

Speculative bubbles in recent oil price dynamics: Evidence from a Bayesian Markov-switching state-space approach

Marc Lammerding, Patrick Stephan, Mark Trede and Bernd Wilfling[†]

23/2012

[†] Department of Economics, University of Münster, Germany

Speculative bubbles in recent oil price dynamics: Evidence from a Bayesian Markov-switching state-space approach

Marc Lammerding, Patrick Stephan, Mark Trede, and Bernd Wilfling*

*Westfälische Wilhelms-Universität Münster, Department of Economics,
Am Stadtgraben 9, 48143 Münster, Germany*

(This version: 15 June 2012)

Abstract

Motivated by repeated spikes and crashes during previous decades we investigate whether the heavily financialized market for crude oil has been driven by speculative bubbles. In our theoretical modeling we draw on the convenience yield approach in order to approximate the fundamental value of the oil price. We separate the oil price fundamental from the bubble component by expressing a standard present-value oil price model in state-space form. We then introduce two Markov-regimes into the state-space representation in order to distinguish between two distinct phases in the bubble process, namely one in which the oil price bubble is a stable process and one in which the bubble explodes. We estimate the entire Markov-switching state-space specification using an econometrically robust Bayesian Markov-Chain-Monte-Carlo (MCMC) methodology. Based on inferential techniques designed for statistically separating both Markov-regimes in the bubble process from each other, we find robust evidence for the existence of speculative bubbles in recent oil price dynamics.

JEL classification: G10, G12, Q40

Keywords: Speculative bubbles; oil price; Markov-switching model; state-space model; Bayesian econometrics

* Corresponding author. Tel.: +49 251 83 25040, fax: +49 251 83 25042. E-mail addresses: marc.lammerding@wiwi.uni-muenster.de (M. Lammerding), patrick.stephan@wiwi.uni-muenster.de (P. Stephan), mark.trede@wiwi.uni-muenster.de (M. Trede), and bernd.wilfling@wiwi.uni-muenster.de (B. Wilfling).

1 Introduction

Between 2001 and mid-2008 the nominal spot price of *West Texas Intermediate* (WTI) crude oil skyrocketed from a level of 20 US-\$/bbl to an all-time high of 147 US-\$/bbl, then collapsed to a low of 30 US-\$/bbl in late 2008, and finally rebounded to a level of 100 US-\$/bbl in late 2011.¹ The economic literature divides the fundamentally justified explanations of these erratic price movements into three groups: (1) oil supply shocks, (2) oil demand shocks driven by global economic activity, and (3) oil-specific demand shocks. An oil supply shock typically results from production disruptions caused by a lack of investments and geopolitical tensions in oil-exporting regions such as the Middle East. By contrast, a general oil demand shock may occur due to unexpectedly strong economic growth in emerging economies, such as China and India, or the surprisingly rapid economic recovery in some countries after the recent global financial crisis. Finally, an oil-specific demand shock may be triggered by the time-varying importance of oil among alternative energy sources and changing expectations of oil fundamentals. Overall, there is now consensus in the literature that oil demand shocks caused by global economic activity were the main drivers of the oil price boom until mid-2008 (e.g. Hamilton, 2009a,b; Kilian, 2009; Kilian and Hicks, 2009; Kilian and Murphy, 2010).

In contrast to these stabilizing fundamentally justified shocks, oil price changes due to financial shocks may lead to speculative bubbles, and are thus considered as destabilizing (Lombardi and Van Robays, 2011). The oil bubble hypothesis is primarily based on the increasing financialization of oil futures markets reflected by the sharp rise in speculative open interest and speculative market shares.² In addition, speculators held

¹WTI, also known as *Texas Sweet Light*, is the reference type of oil in the United States. Similar price dynamics also hold for other oil types like *North Sea Brent*.

²A partial explanation of this phenomenon is that commodities, and oil in particular, have become a new asset class (besides stocks, bonds, and real estate). Turmoil in housing markets worldwide in conjunction with the recent global financial crisis has driven even more investors into commodities. Apart from that, commodity traders and hedge funds were more and more joined by pension funds and commodity index funds during the oil price surge until mid-2008 (Kesicki, 2010).

net long positions during almost the entire oil price surge until mid-2008 betting on further rising prices (CFTC, 2011a). Other anecdotal indicators of a possible oil price bubble until mid-2008 are (1) the combination of low US interest rates and futures prices in contango up to mid-2007 (Garcia, 2006; Tokic, 2010), (2) the large deviation of indexed oil prices from the indexed market valuation of major oil companies (Khan, 2009), and (3) the disappearance of a solid long-run relationship between the oil price and international stock markets during recent years (Miller and Ratti, 2009).

In contrast to this anecdotal evidence, empirical investigations are far less conclusive. On the one hand, there is evidence that the positive shape of the oil futures curve decisively contributed to the rise of the real oil price until late 2006 (Kaufmann et al., 2008). Moreover, destabilizing positive feedback traders are blamed for having subdued the stabilizing influence of fundamentalists, what, in conjunction with weak mean reversion, might have caused the substantial overshooting of the oil price (Cifarelli and Paladino, 2010; Reitz and Slopek, 2009; Tokic, 2011). However, pertaining to the price discovery mechanism, oil futures markets do not unambiguously lead spot markets but are rather interrelated with them, implying that the destabilizing impact of financial shocks may be less severe for entities interested in physical oil (Bekiros and Diks, 2008; Kaufmann and Ullman, 2009; Schwarz and Szakmary, 1994; Silvapulle and Moosa, 1999). Apart from that, speculative activity does not precede oil price movements but rather responds to them, characterizing speculators as trend-extrapolating chartists rather than as key price drivers (CFTC, 2008; IMF, 2005, 2006; Sanders et al., 2004). And finally, despite increasing speculation and betting on rising prices, the oil market was in backwardation between mid-2007 and mid-2008, and oil inventories did not increase substantially, a precondition that would have had to be met for speculators to exert persistent influence on the oil price (Kesicki, 2010).³

³As suggested by Hamilton (2009a), it is possible that speculative trading increases prices even without any inventory accumulation given that the short-run price elasticity of oil (and gasoline) demand is zero. However, Kilian and Murphy (2010) state that this was by far not the case during the oil price surge until mid-2008.

In view of this ambiguity, this paper tries to provide solid statistical inference about the existence of oil price bubbles. Until now, econometric testing for speculative bubbles has mainly been focussing on (US) stock markets. Gürkaynak (2008) provides a recent survey of econometric methods used for detecting asset-price bubbles. This survey includes the well-known variance bounds tests, West's two-step tests, (co)integration-based tests as well as the concept of intrinsic bubbles, and methods treating bubbles as an unobserved variable. By contrast, up to date little effort has been made to identify speculative bubbles in markets for commodities in general and for oil in particular, a shortfall regrettable for various reasons. Since oil is an important input factor of many products, increasing oil prices may cause recessions, bearish stock markets, and inflationary pressure (Kilian, 2008). As a result, overshooting oil prices may lead central banks to erroneously adjust monetary policy.⁴ And finally, misaligned oil prices may expose market participants to large financial losses once the bubble bursts.

