THEMATIC AREA: Quantitative Methods Applied to Market Dynamics
DOI: https://doi.org/10.12660/fgvric.v7.n1.93646

A SECOND-ORDER ORDINARY DIFFERENTIAL EQUATION FOR STUDYING PRICE DYNAMICS IN MARKET EQUILIBRIUM

Name of author: Nathan Vanoukia
Name of the advisor: Patricia Leal da Cunha

ABSTRACT
This paper studies Brent crude oil price dynamics using a second-order ordinary differential equation framework, with a geopolitical shock represented by a conflict dummy. Monthly Brent price data from January 2003 to January 2026 are used to calibrate the model through two approaches: a continuous-time nonlinear optimization method and a discrete-time least-squares method derived from finite-difference approximations. The continuous approach is theoretically appealing but proved numerically unstable: the optimization converged to the imposed parameter bounds, so its parameters are not identified and are reported only as a numerical diagnostic. The discrete approach is computationally stable and yields a positive discriminant, which formally corresponds to an overdamped regime. The calibrated system, however, lies close to the boundary of stability. The slow characteristic root is approximately -0.07, implying a weak restoring force and a half-life of adjustment of roughly ten months, while the fast root falls outside the range in which the monthly finite-difference approximation is reliable and is therefore not interpreted. The paper concludes that the framework provides a useful structured description of price adjustment, but that its quantitative results should be read as illustrative rather than conclusive.

Keywords: second-order differential equations; characteristic roots; damped adjustment; Brent crude oil; price dynamics.

RESUMO
Este artigo analisa a dinâmica dos preços do petróleo Brent utilizando um modelo de equações diferenciais ordinárias de segunda ordem, incorporando um choque geopolítico representado por uma variável dummy de conflito. Foram utilizados dados mensais dos preços do Brent, de janeiro de 2003 a janeiro de 2026, para calibrar o modelo por meio de duas abordagens: um método de otimização não linear em tempo contínuo e um método de mínimos quadrados em tempo discreto, derivado de aproximações por diferenças finitas. A abordagem contínua, embora teoricamente atraente, apresentou instabilidade numérica: a otimização convergiu para os limites impostos aos parâmetros, impedindo sua identificação. Assim, seus resultados são apresentados apenas como diagnóstico numérico. A abordagem discreta demonstrou estabilidade computacional e produziu um discriminante positivo, formalmente correspondente a um regime superamortecido. Entretanto, o sistema calibrado encontra-se próximo ao limite de estabilidade. A raiz característica lenta é aproximadamente -0,07, indicando uma força restauradora fraca e uma meia-vida de ajustamento de aproximadamente dez meses. Já a raiz rápida situa-se fora do intervalo em que a aproximação mensal por diferenças finitas é confiável, não sendo, portanto, interpretada. Conclui-se que o modelo oferece uma descrição estruturada útil do ajustamento dos preços, embora seus resultados quantitativos devam ser interpretados como ilustrativos, e não conclusivos.

Palavras-chave: equações diferenciais de segunda ordem; raízes características; ajustamento amortecido; petróleo Brent; dinâmica de preços.

1 INTRODUCTION

Oil prices play a central role in market equilibrium, inflation, transportation costs, and global production decisions. Because oil is both a commodity and a strategic input, its price responds not only to standard supply-and-demand conditions but also to geopolitical disruptions. Previous research shows that oil-price shocks can have important macroeconomic effects and that not all oil shocks have the same causes or consequences (HAMILTON, 2009; KILIAN, 2009). One recent example is the Russia–Ukraine conflict, which produced a major shock to global energy markets and increased interest in dynamic models capable of describing how prices adjust after geopolitical disruptions (INTERNATIONAL ENERGY AGENCY, 2023).

This paper studies Brent crude oil price dynamics using a second-order ordinary differential equation (ODE). The objective is not only to estimate a dynamic adjustment process, but also to interpret whether the estimated system behaves in an overdamped, critically damped, or underdamped way. In other words, the paper asks whether oil prices appear to return to equilibrium smoothly or with oscillations after a shock.

Two empirical approaches are used. The first estimates the continuous-time ODE directly by nonlinear optimization. The second discretizes the model and estimates the resulting difference equation by ordinary least squares. Comparing the two approaches is useful because the continuous specification is theoretically elegant, whereas the discrete specification is often more stable computationally. At the same time, discretization introduces approximation error. Using both approaches therefore shows whether the continuous model can be calibrated directly from monthly data, or whether the discrete approximation is needed in practice.

The remainder of the paper is organized as follows. Section 2 reviews the mathematical foundations of second-order ODEs. Section 3 presents the methodology, including the data, the conflict dummy, the continuous model, the discretized model, and the estimation strategies. Section 4 reports the empirical results. Section 5 discusses the economic interpretation of the results, relates them to the literature on estimating continuous-time models, and sets out the limitations of the model. Section 6 concludes.

