

Ellwanger, Reinhard

Working Paper

On the tail risk premium in the oil market

Bank of Canada Staff Working Paper, No. 2017-46

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Ellwanger, Reinhard (2017) : On the tail risk premium in the oil market, Bank of Canada Staff Working Paper, No. 2017-46, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2017-46>

This Version is available at:

<https://hdl.handle.net/10419/197852>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Staff Working Paper/Document de travail du personnel 2017-46

On the Tail Risk Premium in the Oil Market



by Reinhard Ellwanger

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Bank of Canada Staff Working Paper 2017-46

November 2017

On the Tail Risk Premium in the Oil Market

by

Reinhard Ellwanger

International Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
rellwanger@bankofcanada.ca

Acknowledgements

I thank Peter Reinhard Hansen, Christiane Baumeister, Lutz Kilian, Marek Raczko, Lai Xu, Victor Todorov, Peter Christoffersen, Antoni Targa Barrera, Gabriela Galassi, Robert Goodhead, Ivan Shaliastovich (discussant at the NBER Economics of Commodity Markets Meeting), Hui Chen (discussant at the Midwest Finance Association Meeting) and Emil Siriwardane (Discussant at the NYU Stern Volatility Conference); the conference and workshop participants at the “Measuring and Modeling Financial Risk with High Frequency Data” conference at the EUI, the 10th BMRC-DEMS conference in London, the 2014 workshop at CREATES, the CFE conference 2015 in London, the MEG 2015 in St. Louis, the MFA conference 2016 in Atlanta, the Annual Volatility Conference at NYU Stern 2016, and the NBER Economics of Commodity Markets meeting 2016; and the seminar participants at the Bank of Norway, UC Davis, the Bank of Canada, Mannheim University, the Bank of England and the University of St. Andrews. This paper was previously circulated under the title “Driven by Fear? The Tail Risk Premium in the Crude Oil Futures Market.”

Abstract

This paper shows that changes in market participants' fear of rare events implied by crude oil options contribute to oil price volatility and oil return predictability. Using 25 years of historical data, we document economically large tail risk premia that vary substantially over time and significantly forecast crude oil futures and spot returns. Oil futures prices increase (decrease) in the presence of upside (downside) fears in order to allow for smaller (larger) returns thereafter. This increase (decrease) is amplified for the spot price because of time varying-benefits from holding physical oil inventories that work in the same direction. We also provide support for view that that time variation in the relative importance of oil demand and supply shocks is an important determinant of oil price fluctuations and their interaction with aggregate outcomes. However, the option-implied tail risk premia are not spanned by traditional macroeconomic and oil market uncertainty measures, suggesting that time-varying oil price fears are an additional source of oil price volatility and predictability.

Bank topics: Asset pricing; Econometric and statistical methods; Financial markets
JEL codes: C53, C58, D84, E44, G12, G13, Q43

Résumé

Cette étude montre que les variations de la peur d'événements extrêmes parmi les acteurs des marchés, inférées à partir des options sur le pétrole brut, contribuent à expliquer la volatilité des cours pétroliers et aident à prévoir les rendements des instruments du marché pétrolier. En nous fondant sur des données couvrant une période de 25 ans, nous trouvons que les primes élevées liées au risque extrême fluctuent sensiblement dans le temps et permettent de prévoir pour une large part les cours à terme et au comptant du pétrole brut. Les cours à terme du pétrole augmentent (diminuent), dans un contexte influencé par la peur d'une évolution haussière (baissière) des prix, pour aboutir à des rendements plus faibles (ou plus hauts). Dans le cas des cours au comptant, cette hausse (ou baisse) est amplifiée par les avantages (variables dans le temps) associés à la détention de stocks réels de pétrole qui évoluent dans le même sens. Nous corroborons l'idée que la variation temporelle, en ce qui concerne l'ampleur relative des chocs de demande et d'offre de pétrole, joue un rôle important dans les fluctuations des cours pétroliers et leurs interactions avec les rendements agrégés. Les primes de risque extrême tirées des options ne sont toutefois pas approximées par les indicateurs macroéconomiques ni par les mesures de l'incertitude sur le marché pétrolier, ce qui donne à penser que la peur, variable dans le temps, d'une évolution défavorable des prix du pétrole est une autre source de volatilité pour les cours pétroliers et d'information pour la prévision de ces cours.

Sujets : Évaluation des prix; Méthodes économétriques et statistiques; Marchés financiers
Codes JEL : C53, C58, D84, E44, G12, G13, Q43

1 Introduction

The large movements in oil prices and the influx of financial investors over the last decade have increased the interest in understanding these fluctuations and forecasting market developments. Derivatives markets constitute a rich source of information about expectations of market participants and their willingness to bear risks. In particular, there is now mounting evidence that the pricing of tail risk as manifested in equity-index options conveys important information on the dynamics of stock returns and equity premia (see, e.g., Bollerslev and Todorov 2011b; Bollerslev, Todorov, and Xu 2015; Andersen, Fusari, and Todorov 2016; Feunou, Jahan-Parvar, and Okou 2017). These studies document a sizeable premium for expected downside volatility—especially for the expected volatility arising from large negative jumps—that is a potent predictor of future excess returns. In a similar fashion, this paper explores option-implied oil market tail risks and their role in explaining oil price volatility, return predictability and the relationship between oil prices and aggregate outcomes. While, naturally, the compensation for large upward jumps tends to be negligible in equity-index markets, there might be a much more prominent role for the upside premium in the oil market, since high oil prices have frequently been associated with poor aggregate outcomes.¹ Our sample of historical oil options prices extends back to the late 1980s and comprises two global recessions, as well as several episodes of oil supply shortfalls, which makes the oil market a particularly interesting case to explore the role of tail premia.

The paper makes several contributions to the existing literature. First, we present novel estimates of oil market tail risk based on the left jump and right jump variation premia embedded in crude oil futures and options. Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and Xu (2015) 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 find that, for the oil market, these jump risk premia are economically large and vary significantly over well-documented episodes of oil supply and demand uncertainty.

Second, we show that the tail variation measures have predictive power for crude oil futures returns that is not contained in traditional oil price predictors: oil futures prices increase (decrease) in the presence of upside (downside) fears in order to allow for smaller (larger) returns thereafter. This finding corroborates existing evidence for return predictability through option-implied tail risk measures in equity markets (Bollerslev, Todorov, and Xu 2015; Andersen, Fusari, and Todorov 2016), and for a time-varying risk premium in the oil market (Baumeister and Kilian 2014). For storable commodities with a futures market, any

¹See, e.g., Hamilton (2009) and Gao et al. (2016).

predictability of *spot* returns depends on the interaction of the futures risk premium and the net value of holding physical oil inventories. We show that in the oil market, the price increase (decrease) associated with a higher upside (downside) tail risk premium and subsequent decline is amplified for the spot price because of time varying benefits from holding physical oil inventories that work in the same direction. Much of this predictability can be attributed to an oil price “fear” index, defined as the difference between the upside tail variation premium and downside tail variation premium. As in Bollerslev and Todorov (2011b), who argue that the fear index computed from stock market index options and futures represents time-varying risk aversion, rather than time-varying risk, we provide evidence that the oil market fear index is associated with the effective risk aversion of oil market participants, rather than expected downside or upside volatility. In contrast to previous studies that have documented structural breaks in oil market risk premia with the onset of the financialization of commodity markets in the mid-2000s, the dynamics and predictive power of this fear index appear to remain remarkable stable over time.²

Finally, to put our jump risk measures in perspective, we also investigate the link between the oil tail premia, aggregate uncertainty and stock market returns. While the oil fear index, i.e., the relative importance of upside and downside tail risk premia, is a strong predictor of oil prices, there is little evidence of a linear relationship with stock market returns. Instead, a relatively high upside tail risk premium seems to be associated with uncertainty about oil supply, while a relatively high downside tail risk premium seems to be associated with macroeconomic uncertainty: intuitively, both states represent adverse economic outcomes, leading to a U-shaped relationship between the oil fear index and stock market returns. Accordingly, higher oil prices are associated with lower stock market returns in the wake of a higher fear index, but with higher stock market returns in the wake of a lower fear index. While these results are broadly consistent with the general equilibrium frameworks of Hitzemann (2016) and Ready (2017), the oil fear measures appear not to be spanned by aggregate uncertainty or traditional measures of oil price uncertainty.³ Instead, these measures seem to provide more accurate measures of directional oil price fear that go beyond changes captured by fundamentals.

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

²See, e.g., Hamilton and Wu (2014) and Christoffersen and Pan (2017a).

³These results are also consistent with Gao et al. (2016), whose results suggest that oil price volatility predicts macroeconomic outcomes, but not oil spot returns. They are also supported by the recent findings of Christoffersen and Pan (2017b), who document a strong U-shape in the state price density on oil returns.

temporal variation in risk aversion (Bollerslev, Tauchen, and Zhou 2009; Bollerslev, Todorov, and Xu 2015; Andersen, Fusari, and Todorov 2016). 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 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 of 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 return dynamics.

