^\pm u}{\sigma^\pm}\right)^{-1} \quad \text{and} && (40)\\ \phi_2^\pm(u) &= \frac{1}{(\xi^\pm)^2}ln\left(1+\frac{\xi^\pm u}{\sigma^\pm}\right) - \frac{u}{\sigma^\pm}\left(1+\frac{1}{\xi}\right)\left(1+\frac{\xi^\pm u}{\sigma^\pm}\right)^{-1} && (41) \end{aligned}$$

are the scores associated with the log-likelihood function of the generalized Pareto distribution. Facing the trade off between a sufficient number of observations of medium and large jumps on the one hand, and the approximation of the jump tails by the generalized Pareto distribution on the other, our choice of $tr^\pm$ corresponds to a jump in the log price of $\pm 1.2\%$. In total, we detect 3266 (3756) positive (negative) jumps, out of which 134 (198) are above the threshold. The large jumps, displayed in figure (4), seem to cluster and the occurrences of positive and negative jumps appear relatively symmetric.

The parameter estimates for our specification of the statistical large jumps' dynamics,

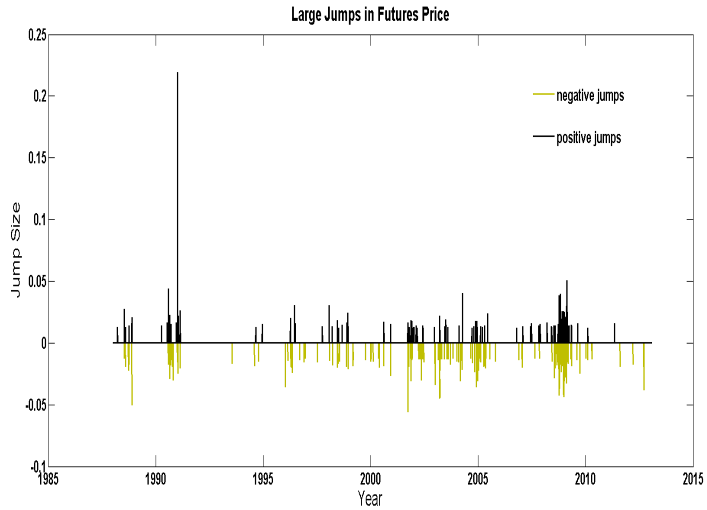

Figure 4: Detected large jumps in the log price. The threshold is ±1.2%, the estimates are based on 5-minute intraday futures data from 1988 to 2013.

presented in table (5) provide further evidence for this symmetry.

Left Tail			Right Tail		
Parameter	Estimate	St. Error	Parameter	Estimate	St.Error
ξ^-	0.3107	0.0707	ξ^+	0.3603	0.1054
$100 \cdot \sigma^-$	0.4317	0.0429	$100 \cdot \sigma^+$	0.3446	0.0457
α_0^-	-0.1194	0.0059	α_0^+	-0.0906	0.0046
α_1^-	11.7926	0.4172	α_1^+	8.7404	0.3248

Table 5: Estimates for $\mathbb{P}$ tail parameters. The estimates are based on 5-minute intraday futures data from 1988 to 2013.

