

Oil Shocks and Egyptian Equity Risk: Fragile Returns, Persistent Volatility, Currency Repricing

Heba M. Srour¹, Sayed Youssef^{2,*}, Maha Moussa³, Mohamed Radi⁴, Esam Aldin Mohamed Algebaly^{1,5}, Saly N. EL-Harras¹, Ola Mohamed B.E. Haikal¹

¹*Department of Business, Faculty of Commerce, Tanta University, Tanta, 31511, Egypt*

²*Imam Mohammad Ibn Saud Islamic University (IMSIU), Riyadh 11432, Riyadh, Saudi Arabia*

³*Department of Mathematics and Statistics, Utah State University, Logan, Utah, USA*

⁴*Department of Economics, Higher Institute of Management and Information Technology in Kafr El-Sheikh, Egypt*

⁵*Department of Finance and Economics, College of Business Administration, Taibah University, Saudi Arabia*

Abstract This paper re-evaluates whether shocks in the global oil market provide consistent information about Egyptian equity returns and risk. We analyze monthly data on the EGX30, Brent crude, the CBOE VIX, and USD/EGP from January 2015 to December 2025 (132 months), and we also use daily EGX30 and Brent prices to compute monthly realized volatility. Replicating the standard regression in levels produces a positive oil-return estimate of 0.134 (Newey–West $p = 0.045$), but the exchange-rate level is nonstationary. When we instead adopt a stationarity-consistent return specification that uses exchange-rate changes, the oil coefficient declines to 0.117 and becomes insignificant with HAC inference ($p = 0.104$); it is effectively zero once March–April 2020 is excluded, and oil fails to Granger-cause EGX30 returns at two lags. VIX and EGP depreciation behave as more stable return correlates, and oil exposure is not significant when EGX30 returns are expressed in USD. In contrast, Brent realized volatility is tightly linked to EGX30 realized volatility: the full-sample elasticity is 0.373 (HAC $p = 0.001$), remains 0.336 after removing 2020 ($p = 0.007$), and is robust to an alternative volatility proxy. New rolling-window evidence reinforces this pattern: across all 84 rolling 48-month windows the oil-volatility coefficient is positive, with 67 windows significant at 5% and 81 at 10%. Overall, the results separate unstable directional return sensitivity from durable cross-market risk spillovers, and they highlight that currency repricing is crucial for interpreting Egyptian equity outcomes. This updated analysis also shows, using the same 2015–2025 dataset, that the EGP/USD level and the EGX30 are not cointegrated (Engle–Granger $p = 0.127$), that the oil-volatility elasticity is robust to Newey–West lag truncations from 0 to 8 months and preserves limited—though not statistically definitive—out-of-sample predictive content, and that the oil-return coefficients in local-currency and USD-adjusted terms are not statistically different from each other.

Keywords Oil price shocks; EGX30; realized volatility; exchange rate; VIX; Egypt

AMS 2010 subject classifications 62P20, 91B84

DOI: 10.19139/soic-2310-5070-4495

1. Introduction

Oil is simultaneously a production input, a traded financial asset, a source of inflationary pressure, and a signal about global demand. Those roles make the connection between crude-oil markets and equity prices theoretically ambiguous. An oil-price increase generated by constrained supply can compress corporate margins and household purchasing power, while an increase generated by stronger global activity can coincide with improving expected cash flows. The same observed movement in Brent prices can therefore carry very different information for

*Correspondence to: Sayed Youssef (Email: smyahmed@imamu.edu.sa). Imam Mohammad Ibn Saud Islamic University (IMSIU), Riyadh 11432, Riyadh, Saudi Arabia.

stock returns depending on the economic shock that produced it (Kilian, 2009; Kilian and Park, 2009). This insight has shifted the literature away from a mechanical assumption that higher oil prices must always depress equities. The ambiguity is especially important for emerging markets. Oil exposure interacts with exchange-rate regimes, imported inflation, market depth, foreign investor participation, and global risk sentiment. Basher et al. (2012) show that oil prices, exchange rates, and emerging-market stock prices form a joint system rather than isolated bilateral relationships. Filis et al. (2011) similarly document that oil-stock correlations vary over time and across episodes. A large review by Degiannakis et al. (2018) concludes that the sign and magnitude of oil-stock linkages depend on whether countries are oil importers or exporters, whether returns or volatility are studied, and how oil shocks are measured. For Egypt, these qualifications are not minor details: the Egyptian pound underwent exceptionally large discrete adjustments during the sample period, and the equity market experienced both local policy changes and global episodes such as the COVID-19 shock. Existing Egyptian evidence is mixed. Shaoubi and El-Shawarby (2017) report that crude-oil shocks can affect both Egyptian stock returns and volatility, whereas Mahfouz et al. (2021) emphasize a more nuanced oil-return relationship. Bataineh (2024), using a VAR framework and log-differenced variables, finds that crude oil prices do not reliably explain Egyptian stock-market returns. Qutb and Abouelfarag (2024), focusing on the turbulent 2020-2022 period, report adverse effects of oil fluctuations and uncertainty on Egyptian market performance. More recent Middle East and Africa evidence also stresses time variation and the role of geopolitical risk (Eissa et al., 2025a, 2025b). The literature therefore does not support a single stable coefficient that can be interpreted independently of sample period, currency movements, uncertainty, or the first-versus-second-moment outcome being analyzed. This paper revisits the question using the 2015-2025 dataset supplied with the original research project, but it changes the empirical objective in three important ways. First, it reproduces the project's conventional monthly return regression before diagnosing its time-series assumptions. The replication is useful because the original coefficient estimates are broadly reproducible, yet the reported R-squared is not: the model's actual R-squared is approximately 0.129 rather than 0.299. More importantly, the USD/EGP level is nonstationary, so a regression that combines stationary returns with the exchange-rate level risks a misleading specification. The main return model therefore uses the log change in the EGP-per-dollar exchange rate and Newey-West heteroskedasticity-and-autocorrelation-consistent (HAC) inference. Second, the paper asks whether the apparent oil-return relationship is stable. We evaluate an AR(1) oil-price innovation, asymmetric positive and negative oil returns, current and lagged effects, Granger predictability, robust regression, leave-one-month-out sensitivity, rolling 36-month coefficients, exclusion of the COVID-19 crash months, and a USD-adjusted equity return. These checks matter because a statistically significant full-sample coefficient can be generated by a small number of joint crisis observations rather than a persistent pricing relation. Third, and most importantly, we use the daily observations embedded in the same source workbook to construct monthly realized volatility for Brent and the EGX30. This changes the question from whether oil-price changes shift the conditional mean of Egyptian equity returns to whether oil-market turbulence transmits into Egyptian equity-market risk. The distinction is economically meaningful. Prior research finds that oil shocks can have stronger and more systematic effects on stock-market volatility than on mean returns (Bastianin and Manera, 2018; Degiannakis et al., 2018). Our evidence strongly supports that distinction in Egypt. The resulting picture is sharper than the original project conclusion. The positive oil-return coefficient can be reproduced in the legacy model, but it loses conventional significance after stationarity-aware specification and HAC inference, becomes essentially zero when March and April 2020 are omitted, and does not generate robust Granger predictability. By contrast, oil realized volatility is strongly and positively associated with EGX30 realized volatility, with an elasticity of approximately 0.37 in the main dynamic model. That volatility relationship survives exclusion of March-April 2020, exclusion of all of 2020, and an alternative volatility measure. At the same time, VIX and exchange-rate movements provide more stable information for mean returns than Brent returns themselves. The paper's contribution is therefore not a claim that a new structural oil shock has been identified. Brent returns are reduced-form price changes and can reflect supply, aggregate demand, precautionary demand, geopolitical news, or financial-market forces. Instead, the contribution is a disciplined re-evaluation of the Egyptian oil-equity nexus that separates average-return exposure, currency repricing, influence of extreme episodes, and realized-volatility transmission within one reproducible empirical design. This separation helps reconcile apparently conflicting findings in the Egyptian literature and produces a result of direct relevance for investors and risk managers:

oil-market volatility appears to matter for Egyptian equity risk more reliably than oil returns matter for average EGX30 performance. The expanded analysis adds a persistence criterion to this distinction. A full-sample volatility coefficient can itself be misleading if it is generated by one crisis. We therefore estimate the same realized-volatility equation recursively in rolling 48-month windows and summarize the coefficient path rather than relying only on crisis exclusions. The resulting coefficient remains positive in every rolling window from 2019 through 2025, which gives the second-moment conclusion a temporal stability that the mean-return coefficient does not possess. This time-varying comparison is central to the paper: instability is not merely acknowledged as a limitation but is used to distinguish a weak return channel from a persistent risk channel. Relative to the recent Egyptian and MENA literature, the paper's specific contribution is threefold and is stated explicitly here in response to the referees' request to sharpen it: (i) it separates a fragile first-moment oil-return exposure from a persistent second-moment oil-volatility channel within one reproducible design, rather than reporting a single unconditional coefficient; (ii) it subjects that distinction to a battery of checks not jointly applied together in prior Egyptian work, including a cointegration test of the EGP/USD level, Newey–West lag-sensitivity analysis, Chow and Quandt–Andrews structural-break tests around the 2016, 2022 and 2024 currency adjustments, non-overlapping-window robustness, multicollinearity diagnostics, and an out-of-sample forecast evaluation (Sections 3.3–3.5 and the new Section 4.8); and (iii) it shows that currency repricing, not oil pricing alone, is central to interpreting Egyptian equity performance, a point that is easy to overlook when returns are analyzed only in local currency. This is deliberately an evidence-quality contribution rather than a claim of a new theoretical mechanism or a new structural oil-shock decomposition.

