Testing and Modeling Speculative Oil Price Bubbles: US and Global Markets

Luiz Eduardo Rocha* Wilfredo Leiva Maldonado†

March 23, 2026

Abstract

This paper employs the model proposed by Scheinkman and Schechtman (1983), applied to oligopolistic competition among producers of storable goods, to establish the fundamental value of oil prices. Based on this value, we test for speculative bubbles in both Global and US market prices between July 1988 and May 2025. We examine explosive, multiple, periodic, and intrinsic bubbles, ultimately analyzing the dynamics of speculative bubbles under two unobservable regimes. Our findings confirm the existence of multiple bubbles in both markets, while periodic bubbles were found only in the US market. Given the presence of multiple bubbles, the regime-switching dynamics proved robust and consistent with periods where commodity prices deviated from their fundamental value at geometric rates. These results support the use of oil inventories as a key fundamental indicator for oil pricing.

Keywords: WTI oil price; Brent oil price; Rational Bubbles; Fundamental value estimation.
JEL Classification codes: C51, Q41.

*University of São Paulo. E-mail: luizrocha99@usp.br
†Corresponding author. University of São Paulo. E-mail: wilfredo.maldonado@usp.br. This author acknowledges financial support from FAPESP (grant 2023/07274-0), CNPq (308815/2019-0, 402949/2021-8), and FIPE.

1 Introduction

The dynamics followed by oil prices in the United States and globally have shown significant fluctuations, maintaining values at distinct levels for extended periods. Throughout this century, the price per barrel has fluctuated, reaching a maximum value of USD 143.95 in July 2008 (during the American subprime crisis) and a minimum value of USD 9.10 in April 2020 (during the COVID-19 crisis). The difficulty of forecasting and defining the regimes to which these prices adjust makes the analysis of this dynamic both challenging and important, as it involves a commodity that is essential to the functioning of world economies.

The financialization of oil in the early 1980s, which emerged with the launch of West Texas Intermediate (WTI) futures contracts on the New York Mercantile Exchange (NYMEX), introduced greater volatility to the product's price. However, it was in the first decade of the 2000s that major fluctuations began to emerge following the opening of the market to the massive use of "non-commercial" capital. This was further compounded by the deregulation of derivative markets contained in the Commodity Futures Modernization Act (CFMA); by 2008, this caused the trading volume of "paper oil" in financial markets to exceed the value of the global physical oil supply by more than 20 times.

To explain oil value, the pioneering work of Pindyck (1993) adapts the standard present value model to the valuation of storable commodities. He utilizes the convenience yield as the commodity equivalent of a stock's dividends, an approach adopted by many subsequent studies in their analyses. Hamilton (2009) analyzes various theories on oil price formation to explain the high prices observed in 2008. He concludes that commodity speculation, global demand, production limitations, the OPEC monopoly, and scarcity rent may have impacted the observed peaks to varying degrees. Kaufmann and Ullman (2009) investigate the origin of variations in crude oil prices and their transmission to spot and futures prices across different regions of the world. Their results suggest that market fundamentals triggered a long-term increase in oil prices, which was subsequently exacerbated by speculation. In the same year, Kilian (2009) proposed a structural VAR (Vector Autoregression) model, decomposing the real price of crude oil into supply shocks, shocks in global demand for all industrial commodities, and demand shocks specific to the global crude oil market. In doing so, he corrects a common approach used to assess the response of macroeconomic aggregates to exogenous changes in oil prices, which assumes that the price varies while all other variables remain constant.

Despite these early studies having identified fundamental factors in oil price formation, subsequent research emerged analyzing the existence of speculative elements influencing the commodity's value. Kaufmann (2011) argues that the spike and collapse of oil prices in 2007-2008 were consequences of both changes in market fundamentals (Chinese oil demand) and speculative pressures (indicated by a prolonged break in the cointegration relationship between spot prices and long-term futures contracts). Lammerding et al. (2013) utilize the convenience yield approach to approximate the fundamental value of oil prices and test for the presence of speculative bubbles, assuming they follow a two-state Markov process (considering stable and explosive states). They find robust evidence of the existence of speculative bubbles in oil price dynamics during the 2005–2011 period. In a similar vein, Shi and Arora (2012) model oil price bubble dynamics using three distinct frameworks: the three-regime model of Brooks and Katsaris (2005) and the two- and three-regime variants of Schaller and van Norden (2002). They find that the latter modeling offers a better fit to the data and identifies a bubble that lasted from 2008 until early 2009.

In contrast, the study by Alquist and Gervais (2013) concludes that macroeconomic fundamentals, rather than financial speculation, are sufficient to explain the oil price increases between 2003 and 2008. These are related to the rising demand for crude oil from emerging Asia, particularly China and India. Similarly, Kruse-Becher (2025) recently argues that misspecified and poorly measured fundamentals can compromise econometric tests for bubbles. Using the concept of rotated expectations – an econometric approach that alternates between market-based expectations and the assumption that future prices are equal to month-end spot prices – the author finds no evidence of speculative bubbles.

An extremely useful methodology for detecting speculative bubbles is the one developed by Phillips et al. (2015a,b). Both papers provide a recursive testing procedure and a dating algorithm useful for detecting multiple bubble events. The authors propose a right-tailed ADF (Augmented Dickey-Fuller) unit root test with a rolling window and a double-sup window selection criterion, known as the Supremum ADF and Generalized Supremum ADF (SADF and GSADF) tests. This SADF and GSADF methodology was utilized by Umar et al. (2021), alongside a set of econometric tools – including breakpoint unit root tests and probability-based detection mechanisms – to identify bubble events and explore their underlying factors. Similarly, Khan et al. (2021) employ the GSADF test to analyze the determinants of bubble formation in crude oil prices. Their results reveal that global economic expansion, supply-demand imbalances, US dollar depreciation, and OPEC oversupply are the primary drivers. Furthermore, US shale oil production and low demand from emerging economies are identified as leading factors contributing to these price bubbles.

This work contributes to the literature on speculative bubbles in two aspects: one theoretical and the other applied. We adapt the model developed by Scheinkman and Schechtman (1983), which analyzes oligopolistic competition in a durable goods market, to the oil production and marketing sector. Through this adaptation, we are able to derive a fundamental value for this commodity's price that has not yet been proposed in the literature: oil inventories. Specifically, we view the oil production and trade market as an oligopolistic market in which firms must strategically decide on both production levels and the inventories they should maintain. As a result, we find that the fundamental price is a multiple of total oil inventories. Subsequently, we perform two empirical exercises using data from the global and U.S. oil markets. The first tests for the existence of multiple, recurrent, and intrinsic bubbles in these markets. The second models the dynamics of the influence of speculative bubbles on oil price innovations (the difference between the spot price and the corresponding futures price). In this latter exercise, we utilize the modeling developed by Maldonado et al. (2012), which also allows for testing the formation of rational expectations in the oil market.

The results point to the existence of multiple and recurrent bubbles, but not intrinsic bubbles. We also find a robust dynamic (well-fitted to the data) that explains the influence of speculative bubbles on oil price innovations and, consequently, on futures prices. Finally, the hypothesis of rational expectations in the formation of the commodity's futures price is rejected.

The article is divided into six sections. In Section 2, we present the theoretical model used to derive the fundamental value of the oil price. In Section 3, we discuss the data to be used for empirical analysis. The tests for multiple, recurrent, and intrinsic bubbles are presented in Section 4. In Section 5, we present and estimate the non-linear two-regime model for the dependence of oil price innovations. We summarize our conclusions in Section 6.

2 Model

We begin this section with a description of the theoretical model used to define the fundamental price of oil. The model is an adaptation of the one presented by Scheinkman and Schechtman (1983), which analyzes competition within an oligopolistic market of firms that produce and sell a durable good in an infinite-horizon economy. Our adaptation is described below.

There are $N \geq 2$ oligopolistic firms competing in an infinite-horizon economy, producing and supplying a storable good. The market demand is given by $P = P(v)$, where $v \geq 0$ is the total quantity demanded (sold) and $P \geq 0$ is the unit price of the product. $P(\cdot)$ is a strictly decreasing and continuously differentiable function, and we will suppose that $P(v)v$ is a concave function. In each period $t \geq 1$, each firm $n \in \{1, 2, \cdots, N\}$ holds an initial stock $x_{n,(t-1)}$, and chooses its production $q_{n,t}$ and sales $v_{n,t}$ to maximize its total discounted profit. In addition to production costs, the firm also incurs storage costs. Let $c(q_{n,t})$ and $\phi(x_{n,t})$ be the costs incurred by the firm to produce and store the quantities $q_{n,t}$ and $x_{n,t}$ respectively. Both functions c and ϕ are continuously differentiable, strictly increasing, and concave functions. Firms could have heterogeneous production and storage costs; however, such a specification does not significantly alter the model's main results. Finally, $\beta \in (0, 1)$ is the discount factor.

Under these assumptions, the problem for firm n consists of choosing a sequence $\{q_{n,t}, v_{n,t}, v_{n,t}\}_{t \geq 0}$ to maximize its objective function, given its initial inventory $x_{n,(-1)}$:

$$\max \mathbb{E}\left[\sum_{t=0}^{+\infty} \beta^t \{P(v_t)v_{n,t} - c(q_{n,t}) - \phi(x_{n,t})\}\right] \tag{1}$$

subject to the restrictions,

$$v_{n,t} + x_{n,t} = q_{n,t} + x_{n,t-1}$$
$$P_t \equiv P(v_t) = P\left(\sum_{n=1}^{N} v_{n,t}\right).$$

The first constraint in problem (1) represents the balance between the total available quantity of the good (right side) and its allocation between sales and the new stock level (left side). The second constraint is simply the definition of the equilibrium price at time t.

The intratemporal and intertemporal equations for problem (1) are, respectively:

$$c'(q_{n,t}) = P'(v_t)v_{n,t} + P(v_{n,t}) \tag{2}$$

$$P'(v_t)v_{n,t} + P(v_t) + \phi'(x_{n,t}) = \beta\mathbb{E}_t\left[P'(v_{t+1})v_{n,(t+1)} + P(v_{n,(t+1)})\right]. \tag{3}$$

Defining the market share of firm n as $s_{n,t} = v_{n,t}/v_t$ and the price elasticity of demand as $\varepsilon = v \times P'(v)/P(v)$,[1] we can rewrite the intertemporal equation (3) as follows,

$$P(v_t)\left[1 - |\varepsilon|^{-1} \times s_{n,t}\right] + \phi'(x_{n,t}) = \beta\mathbb{E}_t\left[P(v_{t+1})\left[1 - |\varepsilon|^{-1} \times s_{n,t}\right]\right]$$

Summing the above equation over $n = 1, 2, \cdots, N$, calling $P_t = P(v_t,)$ and solving for P_t, we obtain the following expression for the price dynamics of the good:

$$P_t = \beta\mathbb{E}_t\left[P_{t+1}\right] + \frac{\sum_{n=1}^{N}\phi'(x_{n,t})}{|\varepsilon|^{-1} - N}. \tag{4}$$

Equation (4) is a standard equation characterized by a multiplicity of solutions; in particular, it admits price bubble solutions. The transversality condition ensures a unique solution. Consequently, testing for commodity price bubbles amounts to verifying whether the transversality condition holds in the market.