In order to evaluate the oil bubble hypothesis more accurately, we draw on the standard present-value model for stocks, but adapt it to the oil market (Pindyck, 1993). In this context, the fundamental value of oil is defined as the sum of discounted oil dividends, which in turn are approximated by the benefits the holder of the physical commodity experiences in contrast to the owner of a futures contract written on the respective asset. These benefits that inventories provide, including the ability to smooth production, avoid stockouts, and facilitate the scheduling of production and sale, are termed *convenience yield*, and bear a particular meaning in the context of agricultural and energy commodities.⁵ Using the relationship between oil prices and

⁴The ECB (2008), for example, justifies its widely criticized decision to reinforce its restrictive monetary stance in mid-2008 by stating: "At its meeting on 3 July 2008, the Governing Council of the ECB decided (...) to raise the minimum bid rate on the main refinancing operations of the Eurosystem by 25 basis points to 4.25%. (...) The Governing Council's decision was taken (...) to counteract the increasing upside risks to price stability over the medium term. (...) These risks include notably the possibility of further increases in energy and food prices."

⁵Apart from the convenience yield approach, only a few (but less elaborated) alternative methods have been put forward to determine the fundamental value of commodities. Reitz and Slopek (2009) and Reitz et al. (2009), for example, use Chinese oil imports to obtain a rough proxy for the fundamentally justified oil price.

oil dividends, we then follow Wu (1995, 1997), and establish a state-space framework from which we extract the bubble component as an unobservable variable. In line with Al-Anaswah and Wilfling (2011), we additionally assume the bubble to evolve over time as a two-state Markov-switching process. By this econometric specification, we aim at separating two distinct bubble states from each other, namely one in which the bubble evolves over time as a stable process and one in which the bubble exhibits explosive dynamics. In order to obtain robust estimation results, we resort to a Bayesian approach and implement a fully-fledged Markov-Chain-Monte-Carlo (MCMC) estimation framework.

In our empirical analysis we find convincing evidence of at least two bubble periods in the oil market, namely (1) between the end of 2004 and mid-2008 and, after a strong correction of oil prices, (2) between mid-2008 and April 2011. Given the relevance of the oil price for the real economy, stock markets, and monetary policy, our results may indicate the need for more efficient regulation of financial investors in oil-derivative markets.

The remainder of this paper is organized as follows. Section 2 presents our convenience yield approach, and briefly reviews the standard present-value model which we transform into a state-space representation and enrich by a Markov-switching specification. Section 3 establishes our MCMC estimation framework. Section 4 describes the data, and presents our empirical results. Section 5 offers some policy implications and ideas for future research.

2 Economic modeling and econometric specification

2.1 The convenience yield and the present-value model

Following Pindyck (1993), we use futures prices to measure the convenience yield of actively traded storable commodities drawing on the so-called cost-of-carry equation.

In the absence of arbitrage, the (capitalized) flow of convenience yield net of storage costs from date t to T per unit of commodity, $d_{t,T}$, is given by

$$d_{t,T} = P_t(1 + r_t(T - t)/365) - F_{t,T}, \quad (1)$$

where P_t denotes the spot price, r_t is the (annualized) risk-free interest rate, and $F_{t,T}$ is the futures price for delivery on date T .⁶ Dividing $d_{t,T}$ by the time to delivery then leads to the standardized convenience yield d_t . Eq. (1) states that in equilibrium the futures price must equal the spot price (adjusted to the opportunity costs) and the benefits of holding the physical commodity. In other words, investing borrowed money only and taking no risk necessarily leads to a terminal wealth of zero.⁷ From an economic point of view, the convenience yield approach can be interpreted as a highly reduced supply and demand model, which allows using daily data, a major advantage compared to many previous studies on speculative bubbles.

Next, we briefly review the standard present-value model with time-varying expected returns as described, among others, in Campbell et al. (1997, Chapter 7). While in its original form the model is designed to explain stock price behavior, we follow Pindyck (1993), and use it to describe the dynamics of commodity prices. To this end, we replace the original variables representing the log stock price and the log dividend by the log commodity price and the standardized convenience yield as defined above.

In order to build up the model, we consider the relationship between the current log commodity price, the next period's expected log commodity price, and the convenience

⁶It should be noted that we compute our convenience yield measure by assuming arbitrage-free markets at each point in time. We are aware of the fact that this assumption may not be continuously satisfied in reality (Lombardi and Van Robays, 2011). However, owing to the fact that no alternative method for determining the fundamental value of commodities has been established so far in the literature, we retain the no-arbitrage condition.

⁷Besides Eq. (1) more sophisticated unobserved-components models have been put forward in the literature to approximate the convenience yield. Schwartz (1997), for example, presents a two-factor model by specifying the change of the logarithmic spot price and the convenience yield rate as a geometric Brownian motion and an Ornstein-Uhlenbeck process, respectively, the parameters of which he estimates via the Kalman-filtering technique. For our data, however, we find that the time series resulting from Schwartz' more complex procedure are qualitatively indistinguishable from those obtained by Eq. (1).

yield under rational expectations:

$$q = \kappa + \psi E_t(p_{t+1}) + (1 - \psi)d_t - p_t, \quad (2)$$

where q is the required log gross return rate, $E_t(\cdot)$ is the mathematical expectation operator conditional on all information available at date t , $p_t \equiv \ln(P_t)$ is the log commodity price at the end of period t , d_t is the convenience yield which the owner of the commodity experiences between t and $t + 1$, and $\kappa = -\ln(\psi) - (1 - \psi)\ln(1/\psi - 1)$ and $0 < \psi < 1$ are parameters of linearization.

Eq. (2) constitutes a linear expectational difference equation for the log commodity price which we routinely solve forward by repeatedly substituting out future prices and by using the law of iterated expectations to eliminate future-dated expectations. Imposing the transversality condition

$$\lim_{i \rightarrow \infty} \psi^i E_t(p_{t+i}) = 0,$$

we obtain the unique no-bubble solution to Eq. (2):

$$p_t^f = \frac{\kappa - q}{1 - \psi} + (1 - \psi) \sum_{i=0}^{\infty} \psi^i E_t(d_{t+i}). \quad (3)$$

Eq. (3) represents the well-known present-value relation stating that the log commodity price is equal to the present-value of expected future convenience yields out to the infinite future. However, it is important to note that from a mathematical point of view the above transversality condition may not be satisfied. Thus, the no-bubble solution p_t^f represents only a particular solution to the difference equation (2), while its general solution has the form:

$$p_t = p_t^f + B_t, \quad (4)$$

with the process $\{B_t\}$ satisfying the homogeneous difference equation:

$$E_t(B_{t+i}) = \frac{B_t}{\psi^i} \quad \text{for } i = 1, 2, \dots \quad (5)$$

(e.g. Cuthbertson and Nitzsche, 2004, pp. 397-401).

Obviously, the general solution in eq. (4) consists of two components. First, the no-bubble solution p_t^f only depends on the convenience yield, and therefore represents the market-fundamental solution. Second, events extraneous to the market may drive the mathematical entity B_t which we thus refer to as the rational speculative bubble component.

In order to circumvent nonstationarity problems, we express the model in first-difference form which, by virtue of Eqs. (3) and (4), is given by

$$\Delta p_t = \Delta p_t^f + \Delta B_t = (1 - \psi) \sum_{i=0}^{\infty} \psi^i [E_t(d_{t+i}) - E_{t-1}(d_{t+i-1})] + \Delta B_t. \quad (6)$$

Following Wu (1995, 1997), we also assume that the convenience yield may contain a unit root but that we can approximate the convenience yield process $\{d_t\}$ by an autoregressive integrated moving average (ARIMA) process. In particular, we assume an ARIMA($h, 1, 0$)-process of the form

$$\Delta d_t = \mu + \sum_{j=1}^h \phi_j \Delta d_{t-j} + \delta_t, \quad (7)$$

with $\delta_t \sim N(0, \sigma_\delta^2)$ denoting a Gaussian white-noise error term, in which we have to estimate the autoregressive order h from the data.