2 THEORETICAL BACKGROUND: SECOND-ORDER ODES AND DYNAMIC REGIMES

This section introduces the mathematical ideas used later in the paper. The objective is not to provide a complete review of differential equations, but to explain the specific concepts needed to analyze the Brent crude oil price model. In particular, the section focuses on second-order linear ordinary differential equations, characteristic roots, and the interpretation of overdamped, critically damped, and underdamped adjustment regimes.

2.1 Second-Order Linear Differential Equations

A second-order ordinary differential equation involves the second derivative of an unknown function. In the context of this paper, the unknown function is the Brent crude oil price $P(t)$, where t represents time. A second-order model is useful because it can describe not only the direction of price movement, but also the acceleration or deceleration of price adjustment.

A general second-order linear ODE can be written as

$$A(x)y''(x) + B(x)y'(x) + C(x)y(x) = f(x),$$

where $y(x)$ is the unknown function, $y'(x)$ is its first derivative, and $y''(x)$ is its second derivative. The functions $A(x)$, $B(x)$, $C(x)$, and $f(x)$ are given, with $A(x) \neq 0$.

The equation is called homogeneous when the forcing term is zero:

$$A(x)y''(x) + B(x)y'(x) + C(x)y(x) = 0.$$

It is called non-homogeneous when the forcing term is not zero:

$$A(x)y''(x) + B(x)y'(x) + C(x)y(x) = f(x), \qquad f(x) \not\equiv 0.$$

This distinction is important for the present paper because the Brent price model contains both an internal adjustment mechanism and an external shock term. The internal adjustment mechanism is studied through the homogeneous equation, while the conflict dummy acts as a non-homogeneous forcing term.

2.2 Characteristic Equation and Root Cases

For a constant-coefficient homogeneous second-order differential equation

$$Ay''(t) + By'(t) + Cy(t) = 0, \qquad A \neq 0,$$

the standard characteristic-equation method assumes a solution of the form $y(t) = e^{rt}$ (BOYCE; DIPRIMA, 2001; THE OPEN UNIVERSITY, 2013). Since $y'(t) = re^{rt}$ and $y''(t) = r^2e^{rt}$, substitution gives

$$Ar^2e^{rt} + Bre^{rt} + Ce^{rt} = 0.$$

Factoring out e^{rt} gives

$$e^{rt}(Ar^2 + Br + C) = 0.$$

Since $e^{rt} \neq 0$, the characteristic equation is

$$Ar^2 + Br + C = 0.$$

The roots are

$$r = \frac{-B \pm \sqrt{B^2 - 4AC}}{2A},$$

where the discriminant is

$$D = B^2 - 4AC.$$

The three possible cases are summarized below.

Root cases for a second-order homogeneous ODE

Condition	Root type	General solution
$D > 0$	$r_1 \neq r_2$, real	$y(t) = C_1e^{r_1t} + C_2e^{r_2t}$
$D = 0$	$r_1 = r_2 = r$, repeated	$y(t) = (C_1 + C_2t)e^{rt}$
$D < 0$	$r = \alpha \pm i\beta$, complex	$y(t) = e^{\alpha t}[C_1\cos(\beta t) + C_2\sin(\beta t)]$

In the distinct real root case, the solution is a sum of two exponential terms and usually represents non-oscillatory adjustment. In the repeated root case, the solution includes the additional factor t, producing the critically damped case. In the complex root case, the sine and cosine terms create oscillations, while the exponential factor $e^{\alpha t}$ controls whether those oscillations grow or decay over time.

2.3 Application to Brent Price Dynamics

The root cases in Section 2.2 are used to classify the Brent crude oil price model. After rewriting the model in terms of deviations from equilibrium, the homogeneous equation becomes

$$y'' + ay' + by = 0.$$

This corresponds to the general form $Ay'' + By' + Cy = 0$, with $A = 1$, $B = a$, and $C = b$. Therefore, the general discriminant $D = B^2 - 4AC$ becomes

$$\Delta = a^2 - 4b.$$

If $\Delta > 0$, the roots are real and distinct, and the price adjustment is overdamped, meaning that prices return toward equilibrium without oscillations. If $\Delta = 0$, the roots are repeated and the adjustment is critically damped, representing the boundary between smooth and oscillatory convergence. If $\Delta < 0$, the roots are complex and the adjustment is underdamped, meaning that prices may overshoot and oscillate around equilibrium before stabilizing.

Thus, the estimated values of a and b determine whether Brent prices adjust smoothly or with oscillations. This classification is central to the paper because it connects the mathematical roots of the model to the economic interpretation of oil price adjustment.