Therefore, solving Equation (4) forward allows us to decompose the price into its fundamental value and a bubble component, as follows:

$$P_t = \sum_{k=0}^{+\infty}\beta^k\left(\frac{\mathbb{E}_t\left[\sum_{n=1}^{N}\phi'(x_{n,t})\right]}{|\varepsilon|^{-1} - N}\right) + \lim_{T\to+\infty}\beta^T\mathbb{E}_t\left[P_{t+T}\right] \equiv P_t^f + B_t. \tag{5}$$

The term $B_t = \lim_{T\to+\infty}\beta^T\mathbb{E}_t\left[P_{t+T}\right]$ in Equation (5) corresponds to the bubble component of the price and satisfies $\mathbb{E}_t[B_{t+T}] = \beta^{-T}B_t$ Therefore, the explosive behavior of the price bubble arises whenever $B_t > 0$. The existence of such bubbles in the oil price and the modeling of their dynamics will be done in Sections 4 and 5 respectively.

It is important to note that the product's fundamental price P_t^f, given in Equation (5) is well-defined as long as $N < |\varepsilon|^{-1}$. For the oil market, Caldara et al.

[1]We will suppose that the price elasticity of demand is constant.

(2019) provide several estimates for the price elasticity of demand for this commodity based on exogenous episodes of significant drops in global oil production. Their study performs three estimations of the absolute value of this parameter, finding values that range between 0.017 and 0.08. Thus, if the value of N is less than $1/0.08 = 12.5$, the fundamental price will be well-defined – a condition that is satisfied for the oil market.

To derive a closed-form expression for the fundamental value of the commodity, we impose the following condition.

Assumption 1: The storage cost is given by $\phi(x) = (A/2)x^2$. Furthermore, the aggregate stock $x_t = \sum_{n=1}^{N} x_{n,t}$ follows a random walk in logarithms:

$$\ln x_t = \mu + \ln x_{t-1} + \varepsilon_t \tag{6}$$

where $\varepsilon_t \sim N(0, \sigma^2)$.

The assumption of a quadratic storage cost captures the increasing marginal costs associated with physical storage constraints and warehouse congestion. Furthermore, quadratic costs ensure the strict concavity of the firm's objective function, providing mathematical tractability and a well-defined unique interior solution for the optimal stock levels. In addition, the presence of a unit root in the $\{\ln x_t\}$ series is easily verified using the available data.

Therefore, using Assumption 1 we can express the fundamental price as a linear function of the oil stocks:

$$P_t^f = \sum_{k=0}^{+\infty} \beta^k \left(\frac{\mathbb{E}_t \left[\sum_{n=1}^{N} \phi'(x_{n,t}) \right]}{|\varepsilon|^{-1} - N} \right) = \frac{A}{|\varepsilon|^{-1} - N} \sum_{k=0}^{+\infty} \beta^k \mathbb{E}_t [x_{t+k}] \tag{7}$$

Furthermore, Equation (6) yields $x_{t+1} = x_t \exp\{\mu + \varepsilon_{t+1}\}$, which implies that the expected value of future stocks satisfies $\mathbb{E}_t[x_{t+k}] = x_t \exp\{k\mu + (k\sigma^2/2)\}$. Substituting this into the expression above, we obtain:

$$P_t^f = \left(\frac{A}{(|\varepsilon|^{-1} - N)(\exp\{-(\ln(\beta) + \mu + \sigma^2/2)\} - 1)} \right) x_t \equiv \kappa x_t, \tag{8}$$

provided that $\ln(\beta) + \mu + \sigma^2/2 < 0$.

In summary, when modeling a commodity market (oil) under oligopolistic competition – where firms face both production and storage costs and maximize the present value of all future profits – the product's fundamental price (the market price minus any potential bubble component) is proportional to the aggregate market stocks. The hypotheses used to derive this result are: quadratic storage costs and a log-stock series that follows a random walk with drift.

3 Data

To test for the existence of oil price bubbles and subsequently model their dynamics, this study uses two primary datasets on oil prices and inventories. The first pertains to the U.S. market, comprising West Texas Intermediate (WTI) spot prices and U.S. petroleum inventories, which include strategic reserves and other liquids. The second pertains to the global market, using Brent spot prices and Organisation for Economic Co-operation and Development (OECD) petroleum inventories as a proxy for global stocks, given the lack of comprehensive world-wide data.

In addition, one-month oil futures prices are obtained from contracts traded on the New York Mercantile Exchange (NYMEX) for WTI and on the Intercontinental Exchange ICE Futures Europe for Brent.

Nominal oil prices are deflated using the monthly U.S. Consumer Price Index (CPI) for All Urban Consumers, sourced from the Federal Reserve Economic Data (FRED) database.

All data are at a monthly frequency, covering the period from July 1988 to May 2025, and were obtained from the U.S. Energy Information Administration (EIA).[2].

Table 1 shows some descriptive statistics on the data, while Figure 1 illustrates the time evolution of spot prices and inventories for the OECD and the United States. The table reports summary measures for inventories, spot prices, and futures prices in both the global (OECD) and U.S. markets. On average, OECD inventories are substantially larger than U.S. inventories, reflecting the broader geographic coverage of the global series. Price and futures series display similar central tendencies across markets, with means around 27–28 dollars per barrel and comparable standard deviations, indicating similar levels of volatility over the sample period.

Table 1: Descriptive Statistics

	Global			United States		
Stat.	OECD Stocks	OECD Price	OECD Futures	US Stocks	US Price	US Futures
Mean	4077.03	27.94	28.16	1708.01	27.34	27.41
Median	4038.64	24.45	25.05	1661.25	23.85	23.92
Minimum	3605.89	7.08	7.55	1482.06	7.66	8.10
Maximum	4745.15	72.10	76.20	2099.43	72.95	76.29
Std. Dev.	265.50	14.49	14.59	144.95	12.79	12.90
JB	22.07***	57.07***	56.14***	54.24***	59.53***	59.06***

Notes: Oil spot and futures prices for both the Global (Brent) and U.S. (WTI) markets are denominated in U.S. Dollars per barrel, while inventory levels for the OECD and U.S. are measured in millions of barrels (MMb). This table reports descriptive statistics (sample mean, median, minimum, maximum, standard deviation, and the Jarque-Bera test statistic with H_0 of normality). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

[2]https://www.eia.gov/petroleum/.

Figure 1: US and OECD spot prices and inventories

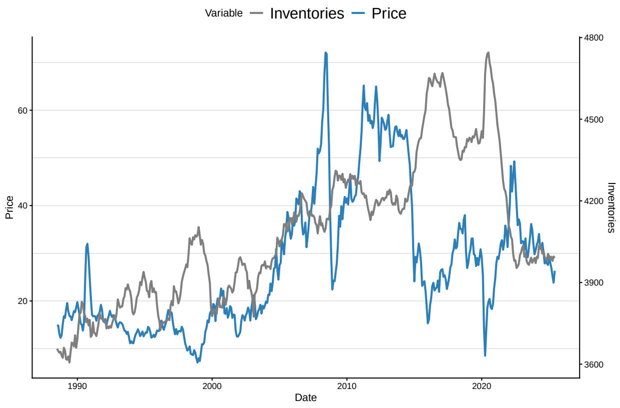

(a) OECD

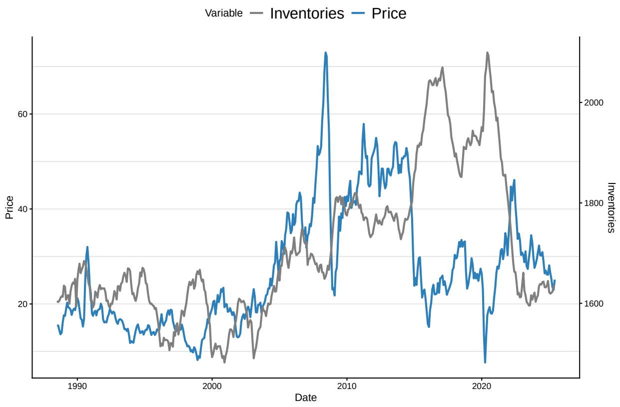

(b) US

4 Testing Rational Bubbles

Testing for the existence of speculative bubbles in an asset's price is equivalent to testing whether the observed price exhibits deviations from the proposed fundamental value that grow at a geometric rate over time. The literature presents various types of speculative bubbles and specific tests for each; consequently, evidence for one or more types of bubbles may exist within the same dataset. For this reason, in this section, we proceed to test for four types of bubbles: explosive, multiple, periodically collapsing, and intrinsic. In each of the following subsections, we describe the respective tests and the results of their applications.

4.1 Explosive Bubbles

The explosive bubble test follows a sequence of steps originally proposed by Diba and Grossman (1984) and Diba and Grossman (1988) to assess whether asset prices deviate from their fundamental values in a way that is consistent with the presence of rational explosive bubbles. The intuition is that if prices are determined solely by fundamentals, they should not exhibit explosive autoregressive behavior — that is, their stochastic process should not grow at a rate faster than the market discount rate. Thus, the detection of explosive behavior in prices, particularly when fundamentals do not share the same behavior, is a strong indicator of the presence of a bubble.

The testing strategy proceeds in three stages. First, unit root tests are applied to both the price and fundamental series to verify whether they share the same order of integration. Second, cointegration tests examine whether prices and fundamentals maintain a stable long-run equilibrium relationship, since a rejection of cointegration implies that prices diverge from fundamentals in a manner inconsistent with a purely fundamental-driven process. Third, the stationarity and explosiveness of the ratio $y_t = P_t/x_t$ are assessed using both the standard ADF and the right-tailed ADF (RTADF) tests: under the null hypothesis of no bubble, y_t should be stationary; under the alternative, the bubble component B_t grows explosively, causing y_t to diverge.

Initially, we conducted Augmented Dickey-Fuller (ADF) unit root tests to verify the order of integration of both the oil price series and the aggregate inventory series. The general model from which we depart consists of the regression proposed by Dickey and Fuller (1979) and Dickey and Fuller (1981):

$$\Delta y_t = \alpha + \beta t + \gamma y_{t-1} + \sum_{i=1}^{k} \psi^i \Delta y_{t-i} + \nu_t \tag{9}$$

where y_t represents the series under analysis - either oil prices or inventories - α and β are the drift and trend constants, respectively, ψ^i are the coefficients of the i lagged differences, γ is the parameter of interest, and ν_t is a white noise error term. The null hypothesis $H_0 : \gamma = 0$ is tested against the one-sided alternative $H_1 : \gamma < 0$, where rejection of the null indicates the absence of a unit root and, consequently, stationarity of the series.