The rest of this paper is structured as follows: the next section provides a formal definition of the tail risk variation measures, and an intuition for the relationship between the tail variation measures and oil return predictability; section 3 describes the empirical implementation, the data and the empirical properties of our estimates; section 4 presents the forecasting results; section 5 discusses the interaction between oil market and aggregate uncertainty; 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 in the oil market. This setup follows Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and Xu (2015); instead of considering the aggregate stock market, however, 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 jump diffusion process

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

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.⁴

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

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

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 risk 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

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

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

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

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

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

⁴The compensator $v_t^{\mathbb{P}}(dx)dt$ ensures that the jump measure $\tilde{\mu}(dt, dx)$ is a martingale.

and large negative jumps:

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

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 $[t, T]$:

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

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

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

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 define $RJV_{t,T}^{\mathbb{P}}(k_t)$ and $LJV_{t,T}^{\mathbb{P}}(k_t)$ as the predictable component of the quadratic variation arising through large positive and large negative jumps under the \mathbb{P} -measure

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

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

$$RJV_{t,T}^{\mathbb{Q}}(k_t) = \int_t^T \int_{x>k_t} x^2 v_s^{\mathbb{Q}}(dx) ds, \quad LJV_{t,T}^{\mathbb{Q}}(k_t) = \int_t^T \int_{x<-k_t} x^2 v_s^{\mathbb{Q}}(dx) ds. \quad (9)$$

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

$$RJVP(k_t) \equiv \frac{1}{T-t} (E_t^{\mathbb{P}}(RJV_{t,T}^{\mathbb{P}}(k_t)) - E_t^{\mathbb{Q}}(RJV_{t,T}^{\mathbb{Q}}(k_t))), \quad (10)$$

while the part due to large negative jumps is

$$LJVP(k_t) \equiv \frac{1}{T-t} (E_t^{\mathbb{P}}(LJV_{t,T}^{\mathbb{P}}(k_t)) - E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{Q}}(k_t))). \quad (11)$$

The results provided by Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and

Xu (2015) indicate that for equity index options, the difference between $LJVP(k_t)$ and $RJVP(k_t)$ is associated with investors' attitudes towards risks. Intuitively, this difference reflects the special compensation demanded by investors for bearing jump risk even if the investment opportunity set does not change over time, and thus seems to reflect market *fears*.

In this paper we define an oil market fear index in a similar fashion:

$$FI_t(k_t) \equiv LJVP(k_t) - RJVP(k_t). \quad (12)$$

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 $FI_t(k)$.

2.2 Tail Risk Premia and Oil Return Predictability

There exists both theoretical and empirical evidence that the pricing of downside jump risk in the aggregate equity market is associated with large risk compensations in terms of expected returns (Bollerslev, Todorov, and Xu 2015). In a similar fashion, we would expect a relatively high downside premium $LJVP(k_t)$ to be associated with higher returns of the underlying security, namely oil *futures*, and a relatively high upside premium $RJVP(k_t)$ with lower returns. It is less obvious, however, how these tail premia should be related to expected changes in the *spot* price of oil, given that even the price of relatively close-to-maturity futures can differ substantially from spot prices.⁵ In order to provide some intuition on this relationship, we rely on a common decomposition of the expected change in the spot price into a futures risk premium and the (net) convenience yield.

The first part of this decomposition is based on the 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

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

Equation (13) reflects the non-arbitrage condition that the price of a futures contract has

⁵Historically, WTI spot oil prices have differed by as much as 15% from the nearby futures price.

to be equal to the expected spot price discounted by the premium associated with holding the futures contract. Much of the existing evidence 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., Bessembinder 1992; Hamilton and Wu 2014; Baumeister and Kilian 2014).

The second part of the decomposition relies 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:

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

where $CY_{t,T}$ is the net (of storage costs and interest rate outlays) equilibrium convenience yield in annualized terms. Thus, equation (14) defines the net convenience yield implicitly as the negative of the futures price over the current spot price. The 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.⁶

Combining equations (13) and (14), we can relate expected (log) spot price returns to the futures risk premium and the convenience yield, up to a scale factor depending on the time to maturity, by

$$E_t(r_{t,T}) = frp_{t,T} - cy_{t,T}, \quad (15)$$

where $r_{t,T}$ denotes log spot price returns, and $frp_{t,T}$ and $cy_{t,T}$ stand for the (appropriately scaled) log futures risk premium and the log of the net convenience yield.

Importantly, equation (15) shows that *any* form of oil spot return predictability, in a mean-squared prediction sense, is determined by the interaction of the futures risk premium and the net value of holding physical oil inventories. Although we remain agnostic about the precise determinants of the net convenience yield, existing explanations suggest a strong link between oil price uncertainty, expectations and the convenience yield: Pindyck (2004) relates the convenience yield to oil price volatility; Milonas and Thomadakis (1997) interpret the convenience yield as a call option on a futures contract; Byun (2017) relates the convenience yield directly to expected spot prices by modelling inventories as an input to next period's refined oil production. In the empirical section we confirm this conjecture and provide empirical evidence for a strong positive relationship between the oil market fear index and the net convenience yield. Given that a higher (lower) fear index is naturally associated with lower (higher) futures returns, equation (15) suggests that the index should predict oil

⁶Such benefits can arise for a variety of reasons such as temporary stock-outs and associated price spikes or convex adjustment costs.

spot returns as well. Moreover, conditionally on a higher (lower) oil market fear index, the expected decrease (increase) in the spot price should be larger than the expected changes in the futures contract.

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 approximated by a portfolio of out-of-the-money (OTM) put and call options. For our calculation we follow the methodology for the Chicago Board Options Exchange Volatility Index (VIX).⁷

Our specification of the jump tails follows Bollerslev, Todorov, and Xu (2015). In particular, the jump distribution and intensity under the \mathbb{Q} -measure are based on the semi-parametric model

$$v_t^{\mathbb{Q}}(dx) = \left(\phi_t^+ \times e^{-\alpha_t^+ x} \mathbf{1}_{(x>0)} + \phi_t^- \times e^{-\alpha_t^- |x|} \mathbf{1}_{(x<0)} \right) dx. \quad (16)$$

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 ϕ^\pm and the shape parameters α^\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 $k \uparrow \infty (k \downarrow -\infty)$,

$$\frac{e^r O_{t,T}(K)}{(T-t)F_{t,T}} \approx \frac{\phi_t^\pm e^{k(1 \pm \alpha_t^\pm)}}{\alpha_t^\pm (\alpha_t^\pm \pm 1)}, \quad (17)$$

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

From equation (17) it follows that the ratio of two OTM options does not depend on ϕ_t^\pm , leading to the natural estimator suggested by Bollerslev and Todorov (2014):

$$\hat{\alpha}_t^\pm = \operatorname{argmin}_{\alpha^\pm} \frac{1}{N_t^\pm} \sum_{i=1}^{N_t^\pm} \left| \log \left(\frac{O_{t,\tau}(k_{t,i})}{O_{t,\tau}(k_{t,i-1})} \right) (k_{t,i} - k_{t,i-1})^{-1} - (1 \pm (-\alpha^\pm)) \right|, \quad (18)$$

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 $0 < |k_{t,1}| < \dots < |k_{t,N_t^\pm}|$. In

⁷See the white paper on the CBOE website for details regarding the VIX methodology.

practice we will pool options such that t refers to a given month, which implicitly assumes that α^\pm is approximately constant during this period.

For a given α^\pm , we then use equation (17) to estimate

$$\hat{\phi}_t^\pm = \operatorname{argmin}_{\phi^\pm} \frac{1}{N_t^\pm} \sum_{i=1}^{N_t^\pm} \left| \log \left(\frac{e^{r\tau} O_{t,\tau}(k_{t,i})}{(T-t)F_{t-,\tau}} \right) + (\pm \hat{\alpha}_t^\pm - 1)k_{t,i} + \log(\hat{\alpha}_t^\pm \mp 1) + \log(\hat{\alpha}_t^\pm) - \log(\phi^\pm) \right|. \quad (19)$$

From the definition of the tail risk premia in equations (10) and (11), and our assumptions for the large jumps dynamics in (16), it follows that for time to maturity $T-t$ and threshold k_t ,

$$RV_{t,T}^{\mathbb{Q}} = (T-t)\phi_t^+ e^{-\alpha_t^+ k_t} (\alpha_t^+ k_t (\alpha_t^+ k_t + 2) + 2) / (\alpha_t^+)^3 \quad \text{and} \quad (20)$$

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

The \mathbb{Q} tail measures are then computed by replacing the population quantities in (20) by their respective 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.⁸ 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 let RV_t denote the realized variation on day t , which is consistently estimated by summing the squared intraday returns:

$$RV_t \equiv \sum_{i=1}^n (\Delta_{t,i}f)^2 \xrightarrow{p} \int_{t+\pi_t}^{t+1} \sigma_s^2 ds + \int_{t+\pi_t}^{t+1} \int_{\mathbb{R}} x^2 \mu(ds, dx). \quad (22)$$

Realized jumps under the statistical measure are estimated using the threshold technique 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.⁹

Due to the lack of observations for sufficiently large jumps under the statistical measure,

⁸Although 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.