3 METHODOLOGY

3.1 Data and Source

The empirical analysis uses monthly Brent crude oil prices obtained from the FRED database and imported into Python from an Excel file. The series used is *Global price of Brent Crude* (POILBREUSDM), reported monthly in U.S. dollars per barrel (FEDERAL RESERVE BANK OF ST. LOUIS, 2026).

The sample begins in January 2003 and ends in January 2026, giving 277 monthly observations. Starting in 2003 provides a sufficiently long pre-conflict period before the Russia–Ukraine shock and also includes several important oil-market episodes, such as the 2007–08 oil price shock, the 2014–16 oil price decline, the COVID-19 shock, and the post-2022 conflict period. This makes the sample useful for studying both long-run price movements and the later conflict period.

The date variable was converted into a time-series format, the observations were sorted chronologically, and missing values were removed. Computations were carried out in Python with NumPy, SciPy, pandas, statsmodels and Matplotlib (HARRIS et al., 2020; VIRTANEN et al., 2020; MCKINNEY, 2010; SEABOLD; PERKTOLD, 2010; HUNTER, 2007; PYTHON SOFTWARE FOUNDATION, 2026).

3.2 Construction of the Conflict Dummy

To represent the geopolitical shock associated with the Russia–Ukraine conflict, a binary conflict variable was created. The dummy takes the value 0 before February 2022 and 1 from February 2022 onward. February 2022 is used as the break date because the full-scale invasion of Ukraine began during that month.

Formally, the conflict dummy is defined as

$$\text{conflict}_t = \begin{cases} 0, & \text{if } t < \text{February 2022}, \\ 1, & \text{if } t \geq \text{February 2022}. \end{cases}$$

This variable is used as the forcing term in the continuous model and as an explanatory variable in the discretized OLS model.

The conflict is singled out as a forcing term because it is a geopolitical shock to the oil market rather than a cyclical one. Kilian (2009) shows that oil price shocks differ according to whether they originate in supply disruptions, in global demand, or in precautionary demand for oil. A war involving one of the world's largest oil and gas exporters combines a potential supply disruption with a precautionary component, and it produced a major shock to global energy markets (INTERNATIONAL ENERGY AGENCY, 2023). By contrast, the 2007–08 price increase is attributed mainly to strong demand meeting stagnant production (HAMILTON, 2009).

This choice of break date has a cost. Because February 2022 falls late in the sample, the conflict period contains only 48 of the 277 monthly observations, against 229 before the break. The dummy is therefore not used to compare two periods of similar length. It enters the model as a forcing term that allows the price dynamics to shift after the invasion, and its coefficient should be read with corresponding caution. Other major episodes in the sample are not modeled separately, a limitation discussed in Section 5.4.

3.3 Continuous Model Specification

The model used in this paper is based on the idea that Brent crude oil prices adjust over time around a long-run equilibrium level. This type of dynamic structure is inspired by second-order differential equation models, where the second derivative captures acceleration or deceleration in the adjustment process. Second-order equations are commonly used to describe systems that move around an equilibrium and may display damping or oscillatory behavior (BOYCE; DIPRIMA, 2001; THE OPEN UNIVERSITY, 2013).

Let $P(t)$ denote the Brent crude oil price at time t, and let $\bar{P}$ denote the long-run equilibrium price. The continuous-time model is specified as

$$\frac{d^2P}{dt^2} = -a\frac{dP}{dt} - b(P - \bar{P}) + c\,\text{conflict}(t).$$

Each term in Equation [eq:cont_model] has an economic interpretation. The term $\frac{dP}{dt}$ represents the speed of price adjustment, while $\frac{d^2P}{dt^2}$ represents the acceleration or deceleration of price changes. The parameter a is a damping parameter. It measures resistance to rapid price changes and captures the idea that markets do not usually adjust infinitely fast.

The parameter b measures the strength of the restoring force toward equilibrium. When $P(t) > \bar{P}$, the term $-b(P - \bar{P})$ is negative, pushing the price downward. When $P(t) < \bar{P}$, the term $-b(P - \bar{P})$ is positive, pushing the price upward. Therefore, this term represents the tendency of prices to move back toward the long-run equilibrium level.

The parameter c measures the effect of the geopolitical conflict shock. If $c > 0$, the conflict dummy increases upward pressure on the price dynamics. If $c < 0$, the estimated effect is negative. However, the sign and significance of this parameter must be interpreted empirically.

This specification is economically more realistic than a model with only $-bP$, because it does not imply that oil prices return to zero when there is no conflict. Instead, prices are assumed to adjust around a nonzero long-run equilibrium level $\bar{P}$. The classification of the model's adjustment behavior through characteristic roots follows the standard theory of second-order linear differential equations (BOYCE; DIPRIMA, 2001).