8 Appendix: Data and Empirical Measures

Tail Risk and Traditional Volatility Measures

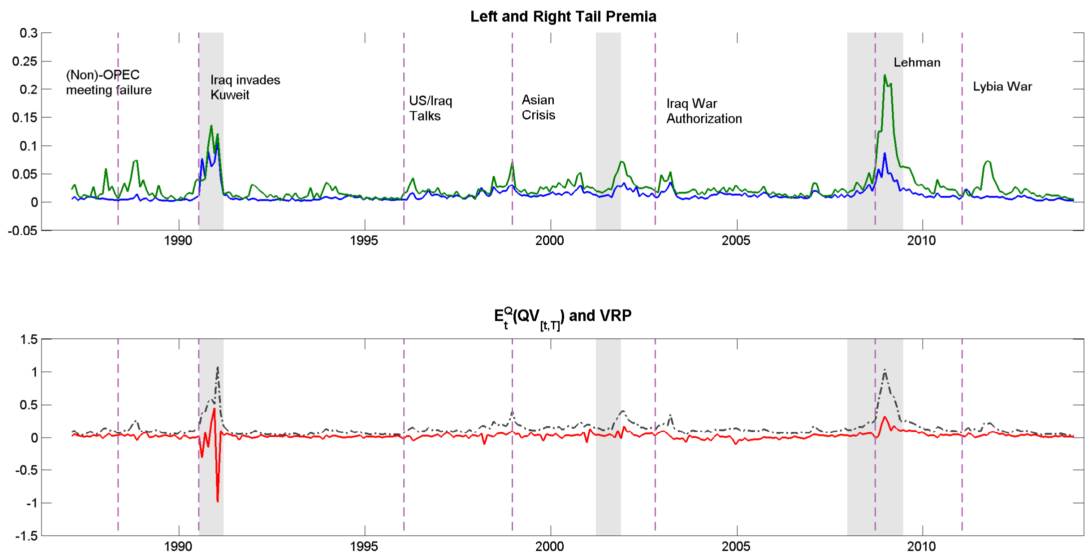

Figure 5: The figure on top displays the left tail (green) and right tail (blue) variation measures under $\mathbb{Q}$. The figure on the bottom displays the traditional measures $E_t^{\mathbb{Q}}(QV_{t,T})$ (dashed line) and $VRP_t \equiv E_t^{\mathbb{Q}}(QV_{t,T}) - RV_{t,T}$ (red line). All measures are presented in annualized form. The tail measures are computed for the threshold $k_t = 3\times$ BS implied volatility. Shaded areas represent NBER recessions.

Robustness to the choice of k_t

Table (6) displays the estimated tail risk premia for a larger thresholds, $k_t = 4 \times$ the at-the-money Black-Scholes implied volatility and $k_t = 6.8 \times$ the at-the-money Black-Scholes implied volatility.

	Obs.	Mean	S.d.
$k_t = 4 \times$ ATM Black-Scholes implied volatility			
$LJV^{\mathbb{Q}}_{t,T}$	324	0.0113	0.0130
$RJV^{\mathbb{Q}}_{t,T}$	324	0.0032	0.0041
$k_t = 6.8 \times$ ATM Black-Scholes implied volatility			
$LJV^{\mathbb{Q}}_{t,T}$	324	$1.060 \cdot 10^{-3}$	$1.467 \cdot 10^{-3}$
$RJV^{\mathbb{Q}}_{t,T}$	324	$0.188 \cdot 10^{-3}$	$0.334 \cdot 10^{-3}$

Table 6: Summary Statistics for $LJV^{\mathbb{Q}}_{t,T}$ and $RJV^{\mathbb{Q}}_{t,T}$ based on pooled monthly data; evaluated at $k_t = 4 \times$ and $k_t = 6.8 \times$ Black-Scholes implied volatility.

9 Appendix: Predictability of Futures and Spot Returns

Cross validation

	Futures				Spot			
	$r_{F,t+3}$	$r_{F,t+3}$	$r_{F,t+6}$	$r_{F,t+6}$	$r_{S,t+3}$	$r_{S,t+3}$	$r_{S,t+6}$	$r_{S,t+6}$
Controls	No	Yes	No	Yes	No	Yes	No	Yes
No change forecast	30.5	-	55.0	-	34.3	-	59.7	-
With $RJV^{\mathbb{Q}}_{t,T}$,$LJV^{\mathbb{Q}}_{t,T}$	29.1	26.7	48.9	41.0	33.0	33.6	50.8	57.1
Without $RJV^{\mathbb{Q}}_{t,T}$, $lJV^{\mathbb{Q}}_{t,T}$	29.8	28.1	51.8	46.7	33.9	35.8	52.2	65.2

Table 7: Cross validation statistics for forecasting models with and without $RJV^{\mathbb{Q}}_{t,T}$ and $LJV^{\mathbb{Q}}_{t,T}$. The dependent variable$r_{F,t+3}$ stands for 3 month futures returns, $r_{F,t+6}$ for six month futures returns, $r_{S,t+3}$ for three month spot returns and $r_{S,t+6}$ for six month spot returns. Each return is evaluated out of sample once, with a 7 (13 in the case of the 6 months prediction) out-of-sample window around the corresponding return. The control variables are those described in section 4. Each model is evaluated twice: Once including the the predictors $RJV^{\mathbb{Q}}_{t,T}$ and $RJV^{\mathbb{Q}}_{t,T}$, and once excluding them. A lower statistics indicates a lower out-of-sample MSPE. Values are multiplied by 100. Sample period is 1989 - 2013.

Forecasting empirical jump variations

Variables	(1) Emp. LJV	(2) Emp. LJV	(3) Emp. RJV	(4) Emp. RJV
L.$RV2$	0.00***	0.00***	0.00***	0.00**
	(0.001)	(0.001)	(0.001)	(0.001)
L.$RJV_{t,T}^{\mathbb{Q}}$		-0.01		0.02*
		(0.016)		(0.011)
L.$LJV_{t,T}^{\mathbb{Q}}$		0.01		-0.00
		(0.004)		(0.004)
Obs.	312	312	312	312
Adj. R^2	0.44	0.46	0.27	0.28

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 8: Forecasting results for empirical jumps. Emp. LJV stands for the sum of squared log returns from large negative jumps in the futures price, Emp. RJV stands for the sum of squared log returns from positive negative jumps. $RV2$ is the average monthly realized variance.