⁹The reader is referred to Appendix A for further details regarding the estimation procedure.

it is infeasible to estimate the same flexible jump tail specification as under the \mathbb{Q} -measure. Instead we follow Bollerslev and Todorov (2011b) in assuming that

$$v_t^{\mathbb{P}} = ((\alpha_0^- 1_{x<0} + \alpha_0^+ 1_{x>0}) + (\alpha_1^- 1_{x<0} + \alpha_1^+ 1_{x>0})\sigma_t^2) v^{\mathbb{P}}(x)dx, \quad (23)$$

where $v^{\mathbb{P}}(x)dx$ is the time-invariant distribution of empirical jumps, and the level parameters α_0^{\pm} and α_1^{\pm} relate linearly to the time-varying continuous volatility σ_t^2 .¹⁰

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 separately for each tail as outlined in Appendix A.

3.2 Data Description

Our empirical analysis is based on light sweet crude oil (West Texas Intermediate, short WTI) futures and options.¹¹ Crude oil derivatives are traded in very liquid markets and available historical data go 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 historical end-of-the-day settlement prices from November 1986 to December 2013.¹² WTI crude oil options are quoted for a variety of 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 estimators. 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 anomalies due to the lack of trading volume, which makes it necessary to discard these data and resort to the first back contract

¹⁰The 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 is presented in the following section. Note also that many existing models assume a proportional relationship between the jump intensity and volatility (see, e.g., Doran and Ronn 2008).

¹¹The CME tickers for the futures and options are CL and LO, respectively.

¹²NYMEX 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 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).

for those days.¹³ This final data set thus contains options with maturities between 9 and 40 days.

In order to mitigate potential influences from the diffusive risk, we retain only OTM call (put) options with 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 three cents and those violating the monotonicity condition in the strike dimension. For the estimation of the jump tail parameters we pool all clean deep OTM options for a given calendar month in order to obtain monthly observations.¹⁴

Our futures data comprise five-minute intraday transaction prices from Tickdatamarket from January 1988 to December 2013. After standard data-cleaning procedures our sample of intraday data comprise 6,332 trading days.¹⁵ For the computation of the realized measures we use the part of the day when trading was actively carried out throughout the sample. Thus the first price observation is at 10:00 (CST), and the last price observation at 14:30. This leaves us with 54 return observations for each trading day.

Finally, we match the futures contracts used for the estimation of the realized measures to the corresponding option contracts used for the jump tail estimation under the risk-neutral measures. To keep the subsequent analysis consistent, all descriptive statistics and forecasting results will be based on a sample period from January 1989 to December 2013.

3.3 Empirical Tail Risk Measures

The estimates for the tail variation measures under \mathbb{Q} implied by equation (20) require a choice of the threshold k_t that separates large from small jumps. Similar to Bollerslev, Todorov, and Xu (2015) 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 differ with economic conditions and corresponding market volatility. The main results are presented for the threshold k_t equal to three times the at-the-money Black-Scholes (BS) implied volatility, which corresponds to a median threshold of 25%. However, we also show that qualitatively similar results are obtained for alternative thresholds.

The estimates for FI_t for different choices of k_t are presented in figure (1). Most of the noticeable movements in the series correspond to well-known periods of oil price or aggregate uncertainty, such as the First Gulf War in 1990, the Arab Spring in 2011, the

¹³This “cleaning procedure” is standard; see, e.g., Trolle and Schwartz (2010) or Bollerslev and Todorov (2011b).

¹⁴The monthly pooling ensures a sufficient number of options for estimation throughout the sample. It also has the advantage that potential seasonalities related to the monthly expiry cycle are averaged out.

¹⁵Some days around Christmas, Thanksgiving and July 4th feature irregular trading hours and were discarded.

financial crisis in 2008 and NBER recessions. The graph also shows that the fear indexes derived from different cut-off levels for large jumps display very similar dynamics. Indeed, table (7) in Appendix B confirms that the correlation between the variation measures for different thresholds is almost perfect.

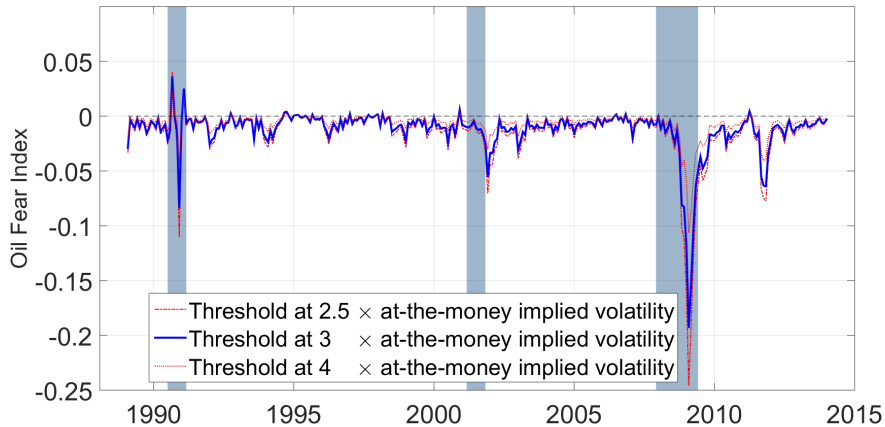


Figure 1: Oil market fear index, $FI_t \equiv RJV_{t,T}^{\mathbb{Q}} - LJV_{t,T}^{\mathbb{Q}}$, in annualized form. Different lines stand for different thresholds for the definition of large jumps in the computation of FI_t . Shaded areas represent NBER recessions. The sample period is 1989:1 to 2013:12.

Similar to results for 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} -measure. The estimates presented in table (1) indicate that, on average, the statistical variation measure for the left tail is more than 200 times smaller than its \mathbb{Q} -measure counterpart. The corresponding ratio for the right tail is larger than 100 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 the \mathbb{Q} -measure.¹⁶ Importantly, this implies that the tail variation premia appear well approximated by the risk-neutral variation measures only:¹⁷

$$RJVP(k_t) \approx -RJV_{t,T}^{\mathbb{Q}} \quad \text{and} \quad (24)$$

$$LJVP(k_t) \approx -LJV_{t,T}^{\mathbb{Q}}. \quad (25)$$

Interestingly, this indicates that changes in the objective jump distributions play a minor

¹⁶In fact, table (6) in Appendix B shows that while the average left and right tail variations under the \mathbb{Q} -measures decrease fast with the choice of k_t , they are still substantial for larger thresholds.

¹⁷A similar approximation empirically holds for the equity market. As noted by Bollerslev, Todorov, and Xu (2015), such an approximation conveniently allows one to discard the \mathbb{P} tail variation measures, which are inevitably estimated with less precision.

	Mean	SD	AR(1)
$LJV_{t,T}^{\mathbb{Q}}$	2.46	2.76	0.87
$RJV_{t,T}^{\mathbb{Q}}$	1.21	1.24	0.75
$LJV_{t,T}^{\mathbb{P}}$	<0.01	<0.01	0.55
$RJV_{t,T}^{\mathbb{P}}$	0.01	0.01	0.76
OILVIX2	13.73	13.08	0.85
VRP	1.33	11.13	0.06

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. OILVIX2 denotes the squared VIX calculated from WTI options, and VRP the variance risk premium calculated from WTI futures and options. The tail variation measures are presented in annualized percentage form. The tail variation measures are evaluated at a time-varying threshold of $k_t = 3 \times$ Black-Scholes implied volatility. The sample period is 1989:1 to 2013:12.

role in explaining the time variation in the size of the tail premia.¹⁸

The parameter estimates for the object variation measures, presented in table (5) in Appendix A, show that the larger estimates for the right tail variation can be attributed to a slower decay of the distribution in the tails, as governed by ξ . Instead, the jump intensities governing the time variation in these measures are similar for both measures. Together with strong time variation of the difference of $RJV_{t,T}^{\mathbb{Q}}$ and $LJV_{t,T}^{\mathbb{Q}}$ documented in figure (1), 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 are much larger than the premia for the right tail (table 1). This is consistent with a positive average futures risk premium and commodity futures markets being in normal “backwardation,” an idea first put forth by Keynes. 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.¹⁹

It is interesting to compare the left and right tail variation measures to the total variation measures. The implied total variation $QV_{t,T}^{\mathbb{Q}}$ is computed using all OTM options on a given day. The computation of the realized total variation also accounts for overnight returns. The contribution of the squared overnight returns to the entire daily observation, π_t , is about

¹⁸This is similar for the equity index market; see, e.g., Bollerslev and Todorov (2011a), Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and Xu (2015).