3.4 Homogeneous Equation and Characteristic Roots

To classify the adjustment behavior implied by the continuous model, the homogeneous version is considered by setting $\text{conflict}(t) = 0$. Starting from Equation [eq:cont_model], this gives

$$\frac{d^2P}{dt^2} = -a\frac{dP}{dt} - b(P - \bar{P}).$$

Rearranging,

$$\frac{d^2P}{dt^2} + a\frac{dP}{dt} + b(P - \bar{P}) = 0.$$

Let

$$y(t) = P(t) - \bar{P}.$$

Since $\bar{P}$ is treated as constant, $y'(t) = P'(t)$ and $y''(t) = P''(t)$. Therefore, the homogeneous equation becomes

$$y'' + ay' + by = 0.$$

Following the characteristic-equation method developed in Section 2.2, the corresponding characteristic equation is

$$r^2 + ar + b = 0.$$

Using the quadratic formula,

$$r = \frac{-a \pm \sqrt{a^2 - 4b}}{2}.$$

The corresponding discriminant is

$$\Delta = a^2 - 4b.$$

This discriminant is the Brent price model version of the general discriminant $D = B^2 - 4AC$, with $A = 1$, $B = a$, and $C = b$.

The sign of Δ determines the adjustment regime. As summarized in Table 1, $\Delta > 0$ corresponds to two distinct real roots and an overdamped regime, $\Delta = 0$ corresponds to a repeated real root and a critically damped regime, and $\Delta < 0$ corresponds to complex conjugate roots and an underdamped regime. In economic terms, this classification determines whether Brent prices return toward equilibrium smoothly or with oscillations.

Having established how the continuous model can be classified through characteristic roots, the next step is to transform the model into a discrete-time form that can be estimated using monthly data.

3.5 Discretization of the Model

The continuous-time model is theoretically useful, but the Brent price data are observed monthly. Therefore, the differential equation must be transformed into a discrete-time equation that can be estimated using monthly observations.

Let P_n denote the Brent crude oil price in month n. With a general time step Δt, the first derivative can be approximated by the forward difference

$$\frac{dP}{dt}(t_n) \approx \frac{P_{n+1} - P_n}{\Delta t}.$$

The second derivative can be approximated by the second finite difference

$$\frac{d^2P}{dt^2}(t_n) \approx \frac{P_{n+1} - 2P_n + P_{n-1}}{(\Delta t)^2}.$$

Because the data are monthly, one period corresponds to one month. Therefore, the time step is normalized to

$$\Delta t = 1.$$

Using $\Delta t = 1$, the finite-difference approximations become

$$\frac{dP}{dt}(t_n) \approx P_{n+1} - P_n$$

and

$$\frac{d^2P}{dt^2}(t_n) \approx P_{n+1} - 2P_n + P_{n-1}.$$

Starting from the continuous model for Brent crude oil prices,

$$\frac{d^2P}{dt^2} = -a\frac{dP}{dt} - b(P - \bar{P}) + c\,\text{conflict}(t),$$

substituting the finite-difference approximations gives

$$P_{n+1} - 2P_n + P_{n-1} = -a(P_{n+1} - P_n) - b(P_n - \bar{P}) + c\,\text{conflict}_n.$$

Expanding the right-hand side gives

$$P_{n+1} - 2P_n + P_{n-1} = -aP_{n+1} + aP_n - bP_n + b\bar{P} + c\,\text{conflict}_n.$$

Collecting the P_{n+1} terms on the left side gives

$$(1 + a)P_{n+1} = (2 + a - b)P_n - P_{n-1} + b\bar{P} + c\,\text{conflict}_n.$$

Dividing by $1 + a$, the discrete model becomes

$$P_{n+1} = \theta + \alpha P_n + \beta P_{n-1} + \gamma\,\text{conflict}_n,$$

where

$$\theta = \frac{b\bar{P}}{1 + a}, \qquad \alpha = \frac{2 + a - b}{1 + a}, \qquad \beta = -\frac{1}{1 + a}, \qquad \gamma = \frac{c}{1 + a}.$$

This discrete equation is the empirical version of the continuous ODE. It shows that next month's Brent price depends on the current price, the previous month's price, the conflict dummy, and a constant term generated by the long-run equilibrium price $\bar{P}$. Therefore, the OLS estimation must include an intercept.

3.6 Estimation Strategies

Two estimation strategies are used in this paper.

First, the continuous-time ODE is estimated by nonlinear optimization. This approach is theoretically more elegant because it remains close to the differential-equation framework. However, in practice it can be numerically unstable, especially when the optimization procedure produces parameter values that lead to explosive simulated paths or non-finite numerical results.

Second, the discrete-time model is fitted by least squares after discretization. Because the model includes the long-run equilibrium term $\bar{P}$, the discretized equation contains a constant term. Therefore, the OLS specification includes an intercept:

$$P_{n+1} = \theta + \alpha P_n + \beta P_{n-1} + \gamma\,\text{conflict}_n.$$

This approach is more computationally stable and easier to interpret, but it introduces approximation error because it replaces continuous derivatives with finite differences.