To identify the order of integration of the WTI and Brent crude oil price series, as well as the U.S. and OECD crude oil inventory series, we follow the sequential testing procedure proposed by Enders (2014). This procedure applies the ADF test starting from the most general specification - including both a trend component and an intercept - and progressively moves toward simpler models, ensuring that the appropriate deterministic components are included at each stage. Lag length selection follows the Modified Akaike Information Criterion (MAIC), as proposed by Ng and Perron (2001), which provides more reliable lag selection in the presence of highly persistent or near-unit-root processes.

The results of the unit root tests for the series in levels and first differences are reported in Table 2 below.

Table 2: Unit root test results

Variable	Model	ADF Level	ADF First Difference	Integration
Panel A: U.S.				
WTI	*drift + trend*	-2.40	-14.34***	$I(1)$
WTI	*drift*	-2.21	-14.36***	$I(1)$
WTI	-	-0.81	-14.37***	$I(1)$
Stocks	*drift + trend*	-1.94	-4.73***	$I(1)$
Stocks	*drift*	-1.75	-4.73***	$I(1)$
Stocks	-	-0.22	-4.74***	$I(1)$
Panel B: Global				
Brent	*drift + trend*	-2.26	-14.60***	$I(1)$
Brent	*drift*	-2.06	-14.62***	$I(1)$
Brent	-	-0.79	-14.64***	$I(1)$
Stocks	*drift + trend*	-1.99	-6.83***	$I(1)$
Stocks	*drift*	-2.07	-6.79***	$I(1)$
Stocks	-	0.18	-5.57***	$I(1)$

Notes: This table reports the results for the Augmented Dickey-Fuller (ADF) unit root tests. The column Model indicates the deterministic components included in the test regression: *drift + trend*, *drift* only, or none (-). Level and First Difference columns report the test statistics for the series in their level and differenced forms, respectively. Integration column indicates the order of integration for each series. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level, indicating the rejection of the null hypothesis of a unit root.

The unit root test results for the price and inventory series in levels indicate the presence of a unit root across both markets and at all stages of the sequential testing procedure. For the series in first differences, however, the null hypothesis of a unit root is rejected at all stages, supporting the alternative of stationarity. These results allow us to conclude that both oil spot prices and inventories are integrated of order one in both markets.

Next, the cointegration tests of Engle and Granger (1987) and Johansen (1995) will be used to verify the existence of a long-run relationship between the oil price and inventory series. If the series share a long-run equilibrium relationship, this constitutes evidence that prices do not systematically deviate from their fundamentals, allowing us to rule out the presence of explosive bubbles.

The Engle and Granger (1987) cointegration test is conducted based on the auxiliary regression

$$P_t = \beta_0 + \beta_1 x_t + \mu_t \tag{10}$$

which represents the long-run equilibrium relationship between oil prices and inventories. The null hypothesis of no cointegration is evaluated by applying the Phillips and Perron (1988) unit root procedure to the estimated residuals $\{\hat{\mu}_t\}_{t=1}^{T}$, where the resulting test statistic is compared against the critical values tabulated by Phillips and Ouliaris (1990).

To complement the cointegration analysis, we additionally employ the procedure of Johansen (1995). The lag order of the VAR model in its error correction representation is initially selected based on the Bayesian Information Criterion

(BIC). Starting from this initial lag length, alternative specifications with additional lags are examined using Lagrange Multiplier (LM) tests for residual autocorrelation. For both the US and global oil markets, a lag order of 1 is selected. Finally, the trace statistic and the maximum eigenvalue statistic are computed and compared against the critical values tabulated by Osterwald-Lenum (1992).

Results for both cointegrations tests are presented in Table 3.

Table 3: Cointegration test results

		Engle-Granger		Johansen				
Market	Cointegration	Stat.	90% Crit. Val.	H_0	λ_{trace}	90% Crit. Val.	λ_{eigen}	90% Crit. Val.
Global	No	-2.86	-2.57	$r = 0$	14.02	15.66	10.89	12.91
				$r \leq 1$	3.12	6.50	3.12	6.50
U.S.	No	-2.93	-2.57	$r = 0$	13.75	15.66	11.82	12.91
				$r \leq 1$	1.94	6.50	1.94	6.50

Notes: This table reports results for the Engle-Granger and Johansen cointegration tests between oil prices and aggregate inventories. Engle-Granger column reports the ADF statistic on the residuals of the cointegrating regression (10). The column Johansen reports the Trace (λ_{trace}) and Maximum Eigenvalue (λ_{eigen}) statistics, where r denotes the number of cointegrating vectors. Cointegration is determined to be present only if both testing frameworks reject the null hypothesis of no cointegration at the 10% significance level.

Although the test statistics from the Engle and Granger (1987) procedure reach the tabulated critical value at the 10% significance level, we require, in this work, that both tests jointly indicate the presence of cointegration in order to conclude that the series are cointegrated. Since the test statistics from the Johansen (1995) procedure - both the trace and maximum eigenvalue statistics - do not reach their respective tabulated critical values, we fail to reject the null hypothesis of no cointegrating vectors. Therefore, the series do not cointegrate in either market, meaning that no long-run equilibrium relationship is found between oil prices and inventories. This result is consistent with the possible presence of explosive bubbles in oil prices.

The final stage consists of evaluating the stationarity and explosiveness of the oil spot prices and price-to-inventory ratios in both markets. We first estimate model (9) to test $H_0 : \gamma = 0$ against the left-sided alternative $H_1 : \gamma < 0$, which corresponds to the standard ADF test for stationarity. Still based on model (9), the explosiveness of these processes is then assessed using the right-tailed ADF (RTADF) test, which maintains the same null hypothesis of $H_0 : \gamma = 0$ but shifts the alternative to the right side, $H_1 : \gamma > 0$, thereby testing for an autoregressive root greater than unity.

Failure to reject the null hypothesis of the standard ADF test combined with failure to reject the null of the RTADF test implies that the spot prices and prices-to-inventory ratios are stationary, which constitutes evidence against the presence of explosive bubbles. Conversely, rejection of the null hypothesis of the RTADF test provides evidence that these processes are explosive, indicating the presence of an explosive bubble episode.

Table 4: ADF and Right-tailed ADF test results

	ADF		RTADF	
Variable	Stat.	90% Crit. Val.	Stat.	90% Crit. Val.
Brent	-2.26	-3.13	-2.01	-0.37
WTI	-2.40	-3.13	-2.15	-0.37
Ratio (Brent)	-2.41	-3.13	-2.07	-0.37
Ratio (WTI)	-2.62	-3.13	-2.23	-0.37

Notes: This table reports the results for the standard Augmented Dickey-Fuller (ADF) test and the Right-tailed ADF (RTADF) test based on Equation (9). ADF tests the null of a unit root ($H_0 : \gamma = 0$) against the stationary alternative ($H_1 : \gamma < 0$). RTADF tests the same null against the explosive alternative ($H_1 : \gamma > 0$). Rejection of the RTADF null occurs if the test statistic is greater than the critical value. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

The results reported in Table 4 indicate that, for all series and in both markets, the standard ADF test statistics do not reach the tabulated critical value at the 10% significance level, meaning that we fail to reject the null hypothesis of a unit root. This confirms that spot oil prices and price-to-inventory ratios are non-stationary processes, consistent with the unit root findings reported in the previous stage.

Turning to the RTADF test, none of the test statistics exceed the 90% critical value of -0.37, so we also fail to reject the null hypothesis against the explosive alternative in any of the series. Taken together, the absence of explosive behavior in both spot prices and price-to-inventory ratios - for the US and global oil markets - constitutes evidence against the presence of explosive bubble episodes over the sample period analyzed. Although we confirm the absence of explosive bubbles in oil prices and price-to-inventory ratios, the presence of unit root may be consistent with different types of bubble processes.

4.2 Multiple Bubbles

To test for the presence of multiple bubbles, we use the methodology developed by Phillips et al. (2011) and Phillips et al. (2015a,b), which we summarize below.

They consider the following moving-window ADF regression for an interval starting in a fraction r_1 of the total sample size T and ending in a fraction r_2 of T:

$$y_t = \mu + \delta y_{t-1} + \sum_{i=1}^{k} \psi_{r_w}^{i} \delta y_{t-i} + \varepsilon_t, \tag{11}$$

where y_t denotes a time series (either the oil spot price or the price/stock ratio), $r_w = r_2 - r_1$ is the window size, $k \in \mathbb{N}_+$ is the number of lags. The terms μ, δ and ψ^i (for $i = 1, \ldots, k$) represent the regression coefficients. The null hypothesis is that of a unit root, $H_0 : \delta = 1$, against the alternative of explosive behavior $H_1 : \delta > 1$. The resulting test statistic is computed as $ADF_{r_1}^{r_2}$.

Given this regression, the initial suggestion of Phillips et al. (2011) is to estimate this statistic by successively expanding the window size r_w from an initial value r_0 up to 1, which would be the total fraction of the sample. By setting $r_1 = 0$, the supremum of the sequence of statistics $ADF_0^{r_2}$, called *SADF*, is defined by:

$$SADF(r_0) = \sup_{r_2 \in [r_0, 1]} \left\{ ADF_0^{r_2} \right\}. \tag{12}$$

Phillips et al. (2015a) proposed an extension of the *SADF* to better suit the testing mechanism in series with multiple periods of explosive behavior. Specifically, starting with the minimum window size $r_0 \in (0,1)$, estimate the regression (11) for all possible subsamples of size r_0 or larger, while allowing both the starting and the ending point to vary in the ranges $r_1 \in [0, r_2 - r_0]$ and $r_2 \in [r_0, 1]$. They define as *GSADF* the resulting statistic of the supremum of the sequence of statistics $ADF_{r_1}^{r_2}$ in this double recursive procedure, or:

$$GSADF(r_0) = \sup_{\substack{r_2 \in [r_0, 1] \\ r_1 \in [0, r_2 - r_0]}} \left\{ ADF_{r_1}^{r_2} \right\}. \tag{13}$$

If the null hypothesis of no bubbles is rejected, we may proceed to record their origination and collapse periods. Both the *SADF* and *GSADF* methodologies allow the identification of exuberance episodes, the first based on the procedure of Phillips et al. (2011) and the latter on that of Phillips et al. (2015a). In the latter, the authors suggest employing a double recursive procedure with a flexible window, refining the approach proposed in Phillips et al. (2011). The idea is to perform the *SADF* procedure described above, but on a backward expanding sample sequence. Accordingly, we compute the sequence of statistics $ADF_{r_1}^{r_2}$ specifically fixing the end point of each window as r_2 while allowing the starting point to vary from 0 to $r_2 - r_0$, with the supremum of this sequence of statistics being defined as the backward *SADF*:

$$BSADF_{r_2}(r_0) = \sup_{r_1 \in [0, r_2 - r_0]} \left\{ ADF_{r_1}^{r_2} \right\}. \tag{14}$$

Finally, we can compute the episodes of exuberance by comparing the values of the *BSADF* statistic series with the critical values of the *SADF* test. Defining the bubble origination date T_{r_e} as the first observation in which the *BSADF* test statistic exceeds the *SADF* critical value, and the termination date T_{r_f} as the first observation at which it falls below that value, the points r_e and r_f can be computed as follows:

$$\hat{r}_e = \inf_{r_2 \in [r_0, 1]} \left\{ r_2 : BSADF_{r_2}(r_0) > scv_{r_2}^{\alpha} \right\},$$
$$\hat{r}_f = \inf_{r_2 \in [\hat{r}_e, 1]} \left\{ r_2 : BSADF_{r_2}(r_0) < scv_{r_2}^{\alpha} \right\},$$

where $scv_{r_2}^{\alpha}$ is the $100(1-\alpha)\%$*SADF* critical value for a significance level α and sample size r_2 of T.