In what follows, it is convenient to express the autoregressive process (7) in companion form. Defining the $(h \times 1)$ vectors

$$\mathbf{y}_t = (\Delta d_t, \Delta d_{t-1}, \dots, \Delta d_{t-h+1})', \quad \mathbf{u} = (\mu, 0, 0, \dots, 0)', \quad \boldsymbol{\nu}_t = (\delta_t, 0, 0, \dots, 0)',$$

and the $(h \times h)$ matrix

$$\mathbf{A} = \begin{pmatrix} \phi_1 & \phi_2 & \phi_3 & \dots & \phi_{h-1} & \phi_h \\ 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 & 0 \end{pmatrix},$$

we may write Eq. (7) in the form

$$\mathbf{y}_t = \mathbf{u} + \mathbf{A}\mathbf{y}_{t-1} + \boldsymbol{\nu}_t. \quad (8)$$

Based on this representation, it follows from Campbell and Shiller (1987) that the solution to our commodity price model (6) results from the formula

$$\Delta p_t = \Delta d_t + \mathbf{m}\Delta \mathbf{y}_t + \Delta B_t, \quad (9)$$

where $\mathbf{m}$ is the $(h \times 1)$ vector

$$\mathbf{m} = \mathbf{g}\mathbf{A}(\mathbf{I} - \mathbf{A})^{-1}[\mathbf{I} - (1 - \psi)(\mathbf{I} - \psi\mathbf{A})^{-1}], \quad (10)$$

with the $(h \times 1)$ vector $\mathbf{g} = (1, 0, 0, \dots, 0)'$ and $\mathbf{I}$ symbolizing the $(h \times h)$ identity matrix.

In line with Wu (1995, 1997), we also assume a linear bubble process $\{B_t\}$. Hence, Eq. (5) implies

$$B_t = (1/\psi)B_{t-1} + \eta_t, \quad (11)$$

where we assume the innovation process $\{\eta_t\}$ to be i.i.d. $N(0, \sigma_\eta^2)$. Additionally, we assume that η_t is uncorrelated with the convenience yield innovation δ_t from Eq. (7).

2.2 Basic state-space representation

When estimating the commodity price equation (9), we are faced with the problem that the bubble component $\{B_t\}$ is unobservable. To sidestep this issue, we closely follow the lines of Wu (1995, 1997) and Al-Anaswah and Wilfling (2011) by first expressing our present-value model from above in state-space form, and then using the Kalman filter to estimate the unobservable oil price bubble $\{B_t\}$.

Let $\boldsymbol{\beta}_t$ be an $(n \times 1)$ vector of unobserved variables referred to as state variables, and $\mathbf{g}_t$ and $\mathbf{z}_t$ ($m \times 1$) and $(l \times 1)$ vectors of observable variables referred to as input and output variables, respectively. Then, we can write the state-space model as

$$\boldsymbol{\beta}_t = \mathbf{F}\boldsymbol{\beta}_{t-1} + \boldsymbol{\xi}_t, \quad (12)$$

$$\mathbf{z}_t = \mathbf{H}\boldsymbol{\beta}_t + \mathbf{D}\mathbf{g}_t + \boldsymbol{\zeta}_t, \quad (13)$$

where $\boldsymbol{\xi}_t$ and $\boldsymbol{\zeta}_t$ are $(n \times 1)$ and $(l \times 1)$ vectors of disturbances, respectively, and $\mathbf{F}$, $\mathbf{H}$, and $\mathbf{D}$ are constant real matrices of conformable dimensions. We assume that the

disturbance vectors $\boldsymbol{\xi}_t$ and $\boldsymbol{\zeta}_t$ are serially uncorrelated, uncorrelated with each other, and that

$$E(\boldsymbol{\xi}_t) = \mathbf{0}, \quad E(\boldsymbol{\zeta}_t) = \mathbf{0},$$

$$E(\boldsymbol{\xi}_t \boldsymbol{\xi}_t') = \boldsymbol{\Omega}, \quad E(\boldsymbol{\zeta}_t \boldsymbol{\zeta}_t') = \mathbf{R}.$$

Eqs. (12) and (13) are known as the transition and the measurement equation.

Basically, our present-value model established above consists of the following three components: the ARIMA($h, 1, 0$) convenience yield process $\{\Delta d_t\}$ from Eq. (7), the commodity price process $\{\Delta p_t\}$ from Eq. (9), and the bubble process $\{B_t\}$ from Eq. (11). It is straightforward to verify that we can express the entire present-value model in state-space form as follows:

$$\boldsymbol{\beta}_t = (B_t, B_{t-1})', \quad \mathbf{z}_t = (\Delta d_t, \Delta p_t)', \quad \mathbf{g}_t = (1, \Delta d_t, \Delta d_{t-1}, \Delta d_{t-2}, \dots, \Delta d_{t-h})',$$

$$\boldsymbol{\xi}_t = (\eta_t, 0)', \quad \boldsymbol{\zeta}_t = (\delta_t, 0)',$$

$$\mathbf{F} = \begin{pmatrix} 1/\psi & 0 \\ 1 & 0 \end{pmatrix}, \quad \mathbf{H} = \begin{pmatrix} 0 & 0 \\ 1 & -1 \end{pmatrix}, \quad (14)$$

and

$$\mathbf{D} = \begin{pmatrix} \mu & 0 & \phi_1 & \phi_2 & \dots & \phi_{h-1} & \phi_h \\ 0 & (1 + m_1) & (m_2 - m_1) & (m_3 - m_2) & \dots & (m_h - m_{h-1}) & -m_h \end{pmatrix}, \quad (15)$$

where m_i is the i th component of the $(h \times 1)$ vector $\mathbf{m}$ defined in Eq. (10). The covariance matrices $\boldsymbol{\Omega}$ and $\mathbf{R}$ are given by:

$$\boldsymbol{\Omega} = \begin{pmatrix} \sigma_\eta^2 & 0 \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad \mathbf{R} = \begin{pmatrix} \sigma_\delta^2 & 0 \\ 0 & 0 \end{pmatrix}. \quad (16)$$

To sum up, our state-space representation treats the commodity price bubble as an unobservable state variable, and specifies two transition and two measurement equations. Both transition equations represent the bubble process (11), while the first measurement equation represents the convenience yield process (7), and the second measurement equation the commodity price process (9).

2.3 State-space representation with Markov-switching

We now introduce two distinct Markov-regimes into the basic state-space model from above. By this approach, we aim at econometrically separating two distinct states in the data-generating process of the bubble from each other, namely one regime in which the bubble is a stable process and a second regime in which the bubble process is explosive. In what follows, we refer to the stable regime as regime 1, and to the exploding regime as regime 2. Using this interpretation of the distinct regimes, we identify a collapse of the bubble by the transition from regime 2 to regime 1. Our formal exposition closely follows Kim and Nelson (1999, Chapter 5).