Using both approaches is useful because it shows whether the continuous model can be calibrated directly, or whether the discrete approximation is needed in practice. The continuous model preserves the theoretical structure of the ODE, while the discrete model provides a practical empirical implementation using monthly data.

Throughout this paper, the coefficients obtained from the discretized equation are treated as a calibration of the model rather than as an inferential exercise. The objective is to obtain plausible numerical values for the structural parameters a, b, and c so that the adjustment regime implied by the characteristic roots can be examined and interpreted. Formal statistical inference on these coefficients would require addressing the time-series properties of the price level, in particular the question of stationarity, which lies outside the scope of the present paper and is left for future work.

4 RESULTS

4.1 Descriptive Statistics

Table 2 reports descriptive statistics for the Brent price series.

Descriptive statistics

Variable	Mean	Std. Dev.	Min	Max	Observations
Brent price	72.28	25.26	25.04	133.59	277

The conflict dummy equals 1 in 48 months and 0 in the remaining 229.

4.2 Continuous Calibration and Numerical Diagnostic

The continuous approach seeks the parameters a, b, and c that minimize the mean squared error between the observed Brent price series and the path simulated from the differential equation.

The optimization converged, but the resulting values $a = 2.0000$, $b = 2.0000$, and $c = -50.0000$ coincide exactly with the bounds imposed on the search region. A solution located at the corner of the feasible set indicates that the objective function did not attain an interior minimum. These values are therefore not identified by the procedure and should not be interpreted as estimates of the structural coefficients. They are reported here as a numerical diagnostic rather than as a result.

Two features of the problem account for this behavior. First, the objective function requires the differential equation to be integrated over the full sample for every candidate parameter vector, so small changes in a or b can produce explosive simulated paths and non-finite values of the objective. During implementation, the unrestricted optimization did produce overflow, which is why bounds and a penalty for invalid parameter combinations were introduced. Second, the resulting objective surface is nearly flat in some directions and very steep in others. The mean squared error of approximately 809.89 confirms the poor fit, which is also visible in Figure 1.

This outcome is informative in itself. It shows that the theoretical appeal of the continuous specification does not translate directly into a tractable numerical problem, and it motivates the discretized approach developed in the next section.

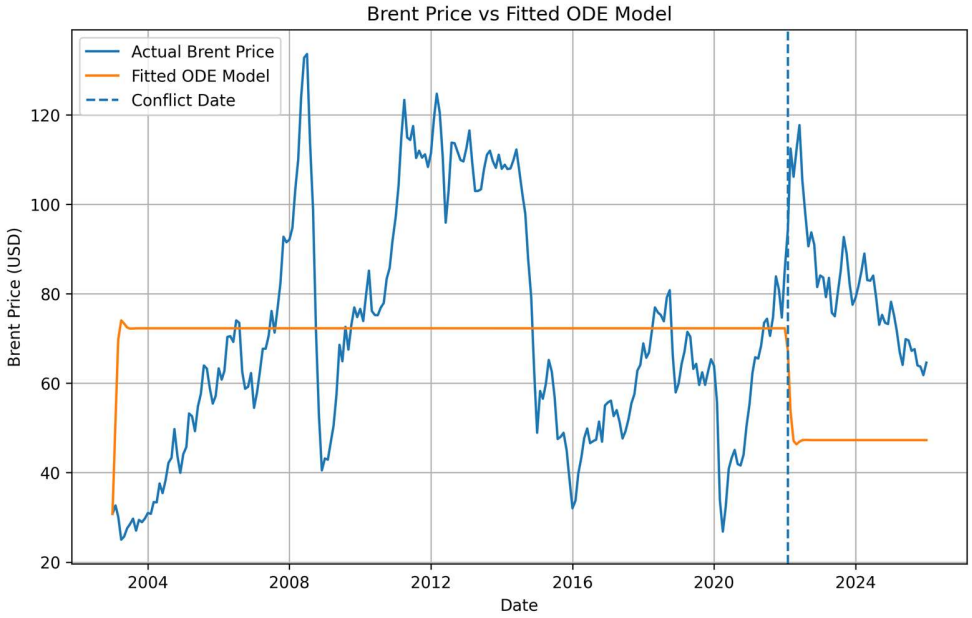
Output of the continuous optimization. All three parameters reached the imposed bounds; the values are not identified and are reported as a diagnostic.

Parameter	Value
a	2.0000
b	2.0000
c	-50.0000
MSE	809.8945
$\bar{P}$	72.2787

4.3 Continuous Fitted Path

Using the values returned by the optimization, the ODE was simulated from the initial Brent price and the initial monthly price change. Figure 1 compares the simulated path with the observed Brent price series. The fitted path is relatively flat and does not

capture the full volatility of Brent prices, which reinforces the interpretation that the continuous model is limited in empirical fit.