¹⁹Hitzemann (2016) provides an alternative explanation for a positive average oil futures risk premium based on macroeconomic long-run shocks.

50% on average.²⁰ We compute the daily series by an appropriate scaling of the intraday realized variance RV_t and obtain the monthly series by averaging over days.²¹ 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.²² Over our sample period, VRP_t is about -1.3%. This is lower than the estimate of Trolle and Schwartz (2010), who report an average premium of almost -3% over the 1996–2006 sample period, but similar to Pan and Kang (2013), who report a premium of -1.65% for a sample period slightly shorter than ours.

For the benchmark threshold of $k_t = 3 \times$ BS-implied volatility, the sum of the two tail premia is on average larger than 3% and thus substantially larger than the absolute value of the average VRP_t . In fact, $RJV_{t,T}^{\mathbb{Q}}$ by itself is of a similar magnitude as VRP_t ; and $LJV_{t,T}^{\mathbb{Q}}$ is even larger. 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 a positive premium on average.²³

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 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 risk premium due to large jumps is 3% over our sample period. In comparison, the average (continuously compounded) total futures risk premium computed from the first back contract is 7.8%, suggesting that one-fourth of the short maturity futures risk premium can be attributed to tail risk.

Empirically, we also find a strong and statistically significant correlation of about almost 0.4 between our measure of FI_t and the log of the short-term net convenience yield measure as implied by equation (14). Figure (2) shows that a higher fear index is systematically associated with a positive net convenience yield (or downward sloping futures curve), and

²⁰This 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.

²¹Since there is no consensus on the proper scaling method in the existing literature, we also experimented with different forms of 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.

²²This 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.

²³Table (6) in Appendix B shows that the magnitudes of $LJV_{t,T}^{\mathbb{Q}}$ and $RJV_{t,T}^{\mathbb{Q}}$ are strongly decreasing in the choice of the threshold, but they are still jointly larger than the average variance risk premium for a threshold of $k_t = 4 \times$ the at-the-money Black-Scholes implied volatility, corresponding to an average threshold of 35%.

a lower fear index with a negative net convenience yield (or upward sloping futures curve). The intuition provided by equation (15) thus suggests that the current spot price rises above (declines below) the expected spot price in the wake of large upside (downside) fears. We address this conjecture more carefully in the next section, where we discuss the forecasting properties of our indicators.




Figure 2: Average annualized net convenience yield for different quintiles of FI_t . The net convenience yield is measured as the negative of the price of the corresponding futures contract over the spot price. The averages are based on 60 observations determined by the respective quintile over the period from 1989:1 to 2013:12.

4 Predictability of Futures and Spot Returns

4.1 Predictive Regression Framework and Control Variables

This section presents new evidence of oil price predictability through the tail variation measures. If the premia capture oil market agents' attitudes towards 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

$$r_{j,t+i} = \beta_{0,i} + \beta_{1,i} \cdot LJV_{t,T}^{\mathbb{Q}} + \beta_{2,i} \cdot RJV_{t,T}^{\mathbb{Q}} + Controls_t' \cdot \beta_{3,i} + \epsilon_{t,i}, \quad j = \{S, F\}, \quad (26)$$

where $r_{S,t+i}$ is the i th-month-ahead spot return and $r_{F,t+i}$ is the i th-month-ahead futures return, while $LJV_{t,T}^{\mathbb{Q}}$ and $RJV_{t,T}^{\mathbb{Q}}$ represent the left and the right tail premia, respectively, approximated by the respective \mathbb{Q} -measure variations.²⁴ The vector $Controls_t$ includes both

²⁴Details on the construction of spot and futures returns can be found in Appendix C.

macroeconomic-financial variables specific to the crude oil market that are potential predictors of commodity spot and futures returns (see, e.g., Bessembinder and Chan 1992; Hong and Yogo 2012). We employ overlapping regressions in order to enhance the efficiency of our estimates, using robust Newey-West standard errors to account for the autocorrelation in the residuals induced by the overlap. The lag length chosen for the computation of the standard errors is twice the length of the overlap. The baseline forecasts are performed in-sample, while cross-validation techniques are employed for the out-of-sample robustness check.

We first describe the set of oil-market-specific control variables. As suggested by the recent literature, we include the estimated variance premium, VRP_t and the realized volatility. We compute VRP_t as the difference between the lagged expected variation under the risk-neutral and the actual variation measures in period t as suggested in Bollerslev, Tauchen, and Zhou (2009) and as described in the previous section. 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 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 U.S. Commodity Futures Trading Commission 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 sixth back contract over the current spot price.²⁵

We control for macroeconomic conditions by including the short-term interest-rate, computed as the yield of a three-month T-Bill, and the Aruoba-Diebold-Scotti Business Conditions Index published by the Federal Reserve Bank of Philadelphia. Since crude oil prices might be also driven by global rather than U.S.-specific factors, we include the Real Activity Index developed in Kilian (2009). The last control variable is the yield spread, computed as the difference between Moody’s Aaa and Baa corporate bond yields.

4.2 Forecasting Results

We first discuss the forecasting results for crude oil spot returns, presented in columns (1) to (4) of table (2). For all regressions, the coefficients have the expected sign. A relatively high right tail premium is associated with negative futures returns, and a high left tail premium with positive returns. For the model without controls, all coefficients are statistically significant 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 also robust to the inclusion of standard predictors of crude oil prices. The only exception is the noticeable

²⁵We also experimented with other definitions of the slope of the term structure, e.g. the net ratio of the spot price with the third back futures contract and the net ratio of the first two futures contracts, yielding almost identical results.

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 for the three-months-ahead prediction. This indicates a certain degree of correlation of $LJV_{t,T}^Q$ and some of the additional predictor variables added as controls. However, the tail premia are always jointly significant at the 1% significance level, as documented by F-tests of the null hypothesis that both $\beta_{1,i}$ and $\beta_{2,i}$ are zero. The adjusted R^2 is about 6% for the three-month horizon and 12% for the six-month horizon regression, which amounts to about one half of the in-sample predictability associated with all regressors.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$r_{S,t+3}$	$r_{S,t+3}$	$r_{S,t+6}$	$r_{S,t+6}$	$r_{F,t+3}$	$r_{F,t+3}$	$r_{F,t+6}$	$r_{F,t+6}$
$RJV_{t,T}^Q$	-5.3*** (1.031)	-7.2*** (1.348)	-8.1*** (1.709)	-12.3*** (1.777)	-3.8*** (1.124)	-7.4*** (1.482)	-5.6*** (1.786)	-12.6*** (1.792)
$LJV_{t,T}^Q$	2.1*** (0.550)	1.3 (0.937)	4.2*** (0.578)	2.2* (1.324)	1.1** (0.552)	0.9 (0.897)	2.5*** (0.649)	2.0 (1.282)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Adj. R2	0.062	0.152	0.118	0.274	0.030	0.120	0.046	0.230
Observations	300	300	300	300	300	300	300	300
Wald test (p-value)	<0.001	<0.001	<0.001	<0.001	0.002	<0.001	<0.001	<0.001

Newey-West standard errors in parentheses

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

Table 2: Forecasting results for $i = 3$ months and $i = 6$ months. The dependent variable $r_{S,t+i}$ stands for spot returns, and $r_{F,t+i}$ for futures returns. Wald test stands for an F-test of the joint significance of $RJV_{t,T}^Q$ and $LJV_{t,T}^Q$. The estimation period is 1989:1 to 2013:12.

We now describe the forecasting results for crude oil futures returns, presented in columns (5) to (8) of table (2). For all regressions, the coefficients associated with the tail premia have magnitudes similar to those of the futures regressions. Again, the significance of the coefficient for the left tail premium in the three-month regression when the other predictor variables are included is decreased due to the rise in the standard error relative to the model without predictors. For all other regressors, 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^2 s associated with the futures return regressions are somewhat smaller than those for spot returns. At the three-month horizon, the tail measures account for about 3% of the variability in the spot return, while this number increases to almost 5% at the six-month horizon, constituting about one-fourth of the in-sample variability explained by all regressors. Thus the tail risk measures have strong predictability for futures and spot price returns in the expected direction, and, judging from the magnitude of the coefficients and the fit of the regressions, this predictability is somewhat larger for the spot than for the futures price.

This evidence corroborates the idea that fears of 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.