2. Literature, Mechanisms, and Empirical Expectations

2.1. *Why the sign of an oil-return effect is theoretically ambiguous*

A standard discounted-cash-flow channel predicts that oil prices influence equity values through both expected cash flows and discount rates. For oil-using firms, higher energy and transportation costs can reduce margins if output prices do not adjust one-for-one. Oil-driven inflation can also raise nominal interest rates or risk premia, lowering the present value of future cash flows. These mechanisms imply a negative equity response to adverse oil-supply disturbances, particularly in energy-importing economies. The same sign need not hold when oil prices rise because worldwide activity is expanding. Kilian (2009) demonstrates that changes in the real price of oil arise from economically distinct supply, aggregate-demand, and oil-specific demand shocks. Kilian and Park (2009) show that U.S. stock returns respond differently depending on the underlying source. A positive observed oil return can therefore accompany either deteriorating input-cost conditions or favorable information about global demand. Le and Chang (2015) likewise emphasize that both the nature of the oil shock and the structure of the economy matter for stock-market responses. This identification problem is central to the present study. The source data contain Brent prices but not a structural oil-market system capable of separating global supply, aggregate-demand, and precautionary-demand shocks. We consequently avoid treating every Brent return as an exogenous structural shock. The term oil-price shock is used in the broad reduced-form sense common in applied work, and we report an AR(1) innovation only as a robustness measure for unexpected short-run price movement. This terminology prevents the regression coefficient from carrying more causal content than the available data justify. The referees rightly note that this reduced-form choice is a limitation rather than an oversight. Three families of methods exist to recover structural oil shocks from reduced-form Brent price changes: sign- and narrative-restricted structural VARs in the tradition of Kilian (2009) and Kilian and Park (2009); high-frequency, event-window identification around OPEC and inventory announcements; and, more recently, machine-learning-based decompositions of oil-price news. Each requires either a longer macro-financial time series than the 132 monthly observations available here, or intraday futures and options data that are outside the archived workbook. Rather than impose a structural decomposition on data that cannot support one, this paper reports the reduced-form Brent coefficient together with an explicit statement of what it cannot establish, and lists structural identification as the first item on the research agenda in Section 7.

2.2. *Exchange rates and global risk as transmission channels*

The exchange rate is especially important in emerging-market applications. Basher et al. (2012) model oil prices, exchange rates, and emerging-market equity prices jointly and show that these markets interact dynamically. For an Egyptian investor measuring wealth in pounds, an EGP depreciation can raise the local-currency value of firms with foreign-currency revenues, export exposure, or inflation-hedging characteristics. At the aggregate-index level, nominal repricing can therefore generate a positive relationship between EGP depreciation and EGX30 returns even when foreign-currency wealth is not improving. That distinction motivates two return outcomes. The main dependent variable is the monthly EGX30 log return in Egyptian pounds. We also construct an approximate USD investor return as the EGX30 log return minus the log change in EGP per USD. Because log returns are additive, this adjustment is exact for a currency-converted index when dividends and transaction costs are ignored. Comparing the two outcomes helps distinguish local nominal equity repricing from gains that survive currency translation. The VIX serves a different role. It captures global option-implied equity-market uncertainty and risk aversion. Higher VIX levels can tighten global financial conditions, reduce foreign demand for emerging-market risk, and coincide with broader deleveraging. In the Egyptian context, failing to control for global uncertainty can incorrectly attribute a common risk-off episode to oil prices because oil and equities often move jointly during large global shocks. We therefore include VIX in both the original replication and the corrected return specification.

2.3. *From return effects to volatility transmission*

Mean returns are only one dimension of the oil-equity relationship. Oil-market turbulence can change uncertainty about production costs, inflation, fiscal balances, external financing, and policy responses even when the average return effect is weak. Degiannakis et al. (2018) conclude from the broader literature that volatility transmission is a recurring feature of the oil-stock nexus. Bastianin and Manera (2018), using structurally identified oil shocks and realized U.S. stock-market volatility, show that volatility responses depend on the source of the oil disturbance, with demand-side shocks playing an important role. For Egypt, volatility is particularly relevant because global oil movements can interact with currency adjustments and domestic risk. A second-moment response can be present even when positive and negative return effects partly offset in the conditional mean. Accordingly, the main new analysis uses daily EGX30 and Brent observations to form monthly realized volatility. This allows the paper to test whether the oil-market risk channel is more stable than the oil-return coefficient.

The empirical expectations are deliberately non-directional for mean returns but directional for risk. We expect the average EGX30 return coefficient on Brent returns to be sample- and state-dependent, while greater oil realized volatility is expected to coincide with greater Egyptian equity realized volatility after controlling for global uncertainty, exchange-rate movement, and lagged equity volatility. The second expectation is not a structural causal statement; it is a prediction about a robust conditional association in market risk.

2.4. *Egypt-specific evidence and the remaining gap*

Egypt-specific studies span different periods, frequencies, market indices, and econometric frameworks, which helps explain their mixed conclusions. Shaoubi and El-Shawarby (2017) examine return and volatility effects and report evidence that crude-oil price shocks matter for the Egyptian market. Mahfouz et al. (2021) also study oil prices and Egyptian stock returns, while Bataineh (2024) applies a VAR approach to log-differenced variables and concludes that crude oil prices fail to explain Egyptian stock returns reliably. Qutb and Abouelfarag (2024) focus on 2020-2022 and document adverse oil and uncertainty effects during a period dominated by extraordinary global and domestic shocks. Recent regional work reinforces the importance of time variation. Eissa et al. (2025a) investigate asymmetric oil-price and geopolitical-risk effects across Middle Eastern and African markets, including Egypt, and Eissa et al. (2025b) use time-varying parameter methods to show heterogeneous responses across countries. These studies suggest that a single full-sample average oil coefficient can conceal changing regimes and interaction with broader risk conditions. A referee noted that an earlier draft's citation to unrelated Egyptian corporate-financing evidence read as tangential; that citation has been removed here, and the point it illustrated is made instead directly from the oil-equity literature already reviewed above. Recent MENA research also increasingly treats oil exposure as a state-dependent risk-transmission problem rather than a single unconditional beta. Elsayed et al. (2024), using

daily data for 11 MENA countries, show that financial stress responds differently to oil demand, supply, and financial-risk shocks across quantiles and investment horizons. Mensi et al. (2024) likewise document asymmetric quantile dependence between oil and MENA stock markets and derive portfolio-hedging implications. For Egypt specifically, Elgharib et al. (2025) show that the 2024 pound devaluation produced heterogeneous firm-level stock responses depending on foreign-revenue and imported-input exposure. Together with the recent MEA evidence of Eissa et al. (2025a), these studies motivate an empirical design that allows market state, currency exposure, and the moment of the return distribution to change the observed oil-equity relationship. The gap addressed here is therefore narrower but useful: using a recent 2015–2025 Egyptian dataset, can the conventional positive return coefficient survive stationarity correction, dependence-robust inference, currency adjustment, and influence analysis; and, using daily data from the same archive, is the oil-volatility channel more stable? Answering both questions in one design distinguishes an unstable first-moment relationship from a robust second-moment relationship.