We begin with the state-space representation of a dynamic system consisting of the transition equation (12) and the measurement equation (13). Additionally, we now allow the parameters in the matrices $\mathbf{F}$, $\mathbf{H}$, $\mathbf{D}$, $\mathbf{\Omega}$, and $\mathbf{R}$ to switch between two distinct regimes. We denote this switching-property by writing the state-space model from above compactly as

$$\boldsymbol{\beta}_t = \mathbf{F}_{S_t} \boldsymbol{\beta}_{t-1} + \boldsymbol{\xi}_t, \quad (17)$$

$$\mathbf{z}_t = \mathbf{H}_{S_t} \boldsymbol{\beta}_t + \mathbf{D}_{S_t} \mathbf{g}_t + \boldsymbol{\zeta}_t, \quad (18)$$

$$\begin{pmatrix} \boldsymbol{\xi}_t \\ \boldsymbol{\zeta}_t \end{pmatrix} \sim N \left(0, \begin{pmatrix} \mathbf{\Omega}_{S_t} & 0 \\ 0 & \mathbf{R}_{S_t} \end{pmatrix} \right), \quad (19)$$

where the subscript $S_t \in \{1, 2\}$ indicates that the parameters in the matrices are governed by an unobservable two-state random variable determining the specific regime the parameters are in at date t . We specify the probabilistic nature of the regime-indicator S_t by a first-order Markov process with constant transition probabilities $p_{ij} = \Pr[S_t = j | S_{t-1} = i]$ which we collect in the matrix

$$\boldsymbol{\Pi} = \begin{pmatrix} p_{11} & 1 - p_{22} \\ 1 - p_{11} & p_{22} \end{pmatrix}. \quad (20)$$

In the next section we establish a Bayesian approach to estimating all parameters contained in the Eqs. (17) to (20).

3 Model adjustments and MCMC estimation

3.1 Model adjustments

We estimate the model in Eqs. (17) to (20) by Bayesian MCMC methods.⁸ Let $\mathbf{S}$ be the vector of regimes, $\boldsymbol{\beta}$ the vector of states, and $\mathbf{z}$ the two-dimensional process $\mathbf{z} = \begin{pmatrix} \Delta d \\ \Delta p \end{pmatrix}$. We thus have:

$$\mathbf{S} = (S_1, \dots, S_T),$$

$$\boldsymbol{\beta} = (\beta_1, \dots, \beta_T),$$

$$\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_T).$$

However, before starting the estimation, we have to tackle two problems. First, convenience yields may become negative. This is in contrast to the investigation of stock price bubbles where both the stock prices and the dividends are always positive. As a result, a gap might open between the fundamental value and the bubble term modeled as an AR(1)-process with regime-switching parameters $1/\psi_{S_t}$ and σ_{η, S_t} . We account for this gap by including an additional parameter, γ_{S_t} , in the bubble specification:

$$B_t - \gamma_{S_t} = \frac{1}{\psi_{S_t}} \cdot (B_{t-1} - \gamma_{S_t}) + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_{\eta, S_t}). \quad (21)$$

With this adjustment, the elements of the transition and measurement equations become

$$\boldsymbol{\beta}_t = \begin{pmatrix} B_t \\ B_{t-1} \\ 1 \end{pmatrix}, \quad \mathbf{F}_{S_t} = \begin{pmatrix} \frac{1}{\psi_{S_t}} & 0 & \gamma_{S_t}(1 - \frac{1}{\psi_{S_t}}) \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \boldsymbol{\xi}_t = \begin{pmatrix} \eta_t \\ 0 \\ 0 \end{pmatrix}, \quad (22)$$

and

$$\mathbf{H}_{S_t} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & -1 & 0 \end{pmatrix}. \quad (23)$$

All other matrices and vectors remain unchanged. The parameters γ_1 and γ_2 capture those components of the fundamental value which are not accounted for by the convenience yield.

⁸For the implementation we use the `MCMCpack` and `dlm` packages available for the statistical software R.

A second problem arises from the fact that our convenience yield data are rather volatile. There are two distinct sources for high volatility. First, it may result from our use of daily data which are known to be noisy in many real-world situations. Second, the convenience yield is modeled as a function of spot and futures prices and the risk-free interest rate, all of which typically exhibit high volatility on their own. However, since the purpose of the fundamental value is to describe the basic movement of the asset price, we decide to smooth our convenience yield data by applying local polynomial regression-fitting.⁹

With this modifications in mind, we can proceed to the MCMC method's sampling process. The distribution we wish to sample from is the joint posterior distribution

$$\pi(\mathbf{P}, \boldsymbol{\beta}, \mathbf{S}, p_{11}, p_{22} | \mathbf{z} = (\Delta d, \Delta p)). \quad (24)$$

where $\mathbf{P} = (\mu_1, \mu_2, \delta_1, \delta_2, \eta_1, \eta_2, \psi_1, \psi_2, \phi_{1,1}, \dots, \phi_{h,1}, \phi_{1,2}, \dots, \phi_{h,2})$ collects the model's parameters. We use the Gibbs sampler in order to generate a sample from (24).

3.2 Description of the Gibbs sampler steps

The model parameters can be split into four groups: States $\boldsymbol{\beta}$, parameters of the model equations $\mathbf{P}$, regimes $\mathbf{S}$, and Markov regime-switching probabilities p_{11}, p_{22} . For each group, we set up a Gibbs sampler step:

- (1) Sample from $\pi(\boldsymbol{\beta} | \mathbf{z}, \mathbf{P}, \mathbf{S}, p_{11}, p_{22})$.
- (2) Sample from $\pi(\mathbf{P} | \mathbf{z}, \boldsymbol{\beta}, \mathbf{S}, p_{11}, p_{22})$.
- (3) Sample from $\pi(\mathbf{S} | \mathbf{z}, \boldsymbol{\beta}, \mathbf{P}, p_{11}, p_{22})$.
- (4) Sample from $\pi(p_{11}, p_{22} | \mathbf{z}, \boldsymbol{\beta}, \mathbf{P}, \mathbf{S})$.

Note that not all parameters listed in the conditioning sets are necessary, and may be omitted at some point in the estimation process. For instance, the conditional posterior

⁹Details of our smoothing procedure are available upon request.

distribution of p_{11} and p_{22} does not depend on β and $\mathbf{P}$. The order of the four Gibbs steps is arbitrary as long as one iteration of the Gibbs sampler includes each step once and only once. The algorithm can be summarized as follows:

- (1) Sampling from the conditional distribution of the states $\pi(\beta|\mathbf{z}, \mathbf{S}, \mathbf{P})$ requires a Forward-Filtering-Backward-Sampling (FFBS) algorithm as in the standard Kalman filter with time-varying parameters. We declare the parameters that are subject to regime switches as time-varying with the two possible values determined by the regimes. For efficient programming, we use the `d1m` package of `R`.¹⁰ The FFBS methods provided by the `d1m` package can be used to generate a sample of the process β . Concerning starting values, we use arbitrarily chosen values for all parameters, and a randomly generated regime process.
- (2) The second step is sampling of the parameters $\mathbf{P}$. Notice that if both the regime process $\mathbf{S}$ and the state process β are given, the measurement equation (18) and the transition equation (21) become independent and thus can be handled separately. As to the measurement equation, we sample the parameters $\mu_1, \mu_2, \phi_{1,1}, \dots, \phi_{h,1}, \phi_{1,2}, \dots, \phi_{h,2}$, and $\sigma_{\delta,1}, \sigma_{\delta,2}$ of the dividend process

$$\Delta d_t = \mu_{S_t} + \phi_{1,S_t} \Delta d_{t-1} + \dots + \phi_{h,S_t} \Delta d_{t-h} + \delta_{t,S_t}, \quad \delta_{t,S_t} \sim \mathcal{N}(0, \sigma_{\delta,S_t}^2). \quad (25)$$

The price process does not have any parameters to be estimated.