Obviously, a high in-sample R^2 is not always indicative of a good forecasting model since adding additional regressors will mechanically increase the in-sample fit. In order to safeguard against in-sample overfitting, we perform a cross-validation exercise to show that the forecasting results hold also out-of-sample. Instead of choosing an arbitrary cut-off for a pseudo out-of-sample exercise, the 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.²⁶ This is repeated until each return observation is treated “out-of-sample” once. The cross-validation statistics are based on the mean squared prediction error; hence lower values indicate a better fit. In order to evaluate the out-of-sample performance of our predictors, we estimate the forecasting model presented in equation (26) for three different sets of regressors: the tail variation measures only, the tail variation measures and controls, and the control variables only.

The results are displayed in table (8) in Appendix C. Including the tail premia as predictive regressors reduces the cross-validation statistics substantially for all specifications, indicating that the forecasting power of the premia is not only due to in-sample properties. The out-of-sample mean squared prediction error (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 no-change forecast. In contrast, some of the models including control variables display a higher out-of-sample MSPE than the no-change forecast, suggesting that the R^2 related to these variables is driven at least in part by in-sample overfitting. The evidence is even stronger for the models that rely on the control variables only: these models perform worse than a no-change forecast out-of-sample.

The stability of the predictive relationship of $LJV_{t,T}^{\mathbb{Q}}$ and $RJV_{t,T}^{\mathbb{Q}}$ is remarkable, since our monthly estimates are based on non-overlapping data over a long time span in which the oil market has experienced profound changes. In contrast to previous studies that have documented structural breaks in oil market risk premia with the onset of the financialization of commodity markets in the mid-2000s, we do not find any evidence that the predictive power of the option-implied tail risk measures has changed over this period.²⁷

²⁶We ensure 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.

²⁷Such a break is documented by Hamilton and Wu (2014) and Christoffersen and Pan (2017a), among others.

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 of whether the premia reflect compensation for potential risks or rather describe oil market participants' attitudes towards risks. From the definition of FI_t , it follows that for negligible or approximately symmetric jump tails under the \mathbb{P} -measure,

$$\begin{aligned} FI_t &= \frac{1}{T-t} [(E_t^{\mathbb{P}}(LJV_{t,T}^{\mathbb{P}}) - E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{Q}})) - (E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{P}}) - E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{Q}}))] \\ &\approx \frac{1}{T-t} [E_t^{\mathbb{Q}}(LJV_{t,T}^{\mathbb{Q}}) - E_t^{\mathbb{Q}}(RJV_{t,T}^{\mathbb{Q}})]. \end{aligned} \quad (27)$$

Bollerslev, Todorov, and Xu (2015) point out that under these conditions, 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 for oil market fears. Further corroboration of the idea that the tail measures represent oil market participants' fear, rather than a rational prediction of future variations, is provided in table (9) presented in Appendix C. The table presents forecasting regressions for the relative importance of the realized upside variance vis-à-vis the realized downside variance. We consider several specifications—based on the realized variance [of all returns, jump returns, and large jump returns—for one- and three-months-ahead predictions. The regressions show no evidence that the asymmetry in the premia, captured by FI_t , is systematically related to the expected relative importance of upside or downside variations.

In contrast, the forecasting results shown in table (3) indicate that the return predictability through our novel predictors is mainly driven by its asymmetry. Here we use the six-month-ahead forecasting regression and decompose the tail measures into FI_t and a total component, $RJV_{t,T}^{\mathbb{Q}} + LJV_{t,T}^{\mathbb{Q}}$. 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 relative large upside fears are associated with decreasing spot price returns. The R^2 for this regression is almost 9%, while the R^2 for the regression with the total component of tail risk is only about 1% and yields a statistically insignificant coefficient. Taken together, these results indicate that time variation in the effective risk aversion of oil market participants—as measured by FI_t —is an important driver of the short- and medium-run fluctuations in the price of oil.

4.4 VAR Estimates of the Impact of the Oil Fear Index

In this section, we compare our forecasting results to estimates from VAR models, where we include FI_t as a directional measure of oil market uncertainty. The VAR framework has

Variables	(1)	(2)	(3)	(4)
	$r_{S,t+6}$	$r_{S,t+6}$	$r_{S,t+6}$	$r_{S,t+6}$
$RJV_{t,T}^Q$	-8.1*** (1.709)			
$LJV_{t,T}^Q$	4.2*** (0.578)			
FI_t		-3.4*** (0.614)		-6.2*** (1.023)
$RJV_{t,T}^Q + LJV_{t,T}^Q$			0.8 (0.910)	-1.9** (0.762)
Adj. R2	0.118	0.086	0.011	0.118
Obs	300	300	300	300

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

Table 3: Forecasting results for six-month oil spot returns, denoted by $r_{S,t+6}$. The sample period is 1989:1 to 2013:12.

been useful for structural analysis of oil price shocks and forecasting the real price of oil and has therefore become a benchmark model in the analysis of the oil market (see, among many others, Kilian 2009; Baumeister and Kilian 2012; Kilian and Murphy 2014). Building on the reduced-form version of the model proposed by Kilian and Murphy (2014), the variables included in the benchmark estimation are oil production (in percentage changes), the global real activity proposed by Kilian (2009), changes in crude oil inventories, FI_t , and the nominal price of oil (in percentage changes). To analyze the relationship between FI_t and the spot-futures spread, we also provide evidence from a model that includes the 12-month futures spread in addition to the other variables.

The identification of shocks to FI_t is based on a Cholesky decomposition of the reduced-form variance-covariance matrix with the following ordering: oil production, real economic activity, changes in crude oil inventories, the futures spread (when included) and prices.²⁸ Ordering production, real activity and inventories first ensures that the impact of other fundamentals is already controlled for when measuring the impact of shocks to FI_t on the oil and futures price.²⁹ Given the trade-off between overparameterization and allowing for sufficient lags to account for business cycle effects, we estimate the VAR with 12 lags on a monthly frequency from 1989:1 to 2013:12.

Figure (3) plots the impulse response function of oil spot prices, the six-month futures

²⁸Unlike Kilian and Murphy (2014), we are not interested in identifying all underlying structural shocks; rather, we focus on the effect of surprise changes in FI_t on oil prices and the futures spread. This exercise is similar to the one presented in Bloom (2009), who investigates the effect of stock market uncertainty on macroeconomic outcomes.

²⁹As pointed out by Kilian and Murphy (2014), inventories tend to 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.




Figure 3: Responses of the spot price of oil, the 12-month futures spread and the oil market fear index to a one-standard-deviation shock in the oil market fear index, FI_t . The estimates are based on a Cholesky decomposition with the following order: oil production, real economic activity, change in inventories, FI_t , futures spread, and oil spot returns (left and right panel); oil production, real economic activity, change in inventories, FI_t , futures spread, and oil spot (middle panel). Dashed lines represent 90% asymptotic confidence intervals. The sample period is 1989:1 to 2013:12.

spread and the oil fear index to a shock in the oil fear index. As expected, oil spot prices react with an instantaneous increase and a gradual reversion to previous levels over the following months. The 90% asymptotic confidence intervals are plotted around this, highlighting that this impact is statistically significant over the first months. The six-month futures spread declines on impact, while the reversion to previous levels seems somewhat faster. These results support the idea that a higher (lower) fear index is associated with lower (higher) returns in future periods, and an increase in the net convenience yield. The right-hand panel of figure (3) also shows that the shocks to the fear index revert quickly themselves, consistent with the idea that directional oil price uncertainty drives a temporary wedge between current and expected spot prices.

To gauge the economic significance of this effect, the left panel of figure (4) in Appendix C also depicts the reaction of oil prices to a one-standard-deviation shock to real economic activity. While these shocks tend to be more permanent, their immediate impact on oil prices seems to be smaller than that of oil fear shocks. Interestingly, the response of inventories to fear shocks, presented in the middle panel of figure (4), does not provide any evidence for a systematic reaction of inventories to FI_t . These results 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, oil market fears do not require an immediate response of inventory in order to display discernible effects on prices. Last, we also show the estimated effect of oil price volatility on oil price shocks, which turns out to be largely insignificant.³⁰ Taken together, these results corroborate the previous evidence for the importance of upside and downside fears in explaining short- and medium-term oil price fluctuations.

5 Interaction with the Macroeconomy

5.1 Oil Market and Aggregate Uncertainty

Thus far, we have treated the oil risk factors in isolation from the aggregate asset uncertainty. Bollerslev, Tauchen, and Zhou (2009) have shown that the variance risk premium embedded in stock market index options and futures is a suitable measure for degree of risk aversion in the market. This work has been extended by Bollerslev and Todorov (2011b) and Bollerslev, Todorov, and Xu (2015), 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. A natural question is therefore whether the measures of aggregate risk aversion, as measured by the variance risk premium and tail risk implied by index options and futures, already capture the predictive information embedded in the oil market fear measures.