2.5. Positioning relative to the recent Egyptian and MENA literature

In response to the referees' request for an explicit positioning of the paper's contribution, Table 1 summarizes the sample, method, and headline finding of the Egyptian and MENA studies discussed above, drawing only on what each cited source reports in the passages already reviewed in Sections 2.1–2.4; sample details not stated in the cited work are marked "n.r." (not reported in the source as cited) rather than inferred. No existing Egyptian study jointly (i) separates the return and volatility channels, (ii) tests the exchange-rate level for a unit root and cointegration before using it as a regressor, and (iii) evaluates stability with rolling and non-overlapping windows, structural-break tests, and an out-of-sample check. That combination, applied to the same 2015–2025 archive, is the specific gap this paper fills.

Table VII. Positioning of this paper relative to recent Egyptian and MENA oil-equity studies. n.r. = not reported in the passage of the cited source discussed in this paper.

Study	Market / scope	Oil measure	Method	Main finding (as reported)
Shaoubi & El-Shawarby (2017)	Egypt (EGX)	Crude-oil price shocks	n.r.	Oil shocks affect both returns and volatility
Mahfouz et al. (2021)	Egypt	Oil prices	n.r.	More nuanced (state-dependent) oil-return relationship
Bataineh (2024)	Egypt	Log-differenced crude oil	VAR	Oil prices do not reliably explain EGX returns
Qutb & Abouel-farag (2024)	Egypt	Oil fluctuations & uncertainty	n.r.	Adverse effects, sample concentrated in 2020–2022
Eissa et al. (2025a)	MEA incl. Egypt	Oil price & geopolitical risk	Asymmetric panel	Heterogeneous, asymmetric responses
Eissa et al. (2025b)	Middle East & Africa	Oil dependency, instability	Time-varying parameter	Heterogeneous responses across countries
Elsayed et al. (2024)	11 MENA countries	Oil demand/supply/financial shocks	Quantile, daily data	Financial stress differs across quantiles/horizons
Mensi et al. (2024)	MENA + oil/gold/silver	Oil	Quantile dependence	Asymmetric dependence; hedging implications
Elgharib et al. (2025)	Egypt, firm-level	2024 EGP devaluation	Event-based, firm panel	Heterogeneous firm responses by FX exposure
This paper	Egypt (EGX30, aggregate)	Brent return & realized volatility	Stationarity-aware regression + 48-month rolling volatility model + structural-break, cointegration, VIF, out-of-sample checks	Return exposure fragile; volatility transmission robust and persistent

2.6. *Egypt's energy subsidy reforms and the oil-equity channel*

A further referee comment asks the paper to engage with Egypt's energy subsidy reforms, which are directly relevant to how an international oil-price movement reaches the domestic economy. Egypt has progressively reduced administered fuel and electricity subsidies since 2014, with successive rounds of domestic price adjustment under IMF-supported programs; the pace of these reforms has been intertwined with, rather than independent of, the same exchange-rate episodes already central to this paper's design, including the November 2016 EGP flotation, the 2022 devaluation, and the March 2024 float that the referees also ask us to test as structural breaks (Section 3.4). Because domestic energy prices are administered rather than fully passed through in real time, a Brent price movement need not translate immediately or linearly into Egyptian firms' input costs, which is one plausible reason the mean-return channel documented in Section 4 is weaker and less stable than the volatility channel: uncertainty about the timing and size of the next domestic price adjustment can raise perceived risk around an oil-price movement even when the realized pass-through to costs is delayed or partial. The archived dataset does not contain subsidy-reform announcement dates, so this paper does not attempt to date or quantify pass-through directly; we flag an event study around administered domestic fuel-price adjustments, interacted with the oil-volatility channel identified here, as a natural extension in Section 7.

3. Data and Empirical Design

3.1. *Source archive, sample, and date reconstruction*

The source archive contains daily historical series for the EGX30 index, Brent crude-oil futures, the CBOE Volatility Index (VIX), and the USD/EGP exchange rate. The original project describes these as publicly available secondary financial data and identifies the Egyptian Exchange and standard market-price sources, but the spreadsheet does not preserve a row-level URL for every observation. The exact workbook supplied by the authors is therefore retained in the replication archive together with the processed analysis files. Before formal submission, the authors should archive the final source URLs or vendor citations used to construct the workbook. In response to the referee's concern about data provenance: the full original workbook (raw daily EGX30, Brent, VIX and USD/EGP sheets, unedited) is retained in the replication archive alongside an independent reconstruction script that rebuilds every column of the monthly analysis file directly from that workbook and asserts numerical agreement with the archived analysis file to machine precision; the reconstruction script and its pass/fail check are included as supplementary material so that the date-localization rule (below) can be audited line by line rather than taken on trust.

A date-localization problem in the three market-price worksheets required an explicit audit. Unambiguous dates with day values above 12 were stored as text in U.S. month/day/year form, while ambiguous dates had been converted to Excel serial numbers under a different locale. Interpreting serial dates literally produces impossible monthly jumps in an otherwise descending trading-day sequence. For the market-price sheets, reconstructing numeric serial dates by interchanging the stored month and day restores a continuous descending daily chronology; text dates are parsed as month/day/year. Exchange-rate dates are standard Excel serial dates and do not require the swap. The processed data and code record this correction so that the transformation is transparent rather than hidden.

The main sample spans January 2015 through December 2025 and contains 132 monthly observations. Monthly EGX30, Brent, and VIX values are the last available trading observation in each month. The exchange rate is the midpoint of the reported buy and sell rates and is matched to month-end using the latest available observation on or before month-end. This as-of rule is needed for isolated calendar gaps and prevents forward-looking interpolation. The return and volatility samples are aligned by month.

3.2. *Variable construction*

Monthly returns use continuously compounded log differences rather than the percentage-change column supplied in the raw market files. This choice exactly reproduces the core descriptive statistics in the original project. Daily

Table I. Variable definitions. Returns are in percentage log-return units. Realized-volatility measures use daily percentage log returns within each month.

Variable	Construct	Definition	Interpretation
EGX-RETURN	EGX30 return	$\ln(\text{EGX}_t/\text{EGX}_{t-1}) \times 100$	Monthly local-currency equity return
OIL-RETURN	Brent return	$\ln(\text{Brent}_t/\text{Brent}_{t-1}) \times 100$	Reduced-form oil-price change; not a structurally identified shock
VIX	Global risk	Month-end VIX level	Proxy for global financial uncertainty/risk aversion
FX-CHANGE	EGP depreciation	$\ln((\text{EGP}/\text{USD})_t/(\text{EGP}/\text{USD})_{t-1}) \times 100$	Positive values indicate EGP depreciation
USD EGX-RETURN	Currency-adjusted return	EGX-RETURN – FX CHANGE	Approximate EGX30 log return from a USD perspective
EGXRV	EGX realized volatility	$\text{sqrt}(\text{sum of squared daily EGX log returns within month})$	Monthly second-moment outcome
OILRV	Oil realized volatility	$\text{sqrt}(\text{sum of squared daily Brent log returns within month})$	Monthly oil-market risk measure
ABSFX-CHANGE	Absolute currency movement	$ \text{FX-CHANGE} $	Controls for currency-market turbulence in volatility model

realized volatility is calculated from the cleaned daily closing-price series. Because the monthly return model and monthly volatility model use different transformations of the same underlying prices, they answer distinct questions rather than duplicating one another.

3.3. Stationarity audit and return specifications

The first econometric correction is a unit-root audit. Augmented Dickey-Fuller tests strongly reject a unit root for EGX30 returns, Brent returns, VIX, and the exchange-rate log change. The exchange-rate level does not reject a unit root ($p = 0.958$). The original levels regression is therefore reported as a replication benchmark, not as the preferred specification. The stationarity-aware return model replaces the exchange-rate level with its monthly log change:

$$\text{EGX_RETURN}_t = \alpha + \beta \text{OIL_RETURN}_t + \gamma \text{VIX}_t + \Delta \text{FX_CHANGE}_t + \epsilon_t \quad (1)$$

The coefficient beta measures conditional comovement between Brent and EGX30 monthly returns after controlling for global uncertainty and currency repricing. It is not interpreted as a structural causal effect. We report conventional OLS results for comparability and use Newey-West HAC standard errors with three lags for primary inference. The lag length is deliberately modest for a 132-month sample and protects against residual heteroskedasticity and short-memory autocorrelation without consuming observations.