Observed Brent prices and fitted ODE path.

4.4 Discrete Estimation

The discrete model was estimated by ordinary least squares (OLS) with an intercept. The dependent variable is next month's Brent price, and the explanatory variables are the current price, the previous month's price, and the conflict dummy. The intercept corresponds to the coefficient θ in the discretized model. The full regression uses all available observations after constructing the lagged and forward price variables.

Calibrated coefficients of the discretized model

Variable	Coefficient
θ / Constant	3.103
α / P_n	1.335
β / P_{n-1}	-0.376
γ / Conflict dummy	-0.152

The fitted coefficients are

$$\theta = 3.103, \qquad \alpha = 1.335, \qquad \beta = -0.376, \qquad \gamma = -0.152.$$

The coefficient on the conflict dummy is small and negative. As stated in Section 3.6, these four coefficients are used as a calibration device: they are intermediate discrete-time quantities whose purpose is to recover the structural parameters a, b, and c in the next step, and no inferential claim is attached to them.

4.5 Recovery of Structural Parameters

The estimated coefficients θ, α, β, and γ were transformed back into the structural parameters a, b, and c using the algebraic relationships derived in Section 3.5. Since the intercept θ is related to the long-run equilibrium term $\bar{P}$, the implied equilibrium price can also be recovered from the discrete estimates.

The relationships are:

$$a = -\frac{1}{\beta} - 1, \qquad b = 2 + a - \alpha(1 + a), \qquad c = \gamma(1 + a).$$

Since

$$\theta = \frac{b\bar{P}}{1 + a},$$

the implied long-run equilibrium price is

$$\bar{P} = \frac{\theta(1 + a)}{b}.$$

Using the estimated discrete coefficients, the recovered structural parameters are:

$$a = 1.657, \qquad b = 0.109, \qquad c = -0.404.$$

The implied long-run equilibrium price recovered from the OLS estimates is:

$$\bar{P} = 75.313.$$

This value is close to the sample mean Brent price of approximately 72.28, suggesting that the recovered equilibrium level is economically reasonable.

Recovered structural parameters from the OLS estimates

Parameter	Estimated value
a_{OLS}	1.657
b_{OLS}	0.109
c_{OLS}	-0.404
$\bar{P}_{\text{OLS}}$	75.313

Table 5 reports the structural parameters recovered from the OLS estimates of the discretized model. The estimated damping parameter a_{OLS} is positive, indicating resistance to rapid price changes. The restoring parameter b_{OLS} is also positive, meaning that the model implies adjustment toward a long-run equilibrium price. The recovered long-run equilibrium price is approximately 75.31 dollars per barrel.

4.6 Characteristic Roots Analysis and Adjustment Regime

Using the recovered values of a and b, the discriminant and the characteristic roots were computed in order to classify the adjustment regime.

Characteristic roots and dynamic regime

Quantity	Value
Discriminant $\Delta = a^2 - 4b$	2.307
Root r_1	-0.069
Root r_2	-1.588

Dynamic regime	Overdamped

The discriminant is:

$$\Delta = 2.307.$$

Since $\Delta > 0$, the characteristic equation has two real roots:

$$r_1 = -0.069, \qquad r_2 = -1.588.$$

Two remarks are necessary before these roots are interpreted. First, the finite-difference approximations introduced in Section 3.5 are most reliable when the characteristic time of the system is long relative to the observation interval. With a monthly time step, the slow root satisfies this condition comfortably, since its magnitude is approximately 0.07. The fast root does not, since its magnitude is approximately 1.59. The fast mode therefore falls outside the range in which the monthly discretization is reliable, and only the slow root is interpreted in what follows.

Second, the slow root is close to zero. A characteristic root approaching zero corresponds to the limiting case in which the restoring parameter b vanishes, so that the price no longer returns to a long-run equilibrium. The calibrated system is thus located close to the boundary between convergence and non-convergence. Taking the slow root at face value, the implied half-life of adjustment is

$$\frac{\ln(2)}{|r_1|} \approx 10.1 \text{ months.}$$

Formally, the positive discriminant places the model in the overdamped region of Table 1. Substantively, the more informative statement is that the restoring force is weak and that adjustment, while non-oscillatory, is slow.

5 DISCUSSION

5.1 Comparison Between the Two Approaches

The continuous-time approach is more directly connected to the theoretical ODE framework and is therefore conceptually elegant. However, in the present application it proved numerically unstable and required practical adjustments to prevent overflow and invalid parameter values. The discrete approach, by contrast, is more stable and easier to estimate with conventional regression tools. Its main weakness is that it introduces approximation error through discretization.

Taken together, the two approaches serve complementary purposes. The continuous model preserves the theoretical structure, while the discrete model provides a more robust empirical implementation.