Table (10) in Appendix D shows that the oil and stock market tail variation measures are indeed highly correlated. Both the left tail oil variation measure $LJV_{t,T}^Q$ and the right tail oil variation measure $RJV_{t,T}^Q$ exhibit strong correlations with the stock market variation measures, with correlations ranging between 48% and 66%.³¹ In table (11) in Appendix D we present the forecasting results for the spot price and six-month futures returns, respectively.³² The aggregate fear index and variance risk premium appear to contain some explanatory power for the six-month spot price return regression, with an R^2 of about 5%.³³ However,

³⁰The estimates are based on a model that has the same setup as the benchmark model, but replaces the oil fear index by the option-implied oil price variance described in the previous section. The results are consistent with Gao et al. (2016), who find little evidence for oil return predictability from oil price uncertainty.

³¹The 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 (2015) for the years 1996 to 2013. I would like to thank Lai Xu for providing the data on the S&P 500 fear index.

³²Since these measures are only available for a shorter time span, we did not include the aggregate fear measures $VRP_{t,SPX}$ and $FI_{t,SPX}$ in our previous forecasting exercises.

³³The predictability can be attributed mainly to $VRP_{t,SPX}$, a result that is consistent with the work of Bollerslev, Todorov, and Xu (2015), who show that $VRP_{t,SPX} - FI_{t,SPX}$ is mainly associated with economic uncertainty, while $FI_{t,SPX}$ captures attitudes towards risk.

this effect is completely dominated by the oil-specific tail risk variation measures, implying that the-oil specific measures already entail the information from the aggregate measures that is relevant for the oil market. Likewise, there is no evidence that the stock market tail variation measures predict oil futures returns. These results 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.

5.2 Equity Return Predictability and Equilibrium Interpretation

In addition to these results, we also address the question of 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, and in particular on the inclusion of data during and after the financial crisis. Table (12) 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, although these results are not statistically significant. One of the most notable events prior to the financial crisis in terms of oil fears was the 1990 Gulf War episode, which is clearly identified with supply risks (Alquist and Kilian 2010), along with a period of heightened demand uncertainty during the NBER recessions of the early 1990s and 2000s. Instead, after 2008, there is a clear dominance of downside uncertainty, which coincides with statistically and economically large coefficients suggesting the opposite relationship. 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.

In fact, it is well known that the macroeconomic impact of oil price shocks depends crucially on the underlying drivers of oil price fluctuations, and that high (low) oil prices could be the result of both high (low) oil demand and low (high) oil supply (Kilian 2009). As such, the empirical evidence and general equilibrium models of Hitzemann (2016) and Ready (2017) suggest that positive macroeconomic shocks are associated with both higher aggregate activity and higher oil prices, while the correlation between macroeconomic outcomes (and equity returns) and oil prices is negative after oil supply shocks.

There is also existing evidence that *uncertainty* about future oil supply and demand can have an impact on oil prices and the macroeconomy via the demand for speculative inventories. Indeed, Alquist and Kilian (2010) and Gao et al. (2016) show that oil supply uncertainty is associated with higher oil prices, a higher convenience yield and poor macroeconomic outcomes. Intuitively, macroeconomic uncertainty is associated with lower

oil prices, a lower convenience yield and poor macroeconomic outcomes. The predictability of oil futures and spot returns via the relative importance of the oil market tail variation measures makes these measures natural candidates for indicators of the relative importance of oil supply and macroeconomic uncertainty.

Industry	Low oil fear index			High oil fear index		
	$r_{t,1}$		R2	$r_{t,1}$		R2
Market	0.20**	(0.083)	0.14	-0.17***	(0.050)	0.17
NoDur	0.12*	(0.066)	0.08	-0.16***	(0.033)	0.21
Durbl	0.32*	(0.173)	0.11	-0.25***	(0.077)	0.2
Manuf	0.28**	(0.114)	0.15	-0.18***	(0.062)	0.16
Enrgy	0.32***	(0.068)	0.31	0.12*	(0.070)	0.07
Chems	0.15*	(0.084)	0.08	-0.19***	(0.059)	0.21
BusEq	0.21*	(0.111)	0.08	-0.29***	(0.062)	0.14
Telcm	0.17**	(0.083)	0.09	-0.10	(0.070)	0.05
Utils	0.14**	(0.057)	0.12	-0.06	(0.054)	0.03
Shops	0.1	(0.093)	0.04	-0.26***	(0.048)	0.32
Hlth	0.06	(0.079)	0.02	-0.15***	(0.041)	0.12
Money	0.21*	(0.113)	0.08	-0.26***	(0.053)	0.3
Other	0.23**	(0.101)	0.12	-0.21***	(0.058)	0.2

Robust standard errors in parentheses

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

Table 4: Univariate regressions of one-month stock market returns on one-month oil spot returns, conditional on the highest and the lowest quintile of the oil market fear index, FI_t . $r_{t,1}$ stands for one-month excess returns, “Low oil fear index” stands for the lowest quintile, and “High oil fear index” stands for the highest quintile. Regressions are based on the market and 12-Industry-Portfolio excess returns from Kenneth French’s website. The regressions are based on 60 observations determined by the respective quintile over the period from 1989:1 to 2013:12.

To further corroborate this idea, we also investigate the relationship between oil prices and equity returns *conditional* on the oil market fear index. Table (4) displays the regressions of one-month stock market returns on one-month oil spot returns, with separate regressions for both the lowest and the highest quintile of FI_t . As expected, higher oil prices are associated with lower stock market returns when upside oil price fears dominate downside oil price fears. On the contrary, higher oil prices are associated with higher stock market returns when downside oil price fears dominate. These effects hold for aggregate stock market returns, as well as for major industry portfolios; the only exception is the positive correlation

between the energy portfolio and oil price returns in the wake of larger upside oil price fears. Consistent with the idea that both very large and very small values of FI_t are associated with higher aggregate risk, the average six-month market excess return following a period of either large or small values of FI_t is 4.4%, versus only 2.3% in more tranquil periods.³⁴

While these results have an intuitive interpretation in existing general equilibrium frameworks, we have also shown that the option-implied tail risk measures dominate traditional predictors of oil spot and futures, and appear not to be spanned by aggregate uncertainty or oil-market-specific variables, such as oil price volatility or indicators of demand and supply. Instead, these measures seem to provide more accurate measures of directional oil price fear that go beyond changes captured by other fundamentals.

6 Conclusion

Oil prices are notoriously volatile. We find that the option-implied premia for large jump risks contain important information on market participants' fears that help to explain the large price fluctuations observed in the data. 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&P 500 option-implied tail risk. Instead, the oil uncertainty measures appear 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 increase (decrease) in the presence of upside (downside) fears in order to allow for smaller (larger) returns thereafter. Moreover, these tail risks extend to the physical market for inventories, as a high (low) fear index is associated with downward (upward) sloping futures curve and a high (low) convenience yield. We also provide support for the view that that time variation in the relative importance of oil demand and supply shocks is an important determinant of oil price fluctuations and their interaction with aggregate outcomes. However, the option-implied tail risk premia are not spanned by traditional macroeconomic and oil market uncertainty measures, suggesting that time-varying oil price fears are an additional source of oil price volatility. On a more general level, our results show 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.

³⁴The recent findings of Christoffersen and Pan (2017b), who document a strong U-shape in the state price density on oil returns, are also supportive of our results.

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.
- Andersen, Torben G, Nicola Fusari, and Viktor Todorov (2016). “The Pricing of Tail Risk and the Equity Premium: Evidence from International Option Markets.” *Working Paper*.
- 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*.
- 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 (2014). “Time-varying jump tails.” *Journal of Econometrics* 183.2, pp. 168–180.
- Bollerslev, Tim, Viktor Todorov, and Lai Xu (2015). “Tail risk premia and return predictability.” *Journal of Financial Economics* 118.1, pp. 113–134.
- Byun, Sung Je (2017). “Speculation in Commodity Futures Markets, Inventories and the Price of Crude Oil.” *The Energy Journal* 38.5.
- Carr, Peter and Liuren Wu (2009). “Variance risk premiums.” *Review of Financial Studies* 22.3, pp. 1311–1341.