Sampling the parameters of the transition equation (21) is similiar. We use the `MCMCregress` command in order to draw from the conditional posterior distribution of the parameters of

$$(\beta_t - \gamma_{S_t}) = \frac{1}{\psi_{S_t}} (\beta_{t-1} - \gamma_{S_t}) + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_{\eta,S_t}^2), \quad (26)$$

¹⁰For a detailed documentation see Petris et al. (2009) or the more general overview over filtering techniques by Petris and Petrone (2011). The package provides extensive methods for estimation, filtering, and sampling of dynamic linear models, and can handle time-varying parameters. Note that the `d1m` package does not natively support additional exogenous variables such as the term $\mathbf{D}_{S_t} \mathbf{g}_t$ in Eq. (18). In order to avoid a re-parametrization of the model, we calculate $\mathbf{z}_t^{dlm} = \mathbf{z}_t - \mathbf{D}_{S_t} \mathbf{g}_t$. This poses no problem in the MCMC estimation since all parameters in $\mathbf{D}_{S_t}$ and the complete regime process S_t are known in this step of the Gibbs iteration.

with given states $\boldsymbol{\beta}$ and regimes $S_t \in \{1, 2\}$.¹¹

- (3) In the third step, we sample the regime process $\mathbf{S}$. Sampling from the conditional posterior distribution $\pi(\mathbf{S}|\mathbf{z}, \boldsymbol{\beta}, \mathbf{P})$ is similar to the FFBS algorithm used in step one. As to the forward-filtering part, we perform successive prediction steps $\pi(S_t|z_{1,\dots,t-1}, \beta_{1,\dots,t-1}, \mathbf{P})$ and update steps $\pi(S_t|z_{1,\dots,t}, \beta_{1,\dots,t}, \mathbf{P})$ to arrive at the filtered distributions

$$\pi(S_t|z_{1,\dots,t}, \beta_{1,\dots,t}, \mathbf{P}) \propto \pi(z_t, \beta_t|S_t, \beta_{t-1}, \mathbf{P}) \pi(S_t|z_{1,\dots,t-1}, \beta_{1,\dots,t-1}, \mathbf{P}),$$

where $\pi(z_t, \beta_t|S_t, \beta_{t-1}, \mathbf{P})$ is multivariate normal:

$$\begin{pmatrix} z_t \\ \beta_t \end{pmatrix} | S_t, \beta_{t-1} \sim \mathcal{N}(\mathcal{M}_{S_t}, \Sigma_{S_t}),$$

with expectation vector $\mathcal{M}_{S_t}$ and covariance matrix Σ_{S_t} for regime S_t . We have:

$$\begin{aligned} \mathcal{M}_{S_t} &= \begin{pmatrix} F_{S_t} G_{S_t} \beta_{t-1} + D_{S_t} g_t \\ G_{S_t} \beta_{t-1} \end{pmatrix}, \\ \Sigma_{S_t} &= \begin{pmatrix} H_{S_t} R_{S_t} H_{S_t}' + \Omega_{S_t} & H_{S_t} R_{S_t} \\ R_{S_t} H_{S_t}' & R_{S_t} \end{pmatrix}. \end{aligned}$$

Note that the process $\boldsymbol{\beta}$ is assumed to be known in this step, and can be treated as a constant. The covariance matrix Σ_{S_t} is singular in light of Eq. (9) and given that:

$$\begin{pmatrix} z_t \\ \beta_t \end{pmatrix} = \begin{pmatrix} \Delta p_t & \Delta d_t & \beta_t & \beta_{t-1} & 1 \end{pmatrix}'.$$

Considering the subvector $\begin{pmatrix} \Delta d_t \\ \beta_t \end{pmatrix}$ yields the expectation vector and the non-singular covariance matrix

$$\mathcal{M}_{S_t} = \begin{pmatrix} D_{1,S_t} \cdot g_t \\ \frac{1}{\psi_{S_t}} \cdot \beta_{t-1} \end{pmatrix}, \quad \Sigma_{S_t} = \begin{pmatrix} \sigma_{\delta,S_t}^2 & 0 \\ 0 & \sigma_{\eta,S_t}^2 \end{pmatrix},$$

where D_{1,S_t} denotes the first row of D_{S_t} for regime S_t .

¹¹Theoretically, it might happen that the random draw from $\mathbf{S}$ only generates a single regime, e.g. $S_t = 1$ for all t . In this case, the parameters pertaining to the other regime would not be identified, and would have to be drawn from their prior distribution. As there are always regime switches in our empirical application, we do not deal with this problem any further.

Once the forward filtering recursion has reached the final period, we draw from $\pi(S_T|z_{1,...,T}, \beta_{1,...,T}, \mathbf{P})$, and then use the backward sampling recursion

$$\pi(S_t|z_{1,...,t-1}, \beta_{1,...,t-1}, S_{t+1,...,T}, \mathbf{P}) \propto \pi(S_{t+1}|S_t) \cdot \pi(S_t|z_{1,...,t}, \beta_{1,...,t}, \mathbf{P})$$

to sample the regimes. $\pi(S_{t+1}|S_t)$ is given by the transition probabilities of the underlying Markov process, and $\pi(S_t|z_{1,...,t}, \beta_{1,...,t}, \mathbf{P})$ are the filtered probability distributions.¹²

- (4) The last step of the Gibbs iteration provides updates for the transition probabilities p_{11} and p_{22} governing the regime process. If we choose conjugated priors for both p_{11} and p_{22} , i.e. beta distributions, then we only need to calculate the number of switches between regimes 1 and 2 in order to derive the posterior distributions of p_{11} and p_{22} . Draws from the posterior beta distribution are standard.

The draws of each iteration are saved in a global results matrix. After eliminating the burn-in period, we have a sample from the joint distribution of all parameters, states, and regimes. Point estimates can be generated by calculating their sample means. Confidence intervals and standard deviations of single parameters are simply calculated from the parameters' sampled distribution.

4 Empirical analysis

4.1 Data

We use spot and futures prices as well as a proxy for the risk-free interest rate to approximate the convenience yield of oil as given in Eq. (1). Concerning price data, we analyze the two main reference types of oil, WTI (from April 1983 to April 2011)

¹²A detailed derivation of the necessary distributions is available from the authors upon request.

and Brent (from April 1989 to April 2011).¹³ Panel A of Figure 1 shows the daily spot price of WTI and Brent.

[Figure 1 about here]

We also calculate artificial spot prices derived from the shape of the futures curve in order to avoid potential distortions in the actual spot prices due to discounts and premiums which may result from long-standing relationships between buyers and sellers (Pindyck, 1993). Our daily futures prices belong to contracts traded on the New York Mercantile Exchange (NYMEX) for WTI, and on the Intercontinental Exchange (ICE) Futures Europe for Brent. Applying the first-day-of-delivery-month criterion, we always draw on the first-nearby contract and roll over to the second-nearby on the first day of the first-nearby's delivery month. The reason for rolling over sufficiently prior to the expiration of the first-nearby is that the latter runs out of liquidity close to maturity. Alternatively, following the liquidity-peak criterion, we also experiment with rolling over once the second-nearby exhibits a continuously higher open interest than the first-nearby. And finally, we approximate the risk-free rate by the three-month US Treasury bill interest rate, and alternatively also experiment with the Federal funds rate. We obtain all time series from *Thomson Reuters Datastream*. Oil prices are quoted in US-\$/bbl, the interest rates are given in percent p.a.