5.2 Economic Interpretation of the Results

The calibrated parameters describe the shape of adjustment rather than providing formal causal estimates. The parameter a reflects damping, meaning resistance to rapid changes in price. The parameter b captures the strength of the force pushing prices back toward equilibrium. The parameter c measures the direction and size of the conflict forcing term within the simplified model.

These dynamics have a practical counterpart for firms exposed to fuel and energy costs. A weak restoring force combined with a half-life of adjustment of roughly ten months implies that deviations of the Brent price from its long-run level are not corrected quickly: approximately half of a given deviation is still present after ten months. For a firm

managing logistics costs, setting fuel budgets, or defining a hedging horizon, this suggests that observed price deviations should be treated as persistent rather than transitory, and that planning horizons of one or two quarters may be too short to rely on mean reversion. At the same time, the absence of oscillatory behavior in the calibrated model means that the framework offers no support for timing strategies based on predictable price cycles.

5.3 Relation to the Literature on Estimating Continuous-Time Models

Fitting a continuous-time model to discretely sampled data is a long-standing problem in econometrics (BERGSTROM, 1984). The approach followed here replaces derivatives by finite differences and therefore estimates an approximate discrete version of the model, which introduces the discretization error discussed in Section 5.4. A second, well-documented source of error is the finite-sample bias of least-squares estimators in autoregressive models. Since the work of Hurwicz (1950) and Marriott and Pope (1954), it has been known that OLS estimates of autoregressive coefficients are biased toward zero in finite samples, and that the bias grows as the process becomes more persistent. In continuous-time models with a linear drift, the same mechanism leads standard estimators to overstate the speed of mean reversion. Phillips and Yu (2005) show that this bias can remain substantial even in large samples and can matter more than the discretization bias, Tang and Chen (2009) derive approximations to its size and propose a bootstrap correction, and Yu (2012) derives an approximation that remains accurate when mean reversion is slow, which is the near-unit-root case relevant here.

This literature bears directly on the present results. In the discretized model, the persistence of the price is measured by the sum $\alpha + \beta$, and the parameter relationships derived in Section 3.5 imply

$$\alpha + \beta = 1 - \frac{b}{1 + a}.$$

The calibrated value $\alpha + \beta \approx 0.959$ is close to one, which is precisely the region in which the least-squares bias is largest. A downward bias in $\alpha + \beta$ translates into an upward bias in the restoring parameter b. If such a bias is present here, the true restoring force is weaker than the calibrated one, the slow root lies even closer to zero, and the half-life of adjustment is longer than ten months. The bias therefore does not overturn the main conclusion of the paper but reinforces it: the calibrated system lies at least as close to the boundary between convergence and non-convergence as reported. Quantifying the effect, for example with the jackknife correction of Phillips and Yu (2005) or the bootstrap correction of Tang and Chen (2009), is left for future work.

5.4 Model Limitations

Several limitations must be acknowledged.

First, the long-run equilibrium price $\bar{P}$ is approximated in a simple way. In this paper, $\bar{P}$ is treated as a constant equilibrium level, but in reality the long-run equilibrium price of oil may change over time due to shifts in global demand, production costs, OPEC decisions, inventories, exchange rates, and energy policy. Therefore, the model improves on the zero-equilibrium problem, but it still simplifies the way equilibrium is determined.

Second, the conflict dummy is a very simple representation of geopolitical shocks. Oil prices are influenced by many other factors, including global demand, supply conditions, inventories, OPEC decisions, exchange rates, and financial speculation. A

binary dummy cannot capture this full complexity. Other major episodes contained in the sample, such as the 2007–08 oil price shock, the 2014–16 price decline, and the COVID-19 shock, are not modeled separately.

Third, the finite-difference approximation introduces numerical approximation error. The first derivative is approximated by a forward difference, while the second derivative is approximated by a centered second difference. Because these approximations have different error orders, the recovered structural parameters inherit approximation bias. A fully centered discretization would be a natural refinement for future work.

Fourth, the reliability of the finite-difference approximation depends on the relationship between the characteristic time of the system and the monthly observation interval. As discussed in Section 4.6, this condition is satisfied by the slow root but violated by the fast root. For this reason, the interpretation is restricted to the slow adjustment mode.

Fifth, the analysis focuses on the homogeneous equation. The particular solution associated with the conflict forcing term is not derived, so the paper characterizes the shape of the adjustment process but does not fully quantify the displacement of the equilibrium caused by the conflict dummy.

6 FINAL CONSIDERATIONS

This paper examined Brent crude oil price dynamics using a second-order differential equation framework with a conflict dummy representing the Russia–Ukraine shock. Two approaches were used to obtain numerical values for the model parameters: a continuous-time nonlinear optimization and a discrete-time least-squares fit derived from finite differences.