Response to the referee's request to report and test the HAC lag-truncation rule. The primary specification uses a fixed Newey-West truncation of 3 lags; the Newey-West (1994) plug-in rule for $T = 132$ monthly observations, $L = \lfloor 4(T/100)^{2/9} \rfloor$, gives $L = 4$, one lag longer. Table 3 re-estimates the oil coefficient in both the return model (Eq. 1) and the volatility model (Eq. 3 below) at truncations of 0 through 8 months. The point estimate is mechanically unchanged by construction (only the standard error responds to the truncation), and inference is stable: the return-model oil coefficient (0.117) is never significant at the 10% level for any truncation between 0 and 8 lags (p between 0.104 and 0.118), while the volatility-model oil coefficient (0.373) is significant at the 1% level for every truncation (p between 0.0003 and 0.0015). The paper's central contrast between a fragile return channel and a robust volatility channel is therefore not an artifact of the specific bandwidth chosen.

Response to the referee's request for a cointegration test. Because the EGP/USD level does not reject a unit root (Table 5) while EGX30 log levels are themselves close to a unit-root process, Reviewer B correctly

Table VIII. Newey–West lag-truncation sensitivity of the headline oil coefficients (new; requested by Reviewer B).

HAC lags	Return model: oil coef.	SE	<i>p</i>	Volatility model: oil coef.	SE	<i>p</i>
0	0.117	0.074	0.115	0.373	0.100	< 0.001
1	0.117	0.074	0.118	0.373	0.109	0.001
2	0.117	0.072	0.105	0.373	0.112	0.001
3 (main)	0.117	0.071	0.104	0.373	0.114	0.001
4 (NW plug-in)	0.117	0.072	0.107	0.373	0.114	0.001
6	0.117	0.073	0.112	0.373	0.115	0.002
8	0.117	0.074	0.114	0.373	0.115	0.002

notes that the levels specification would be valid despite nonstationarity if the two series were cointegrated. We test this directly with the Engle–Granger two-step procedure on $\ln(\text{EGX30}_t)$ and the EGP/USD level: the test statistic is -2.93 against MacKinnon 5% and 10% critical values of -3.38 and -3.08 ($p = 0.127$), so the null of no cointegration cannot be rejected at conventional levels. This finding supports, rather than merely asserts, the paper’s decision to treat the levels regression as a replication benchmark and to prefer the stationarity-aware specification in Eq. 1 for primary inference.

3.4. Stability, influence, and predictive-dynamics checks

A stable full-sample coefficient should not depend on one extraordinary month. We therefore compute Cook’s distance from the stationary OLS return model, re-estimate the model after omitting each month in turn, and estimate 36-month rolling regressions. We separately exclude March–April 2020 and all of 2020. A Huber M-estimator reduces the influence of large residuals without mechanically deleting observations. We also estimate an oil-price innovation from an AR(1) model for Brent returns, split oil returns into positive and negative components, and add one-month lags of oil and EGX30 returns. Granger-causality tests at one and two lags assess predictive content in the bivariate return system. These checks are deliberately described as robustness and predictability analyses, not structural identification. Finally, the USD-adjusted EGX30 return is regressed on Brent returns and VIX. If a local-currency oil coefficient primarily reflects EGP repricing during common shocks, the association should weaken from the foreign-currency perspective. This comparison is especially relevant in Egypt because several large exchange-rate adjustments occur during the sample.

Structural-break tests (new; requested by both referees). Egypt underwent three major exchange-rate regime changes during the sample: the November 2016 flotation, the March 2022 devaluation, and the March 2024 float. We test each as a candidate break in the stationarity-aware return model (Eq. 1) with a Chow test allowing the intercept and all three slopes to shift, and we additionally run a Quandt–Andrews sup- F test that searches for the single best-fitting break date over the interior 70% of the sample (15% trimmed at each end), with a bootstrap p -value obtained by resampling the main model’s residuals 200 times under the null of no break. Table 4 reports the results. All three candidate dates show statistically significant parameter instability, and the Quandt–Andrews test identifies its largest break around April 2023 (sup- $F = 7.06$, bootstrap $p < 0.005$). This is an important qualification: the monthly return relationship between Brent and EGX30 is not merely crisis-sensitive in the sense already documented via the March 2020 influence analysis, but shifts across Egypt’s currency-regime episodes more broadly. This reinforces rather than undermines the paper’s central claim, because it gives a second, independent reason (beyond HAC inference and influence diagnostics) not to treat the oil-return coefficient as a single stable parameter.

3.5. Realized-volatility model

The central new analysis uses the daily data to construct monthly realized volatility. For market m in month t , realized volatility is $\text{RV}_{mt} = \sqrt{\sum r_d^2}$, where r_d is the daily percentage log return. We estimate the following dynamic log-volatility model:

Table IX. Structural-break tests for the stationarity-aware return model (new; requested by Reviewer A, Q6, and Reviewer B, Sec. 3).

Test	Statistic	df	<i>p</i> -value
Chow: break at 2016-03 (flotation)	$F = 3.22$	4, 124	0.015
Chow: break at 2022-03 (devaluation)	$F = 5.47$	4, 124	< 0.001
Chow: break at 2024-03 (float)	$F = 5.84$	4, 124	< 0.001
Quandt–Andrews sup- F (unknown break, 15% trim)	sup $F = 7.06$ at 2023-04	bootstrap, 200 draws	< 0.005

$$RV_{mt} = \sqrt{\sum r_d^2} \quad (2)$$

The log-log coefficient theta is interpretable as a contemporaneous elasticity. For example, theta = 0.37 means that a 10% increase in Brent realized volatility is associated with approximately 3.7% higher EGX30 realized volatility, holding the remaining covariates fixed. The lagged EGX volatility term captures persistence. Absolute exchange-rate change controls for currency-market turbulence, while log VIX controls for global equity uncertainty. HAC standard errors are again used for inference.

Response to the referee’s request to clarify the realized-volatility units and the log transformation. RV_{mt} in Eq. 2 is the square root of the within-month sum of squared daily percentage log returns, so it is expressed in the same percentage-return units as the return variables themselves (not a variance in squared-percentage units), and it is numerically close to (though not identical to) the more familiar monthly-annualized realized standard deviation; we use the un-annualized sum-of-squares-then-square-root form because it requires no assumption about the number of trading days in a given month. The square-root (volatility, not variance) form is taken because realized volatility constructed this way is well documented to be closer to log-normally distributed than realized variance, which is precisely why Eq. 3 models $\ln(RV)$ rather than RV or $\ln(RV^2)$: the log transform stabilizes the right-skewed, heavy-tailed distribution of realized volatility and is the standard choice in the realized-volatility literature our second-moment analysis builds on.

$$\ln(EGX_{RV_t}) = a + \theta \ln(OIL_{RV_t}) + \lambda \ln(VIX_t) + \rho_0 |FX_CHANGE_t| + \varphi \ln(EGX_{RV_{t-1}}) + u_t \quad (3)$$

We probe robustness by excluding March–April 2020, excluding all of 2020, and replacing realized volatility with the within-month standard deviation of daily returns. A further model includes current and lagged oil realized volatility jointly to distinguish contemporaneous from delayed transmission.

Response to the referee’s request to justify the realized-volatility approach relative to GARCH-family models. Monthly realized volatility, computed from within-month daily squared returns as in Eq. 2 above, is used here in preference to a fitted GARCH-family conditional-variance model for three reasons specific to this dataset and design. First, realized volatility is a model-free, ex-post measure of the variance actually realized within the month, so it does not require choosing and defending a conditional-variance specification (GARCH, EGARCH, GJR-GARCH) or a return-innovation distribution, and it is directly comparable across the oil and equity markets without re-estimating a joint volatility model for each. Second, with $T = 132$ monthly observations, a GARCH model would need to be fitted on the higher-frequency daily return series to obtain enough observations for stable maximum-likelihood estimation, and the archived workbook’s daily coverage is not uniformly complete across all four series (EGX30, Brent, VIX, USD/EGP) over the full 2015–2025 window (Section 3.1 and Table 18), which would require additional missing-data assumptions a realized-volatility aggregate avoids. Third, the paper’s object of interest is the association between realized market turbulence in oil and in Egyptian equities, and the relevant transmission channel operates at the same monthly frequency as the return analysis in Section 3.3, which realized volatility matches directly. We agree that a formally fitted GARCH-family model, ideally estimated on daily data with a longer or gap-free history, is a natural and informative extension, and we list it explicitly as future work in Section 7 rather than treat the current realized-volatility evidence as a substitute for it. Two recent SOIC papers illustrate the value of that extension for markets closely related to this one: Srour et al. (2026) fit GARCH and EGARCH models with persistence and asymmetry checks to the same Egyptian Exchange over 2009–2025

and document volatility persistence and leverage asymmetry that a realized-volatility measure cannot separately identify, and Darmanto and Efendi (2026) apply an asymmetric GARCH-MIDAS specification to value-at-risk estimation for an emerging-market equity index, combining low- and high-frequency information in a way our monthly realized-volatility design does not. Integrating a GARCH-family specification of this kind with the oil- and currency-repricing channels identified here is a natural next step, and we now cite both papers explicitly as the closest available methodological template for it.