- Christoffersen, Peter and Xuhui Nick Pan (2017a). “Oil volatility risk and expected stock returns.” *Journal of Banking & Finance*.
- Christoffersen, Peter and Xuhui Nick Pan (2017b). “The State Price Density Implied by Crude Oil Futures and Option Prices.” *Mimeo*.
- 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.
- 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.
- Feunou, Bruno, Mohammad R Jahan-Parvar, and Cedric Okou (2017). “Downside variance risk premium.” *Journal of Financial Econometrics*, nbx020.
- Gao, Lin, Steffen Hitzemann, Ivan Shaliastovich, and Lai Xu (2016). “Oil Volatility Risk.” *Mimeo*.
- Hamilton, James D (2009). “Causes and Consequences of the Oil Shock of 2007–08.” *Brookings Papers on Economic Activity* 1, pp. 215–261.
- 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.
- Hitzemann, Steffen (2016). “Macroeconomic fluctuations, oil supply shocks, and equilibrium oil futures prices.” *SSRN Working Paper*.
- 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.
- Kaldor, Nicholas (1939). “Speculation and economic stability.” *The Review of Economic Studies* 7.1, pp. 1–27.
- 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*.
- Pindyck, Robert S (2004). “Volatility and commodity price dynamics.” *Journal of Futures Markets* 24.11, pp. 1029–1047.
- Ready, Robert C (2017). “Oil prices and the stock market.” *Review of Finance*.

- 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.

Appendix A: Computation of the Jump Variation Measures Under the Statistical Measure

Estimating realized and expected jumps under the statistical measure

This section provides further details on the estimation of the jump properties under the statistical (objective) measure. The estimations are implemented for five-minute intraday oil futures returns and follow those proposed in Bollerslev and Todorov (2011a) and Bollerslev and Todorov (2011b). 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 $\alpha_{t,i}$

$$\sum_{i=1}^n (\Delta_{t,i}f)^2 \xrightarrow{p} \int_{t+\pi_t}^{t+1} \sigma_s^2 ds + \int_{t+\pi_t}^{t+1} \int_{\mathbb{R}} x^2 \mu(ds, dx) \quad \text{and} \quad (28)$$

$$\sum_{i=1}^n (\Delta_{t,i}f)^2 1_{|\Delta_{t,i}f| \leq \alpha_{t,i}} \xrightarrow{p} \int_{t+\pi_t}^{t+1} \sigma_s^2 ds \equiv CV_t. \quad (29)$$

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

$$TOD_j = \frac{\sum_{m=0}^N (\Delta_{m(\pi+n),i}f)^2 1_{|\Delta_{m(\pi+n),i}f| < \bar{\alpha}}}{\sum_{m=0}^N |\Delta_{m(\pi+n),i}f| < \bar{\alpha}} / \frac{\sum_{m=0}^N \sum_{j=1}^n (\Delta_{m(\pi+n),i}f)^2}{\sum_{m=0}^N \sum_{j=1}^n 1_{|\Delta_{m(\pi+n),i}f| < \bar{\alpha}}}, \quad (30)$$

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

$$\alpha_{j,t} = 3\left(\frac{1}{n}\right)^{0.49} \sqrt{CV_{t-n,t}} TOD_j. \quad (31)$$

The dynamics of our empirical jumps in equation (23) require an estimate of $v^{\mathbb{P}}$. This estimate is based on medium- and large-sized jump tails using the extreme value theory 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

$$\bar{v}_{\psi}^{\pm}(x) = \int_x^{\infty} v_{\psi}^{\pm}(u), \quad (32)$$

with $x > 0$ for $\bar{v}_{\psi}^+(x)$, and $x > 1$ for $\bar{v}_{\psi}^-(x)$. Under the assumption that \bar{v}_{ψ}^{\pm} belongs to the

domain of attraction of an extreme value distribution (see Bollerslev and Todorov 2011b), it follows that

$$1 - \frac{\bar{v}_\psi^\pm(u+x)}{\bar{v}_\psi^\pm} \sim G(u; \sigma^\pm, \xi^\pm), \quad u > 0, x > 0, \quad (33)$$

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

$$G(u; \sigma^\pm, \xi^\pm) = \begin{cases} 1 - (1 + \xi^\pm u / \sigma^\pm)^{-1/\xi^\pm}, & \xi^\pm \neq 0, \sigma^\pm > 0 \\ 1 - e^{-u/\sigma^\pm}, & \xi^\pm = 0, \sigma^\pm > 0. \end{cases} \quad (34)$$

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

$$\Theta \equiv [\sigma^\pm, \xi^\pm, \alpha_0^\pm \bar{v}_\psi^\pm(tr^\pm), \alpha_1^\pm \bar{v}_\psi^\pm(tr^\pm)], \quad (35)$$

which can be estimated using the GMM framework suggested by Bollerslev and Todorov (2011a). Denoting the scores associated with the log-likelihood function of the generalized Pareto distribution

$$\phi_1^\pm(u) = -\frac{1}{\sigma^\pm} + \frac{\xi^\pm}{(\sigma^\pm)^2} \left(1 + \frac{1}{\xi^\pm}\right) \left(1 + \frac{1 + \xi^\pm u}{\sigma^\pm}\right)^{-1} \quad \text{and} \quad (36)$$

$$\phi_2^\pm(u) = \frac{1}{(\xi^\pm)^2} \ln \left(1 + \frac{\xi^\pm u}{\sigma^\pm}\right) - \frac{u}{\sigma^\pm} \left(1 + \frac{1}{\xi^\pm}\right) \left(1 + \frac{\xi^\pm u}{\sigma^\pm}\right)^{-1}, \quad (37)$$

we use the moment conditions

$$\frac{1}{N} \sum_{t=1}^N \sum_{i=1}^n \phi_j^\pm(\psi^\pm(\Delta_{t,if}) - tr^\pm) 1_{\psi^\pm(\Delta_{t,if}) > tr^\pm} = 0, \quad j = 1, 2 \quad (38)$$

to estimate parameters governing the left and right tail distributions; and the conditions

$$\frac{1}{N} \sum_{t=1}^N \sum_{i=1}^n 1_{\psi^\pm(\Delta_{t,if}) > tr^\pm} - \alpha_0^\pm \bar{v}_\psi^\pm - \alpha_1^\pm \bar{\psi}^\pm(tr^\pm) CV_t = 0, \quad (39)$$

$$\frac{1}{N} \sum_{t=2}^N \left(\sum_{i=1}^n 1_{\psi^\pm(\Delta_{t,if}) > tr^\pm} - \alpha_0^\pm \bar{v}_\psi^\pm(tr^\pm) - \alpha_1^\pm \bar{v}_\psi^\pm(tr^\pm) CV_t \right) CV_{t-1} = 0, \quad (40)$$

to estimate the time-variation in the jump intensity.

Parameter estimates

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 $\pm 1\%$. In total, we detect 3,610 (4,363) positive (negative) jumps, out of which 256 (347) are above the threshold. The parameter estimates for our specification of the statistical large jumps' dynamics are presented in table (5).

Left Tail			Right Tail		
Parameter	Estimate	(St. Error)	Parameter	Estimate	(St. Error)
ξ^-	0.1670	(0.0688)	ξ^+	0.3069	(0.0735)
$100 \cdot \sigma^-$	0.3417	(0.0297)	$100 \cdot \sigma^+$	0.3002	(0.0284)
α_0^-	-0.1565	(0.0093)	α_0^+	-0.1305	(0.0080)
α_1^-	16.661	(0.6767)	α_1^+	13.432	(0.5809)

Table 5: Estimates for \mathbb{P} -tail parameters based on five-minute intraday oil futures data. The sample period is 1989:1 to 2013:12.

Appendix B: Robustness to the Choice of the Threshold

Tail risk premia for different threshold levels k_t

	$k_t = 2.5 \times \text{ATM-BS vola.}$	$k_t = 3 \times \text{ATM-BS vola.}$	$k_t = 4 \times \text{ATM-BS vola.}$	$k_t = 6 \times \text{ATM-BS vola.}$
$LJV_{t,T}^{\mathbb{Q}}$	3.48 (3.81)	2.46 (2.76)	1.14 (1.37)	0.21 (0.28)
$RJV_{t,T}^{\mathbb{Q}}$	1.89 (1.84)	1.21 (1.24)	0.47 (0.52)	0.06 (0.08)

Table 6: Summary statistics for $LJV_{t,T}^{\mathbb{Q}}$ and $RJV_{t,T}^{\mathbb{Q}}$ for different choices of the time-varying threshold k_t . The tail variation measures are expressed in annualized percentage terms. Standard deviations are in parentheses. *ATM-BS vola.* stands for the at-the-money Black-Scholes implied volatility. The sample period is 1989:1 to 2013:12.

Correlation of tail risk premia for different thresholds

	$LJV_{t,T}^{\mathbb{Q},3}$		$RJV_{t,T}^{\mathbb{Q},3}$		FI_t^3
$LJV_{t,T}^{\mathbb{Q},2.5}$	0.99	$RJV_{t,T}^{\mathbb{Q},2.5}$	0.99	$FI_t^{2.5}$	0.99
$LJV_{t,T}^{\mathbb{Q},3}$	1	$RJV_{t,T}^{\mathbb{Q},3}$	1	FI_t^3	1
$LJV_{t,T}^{\mathbb{Q},4}$	0.99	$RJV_{t,T}^{\mathbb{Q},4}$	0.99	FI_t^4	0.99
$LJV_{t,T}^{\mathbb{Q},6}$	0.94	$RJV_{t,T}^{\mathbb{Q},6}$	0.93	FI_t^6	0.97

Table 7: Correlations among option-implied left jump variation measures ($LJV_{t,T}^{\mathbb{Q}}$), right jump variation measures ($RJV_{t,T}^{\mathbb{Q}}$) and oil fear index (FI_t) for different choices of k_t . Superscripts indicate the multiple of the at-the-money Black-Scholes implied volatility that is used to compute the time-varying threshold used to classify large jumps. The sample period is 1989:1 to 2013:12.