The continuous approach proved numerically intractable. The optimization converged to the boundary of the imposed search region, so its parameters are not identified and were reported only as a diagnostic of the difficulty of fitting differential equations directly to data. The discretized approach was stable and produced a positive discriminant, corresponding formally to an overdamped adjustment regime.

The interpretation of this result requires care, and the paper argues that the regime label is less informative than the magnitude of the roots. The slow characteristic root is close to zero, indicating a weak restoring force and a half-life of adjustment of roughly ten months, so that the calibrated system lies near the boundary between convergence and non-convergence. The fast root falls outside the range of validity of the monthly discretization and was not interpreted.

Several limitations remain. The long-run equilibrium price is treated as constant, although it plausibly varies with global demand, production costs, inventories, and policy. The conflict dummy is a coarse representation of geopolitical disruption, and other major episodes contained in the sample are not modeled. The discretization scheme mixes approximations of different orders, and the analysis is confined to the homogeneous equation.

The contribution of the paper is therefore methodological and illustrative rather than predictive. It shows how the classical theory of second-order linear differential equations can be used to organize and interpret the adjustment of a commodity price, and it identifies precisely where such a model becomes fragile. Future work should derive the particular solution associated with the forcing term, adopt a second-order accurate

discretization, correct the calibrated parameters for finite-sample bias, allow the equilibrium price to vary over time, and address the time-series properties of the price level before any inferential claim is made.

REFERENCES

BERGSTROM, A. R. Continuous time stochastic models and issues of aggregation over time. In: GRILICHES, Z.; INTRILIGATOR, M. D. (ed.). **Handbook of econometrics**. Amsterdam: North-Holland, 1984. v. 2, p. 1145-1212.

BOYCE, W. E.; DIPRIMA, R. C. **Elementary differential equations and boundary value problems**. 7. ed. New York: John Wiley & Sons, 2001.

FEDERAL RESERVE BANK OF ST. LOUIS. **Global price of Brent Crude (POILBREUSDM)**. St. Louis: FRED, 2026. Available at: https://fred.stlouisfed.org/series/POILBREUSDM. Accessed on: 12 Oct. 2026.

HAMILTON, J. D. Causes and consequences of the oil shock of 2007-08. **Brookings Papers on Economic Activity**, v. 2009, n. 1, p. 215-261, 2009.

HARRIS, C. R. et al. Array programming with NumPy. **Nature**, v. 585, p. 357-362, 2020.

HUNTER, J. D. Matplotlib: a 2D graphics environment. **Computing in Science & Engineering**, v. 9, n. 3, p. 90-95, 2007.

HURWICZ, L. Least-squares bias in time series. In: KOOPMANS, T. C. (ed.). **Statistical inference in dynamic economic models**. New York: John Wiley & Sons, 1950. p. 365-383.

INTERNATIONAL ENERGY AGENCY. **The impact of Russia's invasion of Ukraine on global energy markets**. Paris: IEA, 2023. Available at: https://www.iea.org/spotlights/the-impact-of-russia-s-invasion-of-ukraine-on-global-energy-markets. Accessed on: 12 Oct. 2026.

KILIAN, L. Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. **American Economic Review**, v. 99, n. 3, p. 1053-1069, 2009.

MARRIOTT, F. H. C.; POPE, J. A. Bias in the estimation of autocorrelations. **Biometrika**, v. 41, n. 3-4, p. 390-402, 1954.

MCKINNEY, W. Data structures for statistical computing in Python. In: PYTHON IN SCIENCE CONFERENCE, 9., 2010, Austin. **Proceedings** [...]. Austin: SciPy, 2010. p. 56-61.

PHILLIPS, P. C. B.; YU, J. Jackknifing bond option prices. **The Review of Financial Studies**, v. 18, n. 2, p. 707-742, 2005.

PYTHON SOFTWARE FOUNDATION. **Python 3 documentation**. [S. l.]: PSF, 2026. Available at: https://docs.python.org/3/. Accessed on: 12 Oct. 2026.

SEABOLD, S.; PERKTOLD, J. Statsmodels: econometric and statistical modeling with Python. In: PYTHON IN SCIENCE CONFERENCE, 9., 2010, Austin. **Proceedings** [...]. Austin: SciPy, 2010. p. 92-96.

TANG, C. Y.; CHEN, S. X. Parameter estimation and bias correction for diffusion processes. **Journal of Econometrics**, v. 149, n. 1, p. 65-81, 2009.

THE OPEN UNIVERSITY. **Second-order differential equations**. Milton Keynes: The Open University, 2013.

VIRTANEN, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. **Nature Methods**, v. 17, p. 261-272, 2020.

YU, J. Bias in the estimation of the mean reversion parameter in continuous time models. **Journal of Econometrics**, v. 169, n. 1, p. 114-122, 2012.