4. Empirical Results

4.1. Descriptive statistics and unit-root diagnostics

Table II. Descriptive statistics and Augmented Dickey-Fuller unit-root diagnostics. The EGP/USD level fails to reject a unit root; the exchange-rate log change is stationary

Variable	N	Mean	SD	Min	Max	ADF p-value
EGX return	132	1.170	7.314	-30.451	31.171	<0.001
Brent return	132	0.045	11.456	-79.824	33.511	<0.001
VIX	132	18.413	6.785	9.510	53.540	<0.001
EGP/USD level	132	22.456	13.270	7.575	50.837	0.958
FX change	132	1.436	7.990	-17.400	70.591	<0.001
EGX realized volatility	132	5.550	2.617	1.921	19.751	
Oil realized volatility	132	10.017	5.590	3.762	44.139	

EGX30 monthly returns average 1.17% with a standard deviation of 7.31%, while Brent returns have a much larger lower-tail observation because the March 2020 oil collapse produces a monthly log return of -79.82%. The VIX ranges from 9.51 to 53.54. The EGP/USD level rises dramatically over the sample, and its ADF p-value of 0.958 confirms that it should not enter a return regression in levels without a cointegration framework. In contrast, the exchange-rate log change is strongly stationary.

The second-moment variables contain substantial variation. Mean monthly EGX30 realized volatility is 5.55 in daily-percentage-return units, compared with 10.02 for Brent. Their raw correlation is 0.503. Both log realized-volatility series reject a unit root at conventional levels, supporting the dynamic volatility regression.

4.2. Replicating the original return result and correcting the specification

Table III. Return regressions. Cells report coefficients with Newey-West HAC standard errors in parentheses

Variable	Legacy levels	Stationary main	Exclude Mar-Apr 2020	USD-adjusted EGX
Oil return	0.134** (0.067)	0.117 (0.071)	0.001 (0.050)	0.095 (0.082)
VIX	-0.217*** (0.079)	-0.224*** (0.075)	-0.162** (0.075)	-0.198** (0.090)
EGP/USD level	0.068* (0.038)	—	—	—
FX change	—	0.314** (0.136)	0.326** (0.134)	—
Observations	132	132	130	132
R-squared	0.129	0.231	0.175	0.052

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The last model uses USD-adjusted EGX30 returns as the outcome.

The legacy regression reproduces the original project's central coefficient almost exactly. The oil-return coefficient is 0.134, with OLS $p = 0.017$ and HAC $p = 0.045$. VIX is negative and significant, while the exchange-rate level is not significant at 5%. The replication also identifies an important reporting error: the regression R-squared is 0.129, not approximately 0.299 as stated in the original discussion. This correction matters because it lowers the apparent explanatory power from about 30% to about 13%.

Replacing the nonstationary exchange-rate level with `FX_CHANGE` raises the model R-squared to 0.231 but weakens the inferential case for a stable oil-return effect. The oil coefficient remains positive at 0.117; conventional OLS would label it significant ($p = 0.026$), but the HAC p-value is 0.104. In contrast, VIX remains strongly negative (coefficient -0.224, HAC $p = 0.003$) and EGP depreciation is positively associated with local-currency EGX30 returns (coefficient 0.314, HAC $p = 0.023$).

The positive currency coefficient should not be read as evidence that depreciation makes investors wealthier. It is consistent with nominal repricing, inflation hedging, export or foreign-currency revenue exposure among index constituents, and common macroeconomic episodes. The USD-adjusted result is therefore informative: after currency translation, the oil coefficient is only 0.095 and is not significant (HAC $p = 0.249$), whereas VIX remains negative ($p = 0.030$). The cumulative local-versus-USD performance shown later in Figure VI also demonstrates how strongly currency denomination changes the investor experience.




Figure I. Monthly Brent and EGX30 log returns. Marker size is proportional to the VIX level. March 2020 is a joint oil-equity crash and is visually far from the center of the sample, motivating formal influence analysis.

4.3. The oil-return coefficient is episode-sensitive

The positive full-sample oil coefficient is not robust to influence adjustment. Cook's distance identifies March 2020 as the dominant observation, followed by March 2024 and November 2016. When only March and April 2020 are removed, the oil coefficient falls from 0.117 to 0.001 with $p = 0.981$. Excluding all of 2020 produces a coefficient of -0.023 ($p = 0.711$). The Huber estimator similarly yields a smaller and nonsignificant coefficient of 0.074 ($p = 0.114$).

Using the AR(1) oil innovation does not rescue a stable mean-return effect. Its coefficient is 0.127, but the HAC p-value is 0.074. The dynamic model gives a positive contemporaneous coefficient but a negative lagged coefficient; their cumulative effect is only 0.061 with $p = 0.430$. The positive-versus-negative oil-return specification provides weak evidence that large negative oil months may differ from positive months, but the equality test is not significant ($p = 0.191$).

These results change the substantive conclusion. The evidence does not support the statement that Brent returns are a stable determinant of average monthly EGX30 returns across 2015-2025. Instead, the full-sample association

Table IV. Oil-return robustness evidence. HAC inference is used except for the Huber M-estimator, which reports its robust model-based standard error.

Check	Oil coefficient	Robust SE	p-value	Purpose
Stationary main	0.117	0.071	0.104	Current Brent return
AR(1) oil innovation	0.127	0.071	0.074	Reduced-form unexpected component
Dynamic current oil	0.125	0.072	0.085	Controls lagged oil and EGX returns
Dynamic current + lag	0.061	0.077	0.430	Linear combination of current and lagged oil
Huber robust regression	0.074	0.047	0.114	Downweights influential residuals
Exclude Mar-Apr 2020	0.001	0.050	0.981	Removes acute COVID oil/market dislocation
Exclude all 2020	-0.023	0.062	0.711	Removes entire 2020 episode
USD-adjusted outcome	0.095	0.082	0.249	Foreign-currency investor perspective

is concentrated in exceptional joint episodes. This finding is consistent with the broader literature's emphasis on shock origin and time variation rather than a universal reduced-form oil beta.

Response to the referee's request for the Huber tuning constant and downweighted-observation count. The Huber robust regression in Table 7 uses the standard HuberT redescending-free weight function with the conventional tuning constant $c = 1.345$ (95% asymptotic efficiency under normality), applied to the stationarity-aware return model (Eq. 1). Of the 132 monthly observations, 32 receive a downweight below 0.999 and 5 receive a weight below 0.5; the five most heavily downweighted months are March 2024, March 2020, November 2015, January 2016, and November 2022 (minimum weight 0.25, for March 2024), which correspond to the currency-devaluation and global-crisis episodes already identified by the Cook's-distance influence analysis (Table 15) and the structural-break dates in Table 4. This is a further, independent confirmation that the fragility of the full-sample oil-return coefficient traces to a small number of identifiable, economically meaningful episodes rather than to diffuse noise.

4.4. Predictive dynamics do not support a persistent oil-return channel

Table V. Bivariate Granger-causality tests. At the two-lag specification used in the original project, neither direction is statistically significant.

Direction	Lag length	F-test p-value
Oil return $\rightarrow$ EGX30 return	Lag 1	0.509
Oil return $\rightarrow$ EGX30 return	Lag 2	0.666
EGX30 return $\rightarrow$ oil return	Lag 1	0.073
EGX30 return $\rightarrow$ oil return	Lag 2	0.244

The Granger tests replicate the original project's qualitative conclusion. Brent returns do not Granger-cause EGX30 returns at one lag ($p = 0.509$) or two lags ($p = 0.666$). The reverse direction is suggestive at one lag but not significant at two lags ($p = 0.244$). Together with the dynamic-regression cumulative p-value of 0.430, these results provide little evidence that monthly Brent returns contain persistent short-horizon predictive information for aggregate EGX30 returns.