In our economic model, the spot price depends on the convenience yield. According to Eq. (1), the convenience yield is modeled as a function of the spot price. In order to sidestep endogeneity problems, we draw on the artificial (actual) spot price in the economic model if the actual (artificial) spot price is used to calculate the convenience yield. Additionally, we run our analysis for both nominal and real data. Table 1 provides an overview of the alternative models. For illustrative purposes, Panel B of

¹³While the prices of WTI and Brent are typically close together, substantial deviations between both time series have emerged over the last couple of months. In order to check whether this anomaly has an impact on our overall results, we consider both oil types in our empirical analysis.

Figure 1 displays the convenience yield of WTI and Brent computed from the models 1 and 9, respectively.

[Table 1 about here]

4.2 Empirical results

We begin with a description of our data and the priors used in the estimation procedure. Our WTI and Brent datasets consist of 7047 and 5611 observations, respectively. We run a total of $N + R = 25000 + 1000$ Gibbs iterations, and delete the first R results as burn-in phase. We specify the fundamental dynamics (7) as an ARIMA(2,1,0)-process in both regimes. Table 2 summarizes our starting values and prior distributions.

[Table 2 about here]

Next, we turn to the estimation of model 1 for the WTI dataset. Figure 2 displays the posterior distributions of the relevant model parameters. The bold vertical lines represent our point estimates of the respective parameters (that is the mean of the respective posterior distribution).

[Figure 2 about here]

Tables 3 and 4 display the estimates of the parameters from the models 1 to 8 (WTI dataset) and the models 9 to 16 (Brent dataset), respectively. Since our study focusses on the detection of speculative bubbles, we confine ourselves here to discussing only the bubble-relevant regime-specific parameters ψ_1 and ψ_2 , and the transition probabilities p_{11} and p_{22} . It is instructive to recall that according to our bubble specification (21) the parameters ψ_1 and ψ_2 indicate whether a bubble regime is explosive ($\psi_{S_t} < 1$) or stable ($\psi_{S_t} > 1$). For all models estimated in our study, we consistently find $\hat{\psi}_1 > 1$ indicating the stable bubble regime 1, and (except for model 4) $\hat{\psi}_2 < 1$ marking the explosive

bubble regime 2. Analyzing the transition probabilities p_{11} and p_{22} , we typically find the stable bubble regime 1 to be more persistent than the explosive regime 2 (except for model 11).

[Table 3 about here]

[Table 4 about here]

Tables 5 and 6 report the 99% confidence intervals of the bubble parameters ψ_1 and ψ_2 . Evidently, the large majority of the confidence intervals do not contain the value 1 implying that ψ_1 and ψ_2 can be considered as significantly different from 1 at the 1%-level in these cases. Only in models 4 and 8, ψ_2 does not appear to be significantly smaller than 1 (meaning that we do not find statistical significance of an explosive behavior of the bubble process), while in model 14, ψ_1 is not significantly larger than 1 (indicating that the bubble process is not significantly stable at the 1% level).

[Table 5 about here]

[Table 6 about here]

Figure 3 displays the estimated bubble process and the corresponding regime-2 probabilities for model 1.¹⁴ Obviously, the bubble process is negative for most of the first 5,000 observations, then steadily increases, and reaches a peak in mid-2008. After a sharp decline, the bubble process grows again until the end of the sample in April 2011. Pertaining to the regime-2 probabilities, we observe various short-lived switches from the stable bubble regime 1 to the explosive regime 2 during the first 5,000 observations. For the rest of the sample the bubble process appears to stay in the explosive bubble regime 2 with a single interruption observed in mid-2008 reflecting the bursting of the bubble at that time.

¹⁴We define a regime-2 probability as the probability of being in regime 2 at date t given all information contained in $\mathbf{z}_1, \dots, \mathbf{z}_T$.

[Figure 3 about here]

When visually inspecting the results for the 16 models estimated, we find striking synchronicity among all WTI and all Brent regime-2 probability time-series.¹⁵ Owing to this synchronicity, we consolidate the multiple regime-2 probability processes to one single WTI and one single Brent regime process by first taking arithmetic averages of the 8 WTI and the 8 Brent regime-2 probability processes, respectively. We then use these average WTI and Brent regime-2 probability processes to construct a WTI and a Brent regime process by the following rules:

$$S_t^{\text{WTI}} = \begin{cases} 1, & \text{if at time } t \text{ the average WTI regime-2 probability} < 0.5 \\ 2, & \text{if at time } t \text{ the average WTI regime-2 probability} \geq 0.5 \end{cases},$$

and S_t^{Brent} analogously defined.

Figure 4 displays the WTI and Brent regime processes $\{S_t^{\text{WTI}}\}$ and $\{S_t^{\text{Brent}}\}$ during the respective sampling periods. Both panels provide clear-cut evidence that the bubble component was in the explosive regime 2 at the beginning of the 1990s as well as between 2004 and 2011. Interestingly, both panels also exhibit a transition from the explosive bubble-regime 2 to the stable bubble-regime 1 in mid-2008 hinting at the bursting of the oil price bubble.

[Figure 4 about here]

As a final example, Figure 5 illustrates the WTI spot prices and convenience yields computed from model 1. The shaded areas show the periods in which we find for the consolidated regime process $S_t^{\text{WTI}} = 2$ signaling that the WTI bubble process is likely to have been in the explosive regime 2. In line with our expectations, the explosive bubble regime 2 appears to have been in force between the end of 2004 and mid-2008, and again between mid-2008 and April 2011 (the end of our sampling period) providing clear-cut evidence in favor of the existence of speculative bubbles in recent oil prices.

¹⁵All remaining 15 regime-2 probability processes not shown in this study are available upon request.

[Figure 5 about here]

5 Concluding remarks

Motivated by repeated price spikes and crashes over the last years, we investigate whether the heavily financialized market for crude oil has been driven by speculative bubbles. In our theoretical modeling we draw on the convenience yield approach in order to approximate the fundamental value by the sum of discounted future oil dividends. We separate the fundamentally justified part of the oil price from the bubble component by expressing the standard present-value model in state-space form. We then introduce two Markov-regimes into the state-space representation in order to distinguish between stable and explosive phases in the bubble process. In contrast to Al-Anaswah and Wilfling (2011), who estimate a similar Markov-switching state-space approach for stock prices by maximum likelihood techniques, we establish an econometrically more robust Bayesian MCMC methodology in this paper. Based on inferential techniques designed for statistically separating both Markov-regimes from each other, we find robust evidence for the existence of speculative bubbles in recent oil price dynamics.