The initial window r_0 selected in this study follows the recommendation of Phillips et al. (2015a), which is based on the total sample size T and takes the functional form $r_0 = 0.01 + (1.8/\sqrt{T})$. In addition, the 90% critical values for the date-stamping strategy are computed using the bootstrap resampling procedure introduced by Phillips and Shi (2020), to ensure that the statistics converge to the distributions reported in Phillips et al. (2015a).

Given the statistics and the corresponding critical values for each point in time, we can plot these series and establish a visual reference for the periods of explosiveness.

First, we report the results of the GSADF tests in Table 5 for the Brent and WTI oil price series and their respective ratios to inventories.

Table 5: GSADF test results

Variable	Bubble	Stat.	90% Crit. Value
Brent	Yes	3.55***	1.94
WTI	Yes	3.98***	1.94
Ratio (Brent)	Yes	3.27***	1.94
Ratio (WTI)	Yes	3.66***	1.94

Notes: This table reports the GSADF test results based on the recursive procedure of Phillips et al. (2015a). Column Stat. refers to the supremum ADF statistic. Bubble column indicates a rejection of the unit root null hypothesis in favor of the explosive alternative at the 10% level or better. Critical values and p-values are obtained via bootstrap resampling procedure (Phillips and Shi, 2020) with 500 bootstrap repetitions. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

As can be seen in Table 5, all the GSADF test statistics are significant at the 99% level. These results indicate that we can reject the null hypothesis of the presence of a unit root in favor of the alternative, providing evidence of the presence of multiple bubbles in both the U.S. and Global markets.

After identifying the presence of multiple bubbles within the analyzed series, we can proceed to date the specific periods of exuberance. These periods correspond to those in which the computed BSADF test statistics exceed the 90% critical values. The detailed time-stamping results from this procedure are reported in Appendix A. In Figure 2, we overlay the price and ratio series for both markets with these periods to identify how the dynamics of these processes are associated with explosive episodes.

Roughly speaking, an analysis of the exuberance episodes in prices (the first two charts in Figure 2) reveals the emergence of speculative bubbles during short intervals between October 2004 and September 2006, and again for more sustained periods from July 2007 to October 2008. In the first period, accelerated urbanization in China and India drove demand for the commodity to such an extent that global supply struggled to keep pace. In addition, Hurricane Katrina in 2005 damaged the refining infrastructure in the Gulf of Mexico, limiting the supply available for inventory buildup. Furthermore, tensions surrounding Iran's nuclear program and instability in Nigeria increased the product's risk premium. Regarding the later bubbles, the U.S. sub-prime mortgage crisis — which began in July 2007 and peaked in 2008 — prompted many investors to flee equities and real estate in favor of commodities, as they viewed oil as a “safe haven” and a hedge against inflation. This pushed the commodity’s price to a historic high of approximately $147 in July 2008.

On the other hand, the sharp declines observed starting in early 2015 for Brent crude — and by the end of the same year for WTI — were due to the maturation of the shale oil era in the United States. The development of fracking allowed the U.S. to drastically increase production, flooding the global market with millions of additional barrels per day. Consequently, in November 2014, Saudi Arabia and OPEC decided to maintain (or even increase) production levels to defend their market share. In addition to these supply increases, which neutralized price speculation above fundamental value, China transitioned from an industry-led economy to a service-oriented one, thereby reducing its energy consumption intensity. In Section 5, we will see that in all episodes where bubbles burst, the oil price always trends toward our estimated fundamental value.

Figure 2: Episodes of Exuberance (GSADF)

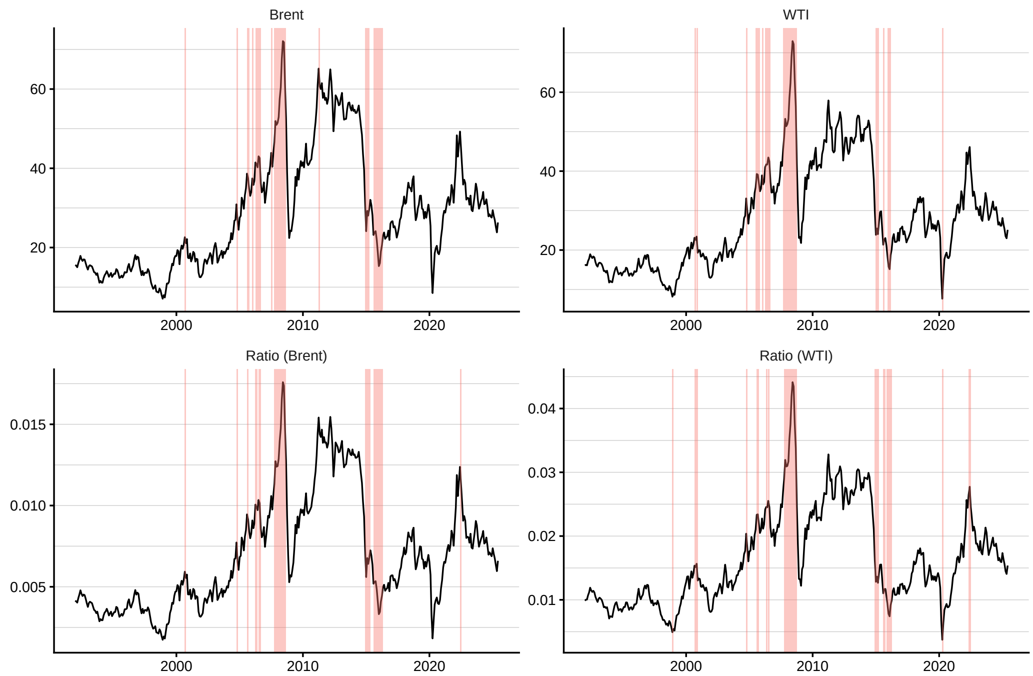

Finally, the last two charts in Figure 2 report the exuberance episodes in the price-to-inventory ratios. This is done to perform a preliminary analysis of whether the fundamental proposed in Equation (8) allows for the elimination of some bubbles observed in prices alone. As the charts show and Table 10 in Appendix A corroborates, many of the bubble episodes (especially the short-term ones) are reduced. This demonstrates that the proposed fundamental value is capable of explaining certain price movements without them being attributed to speculative shifts.

4.3 Periodically Collapsing Bubble

A third type of bubble we will test is the parametric, periodically collapsing bubble proposed by Evans (1991). These bubbles initially grow at a low geometric rate until they reach a size at which speculation intensifies at a higher rate. Finally, after reaching a predefined threshold, they collapse, initiating the dynamics anew.

To test this type of bubble, we employ the Threshold Autoregressive (TAR) and Momentum Threshold Autoregressive (MTAR) models introduced by Enders and Siklos (2001) . These models allow for cointegration testing while accounting for asymmetry in the adjustment of variables toward an equilibrium path. Furthermore, we apply specific tests to identify asymmetries within two forms of long-run adjustments between variables.

To implement these tests, we first estimate the cointegration relation with OLS:

$$P_t = \hat{\beta}_0 + \hat{\beta}_1 x_t + \hat{\mu}_t \tag{15}$$

then, we use the estimated residual deviations relative to a threshold (the attractor) to capture the dynamics of interest. This is reflected in the following error-correction regression:

$$\Delta\mu_t = I_t \rho_1 \mu_{t-1} + (1 - I_t)\rho_2 \mu_{t-1} + \sum_{i=1}^{p-1} \gamma_i \Delta\mu_{t-i} + \nu_t \tag{16}$$

I_t being an indicator function that in the TAR model is defined as:

$$I_t = \begin{cases} 1, \text{if } \mu_{t-1} \geq \tau \\ 0, \text{if } \mu_{t-1} < \tau \end{cases} \tag{17}$$

where τ denotes the threshold taken as a reference for the deviations.

Equation (16) represents an alternative specification of the error-correction model, which incorporates two correction regimes for $\Delta\mu_t$ that allow different adjustment speeds, ρ_1 and ρ_2, for deviations above and below the long-run equilibrium, respectively. For the mechanism to constitute a valid error correction, both ρ_1 and ρ_2 must be less than or equal to zero, with at least one strictly negative. By adopting this more general error-correction specification, it becomes possible to test for the presence of periodically collapsing bubbles in oil prices within a cointegration framework.

The MTAR model is also based on the regression (16), but modifies the indicator function I_t as follows:

$$I_t = \begin{cases} 1, \text{if } \Delta\mu_{t-1} \geq \tau \\ 0, \text{if } \Delta\mu_{t-1} < \tau \end{cases} \tag{18}$$

Enders and Siklos (2001) argue that, for many economic applications, it is convenient to set $\tau = 0$ so that the cointegration vector coincides with the attractor. However, it is also possible to estimate the consistent versions of the TAR and MTAR models (TARC and MTARC) by obtaining a consistent estimate of the threshold value τ. This consistent value is estimated using a selection criterion that relies on the residuals (the first difference of the residuals in the case of the MTAR model) to iteratively estimate regression (16), minimizing the sum of squared errors — as proposed by Enders and Siklos (2001). In this study, we estimate equation (16) using both the consistent threshold value and assuming it equal to zero.

To select the threshold corresponding to the long-run equilibrium of each market, we first use the estimated residuals of the long-run relationship (15) following the methodology of Chan (1993). Then, the estimated residuals μ_t $(\Delta\mu_t)$ are sorted in ascending order, and the highest and lowest 1,5% values are discarded[3]. Among the remaining 97%, the τ that minimizes the sum of squared residuals from the regression estimation (16) is then selected.