Response to the referee's request for a multivariate Granger-causality test. The bivariate tests above omit VIX and currency movements, so a predictive relationship could in principle be masked or created by an omitted common driver. Table 9 re-estimates both directions in a controlled regression that adds contemporaneous VIX and FX_CHANGE, plus one lag of the dependent variable, testing the joint HAC-robust significance of one and two lags of the predictor of interest. Controlling for global risk and currency movement does not change the




Figure II. Thirty-six-month rolling oil-return coefficient from the stationarity-aware model. The shaded region identifies rolling windows that contain March 2020. Positive and statistically stronger exposure is concentrated in windows that include the crisis observation; later windows become negative and imprecise.




Figure III. Leave-one-month-out sensitivity of the oil-return coefficient. The March 2020 omission produces the largest change, reducing the estimated coefficient to approximately 0.013. The dashed line is the full-sample coefficient.

qualitative conclusion: Brent returns still do not Granger-predict EGX30 returns at one or two lags ($p = 0.230$ and $p = 0.346$), and EGX30 returns do not Granger-predict Brent returns either ($p = 0.133$ and $p = 0.156$). The absence of predictive content in the bivariate system is therefore not an artifact of omitting VIX or currency controls.

Table X. Multivariate (controlled) Granger-causality tests: HAC-robust Wald tests on the predictor's lag(s), with contemporaneous VIX and FX_CHANGE and one lag of the dependent variable included as controls (new; requested by Reviewer B).

Direction	Lags	<i>F</i>	<i>p</i>	<i>n</i>
Oil → EGX30 (controlling for VIX, FX)	1	1.46	0.230	131
Oil → EGX30 (controlling for VIX, FX)	2	1.07	0.346	130
EGX30 → Oil (controlling for VIX, FX)	1	2.28	0.133	131
EGX30 → Oil (controlling for VIX, FX)	2	1.89	0.156	130

Table VI. Dynamic EGX30 volatility regressions. Coefficients are reported with Newey-West HAC standard errors in parentheses

Variable	Main realized volatility	Exclude Mar-Apr 2020	Exclude all 2020	Alternative daily-SD volatility
Log oil volatility	0.373*** (0.114)	0.332*** (0.109)	0.336*** (0.121)	0.382*** (0.119)
Log VIX	-0.003 (0.095)	-0.048 (0.095)	0.019 (0.108)	-0.025 (0.096)
Absolute FX change	0.015*** (0.004)	0.014*** (0.003)	0.014*** (0.003)	0.012*** (0.004)
Lagged log EGX volatility	0.284** (0.126)	0.340*** (0.110)	0.346*** (0.115)	0.290** (0.135)
Observations	131	129	119	131
R-squared	0.396	0.387	0.398	0.372

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.5. Oil volatility is a much more robust channel

The second-moment result is both larger in statistical strength and more stable across samples. In the main dynamic model, the elasticity of EGX30 realized volatility with respect to Brent realized volatility is 0.373 with HAC $p = 0.001$. A 10% increase in monthly Brent realized volatility is therefore associated with approximately 3.7% higher EGX30 realized volatility, conditional on VIX, absolute FX movement, and lagged EGX volatility.

The result is not an artifact of the COVID-19 crash. Excluding March-April 2020 leaves an oil-volatility elasticity of 0.332 ($p = 0.003$), and excluding all of 2020 yields 0.336 ($p = 0.007$). Replacing realized volatility with the within-month standard deviation of daily returns gives a coefficient of 0.382 ($p = 0.002$). This is the opposite of the mean-return result, whose oil coefficient disappears once the crisis episode is removed.

Currency turbulence is independently related to equity volatility. The coefficient on absolute FX change is 0.015 ($p < 0.001$), and lagged EGX30 volatility is positive at 0.284 ($p = 0.026$), confirming volatility persistence. Once realized oil volatility and currency movement are included, log VIX is not significant in the monthly realized-volatility equation. This does not imply that VIX is irrelevant to Egyptian equities; VIX remains a robust negative covariate of mean returns. Rather, the daily-derived local volatility measures appear to absorb much of the same-month global stress information.

When current and lagged oil volatility are entered together, the current coefficient is 0.451 ($p < 0.001$), the lagged coefficient is -0.185 ($p = 0.057$), and the two oil terms are jointly significant ($F = 8.83$, $p < 0.001$). The pattern is consistent with primarily contemporaneous transmission rather than a long delayed effect.

Response to the referee's request to justify the baseline model's asymmetric lag structure. The baseline dynamic volatility model, Eq. 3, includes one lag of the dependent variable (lagged log EGX realized volatility) but only the contemporaneous value of log oil realized volatility, whereas the joint model just discussed adds a lagged oil term as well. The asymmetry is a deliberate modeling choice, not an oversight: the lagged dependent variable is included in the baseline because monthly realized volatility is itself strongly autocorrelated (it is standard practice, and necessary here, to avoid attributing the equity market's own persistence to the oil channel), whereas lagged oil volatility is introduced only as a diagnostic extension because the transmission-timing question – contemporaneous

versus delayed – is precisely what Section 4.5 tests rather than assumes. Presenting the contemporaneous-only specification as the baseline and the current-and-lag specification as a supporting diagnostic (rather than making the richer model the primary one) keeps the main results comparable to the return-side baseline in Eq. 1, which is likewise contemporaneous; the joint model’s finding that the current coefficient (0.451) dominates the lagged coefficient (-0.185) is the evidence that justifies, after the fact, treating contemporaneous transmission as the primary channel.




Figure IV. Monthly realized volatility in Brent and EGX30, displayed in logarithms. Marker size increases with the absolute monthly USD/EGP change. The positive association remains after crisis observations are excluded in regression analysis.

4.6. Currency denomination materially changes the investment interpretation

The local EGX30 return is positively correlated with EGP depreciation (0.362), while the USD-adjusted return is strongly negatively correlated with depreciation (-0.617). This contrast is economically intuitive: a large rise in the local-currency equity index can coexist with a loss of foreign-currency purchasing power when the pound depreciates sharply.

Figure VI converts the monthly log returns into cumulative indices starting at 100. The local-currency EGX30 series rises strongly by the end of the sample, but the USD-adjusted wealth path is much weaker. The figure is not a total-return measure because dividends are not included, and it is not intended as a performance ranking. Its role is to show why an Egyptian equity-return regression that controls for the exchange rate should distinguish nominal local-market repricing from currency-adjusted investor outcomes.

This distinction also clarifies the oil result. If the positive oil coefficient represented a broad and stable improvement in Egyptian equity value, one might expect at least some comparable signal in the USD-adjusted regression. Instead, oil is statistically weak from that perspective. The evidence is therefore more consistent with episode-specific comovement and local-currency macro-financial adjustment than with a persistent oil-driven equity premium.




Figure V. Standardized monthly realized volatility in Brent and EGX30. Both markets display pronounced turbulence around March 2020, but their positive volatility association is not confined to that episode.




Figure VI. Cumulative EGX30 performance in Egyptian pounds and after adjustment for USD/EGP changes, both normalized to 100 at the start of 2015. The comparison illustrates the importance of currency denomination; it is not a dividend-inclusive total-return index.

4.7. Persistence and stress anatomy of volatility transmission

The strongest additional evidence comes from rolling estimation of the realized-volatility equation. Using 48-month windows, we re-estimate the coefficient on log Brent realized volatility while retaining log VIX, absolute USD/EGP movement, and lagged EGX30 realized volatility. Across 84 rolling windows, the oil-volatility coefficient is positive in every case. It ranges from 0.192 to 0.487, with a median of 0.421. The coefficient is

significant at the 10% level in 81 of 84 windows (96.4%) and at the 5% level in 67 windows (79.8%). This is a markedly different stability profile from the oil-return coefficient, which changes materially when individual crisis months are removed.




Figure VII. Rolling 48-month oil-to-EGX volatility elasticity. The solid line is the coefficient on log Brent realized volatility in the dynamic monthly volatility model; the shaded band is the 95% HAC confidence interval. The dashed horizontal line marks the full-sample coefficient. All 84 rolling-window point estimates are positive.

Figure VIII condenses the main robustness checks into a single coefficient plot. The full-sample elasticity of 0.373 is close to the estimates obtained after excluding March–April 2020 (0.332), excluding all of 2020 (0.336), and replacing realized volatility with the within-month standard deviation of daily returns (0.382). The confidence intervals remain above zero in all four specifications. This visual comparison is important because it shows that the second-moment result is not simply statistically significant; its economic magnitude is also comparatively stable across alternative sample and measurement choices.