Appendix C: Predictability of Futures and Spot Returns

Construction of oil spot and futures returns

Oil spot returns $r_{S,t+i}$ are constructed from daily WTI spot prices provided by the U.S. Energy Information Administration as the log ratio of the last closing price in month $t+i$ over the last closing price in month i .³⁵

Oil futures returns are based on the settlement price of WTI futures on the last trading day of each month, obtained from Bloomberg. Denoting the price of the first (nearby) contract observed in period t by F_t^1 , and the price of the i th contract by F_t^i , log futures returns are constructed from a single buy-sell strategy of the same contract:

$$r_{F,t+i} = \ln \frac{F_{t+i}^1}{F_t^{i+1}}. \quad (41)$$

Other studies also consider multi-horizon oil futures returns from the total return of a roll strategy where the next out contract is bought at F_t^2 and sold as the front contract at price F_{t+1}^1 :

$$r_{F,t+i} = \sum_{j=1}^i \ln \frac{F_{t+j}^1}{F_{t+j-1}^2}. \quad (42)$$

Although not reported in this paper, constructing returns from a continuous roll strategy yields very similar forecasting results.

³⁵The price data are published at https://www.eia.gov/dnav/pet/pet_pri_spt_s1_d.htm.

Cross validation

Regressors	Spot Returns		Futures Returns	
	$r_{S,t+3}$	$r_{S,t+6}$	$r_{F,t+3}$	$r_{F,t+6}$
$RJV_{t,T}^Q, LJV_{t,T}^Q$ only	0.96	0.86	0.98	0.92
$RJV_{t,T}^Q, JV_{t,T}^Q$ & controls	1.03	0.93	1.04	0.95
controls only	1.06	1.12	1.09	1.15

Table 8: Cross-validation statistics for forecasting models with and without control variables. The dependent variable $r_{S,t+3}$ stands for three-month spot returns, $r_{S,t+6}$ for six-month spot returns, $r_{F,t+3}$ for three-month futures returns and $r_{F,t+6}$ for six-month futures returns. Each return is evaluated out of sample once, with a seven-month (13-month in the case of the six-months-ahead prediction) out-of-sample window around the corresponding return. The control variables are those described in section 4. Values show the out-of-sample MSPE relative to the no-change forecast. The sample period is 1989:1 to 2013:12.

Forecasting empirical jump variations

	(1)	(2)	(3)	(4)	(5)	(6)
	$(Up/Down)_t^{all}$	$(Up/Down)_{t-2,t}^{all}$	$(Up/Down)_t^{jumps}$	$(Up/Down)_{t-2,t}^{jumps}$	$(Up/Down)_t^{large\ jumps}$	$(Up/Down)_{t-2,t}^{large\ jumps}$
FI_{t-1}	-0.01 (0.201)		-0.16 (0.115)		-0.11 (0.123)	
FI_{t-3}		-0.50 (0.503)		-0.66 (0.442)		-0.59 (0.462)
Obs.	300	300	300	300	300	300
Adj. R2	0.00	0.01	0.00	0.01	0.00	0.01

Newey-West standard errors in parentheses

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

Table 9: Forecasting results for the ratio of upside volatility to downside volatility. FI_t is the oil market fear measure defined in section 3. $(Up/Down)_t^{all}$ and $(Up/Down)_{t-2,t}^{all}$ stand for the ratio of the realized upside variance relative to the realized downside variance in months t , and in the three months from $t - 2$ to t , respectively. $(Up/Down)_t^{jumps}$ and $(Up/Down)_t^{large\ jumps}$ stand for the ratio of the realized upside variance to the realized downside variance computed only from returns classified as jumps or large jumps, respectively (see section 3 for details the different classifications of jump and large jump returns). The sample period is 1989:1 to 2013:12.

Additional evidence from VAR estimations




Figure 4: Selected impulse responses for different models. The left panel depicts the reaction of oil spot returns to a one-standard-deviation shock in real economic activity; the middle panel depicts the reaction of inventories to a one-standard-deviation shock to the oil fear index; the right panel depicts the reaction of oil spot returns to a one-standard-deviation shock of the option-implied oil price variance. The left panel is based on a Cholesky-decomposition with the following order: oil production, real economic activity, change in inventories, FI_t , futures spread, and oil spot returns. The middle panel is based on a Cholesky decomposition with the following order: oil production, real economic activity, FI_t , change in inventories, futures spread, and oil spot returns. The right panel is based on a Cholesky decomposition with the following order: oil production, real economic activity, change in inventories, option-implied oil price variance, futures spread, and oil spot returns. Dashed lines represent 90% asymptotic confidence intervals. The sample period is 1989:1 to 2013:12.

Appendix D: Interaction with the Macroeconomy

Relationship with aggregate uncertainty and oil fear

	$RJV_{t,T}^Q$	$LJV_{t,T}^Q$	$FI_{t,SPX}$	$VRP_{t,SPX}$
$RJV_{t,T}^Q$	1.00			
$LJV_{t,T}^Q$	0.80	1.00		
$FI_{t,SPX}$	0.53	0.66	1.00	
$VRP_{t,SPX}$	0.48	0.62	0.37	1.00

Table 10: Correlation between monthly oil risk measures and monthly stock market uncertainty measures. $FI_{t,SPX}$ is the fear index derived from S&P 500 index options, $VRP_{t,SPX}$ the variance risk premium derived from S&P 500 index options and futures as described in Bollerslev, Todorov, and Xu (2015). The sample period is 1996:1 to 2013:8.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	$r_{S,t+6}$	$r_{S,t+6}$	$r_{S,t+6}$	$r_{S,t+6}$	$r_{S,t+6}$	$r_{F,t+6}$	$r_{F,t+6}$	$r_{F,t+6}$	$r_{F,t+6}$	$r_{F,t+6}$
$VRP_{t,SPX}$	3.13 (1.899)		2.64 (2.008)		1.39 (2.389)	1.41 (1.211)		1.56 (1.258)		0.80 (1.981)
$FI_{t,SPX}$		8.49** (3.640)	4.97 (3.730)		0.53 (4.739)		0.58 (3.089)	-1.50 (2.919)		-4.03 (5.387)
$RJV_{t,T}^Q$				-7.94* (4.420)	-7.72* (4.233)					-6.68 (4.315)
$LJV_{t,T}^Q$				4.09*** (0.983)	3.47** (1.591)					2.44** (0.943)
Adj. R2	0.049	0.025	0.053	0.095	0.092	0.006	-0.005	0.002	0.025	0.022
Obs.	206	206	206	206	206	206	206	206	206	206
Wald test 1 (p-value)	-	-	0.02	-	0.84	-	-	0.45	-	0.33
Wald test 2 (p-value)	-	-	-	<0.001	0.04	-	-	-	0.02	0.17

Newey-West standard errors in parentheses

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

Table 11: Forecasting results for six-month oil futures and spot market returns. The dependent variable $r_{S,t+6}$ denotes the six-month oil spot return, $r_{F,t+6}$ the six-month futures excess returns. $RJV_{t,T}^Q$ and $LJV_{t,T}^Q$ 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 index options, proxied through the left tail variation measure as in Bollerslev, Todorov and Xu (2015). $VRP_{t,SPX}$ is the variance risk premium computed from S&P 500 futures and options. Wald test 1 stands for a test of the joint significance of $VRP_{t,SPX}$ and $FI_{t,SPX}$; and Wald test 2 for a test of the joint significance of $RJV_{t,T}^Q$ and $LJV_{t,T}^Q$. The sample period is 1996:1 to 2013:8.

Forecasting aggregate stock market returns

	(1)	(2)	(3)
Sample period	1989-2013	1989-2007	2008-2013
	$r_{Mkt,t+6}$	$r_{Mkt,t+6}$	$r_{Mkt,t+6}$
$RJV_{t,T}^Q$	-1.54 (2.297)	1.67 (1.135)	-13.42*** (4.184)
$LJV_{t,T}^Q$	0.97 (0.623)	-1.08 (1.027)	4.28*** (1.078)
Adj. R2	0.02	0.01	0.35
Observations	300	228	73

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

Table 12: Forecasting results for stock market returns for different sample periods. $r_{Mkt,t+6}$ is the six-month market excess return based CRSP data from Kenneth French's website.