According to Bohl (2003), the models described by equation (16) are able to capture the characteristic dynamics of periodic bubbles, in which rapid price increases relative to dividends are followed by a crash. This pattern is consistent with the presence of asymmetry in the different types of adjustment in (16), such that the corrections $\rho_1\mu_{t-1}$ for deviations μ_t $(\Delta\mu_t)$ above the threshold τ are stronger compared to the corrections $\rho_2\mu_{t-1}$ for deviations below the threshold.

In general, Bohl (2003) also emphasize that MTAR-type models are more suitable for identifying periodically collapsing bubbles due to their compatibility with the rapid adjustments that these bubbles typically exhibit. However, the estimated thresholds for the MTARC models in the highlighted oil markets were negative, which is not consistent with the theoretical model of periodic bubbles proposed by Evans (1991), in which bubble collapses are partial down to a positive value.[4] Adopting a negative τ would also imply a very low tolerance for the expansion of the type of bubble we aim to capture, since even small negative variations in bubble size $\Delta\mu_{t-1}$ could activate the adjustment mechanism that bursts the bubble.

In this regard, we test for the presence of periodically collapsing bubbles using the TAR and TARC adjustments, which have proven to be consistent with the characteristics of the bubbles described by Evans (1991).

To test the null hypothesis of no periodically collapsing bubbles, we follow a three-step procedure: (i) we test for no cointegration ($H_0 : \rho_1 = \rho_2 = 0$) using the F_{coint} statistic; (ii) we test the null of symmetric adjustment ($H_0 : \rho_1 = \rho_2$) using the F_{sym} statistic; and (iii) we verify that the estimates of ρ_1 and ρ_2 are non-positive and satisfy the condition $|\rho_1| > |\rho_2|$. If conditions (i), (ii), and (iii) are satisfied, we find evidence of periodically collapsing bubbles.

Table 6 below reports the estimation results of the TAR and TARC models,

[3]Enders and Siklos (2001) perform this procedure by discarding the top and bottom 15% of μ_t $(\Delta\mu_t)$ values. However, in this study, we adopt a 1,5% cutoff, as it yields a lower sum of squared errors for the processes under analysis and is therefore more consistent with the methodology of Chan (1993).

[4]Bohl (2003) also highlight this feature of periodically collapsing bubbles in Evans (1991), which are always positive and never change sign.

along with the values of the corresponding test statistics for cointegration (F_{coint}) and symmetry (F_{sym})[5].

Table 6: TAR and TARC estimation results

Model	Market	Bubble	τ	ρ_1	ρ_2	Lag	F_{coint}	F_{sym}
TAR	Global	No	0	-0.028**	-0.031*	1	4.55**	0.02
	U.S.	No	0	-0.030**	-0.029*	1	4.66**	0.00
TARC	Global	Yes	30.53	-0.076***	-0.021**	1	7.65***	4.95**
	U.S.	Yes	26.42	-0.080***	-0.019*	1	8.79***	6.77**

Notes: This table reports estimation results for the Threshold Autoregressive (TAR) and Consistent Threshold Autoregressive (TARC) models from Eq. (16). τ is the estimated threshold. ρ_1 and ρ_2 represent the adjustment speeds in the two regimes. F_{coint} is the test for the null of no cointegration. F_{sym} is the test for the null of symmetric adjustment ($H_0 : \rho_1 = \rho_2$). *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Focusing first on the TAR model estimation results, we find that the parameter estimates for the different adjustment regimes, ρ_1 and ρ_2, are negative in all cases. Only for the U.S. market, $|\rho_1|$ is greater than $|\rho_2|$, thus satisfying the preliminary conditions of the test. In addition, the null hypothesis of no cointegration is rejected at the 5% level for both markets. However, the null hypothesis of symmetric adjustment cannot be rejected for either case.

Regarding the results of the consistent TAR (TARC) model, we find that the parameters ρ_1 and ρ_2 are negative in all cases, with $|\rho_1|$ greater than $|\rho_2|$. Furthermore, for both markets, we reject both the null hypothesis of no cointegration and the null hypothesis of symmetric adjustment.

We can summarize the results as follows. It is not possible to provide evidence of periodically collapsing bubbles using the standard TAR model in either the global or the U.S. market. However, the consistent estimation (TARC) for both markets satisfies the necessary conditions to reject the null hypothesis of no periodically collapsing bubbles. Therefore, we do not rule out the existence of this type of bubble in oil prices.

4.4 Intrinsic Bubble

Lastly, we test the presence of intrinsic bubbles as in Froot and Obstfeld (1991). This type of rational bubble differs from the others presented above in the sense that the rapid deviations exhibited by asset prices with respect to the fundamental value are attributed to changes in their fundamentals alone. In the context of a commodity market such as oil, aggregate stock accumulation can generate scarcity in spot markets, which by itself may lead to overreactions in oil prices.

Thus, testing for the presence of intrinsic bubbles in oil markets requires us to write a nonlinear function of aggregate stocks that represents the bubble compo-

[5] Calculations were performed in R using the apt package, version 4.0.

nent in price decomposition (5). The authors define this function as follows

$$B_t \equiv b x_t^{\lambda}. \tag{19}$$

As stated in Section 2, the above component must satisfy the bubble equation, which means that

$$\mathbb{E}_t[B_{t+1}] = \beta^{-1} B_t \leftrightarrow \mathbb{E}_t[b x_{t+1}^{\lambda}] = \beta^{-1} b x_t^{\lambda}. \tag{20}$$

This relation also implies certain conditions that the parameter λ must satisfy. Recalling Assumption 1, it is known that $x_{t+1} = x_t \exp\{\mu + \varepsilon_{t+1}\}$, where $\varepsilon_t \sim N(0, \sigma^2)$, by raising both sides of the equality to the power λ, we have

$$x_{t+1}^{\lambda} = x_t^{\lambda} \exp\{\lambda\mu + \lambda\varepsilon_{t+1}\}$$

and conditioning on the information available at time t:

$$\mathbb{E}_t[x_{t+1}^{\lambda}] = x_t^{\lambda} \exp\left\{\lambda\mu + \frac{\lambda^2\sigma^2}{2}\right\}. \tag{21}$$

Substituting the conditional expectation term $\mathbb{E}_t[x_{t+1}^{\lambda}]$ from (20) into the left-hand side of (21) yields

$$\beta^{-1} = \exp\left\{\lambda\mu + \frac{\lambda^2\sigma^2}{2}\right\}, \tag{22}$$

i.e. λ must be the positive root of the quadratic equation $\lambda\mu + \frac{\lambda^2\sigma^2}{2} + \log\beta = 0$. By summing the fundamental price (8) and the bubble component (19) we get the present value solution for oil prices with an intrinsic bubble:

$$P_t = \kappa x_t + b x_t^{\lambda} \tag{23}$$

where κ and b are both arbitrary constants. Using the inequality $\ln(\beta) + \mu + \sigma^2/2 < 0$ from Section 2 and the quadratic characteristic equation from this Section, we can also infer that λ must exceed 1.

Equation (23) summarizes the main characteristics of intrinsic bubbles presented above. For $b \neq 0$, the transversality condition is violated, allowing solutions such as (5); however, in this case, oil spot prices depend only on the aggregate stocks of the market. In addition, the fact that λ is greater than 1 causes the explosive behavior that leads spot prices to depart from the fundamental value following changes in fundamentals.

Testing for the presence of intrinsic bubbles requires estimating the parameters κ, b and λ from equation (23). Our empirical approach to perform this estimation follows the suggestions of Froot and Obstfeld (1991), in which this equation is divided by x_t and the error term is specified as an $AR(1)$ process, thus avoiding collinearity among regressors:

$$\frac{P_t}{x_t} = c_0 + c_1 (x_t)^{\lambda-1} + \eta_t;\ \eta_t = \rho\eta_{t-1} + \upsilon_t,\ \upsilon_t \sim iid(0, \sigma_{\upsilon}^2) \tag{24}$$

The Maximum-likelihood estimates of (24) are presented in Table 7. Testing the null hypothesis of no intrinsic bubbles is equivalent to testing that $c_0 = \kappa$ and $c_1 = 0$. Meanwhile, the alternative hypothesis of the presence of intrinsic bubbles corresponds to $c_0 = \kappa$ and $c_1 > 0$.

Table 7: Maximum likelihood estimation of (24) results

Market	Bubble	c_0	c_1	$(\lambda - 1)$	ρ	σ_υ
Global	No	0.0049***	0.0012	0.0257***	0.9788***	0.0006
		(0.0011)	(0.0044)	(0.0059)	(0.0085)	(0.0000)
U.S.	No	0.0098***	0.0037	0.0513***	0.9716***	0.0014
		(0.0033)	(0.0065)	(0.0076)	(0.0095)	(0.0000)

*Notes:*This table reports Maximum Likelihood Estimation (MLE) results for Equation (24). Column 2 indicates the outcome of the hypothesis test $H_0 : c_1 = 0$. c_0 and c_1 are the fundamental and bubble coefficients. $(\lambda - 1)$ is the estimated exponent of the nonlinear term. ρ and σ_υ are the $AR(1)$ error parameters. Standard errors are in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

The estimation results presented in Table 7 indicate a lack of empirical support for intrinsic bubbles in both Global and U.S. oil markets. Although the estimated λ parameters satisfy the theoretical condition for convexity ($\lambda > 1$), the bubble coefficients (c_1) are statistically non significant. Consequently, we conclude that while oil prices are highly sensitive to fundamental stock levels, they do not exhibit the specific nonlinear overreactions characterized by the Froot and Obstfeld (1991) intrinsic bubble model. The high value of the AR(1) term (ρ) further suggests that any deviations from fundamentals are better characterized by persistent stochastic shocks rather than a deterministic rational bubble component.

5 Regime-switching Bubble

Although the methodology employed in the previous section for detecting multiple bubbles is well-established in the literature and widely used in applied research, it suffers from two limitations. The first concerns the power constraints of unit root tests; sample size can often be a limiting factor for the conclusions drawn from these tests. Furthermore, these tests are designed to detect geometric deviations from a predefined fundamental value without providing an actual estimate of that value.

Therefore, after applying the bubble detection tests in Section 4 and confirming the presence of multiple and periodically collapsing bubbles in the analyzed oil prices, we proceed to model the bubble dynamics based on oil price innovations, following the methodology proposed by van Norden (1996) and extended by Maldonado et al. (2012).

In that study, the authors modeled bubble size dynamics using a regime - switching model where the dependence of asset price innovations on the bubble size is governed by unobservable regimes. This approach not only provides greater flexibility to the bubble's growth dynamics, but also allows for the estimation of the fundamental value level, which, as shown in Equation (8) of Section 2, contains an indeterminate scale factor. Next, we describe the model derivation and the corresponding estimation procedure.