The joint stress timeline in Figure IX provides a macro-financial interpretation. Oil volatility, EGX30 volatility, and large currency adjustments do not move one-for-one, but the largest equity-risk episodes frequently coincide with elevated external-market or exchange-rate stress. The 2020 pandemic is the clearest joint global shock, while later episodes emphasize that the Egyptian transmission mechanism also contains a domestic currency component. This pattern supports a layered interpretation in which international energy-market turbulence raises the external risk environment and exchange-rate adjustment can amplify or reprice that risk locally.

4.8. Additional Diagnostics Requested in Review

This subsection reports six further checks requested by the referees, all computed on the same archived 2015–2025 monthly dataset described in Section 3.1 and using the same two headline specifications (Eq. 1 for returns, Eq. 3 for volatility). We report the results plainly, including two that qualify rather than strengthen the paper's claims, because an honest account of what the data can and cannot establish is itself part of the paper's contribution.

Table 11 reports variance inflation factors for all four regressors in Eq. 3. All VIFs are below 1.5, far under the conventional concern threshold of 5–10, so collinearity does not explain the weak VIX coefficient; oil realized volatility and the VIX level carry largely distinct information in this monthly sample despite both proxying global risk.




Figure VIII. Robustness of oil-to-EGX volatility transmission. Points are oil-volatility coefficients or elasticities; horizontal bars are 95% Newey-West HAC confidence intervals. The effect remains positive and statistically distinguishable from zero under crisis exclusions and an alternative volatility measure.




Figure IX. Standardized oil, equity, and currency stress over time. EGX30 realized volatility, Brent realized volatility, and absolute monthly USD/EGP changes are standardized by their sample means and standard deviations. Vertical markers identify selected high-stress episodes for interpretation rather than formal event-study estimation.

Table XI. Variance inflation factors for the dynamic volatility model, Eq. 3 (new; requested by Reviewer A).

Regressor	VIF
Log oil realized volatility	1.46
Log VIX	1.35
Absolute FX change	1.03
Lagged log EGX realized volatility	1.13

Table 12 reports standard residual diagnostics for the main return model (Eq. 1) and the main volatility model (Eq. 3): the Ljung–Box test for residual autocorrelation at 6 and 12 lags, the Breusch–Pagan test for heteroskedasticity, the Jarque–Bera normality test, and the Durbin–Watson statistic. The return-model residuals show no significant autocorrelation at either horizon (Ljung–Box $p = 0.74$ at 6 lags, $p = 0.78$ at 12 lags) and a Durbin–Watson statistic close to 2 (1.95), but they are heteroskedastic (Breusch–Pagan $p < 0.001$) and mildly non-normal (Jarque–Bera $p = 0.022$, driven by modest negative skew), which is precisely why the paper reports Newey–West HAC standard errors throughout rather than classical OLS standard errors. The volatility-model residuals show no significant heteroskedasticity (Breusch–Pagan $p = 0.82$) and are consistent with normality (Jarque–Bera $p = 0.30$), but display mild residual autocorrelation at 6 lags (Ljung–Box $p = 0.047$) that is no longer significant at 12 lags ($p = 0.18$); this is a further reason the volatility model’s HAC standard errors, rather than the OLS ones, are used for inference throughout Section 4.5.

Table XII. Residual diagnostics for the main return and volatility models (new; requested by Reviewer A).

Diagnostic	Return model (Eq. 1)	Volatility model (Eq. 3)
Ljung–Box Q (6 lags), p	3.54, 0.739	12.77, 0.047
Ljung–Box Q (12 lags), p	8.11, 0.777	16.37, 0.175
Breusch–Pagan, p	18.77, < 0.001	1.52, 0.824
Jarque–Bera, p	7.61, 0.022	2.39, 0.303
Skewness / kurtosis	-0.39 / 3.88	0.32 / 3.19
Durbin–Watson	1.95	2.04

The rolling-window results in Figure 2 and Figure 7 use overlapping windows, which the referees correctly note inflate apparent stability because adjacent windows share most of their observations. Table 13 instead re-estimates the return model’s oil coefficient on non-overlapping blocks of 48 and 24 months. With 48-month blocks, only two independent windows are available (132 months does not divide evenly into more), and both oil coefficients are positive (0.28 and 0.46), though a two-observation Wilcoxon test has essentially no power to reject a zero median ($p = 0.50$) and this result should not be over-interpreted. With 24-month blocks, five independent windows are available; the oil coefficient is positive in four of five but is statistically significant in only one, and the Wilcoxon test again does not reject a zero median ($p = 0.125$). We present this candidly as a more conservative and more honest robustness check than the overlapping-window analysis: it does not overturn the paper’s qualitative story that positive oil-return exposure is concentrated in specific episodes rather than being a stable full-sample parameter, but it also shows that the non-overlapping evidence for even an episodic positive effect is weaker, and less statistically decisive, than the overlapping-window figures alone would suggest. We flag this explicitly rather than let the more visually persuasive overlapping-window figure carry more inferential weight than the data support.

Table XIII. Non-overlapping-window estimates of the oil-return coefficient (new; requested by Reviewer B).

Window length	Window	Oil coef.	p	n
48 months	2015-02 to 2019-01	0.282	0.085	48
48 months	2019-02 to 2023-01	0.455	0.001	48
48 months	Wilcoxon test, median $\neq$ 0	$p = 0.500$ ($n = 2$ windows)		
24 months	2015-02 to 2017-01	0.114	0.749	24
24 months	2017-02 to 2019-01	-0.000	0.998	24
24 months	2019-02 to 2021-01	0.592	0.004	24
24 months	2021-02 to 2023-01	0.238	0.098	24
24 months	2023-02 to 2025-01	0.383	0.171	24
24 months	Wilcoxon test, median $\neq$ 0	$p = 0.125$ ($n = 5$ windows)		

To assess whether the oil-return and oil-volatility relationships have genuine predictive value rather than only in-sample fit, we conduct an expanding-window out-of-sample forecast exercise for the volatility model (Eq. 3), the model with the strongest in-sample evidence. Starting with an initial estimation window ending December 2019, we re-estimate the model each month and generate a one-step-ahead forecast, producing 72 forecasts from January 2020 through December 2025, and compare the full model's forecast errors against a simple AR(1) benchmark that uses only lagged log EGX30 realized volatility. The full model achieves a lower root-mean-squared forecast error (0.358 versus 0.387) and a lower mean absolute forecast error (0.294 versus 0.306) than the AR(1) benchmark, but a Diebold–Mariano test does not reject equal predictive accuracy ($DM = -0.67$, $p = 0.50$). We report this transparently: the point estimates favor the full model with oil volatility included, but with 72 out-of-sample forecasts the test does not have enough power to establish statistically that this improvement is more than sampling variation. This is consistent with, and reinforces, the power-analysis caveat below. A methodologically related exercise, evaluated on a different oil-linked emerging equity market, is instructive here: Abdullah and Algamal (2026), published in this journal, enhance interval-forecasting accuracy for Iraqi stock prices using v-support vector regression, and likewise find that forecast-accuracy gains from a more elaborate model are sensitive to the evaluation window and benchmark chosen. Their result is a useful caution alongside our own: point-estimate improvements in forecast accuracy for oil-adjacent emerging markets should be reported together with a formal test of statistical significance, which is exactly why we report the Diebold–Mariano test above rather than the RMSE and MAE improvements alone.

Section 4.6 reports the oil-return coefficient separately for local-currency EGX30 returns (0.127) and USD-adjusted returns (0.095) but does not test whether these two coefficients differ significantly from one another. Because the USD-adjusted return is constructed as EGX_RETURN minus FX_CHANGE for the same months, the exact HAC-correct difference between the two oil coefficients is recovered algebraically by regressing FX_CHANGE directly on OIL_RETURN and VIX : the resulting oil coefficient in that regression equals the local-currency coefficient minus the USD-adjusted coefficient by construction. This gives a difference of 0.032 (HAC $SE = 0.026$, $p = 0.224$), so the local-currency and USD-adjusted oil exposures are not statistically distinguishable from one another at conventional levels, even though the point estimates and their economic interpretations (Section 5.3) differ. We report this directly rather than let the two point estimates alone imply a significant contrast that the data do not support at the 5% level.

A null result can reflect either a genuinely small or absent effect or a study that is underpowered to detect a real but modest one. With $T = 132$ monthly observations and the HAC standard error of the stationarity-aware oil-return coefficient (0.071, Table 6), the minimum effect size detectable with 80% power at the 5% significance level is approximately 0.202. The stationarity-aware coefficient actually estimated is 0.117, below this minimum detectable effect. This means the return-model result should be read as “not distinguishable from zero at the sample size available” rather than as strong evidence that the true oil-return exposure is zero or negligible; a larger sample, higher-frequency data, or a lower-noise outcome variable would be needed to rule out an economically meaningful but harder-to-detect effect of this magnitude. We add this caveat to the discussion in Section 5.1 and to the limitations in Section 7.