Given the importance of the oil price for the real economy, stock markets, and monetary policy, our findings may indicate the need for a more efficient regulation of financial investors committed to oil derivative markets. As generally accepted, futures trading is a valuable activity since it improves price discovery, enhances market efficiency, increases market depth and informativeness, and contributes to market completion. However, if these benefits are outweighed by over- and undershooting oil prices caused by speculative trading, regulators should consider implementing effective position limits, as currently executed in the United States (CFTC, 2011b) and at least discussed in Europe. Apart from that, our results contribute to the on-going discussion on whether central banks should actively fight speculative bubbles or just observe their evolutions

and crashes. In addition, given the potential of oil price bubbles to bias the consumer price index upwards, central banks might benefit from putting more focus on core inflation instead of headline inflation in order to avoid erroneous decisions on the course of monetary policy.

As to commodity portfolio management, our approach can be operated in a routine way to make density forecasts for future oil prices, taking into account both the conditional information on the regime probabilities and the uncertainty about the model parameters. Each draw from the forecast density can be generated by first drawing a sample from the joint posterior density (24). Given the realization of all model parameters, states, regimes, and transition probabilities, one can use Eqs. (17) to (20) to draw next period's state variables, regimes, and observations – including the oil price change. Repeating both steps a large number of times results in a distribution of future prices.

Finally, in view of the little explicit testing for speculative bubbles in commodity markets and the on-going academic and political debate on tighter regulation of speculators, scope for future research is given by applying our MCMC Markov-switching state-space approach to other raw material prices which have recently been blamed for over- and undershooting as well. In addition, similarly powerful testing procedures should be used to broaden the empirical evidence of speculative bubbles on markets for raw materials in general and for crude oil in particular.

References

- Al-Anaswah, N., Wilfling, B., 2011. Identification of speculative bubbles using state-space models with Markov-switching. *Journal of Banking and Finance* 35(5), 1073-1086.
- Bekiros, S. D., Diks, C. G. H., 2008. The relationship between crude oil spot and futures prices: Cointegration, linear and nonlinear causality. *Energy Economics* 30(5), 2673-2685.
- CFTC, 2008. Interim report on crude oil. Interagency Task Force on Commodity Markets, Washington, DC: US Commodity Futures Trading Commission.
- CFTC, 2011a. CFTC commitments of traders. Washington, DC: US Commodity Futures Trading Commission.
- CFTC, 2011b. Position limits for derivatives. *Federal Register* 76(17), 4752-4777.
- Cifarelli, G., Paladino, G., 2010. Oil price dynamics and speculation: A multivariate financial approach. *Energy Economics* 32(2), 363-372.
- ECB, 2008. Editorial. *Monthly Bulletin* July 2008, Frankfurt/Main: European Central Bank.
- Garcia, E., 2006. Bubbling crude: Oil price speculation and interest rates. *E-Journal of Petroleum Management and Economics* 1(1).
- Gürkaynak, R. S., 2008. Econometric tests of asset price bubbles: Taking stock. *Journal of Economic Surveys* 22(1), 166-186.
- Hamilton, J. D., 2009a. Causes and consequences of the oil shock of 2007-2008. NBER Working Papers 15002.
- Hamilton, J. D., 2009b. Understanding crude oil prices. *Energy Journal* 30(2), 179-206.
- IMF, 2005. *World economic outlook – Globalization and external imbalances*. Washington, DC: International Monetary Fund, pp. 57-65.
- IMF, 2006. *World economic outlook – Financial systems and economic cycles*. Washington, DC: International Monetary Fund, pp. 15-18.

- Kaufmann, R. K., Dees, S., Gasteuil, A., Mann, M., 2008. Oil prices: The role of refinery utilization, futures markets and non-linearities. *Energy Economics* 30(5), 2609-2622.
- Kaufmann, R. K., Ullman, B., 2009. Oil prices, speculation, and fundamentals: Interpreting causal relations among spot and futures prices. *Energy Economics* 31(4), 550-558.
- Kesicki, F., 2008. The third oil price surge – What’s different this time? *Energy Policy* 38(3), 1596-1606.
- Khan, M. S., 2009. The 2008 oil price “bubble”. Washington, DC: Peterson Institute for International Economics Policy Brief, August 2009.
- Kilian, L., 2008. The economic effects of energy price shocks. *Journal of Economic Literature* 46(4), 871-909.
- Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99(3), 1053-1069.
- Kilian, L., Hicks, B., 2009. Did unexpectedly strong economic growth cause the oil price shock of 2003-2008? CEPR Discussion Paper No. DP7265.
- Kilian, L., Murphy, D., 2010. The role of inventories and speculative trading in the global market for crude oil. University of Michigan: mimeo.
- Kim, C.-J., Nelson, C. R., 1999. State space models with regime switching. Cambridge: MIT Press.
- Lombardi, M. J., Van Robays, I., 2011. Do financial investors destabilize the oil price? Working Paper Series No. 1346, Frankfurt/Main: European Central Bank.
- Miller, J. I., Ratti, R. A., 2009. Crude oil and stock markets: Stability, instability, and bubbles. *Energy Economics* 31(4), 559-568.
- Petris, G., Petrone, S., 2011. State space models in R. *Journal of Statistical Software* 41(4), 1-24.
- Petris, G., Petrone, S., Campagnoli, P., 2009. Dynamic linear models with R. New York: Springer.

- Pindyck, R. S., 1993. The present-value model of rational commodity pricing. *Economic Journal* 103(418), 511-530.
- Reitz, S., Rülke, J. C., Stadtmann, G., 2009. Are oil price forecasters finally right? Regressive expectations toward more fundamental values of the oil price. Discussion Paper Series 1: Economic Studies No. 32/2009, Frankfurt/Main: Deutsche Bundesbank.
- Reitz, S., Slopek, U., 2009. Non-linear oil price dynamics: A tale of heterogeneous speculators? *German Economic Review* 10(3), 270-283.
- Sanders, D. R., Boris, K., Manfredo, M., 2004. Hedgers, funds, and small speculators in the energy futures markets: An analysis of the CFTC's commitments of traders report. *Energy Economics* 26(3), 425-445.
- Schwartz, E. S., 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance* 52(3), 923-973.
- Schwarz, T. V., Szakmary, A. C., 1994. Price discovery in petroleum markets: Arbitrage, cointegration, and the time interval of analysis. *Journal of Futures Markets* 14(2), 147-167.
- Silvapulle, P., Moosa, I. A., 1999. The relationship between spot and futures prices: Evidence from the crude oil market. *Journal of Futures Markets* 19(2), 175-193.
- Tokic, D., 2010. The 2008 oil bubble: Causes and consequences. *Energy Policy* 38(10), 6009-6015.
- Tokic, D., 2011. Rational destabilizing speculation, positive feedback trading, and the oil bubble of 2008. *Energy Policy* 39(4), 2051-2061.
- Wu, Y., 1995. Are there rational bubbles in foreign exchange markets? Evidence from an alternative test. *Journal of International Money and Finance* 14(1), 27-46.
- Wu, Y., 1997. Rational bubbles in the stock market: Accounting for the U.S. stock-price volatility. *Economic Inquiry* 35(2), 309-319.