Suppose that bubble dynamics depend on two unobservable regimes represented by the stochastic process $\{r_t\}_{t\geq 0}$, where $r_t = 0$ represents the bubble collapse regime (C) and $r_t = 1$ represents the bubble survival regime (S). Furthermore, we define $q_{t+1} \equiv \Pr(r_{t+1} = 1)$ as the probability that the process remains in regime S at time $t+1$.

Assumption 2: Given these regimes, there exist two functions u and q satisfying:

$$\mathbb{E}_t[B_{t+1}|r_{t+1} = 0] = u(B_t) \tag{25}$$

$$q_{t+1} = q(B_t) \tag{26}$$

Assumption 2 establishes the relationship between the unobservable regimes and the bubble size. Specifically, Equation (25) governs how the regime affects the evolution of the bubble size, while Equation (26) specifies how the bubble size influences the probability of each regime occurring.

Applying the law of iterated expectations $E_t[B_{t+1}] = q_{t+1}E_t[B_{t+1}|r_{t+1} = 1] + (1-q_{t+1})E_t[B_{t+1}|r_{t+1} = 0]$, and Assumption 2, we can compute the expected bubble size under the survival regime. Specifically, we can isolate the term $E_t[B_{t+1}|r_{t+1} = 1]$ from this identity to obtain:

$$\mathbb{E}_t[B_{t+1}|r_{t+1} = 1] = \frac{\mathbb{E}_t[B_{t+1}]}{q_{t+1}} - \frac{(1-q_{t+1})}{q_{t+1}}\mathbb{E}_t[B_{t+1}|r_{t+1} = 0]$$

By substituting Equations (25) and (26) into the expression above, we get the following:

$$\mathbb{E}_t[B_{t+1}|r_{t+1} = 1] = \frac{\mathbb{E}_t[B_{t+1}]}{q(B_t)} - \left(\frac{1}{q(B_t)} - 1\right)u(B_t).$$

Finally, we can substitute the bubble equation $B_t = \beta\mathbb{E}_t[B_{t+1}]$ into the previous equation and obtain the expected bubble size under the regime S:

$$\mathbb{E}_t[B_{t+1}|r_{t+1} = 1] = \frac{B_t}{\beta q(B_t)} - \frac{u(B_t)}{q(B_t)} + u(B_t). \tag{27}$$

Define oil price innovation as the gain (or loss) realized by a market participant when trading a futures contract. Assuming that futures price of a barrel of oil quoted at time t is equal to the expected future spot price, $\mathbb{E}_t[P_{t+1}]$, the innovation is expressed as:

$$R_{t+1} = P_{t+1} - \mathbb{E}_t[P_{t+1}]. \tag{28}$$

Analogously, we can define fundamental innovation as $R^f_{t+1} = P^f_{t+1} - \mathbb{E}_t[P^f_{t+1}]$. By applying these definitions and the price decomposition $P_{t+1} = P^f_{t+1} + B_{t+1}$, we can establish a relationship between price innovations and bubble size. First, we substitute the price decomposition into the innovation expression given in (28):

$$R_{t+1} = P^f_{t+1} + B_{t+1} - \mathbb{E}_t[P_{t+1}],$$

subsequently, we express the fundamental price in terms of the fundamental innovation,

$$R_{t+1} = R^f_{t+1} + \mathbb{E}_t[P^f_{t+1}] + B_{t+1} - \mathbb{E}_t[P_{t+1}].$$

Since $\mathbb{E}_t[P^f_{t+1}] - \mathbb{E}_t[P_{t+1}] = -\mathbb{E}_t[B_{t+1}] = -\frac{B_t}{\beta}$, we finally obtain:

$$R_{t+1} = R^f_{t+1} + B_{t+1} - \frac{B_t}{\beta}$$

Then, taking the expected value of the innovation conditional on the collapse and survival states and using Equation (25) and (27), we get

$$\mathbb{E}_t[R_{t+1}|r_{t+1} = 0] = \mathbb{E}_t[R^f_{t+1}|r_{t+1} = 0] + u(B_t) - \frac{B_t}{\beta}$$
$$\mathbb{E}_t[R_{t+1}|r_{t+1} = 1] = \mathbb{E}_t[R^f_{t+1}|r_{t+1} = 1] + \frac{B_t}{\beta q(B_t)} - \frac{u(B_t)}{q(B_t)} + u(B_t) - \frac{B_t}{\beta}$$

As in Maldonado et al. (2012), we assume that when prices are aligned with fundamentals, the expected gains from futures market operations are zero, regardless of whether there is information indicating that the speculative bubble is in a collapse or survival regime; that is,

Assumption 3: The regimes do not affect the expected value of the fundamental component such that:

$$\mathbb{E}_t[R^f_{t+1}|r_{t+1} = 0] = \mathbb{E}_t[R^f_{t+1}|r_{t+1} = 1] = 0$$

Thus, given Assumption 3, the expected value of the innovation, conditional on each regime, can be written as:

$$\mathbb{E}_t[R_{t+1}|r_{t+1} = 0] = u(B_t) - \frac{B_t}{\beta} \tag{29}$$

$$\mathbb{E}_t[R_{t+1}|r_{t+1} = 1] = \frac{B_t}{\beta q(B_t)} - \frac{u(B_t)}{q(B_t)} + u(B_t) - \frac{B_t}{\beta} \tag{30}$$

Finally, if we consider a linear function to represent $u(B) = a + bB$ and a logit function for $q(B) = \left[1 + \exp\{\beta_{q0} + \beta_{q2}B^2\}\right]^{-1}$, the right-hand sides of equations (21) and (22) will be a linear function and a combination of linear and quadratic-exponential functions, respectively. In this manner, we obtain a two-regime regression that can be estimated using the available data. We specify the details in the following subsection.

5.1 Empirical Model

To estimate the two-regime model resulting in the previos subsection, let us assume:

$$\varepsilon^S_{t+1} = R_{t+1} - \mathbb{E}_t[R_{t+1}|r_{t+1}=1],\ \varepsilon^S_t \sim N(0,\sigma^2_S)$$
$$\varepsilon^C_{t+1} = R_{t+1} - \mathbb{E}_t[R_{t+1}|r_{t+1}=0],\ \varepsilon^C_t \sim N(0,\sigma^2_C)$$

Thus, we can write (29) and (30) as the following regime-switching regression:

$$R_{t+1} = \begin{cases} \beta^C_0 + \beta^C_1 B_t + \varepsilon^C_{t+1} & \text{with prob. } (1-q(B_t)) \\ \beta^S_{00} + (\beta^S_0 + \beta^S_1 B_t)\exp\{\beta^S_2 B^2_t\} + \varepsilon^S_{t+1} & \text{with prob. } q(B_t) \end{cases} \tag{31}$$

where $q(B) = \left[1+\exp\{\beta_{q0}+\beta_{q2}B^2\}\right]^{-1}$. We can estimate the parameters in (31) using the method of Maximum Likelihood (ML). Therefore, we maximize the following function:

$$\mathbb{L}(\theta)=\sum_{t=1}^{T}\ln\left[(1-q(B_t))\phi\left(\frac{R_{t+1}-\beta^C_0-\beta^C_1 B_t}{\sigma^2_C}\right)+q(B_t)\phi\left(\frac{R_{t+1}-\beta^S_{00}-(\beta^S_0+\beta^S_1 B_t)\exp\{\beta^S_2 B^2_t\}}{\sigma^2_S}\right)\right] \tag{32}$$

where $\phi(\cdot)$ denotes the probability density function of a standard normal distribution. The bubble size B_t is defined as the difference between the spot price of oil and the fundamental price, where the latter is determined by total inventories x_t multiplied by a scale factor κ, namely, $B_t = P_t - \kappa x_t$. Oil price innovations are constructed according to Equation (28), using the prices of oil futures contracts traded at time t and maturing in the following month (f_t) as a proxy for the expected value of P_{t+1}, i.e., $R_{t+1} = P_{t+1} - f_t$.

Thus, in addition to the parameters of the regime-switching regression, it is also possible to jointly estimate the parameter κ from the bubble-size equation, using data on spot and futures oil prices, as well as global and U.S. oil inventories.

Estimation by the Maximum Likelihood method allows us to analyze whether more restricted specifications of the model provide a better fit to the data, using the Likelihood Ratio (LR) test. In addition, we can also use a specification that allows us to test the rational expectations hypothesis in oil futures price formation, following Maldonado et al. (2012).

In this manner, we can test whether the more general model is more appropriate than the following specifications:

- Two linear regimes

  Innovations in both regimes are linear functions of the bubble size. This is equivalent to saying that linear approximations for both regimes are sufficient for modeling the dynamics.

  $\beta^S_{00} = 0;\ \beta^S_2 = 0$

- Linear regimes with equal slopes

In the case of two linear models, the innovation response to changes in bubble size is identical across regimes.

$\beta_{00}^S = 0; \beta_2^S = 0; \beta_1^C = \beta_1^S$

- Single linear regime

  A single-regime linear specification provides a sufficient fit to the data, obviating the need for a two-regime model.

  $\beta_{00}^S = 0; \beta_2^S = 0; \beta_1^C = \beta_1^S; \beta_0^C = \beta_0^S$

- Rational expectations hypothesis

  Test if the expectations regarding the innovations are rational.

  $\beta_{00}^S = 0; \beta_{q2} = \beta_2^S; \beta_1^S + \beta_1^C \exp\{\beta_{q0}\} = 0; \beta_0^S + \beta_0^C \exp\{\beta_{q0}\} = 0$

Table 8 reports the estimation results for the parameters of the regime-switching equation (31).

Table 8: Maximum likelihood estimates of the parameters

Market	κ	Logit		Survival Regime					Collapse Regime		
		β_{q0}	β_{q2}	β_{00}^S	β_0^S	β_1^S	β_2^S	σ_S	β_0^C	β_1^C	σ_C
Global	0.0042***	0.144	0.039**	0.790***	-0.460***	-0.048***	0.053***	0.584*	-0.075	-0.016**	2.101***
	(0.0002)	(0.958)	(0.019)	(0.134)	(0.024)	(0.004)	(0.002)	(0.298)	(0.165)	(0.008)	(0.057)
U.S.	0.0094***	-0.915**	0.011***	0.155	0.035***	-0.002***	0.031***	1.280***	0.238	-0.026*	2.505***
	(0.0005)	(0.416)	(0.002)	(0.110)	(0.002)	(0.0004)	(0.002)	(0.087)	(0.311)	(0.013)	(0.070)

Notes: This table reports MLE results for the regime-switching model in Eq. (31). κ is the constant that defines the level of the fundamental oil price. Logit parameters (β_{q0}, β_{q2}) govern the regime transition probabilities. Survival Regime parameters describe the explosive bubble phase, while Collapse Regime parameters describe the price correction phase. σ_S and σ_C represent regime-specific variance. Standard errors are in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

It is important to note that most parameters are significant at the 95% level or higher. Furthermore, the estimates for the global market and the U.S. market are statistically significantly different. This is not only due to the relative size of inventories within the explanatory variable, but also because innovations respond differently to the bubble size in each market. The parameter β_{q2} has a sign consistent with the prediction of a logit function ($\beta_{q2} > 0$), and it is significantly different from the parameter β_2^S. This indicates an absence of rational expectation formation in the oil futures market, as will be tested further below.