10 Appendix: Interaction with the Macroeconomy

Robustness to Aggregate Uncertainty and Fears

	$RJV_{t,T}^{\mathbb{Q}}$	$LJV_{t,T}^{\mathbb{Q}}$	$FI_{t,SPX}$	$VRP_{t,SPX}$
$RJV_{t,T}^{\mathbb{Q}}$	1.00			
$LJV_{t,T}^{\mathbb{Q}}$	0.85	1.00		
$FI_{t,SPX}$	0.60	0.65	1.00	
$VRP_{t,SPX}$	0.48	0.62	0.37	1.00

Table 9: Correlation between monthly oil risk measures and monthly stock market uncertainty measures. $FI_{t,SPX}$ is the fear index derived from S&P500 index options, $VRP_{t,SPX}$ the variance risk premium derived from S&P500 index options and futures as described in Bollerslev, Todorov, and Xu (2014). The sample period is 1996:1 to 2013:8.

Variables	(1) $r_{S,t+6}$	(2) $r_{S,t+6}$	(3) $r_{S,t+6}$	(4) $r_{S,t+6}$	(5) $r_{S,t+6}$	(6) $r_{F,t+6}$	(7) $r_{F,t+6}$	(8) $r_{F,t+6}$	(9) $r_{F,t+6}$	(10) $r_{F,t+6}$
$VRP_{t,SPX}$	0.03*		0.03		0.01	0.02		0.02		0.00
	(0.019)		(0.020)		(0.020)	(0.012)		(0.012)		(0.017)
$FI_{t,SPX}$,		0.09**	0.05		0.02		0.01	-0.01		-0.02
		(0.036)	(0.037)		(0.049)		(0.031)	(0.030)		(0.060)
$RJV_{t,T}^{\mathbb{Q}}$				-14.98*	-14.33*				-14.88**	-13.91*
				(7.902)	(7.729)				(7.125)	(7.253)
$LJV_{t,T}^{\mathbb{Q}}$				5.72***	4.97***				4.44***	4.35**
				(1.792)	(1.815)				(1.491)	(1.731)
Adj. R^2	0.05	0.03	0.05	0.14	0.13	0.01	0.00	0.01	0.06	0.06
Obs.	210	210	210	210	210	210	210	210	210	210

Newey-West Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 10: Forecasting results for six months oil futures and spot market returns. The dependent variable $r_{S,t+6}$ denotes the six months oil spot return, $r_{F,t+6}$ six months futures excess returns. $RJV_{t,T}^{\mathbb{Q}}$ and $LJV_{t,T}^{\mathbb{Q}}$ are the right tail oil variation measure and left tail oil variation measure. $FI_{t,SPX}$ is the fear index computed from S&P 500 as proxied through the left tail variation measure, suggested in Bollerslev, Todorov, and Xu 2014. $VRP_{t,SPX}$ is the variance risk premium computed from S&P 500 futures and options. The sample period is 1996 - 2013.

Forecasting stock market returns

Variables	(1) 1989 - 2013 $r_{Mkt,t+6}$	(2) 1989 - 07/2008 $r_{Mkt,t+6}$
L6.$RJV_{t,T}^{\mathbb{Q}}$	-1.50	5.78***
	(3.349)	(1.409)
L6.$LJV_{t,T}^{\mathbb{Q}}$	1.36	-3.84***
	(1.075)	(1.344)
Observations	294	229

Newey-West standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 11: Forecasting results for stock market returns. $r_{Mkt,t+6}$ is the six months market excess return using CRSP data.

Forecasting real activity

	(1)	(2)	(3)	(4)
Variables	$r_{IP,2}$	$r_{IP,2}$	$r_{IP,2}$	$r_{IP,2}$
L2.$RJV^{\mathbb{Q}}_{t,T}$		-0.26***	-0.27***	-0.28**
		(0.084)	(0.070)	(0.133)
l2.$r_{IP,2}$	0.35***	0.25***	0.25***	0.24***
	(0.057)	(0.071)	(0.066)	(0.073)
L4.$r_{IP,2}$	0.30***	0.22**	0.22**	0.23**
	(0.115)	(0.094)	(0.094)	(0.093)
L6.$r_{IP,2}$		-0.04	-0.04	-0.04
		(0.085)	(0.078)	(0.075)
L2.$r_{SPX,1}$				0.03***
				(0.010)
L2.$r_{WTI,1}$				0.00
				(0.006)
L2.OILVIX2				0.00
				(0.012)
L2.$LJV^{\mathbb{Q}}_{t,T}$		-0.01		
		(0.038)		
Observations	316	316	316	316
Adj. R^2	0.29	0.34	0.35	0.37

Newey-West standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 12: Forecasting results for industrial production growth. $r_{IP,2}$ is the two month growth rate of US industrial production, $r_{SPX,1}$ is the one month stock market return, $r_{WTI,1}$ is the one month increase in the crude oil spot price, OILVIX2 is the option implied oil market volatility.