Table 1: Overview of the different types of models

Model	Oil Type	Spot price	Convenience yield			Nom./Real
			Spot price	Roll mechanism	Interest rate	
1	WTI	Original	Artificial	First-day-of-delivery-month	US T-bill rate	Nominal
2	WTI	Artificial	Original	First-day-of-delivery-month	US T-bill rate	Nominal
3	WTI	Original	Artificial	Liquidity-peak	US T-bill rate	Nominal
4	WTI	Original	Artificial	First-day-of-delivery-month	Fed funds rate	Nominal
5	WTI	Original	Artificial	First-day-of-delivery-month	US T-bill rate	Real
6	WTI	Artificial	Original	First-day-of-delivery-month	US T-bill rate	Real
7	WTI	Original	Artificial	Liquidity-peak	US T-bill rate	Real
8	WTI	Original	Artificial	First-day-of-delivery-month	Fed funds rate	Real
9	Brent	Original	Artificial	First-day-of-delivery-month	US T-bill rate	Nominal
10	Brent	Artificial	Original	First-day-of-delivery-month	US T-bill rate	Nominal
11	Brent	Original	Artificial	Liquidity-peak	US T-bill rate	Nominal
12	Brent	Original	Artificial	First-day-of-delivery-month	Fed funds rate	Nominal
13	Brent	Original	Artificial	First-day-of-delivery-month	US T-bill rate	Real
14	Brent	Artificial	Original	First-day-of-delivery-month	US T-bill rate	Real
15	Brent	Original	Artificial	Liquidity-peak	US T-bill rate	Real
16	Brent	Original	Artificial	First-day-of-delivery-month	Fed funds rate	Real

Table 2: Starting values and prior distributions

Parameter	Regime 1	Regime 2	Prior distribution
μ	0	0	$\mathcal{N}(0, 0.1^2)$
σ_δ^2	0.15^2	0.15^2	$\Gamma(1, \frac{1}{2})^{-1}$
ψ	1.413	0.95	$\mathcal{N}([1, 0.75], 0.05^2)^{-1}$
σ_η^2	0.45^2	0.6^2	$\Gamma(1, \frac{1}{2})^{-1}$
p_{11}	-	-	$Beta(15, 1)$
p_{22}	-	-	$Beta(15, 1)$
γ	0.2	0.5	$\mathcal{N}(0, 1)$
$\phi_{1, \dots, h}$	random	random	$\mathcal{N}(0, 0.1^2)$

Table 3: Parameter estimates, WTI dataset

Regime	Model 1		Model 2		Model 3		Model 4	
	1	2	1	2	1	2	1	2
μ	-0.012	-0.018	0.011	0.033	-0.213	0.182	0.056	0.083
σ_δ^2	0.024	0.028	0.057	0.173	0.058	0.143	0.032	0.154
ψ	1.393	0.958	1.263	0.791	1.508	0.778	1.260	1.018
σ_η^2	0.188	0.369	0.295	0.427	0.297	0.278	0.187	0.446
p_{11}, p_{22}	0.973	0.713	0.986	0.793	0.980	0.929	0.857	0.785
γ	0.034	-2.816	-0.104	-2.799	0.125	-2.628	-0.009	-2.895
ϕ_1	0.112	-0.247	0.295	-0.556	0.276	-0.103	0.321	-0.093
ϕ_2	0.159	-0.405	0.385	-0.700	-0.195	-0.429	0.323	-0.422
Regime	Model 5		Model 6		Model 7		Model 8	
	1	2	1	2	1	2	1	2
μ	0.033	-0.009	-0.277	-0.120	0.059	-0.335	-0.359	-0.120
σ_δ^2	0.077	0.185	0.049	0.356	0.173	0.228	0.020	0.080
ψ	1.238	0.958	1.255	0.676	1.389	0.884	1.442	0.998
σ_η^2	0.222	0.176	0.281	0.554	0.041	0.385	0.016	0.149
p_{11}, p_{22}	0.992	0.742	0.904	0.738	0.972	0.769	0.942	0.743
γ	-0.071	-2.857	0.042	-2.722	-0.009	-2.760	0.328	-2.628
ϕ_1	0.232	-0.385	0.127	-0.286	0.239	-0.071	-0.038	-0.599
ϕ_2	0.202	-0.592	0.042	-0.352	0.286	-0.606	0.015	-0.333

Table 4: Parameter estimates, Brent dataset

Regime	Model 9		Model 10		Model 11		Model 12	
	1	2	1	2	1	2	1	2
μ	-0.273	-0.116	0.080	-0.037	0.041	-0.225	0.067	0.133
σ_δ^2	0.145	0.122	0.261	0.191	0.219	0.196	0.210	0.439
ψ	1.289	0.859	1.422	0.798	1.144	0.668	1.471	0.854
σ_η^2	0.403	0.487	0.275	0.258	0.280	0.524	0.251	0.368
p_{11}, p_{22}	0.940	0.870	0.969	0.793	0.817	0.848	0.927	0.638
γ	-0.110	0.094	0.020	-0.038	-0.034	-0.128	0.054	0.074
ϕ_1	-0.176	0.077	-0.057	-0.062	0.262	0.166	0.196	-0.051
ϕ_2	0.101	0.274	-0.112	0.061	0.191	-0.200	0.043	-0.257
Regime	Model 13		Model 14		Model 15		Model 16	
	1	2	1	2	1	2	1	2
μ	-0.101	0.071	-0.088	-0.126	-0.057	0.113	-0.018	0.313
σ_δ^2	0.044	0.018	0.083	0.217	0.081	0.001	0.190	0.029
ψ	1.233	0.782	1.078	0.677	1.364	0.757	1.563	0.839
σ_η^2	0.154	0.611	0.185	0.381	0.268	0.106	0.273	0.356
p_{11}, p_{22}	0.964	0.907	0.889	0.883	0.985	0.932	0.922	0.828
γ	0.183	-0.151	0.360	0.009	-0.038	-0.208	-0.173	0.044
ϕ_1	0.117	-0.129	0.040	0.124	-0.293	-0.082	-0.158	-0.272
ϕ_2	0.123	-0.027	0.022	0.082	0.017	0.132	0.011	-0.214

Table 5: 99% confidence intervals, WTI dataset


	Model 1	Model 2	Model 3	Model 4
ψ_1	[1.284 ; 1.480]	[1.152 ; 1.373]	[1.424 ; 1.588]	[1.197 ; 1.346]
ψ_2	[0.921 ; 0.981]	[0.711 ; 0.865]	[0.697 ; 0.854]	[0.901 ; 1.180]
	Model 5	Model 6	Model 7	Model 8
ψ_1	[1.090 ; 1.371]	[1.103 ; 1.392]	[1.275 ; 1.517]	[1.345 ; 1.622]
ψ_2	[0.899 ; 0.987]	[0.653 ; 0.698]	[0.811 ; 0.957]	[0.946 ; 1.051]

Table 6: 99% confidence intervals, Brent dataset

	Model 9	Model 10	Model 11	Model 12
ψ_1	[1.134 ; 1.321]	[1.367 ; 1.506]	[1.093 ; 1.279]	[1.355 ; 1.618]
ψ_2	[0.814 ; 0.899]	[0.742 ; 0.837]	[0.636 ; 0.698]	[0.784 ; 0.921]
	Model 13	Model 14	Model 15	Model 16
ψ_1	[1.174 ; 1.300]	[0.912 ; 1.224]	[1.268 ; 1.593]	[1.477 ; 1.628]
ψ_2	[0.755 ; 0.826]	[0.613 ; 0.738]	[0.697 ; 0.814]	[0.775 ; 0.894]

Figure 1: Spot prices and convenience yields

Panel A



Panel B





Figure 2: Parameter *posterior* distributions

Figure 3: Bubble process and regime-2 probabilities




Figure 4: Average WTI and Brent regime-processes




Figure 5: Spot prices, convenience yields, and average regime-2 probabilities