Regarding the response of innovations to increases in the bubble size, we observe that this response is negative in the collapse regime; that is, increases in the bubble size imply decreases in the gains from oil futures market operations. In the survival regime, the effect of the bubble on the innovations depends on its magnitude. Since $\beta_1^S < 0$ and $\beta_2^S > 0$ in both markets, the effect is positive for large bubble sizes, whereas the opposite holds true for small bubble sizes.

It is important to note that the scale factor κ can help us compare the fundamental prices of both markets. Considering that the global market could have a higher market power compared to the U.S., one might expect that its fundamental price reflects that. This condition is satisfied if $\kappa^{global} \times x_t^{global} > \kappa^{us} \times x_t^{us}$, i.e.

the ratio of scale factors κ exceeds the inverse ratio of the aggregate inventories ($\kappa^{global}/\kappa^{us} > x_t^{us}/x_t^{global}$). Using the estimates from Table 8, the ratio of scale factors is 0.4468, while the ratio of the sample means for inventories from Table 1 is 0.4189. The fact that the scale factor ratio surpasses the inventory ratio confirms that the Global fundamental price is higher, consistent with the greater market significance of the Brent benchmark.

Using the estimated parameters β_{q0} and β_{q2} reported in Table 8, we can plot the corresponding logit functions for each market. Figure 3 illustrates the sensitivity of the survival probability to changes in the bubble size. The results suggest that the survival regime is more resilient to bubble expansion in the U.S. market than in the global market.

Figure 3: Probability of the survival regime

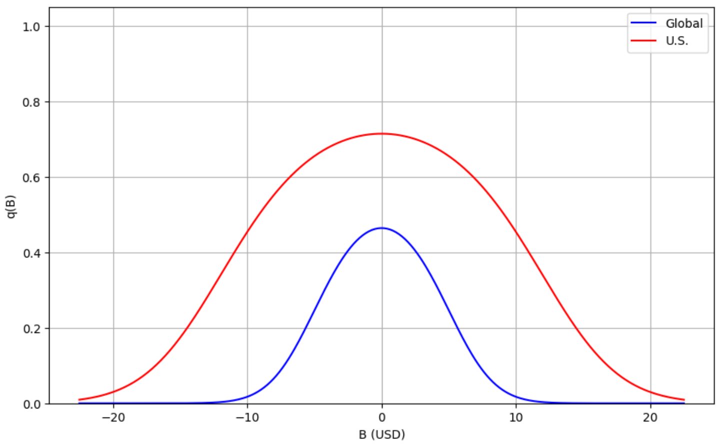

To verify the robustness of the general model relative to specifications with fewer parameters, we apply the log-likelihood ratio test to the aforementioned specifications. The test consists of calculating the log-likelihood with the parameters of the restricted model ($L_{restricted}$) and comparing it with that of the unrestricted model ($L_{unrestricted}$). Thus, $\lambda = -2(L_{restricted} - L_{unrestricted})$ follows a $\chi^2(p)$ distribution, where p is the difference in the number of parameters between the two models.

The test results presented in Table 9 point to the same conclusions in both markets. At a 10% significance level, the specification of a linear model with two regimes cannot be rejected; however, the restrictions of having identical slopes or consisting of only a single regime are indeed rejected. Although a two-regime linear specification could be considered, it would not be useful to testing the rational expectations hypothesis in oil futures pricing. The general model we employ does allow for such a test, and we find that the rational expectations hypothesis cannot be rejected in either market. This leads us to conclude that the informational structure in these markets is sufficient for an accurate one-month-ahead futures price forecast.

Table 9: Likelihood Ratio Test Results

Market	Null hypothesis	Restrictions	Log-likelihood	Restricted log-likelihood	$\chi^2(n)$	p-value
Global	Two linear regimes	2	-891.507	-893.590	4.164	12.4%
	Linear regimes with identical slopes	3	-891.507	-896.508	10.001	1.8%
	One linear regime	5	-891.507	-896.854	10.692	5.7%
	Rational Expectations	5	-891.507	-892.348	1.681	89.1%
U.S.	Two linear regimes	2	-857.386	-859.544	4.316	11.5%
	Linear regimes with identical slopes	3	-857.386	-860.601	10.47	1.4%
	One linear regime	5	-857.386	-862.061	9.351	9.5%
	Rational Expectations	5	-857.386	-861.468	8.164	14.7%

Notes: This table reports likelihood ratio (LR) tests comparing the unrestricted regime-switching specification with alternative restricted models. The LR statistic is computed as $\lambda = -2(L_{restricted} - L_{unrestricted})$, where $L_{unrestricted}$ and $L_{restricted}$ denote the log-likelihood values of the unrestricted and restricted models, respectively. The column "Restrictions" indicates the number of parameter restrictions imposed under the null hypothesis. Under the null, the test statistic is asymptotically distributed as a χ^2 distribution with degrees of freedom equal to the number of restrictions. The reported p-values correspond to this distribution. Rejection of the null hypothesis indicates that the unrestricted regime-switching specification provides a significantly better fit to the data.

Figure 4: Spot and fundamental prices X Episodes of exuberance

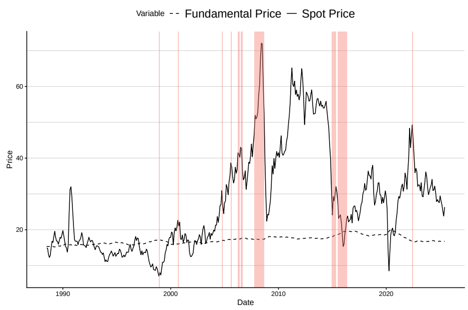

(a) Level series - Global

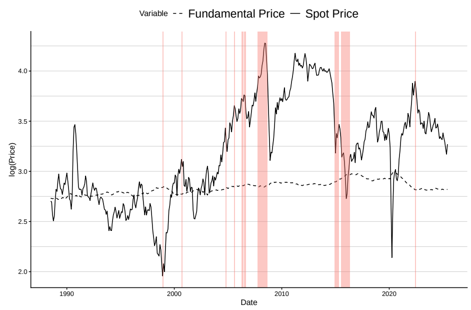

(b) Logarithmic series - Global

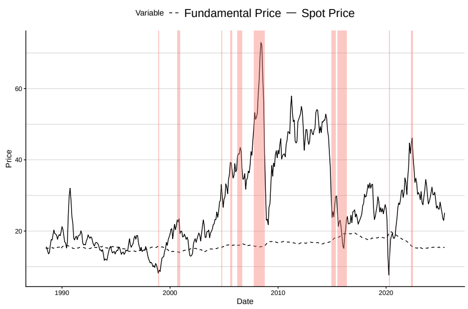

(c) Level series - U.S.

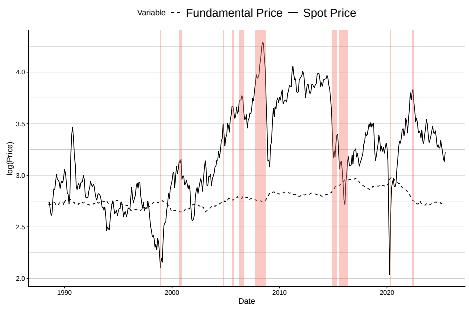

(d) Logarithmic series - U.S.

6 Conclusions

In this study, we utilize the oligopolistic competition model for firms producing a storable good, as presented by Scheinkman and Schechtman (1983), to define a fundamental value for oil prices. Under several simple hypotheses, the model proposes a multiple of the aggregate inventory levels of the commodity as its fundamental value.

Next, using spot price and inventory data for the Global and U.S. markets, we tested four types of speculative bubbles in oil prices: explosive, multiple, recurrent, and intrinsic. Finally, using futures price data, we modeled the joint dynamics between bubble size and price innovation (the difference between the spot price and the futures contract price for the same period).

For the first two types of bubbles, we employed the unit root versus explosive tests developed by Phillips et al. (2011) and Phillips et al. (2015a,b). These tests were applied to both the oil price series and the price-to-inventory ratio series. The results indicated periods of price exuberance, which decrease when using the price-to-inventory ratios; this suggests that including inventories in the fundamentals eliminates several periods previously thought to be driven by speculation.

Given the evidence of multiple bubbles, we proceeded to analyze two alternative dynamics. The first assumes that they follow a periodically collapsing behavior, as proposed by Evans (1991). Using the methodology of Enders and Siklos (2001), we did not reject the hypothesis of the non-existence of such bubbles in either market. Lastly, fitting the dynamics of bubble size as dependent on fundamentals (inventories)—intrinsic bubbles—provided no evidence of their influence on speculative episodes.

Once the existence and types of bubbles were identified, we performed a final analysis to study the influence of bubbles following unobservable regimes on returns in oil futures markets. This analysis allows us not only to estimate the commodity's fundamental price level for each period and market, but also to test futures price formation under the rational expectations hypothesis. In general, the results pointed to a fundamental price in both markets, to which spot prices invariably return whenever a bubble bursts. Additionally, the fundamental price in the Global market is higher than that of the U.S. market, reflecting greater market power on the global stage. Finally, the rational expectations hypothesis in oil futures price formation was not rejected in either market.

Although our research shares similarities with Shi and Arora (2012), it diverges in three key aspects. Our model: *(i)* employs a distinct *proxy* for the fundamental value; *(ii)* eschews linear approximations for the regimes to allow for testing the rational expectations hypothesis; and *(iii)* utilizes a more general specification, leading to the rejection of the linear approximations used in their study.

Certain additional aspects may be analyzed in future research. The heterogeneity of oligopolists across different regions may allow for local dispersion in prices and storage costs. Furthermore, the variability of interest rates offered by U.S. Treasuries and other major blocs over long periods challenges the constant discount factor hypothesis, suggesting it should be considered stochastic and variable. Lastly, the observed bubble dynamics across the entire period suggest that, beyond exponential growth and bursting, phases of bubble maintenance and slow decay should be included. All these aspects must be taken into account when defining local policies for resilience against sudden oil price surges—a commodity that remains indispensable to national economies.

A Time-Stamping Procedure

This appendix presents the detailed time-stamping results obtained from the GSADF procedure. The method identifies periods of explosive behavior by comparing the backward SADF statistics with their corresponding critical values.

Figure 5 reports these results for the series considered in the analysis. The solid blue line represents the backward sup ADF (BSADF) statistic, while the dashed red line corresponds to the 90% critical value obtained from the GSADF procedure. Periods in which the BSADF statistic exceeds the critical value indicate episodes of explosive behavior. These intervals are highlighted by the shaded vertical bands in each panel.

The full results of the GSADF time-stamping procedure are presented in Appendix Table A1, which reports the start, peak, and termination dates of each episode of explosive dynamics for Brent, WTI, and the price–inventory ratios.

Figure 5: Time-Stamping procedure for Oil Prices and Ratios

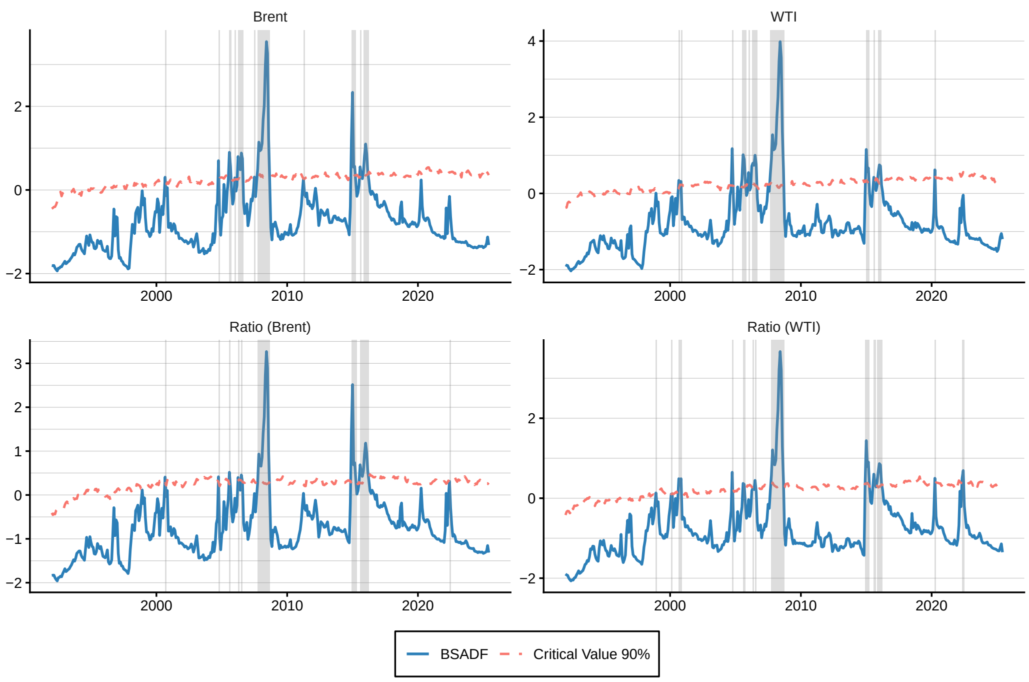

Table 10: GSADF Time-Stamping Results

Start	Peak	End	Duration	Signal
Panel A: Brent				
2000-09	2000-09	2000-10	1	Positive
2004-10	2004-10	2004-11	1	Positive
2005-08	2005-08	2005-10	2	Positive
2006-01	2006-01	2006-02	1	Positive
2006-04	2006-07	2006-09	5	Positive
2007-07	2007-07	2007-08	1	Positive
2007-10	2008-06	2008-09	11	Positive
2011-04	2011-04	2011-05	1	Positive
2014-12	2015-01	2015-04	4	Negative
2015-08	2015-08	2015-09	1	Positive
2015-11	2016-01	2016-04	5	Negative
Panel B: WTI				
2000-09	2000-09	2000-10	1	Positive
2000-11	2000-11	2000-12	1	Positive
2004-10	2004-10	2004-11	1	Positive
2005-07	2005-08	2005-11	4	Positive
2006-01	2006-01	2006-02	1	Positive
2006-04	2006-07	2006-09	5	Positive
2007-09	2008-06	2008-10	13	Positive
2015-01	2015-01	2015-04	3	Positive
2015-08	2015-08	2015-09	1	Positive
2015-12	2016-01	2016-03	3	Negative
2020-04	2020-04	2020-05	1	Positive
Panel C: Brent Price–Inventory Ratio				
2000-09	2000-09	2000-10	1	Positive
2004-10	2004-10	2004-11	1	Positive
2005-08	2005-08	2005-09	1	Positive
2006-04	2006-04	2006-05	1	Positive
2006-07	2006-07	2006-08	1	Positive
2007-10	2008-06	2008-09	11	Positive
2014-12	2015-01	2015-05	5	Negative
2015-08	2016-01	2016-04	8	Negative
2022-06	2022-06	2022-07	1	Positive
Panel D: WTI Price–Inventory Ratio				
1998-12	1998-12	1999-01	1	Positive
2000-02	2000-02	2000-03	1	Positive
2000-09	2000-09	2000-12	3	Positive
2004-10	2004-10	2004-11	1	Positive
2005-08	2005-08	2005-10	2	Positive
2006-05	2006-05	2006-06	1	Positive
2006-07	2006-07	2006-08	1	Positive
2007-10	2008-06	2008-10	12	Positive
2014-12	2015-01	2015-04	4	Negative
2015-08	2015-08	2015-10	2	Positive
2015-11	2016-01	2016-04	5	Negative
2020-04	2020-04	2020-05	1	Positive
2022-05	2022-06	2022-07	2	Positive

Notes: The table reports the periods of explosive behavior identified by the GSADF time-stamping procedure. Start, peak, and end denote the estimated beginning, maximum, and termination of each episode. Duration is measured in months. Positive (negative) signals correspond to upward (downward) explosive dynamics.

References

Alquist, R. and O. Gervais (2013). The role of financial speculation in driving the price of crude oil. *The Energy Journal 34*(3), 35—54.

Bohl, M. T. (2003). Periodically collapsing bubbles in the US stock market? *International Review of Economics & Finance 12*(3), 385–397.

Brooks, C. and A. Katsaris (2005). A three-regime model of speculative behaviour: Modelling the evolution of the s&p 500 composite index. *The Economic Journal 115*(505), 767–797.

Caldara, D., M. Cavallo, and M. Iacoviello (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics 103*, 1–20.

Chan, K. S. (1993). Consistency and limiting distribution of the least squares estimator of a threshold autoregressive model. *The Annals of Statistics 21*(1), 520–533.

Diba, B. and H. Grossman (1984). Rational bubbles in the price of gold. NBER Working Papers 1300, National Bureau of Economic Research, Inc.

Diba, B. T. and H. I. Grossman (1988). Explosive rational bubbles in stock prices? *The American Economic Review 78*(3), 520–530.

Dickey, D. A. and W. A. Fuller (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association 74*(366), 427–431.

Dickey, D. A. and W. A. Fuller (1981). Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica 49*(4), 1057–1072.

Enders, W. (2014). *Applied Econometric Time Series*. Wiley Series in Probability and Statistics. Wiley.

Enders, W. and P. L. Siklos (2001). Cointegration and threshold adjustment. *Journal of Business Economic Statistics 19*(2), 166–176.

Engle, R. F. and C. W. J. Granger (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica 55*(2), 251–276.

Evans, G. W. (1991). Pitfalls in testing for explosive bubbles in asset prices. *The American Economic Review 81*(4), 922–930.

Froot, K. A. and M. Obstfeld (1991). Intrinsic bubbles: The case of stock prices. *The American Economic Review 81*(5), 1189–1214.

Hamilton, J. D. (2009). Understanding crude oil prices. *The Energy Journal 30*(2), 179–206.

Johansen, S. (1995, 12). *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford University Press.

Kaufmann, R. K. (2011). The role of market fundamentals and speculation in recent price changes for crude oil. *Energy Policy 39*, 105—-115.

Kaufmann, R. K. and B. Ullman (2009). Oil prices, speculation, and fundamentals: Interpreting causal relations among spot and futures prices. *Energy Economics 31*, 550–558.

Khan, K., C.-W. Su, M. Umar, and X.-G. Yue (2021). Do crude oil price bubbles occur? *Resources Policy 71*, 101936.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *The American Economic Review 99*(3), 1053–1069.

Kruse-Becher, R. (2025). Let's switch again! testing for speculative oil price bubbles based on rotated market expectations. *Finance Research Letters 78*, 107116.

Lammerding, M., P. Stephan, M. Trede, and B. Wilfling (2013). Speculative bubbles in recent oil price dynamics: Evidence from a bayesian markov-switching state-space approach. *Energy Economics 36*, 491–502.

Maldonado, W. L., O. A. Tourinho, and M. Valli (2012). Exchange rate bubbles: Fundamental value estimation and rational expectations test. *Journal of International Money and Finance 31*(5), 1033–1059.

Ng, S. and P. Perron (2001). Lag length selection and the construction of unit root tests with good size and power. *Econometrica 69*(6), 1519–1554.

Osterwald-Lenum, M. (1992). A note with quantiles of the asymptotic distribution of the maximum likelihood cointegration rank test statistics. *Oxford Bulletin of Economics and Statistics 54*(3), 461–472.

Phillips, P. and S. Ouliaris (1990). Asymptotic properties of residual based tests for cointegration. *Econometrica 58*(1), 165–93.

Phillips, P. C. and S. Shi (2020). Chapter 2 - real time monitoring of asset markets: Bubbles and crises. In H. D. Vinod and C. Rao (Eds.), *Financial, Macro and Micro Econometrics Using R*, Volume 42 of *Handbook of Statistics*, pp. 61–80. Elsevier.

Phillips, P. C. B. and P. Perron (1988). Testing for a unit root in time series regression. *Biometrika 75*(2), 335–346.

Phillips, P. C. B., S. Shi, and J. Yu (2015a). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *International Economic Review 56*(4), 1043–1078.

Phillips, P. C. B., S. Shi, and J. Yu (2015b). Testing for multiple bubbles: Limit theory of real-time detectors. *International Economic Review 56*(4), 1079–1133.

Phillips, P. C. B., Y. Wu, and J. Yu (2011). Explosive behavior in the 1990s nasdaq: When did exuberance escalate asset values? *International Economic Review 52*(1), 201–226.

Pindyck, R. S. (1993). The present value model of rational commodity pricing. *The Economic Journal 103*(418), 511–530.

Schaller, H. and S. van Norden (2002). Fads or bubbles? *Empirical Economics 27*, 335–362.

Scheinkman, J. A. and J. Schechtman (1983). A simple competitive model with production and storage. *The Review of Economic Studies 50*(3), 427–441.

Shi, S. and V. Arora (2012). An application of models of speculative behaviour to oil prices. *Economics Letters 115*(3), 469–472.

Umar, M., C.-W. Su, S. K. A. Rizvi, and O.-A. Lobont (2021). Driven by fundamentals or exploded by emotions: Detecting bubbles in oil prices. *Energy 231*, 120873.

van Norden, S. (1996). Regime switching as a test for exchange rate bubbles. *Journal of Applied Econometrics 11*(3), 219–251.