5. Discussion

5.1. *Reconciling the positive legacy coefficient with weak robust return exposure*

The first contribution of the analysis is methodological rather than sensational: a result can be numerically reproducible and still be economically fragile. The original OLS coefficient on Brent returns is not fabricated; it emerges from the archived data. The problem is that its interpretation changes once the regression is evaluated as a time-series model. The exchange-rate level is nonstationary, HAC inference is materially wider than conventional OLS inference, and the full-sample oil coefficient is dominated by an exceptional crisis observation. These are exactly the circumstances in which a single p-value can overstate the stability of an economic relation.

March 2020 is unusually influential because Brent and EGX30 both experienced extraordinary negative returns as the COVID-19 shock disrupted global demand, financial markets, and oil pricing simultaneously. A positive cross-market coefficient generated by that observation does not mean that a rise in oil prices causes a rise in EGX30 returns. It means that during a common global crash, the two assets moved sharply in the same direction. Once the crisis months are omitted, the oil coefficient is approximately zero. The rolling regressions tell the same story: positive exposure is concentrated in 36-month windows that contain the crisis, and later windows turn negative and imprecise.

This interpretation aligns with the structural-shock literature. Kilian (2009) and Kilian and Park (2009) show that oil-price movements generated by global demand carry different information from movements generated by supply or oil-specific demand. Our data do not identify these components. The correct response is not to label the reduced-form Brent return as an exogenous oil shock, but to state what the model can establish: conditional monthly comovement is unstable and crisis-sensitive. This reading is reinforced, not merely asserted, by the diagnostics in Section 4.8: the non-overlapping-window evidence (Table 13) shows the oil-return coefficient is positive in most but not all independent sub-periods and rarely individually significant, and the power analysis shows the full-sample estimate (0.117) sits below the 0.202 effect size this sample could reliably detect. Taken together, the honest conclusion is that the data neither establish a stable positive return effect nor rule one out; they establish that the naive full-sample coefficient is fragile to influential observations and to the sub-period examined, which is a materially different and more defensible claim.

5.2. *Why the volatility result is stronger*

The volatility result survives the very checks that invalidate a stable mean-return conclusion. Oil realized volatility remains positive and statistically significant after removing the acute March–April 2020 episode, after removing all of 2020, and after changing the volatility proxy. That robustness suggests a broader risk-transmission mechanism rather than a coefficient created by one common crash.

Several mechanisms can produce this pattern. Oil turbulence changes uncertainty about energy costs, inflation, fiscal and external balances, transport costs, and the policy path. Even when those forces have offsetting effects on expected stock returns across firms, they can increase the dispersion of possible outcomes and hence aggregate market volatility. An oil-demand shock associated with strong global activity may support some expected cash flows while still increasing uncertainty; a supply disruption may raise costs and inflation risk; geopolitical oil disturbances can influence both energy prices and investor risk tolerance. Aggregating these heterogeneous mechanisms can leave the conditional mean close to zero while raising the conditional variance.

The finding is consistent with the review evidence summarized by Degiannakis et al. (2018) and with the realized-volatility approach of Bastianin and Manera (2018). The present paper does not identify which structural shock drives the Egyptian volatility response, but it demonstrates that second-moment transmission is empirically more stable in this sample than first-moment exposure. That is a useful empirical distinction for an emerging market in which exchange-rate and uncertainty shocks are substantial. The additional diagnostics in Section 4.8 support this reading from three independent angles: the volatility-model VIF values (Table 11) rule out multicollinearity as the reason the VIX level itself enters insignificantly; the Newey–West lag-sensitivity results (Table 3) show the oil-volatility elasticity remains significant at the 1% level across every HAC truncation tested; and the structural-break tests (Table 4) confirm the underlying return relationship is genuinely episode-dependent rather than merely appearing so because of one estimator choice, which is precisely the contrast the volatility channel does not share.

5.3. *Interpreting the exchange-rate channel*

The exchange-rate findings are not a side result. EGP depreciation has a large positive coefficient in the local-currency return regression and absolute FX movement has a positive coefficient in the volatility equation. Together these results indicate that currency adjustment is part of the market-transmission mechanism. A depreciation can increase the pound value of foreign-currency revenues and tangible assets, encourage nominal repricing, and simultaneously signal macroeconomic stress. These forces can raise local-currency equity returns and equity volatility at the same time.

The USD-adjusted analysis prevents the positive local-currency coefficient from being mistaken for a welfare gain. A foreign investor converting the index back into dollars experiences a much weaker cumulative outcome, and the oil-return coefficient is not significant in that specification. For portfolio interpretation, the relevant unit of account therefore matters. Egyptian investors concerned with domestic nominal wealth and international investors concerned with hard-currency wealth do not face identical return exposures. Section 4.8 formally tests whether the local-currency and USD-adjusted oil coefficients differ; they do not, at conventional significance levels (difference = 0.032, HAC $p = 0.224$). This does not weaken the substantive point made here: the two coefficients still carry different economic interpretations for a local-currency versus a foreign-currency investor even when they are not statistically distinguishable, precisely because the underlying currency-repricing channel identified in this section operates on the mean return level, not only on any gap between the two oil coefficients.

5.4. *Relation to Egyptian and regional evidence*

The results help reconcile prior Egyptian studies rather than simply choosing one side. Bataineh (2024) finds little evidence that crude oil prices explain Egyptian stock returns in a VAR framework. Our stationarity-aware and influence-robust results are consistent with that conclusion for average returns. Qutb and Abouelfarag (2024), by contrast, study a period dominated by pandemic and geopolitical uncertainty and find stronger adverse relationships. Our rolling and interaction diagnostics show why a crisis-focused sample can produce substantially stronger oil-market coefficients than a broader 2015-2025 sample.

Shaoubi and El-Shawarby's (2017) emphasis on both return and volatility effects is particularly relevant to our main contribution. The current results indicate that the volatility dimension is the more robust feature of the recent Egyptian sample. Recent MEA work by Eissa et al. (2025a, 2025b) likewise stresses heterogeneity and time variation. Taken together, the evidence argues against a universal Egyptian oil beta and in favor of models that distinguish shock origin, market state, currency conditions, and the moment of the return distribution under study.

5.5. *Persistent risk transmission versus episodic return co-movement*

The rolling analysis sharpens the paper's interpretation of persistence. A relationship should not be described as robust merely because it remains significant in one full-sample regression. For the mean-return channel, influence diagnostics, crisis exclusions, robust regression, and rolling estimation all show substantial instability. For the volatility channel, however, the coefficient is positive across every 48-month window and remains economically similar under crisis exclusions and an alternative volatility definition. This contrast suggests that oil shocks reach the Egyptian market more consistently through uncertainty and risk repricing than through a stable directional cash-flow effect.

This interpretation is consistent with recent regional evidence that emphasizes state dependence, extreme quantiles, and financial stress rather than a universal oil beta (Elsayed et al., 2024; Mensi et al., 2024; Eissa et al., 2025a). It also fits the institutional characteristics of an emerging market in which global uncertainty, exchange-rate adjustment, imported-input exposure, and foreign-investor risk appetite can all change simultaneously. The practical object of interest is therefore not simply the sign of the Brent return, but the intensity and persistence of the risk environment surrounding oil-market turbulence.

6. Implications for Investors and Policy

For portfolio managers, the primary implication is that Brent returns should not be treated as a stable directional signal for the EGX30. A rule that increases or decreases Egyptian equity exposure mechanically in response to monthly oil returns would be poorly supported by the full set of results. The absence of robust Granger predictability and the instability of the rolling coefficient are especially important for tactical allocation. Oil information may matter conditionally, but the conditional state must be modeled rather than assumed.

Oil volatility, however, is relevant for risk budgeting. The estimated elasticity implies that months with unusually turbulent Brent trading tend to be months with higher EGX30 realized volatility even after controlling for currency movement, VIX, and volatility persistence. Risk managers could therefore use oil-market volatility as an external stress indicator when setting position limits, expected-shortfall scenarios, or liquidity buffers for Egyptian equities. The result does not imply that oil volatility alone forecasts all future EGX risk, but it provides a more robust contemporaneous risk signal than the average oil return.

The rolling estimates make this implication more operational. Because the oil-volatility coefficient remains positive across all 48-month windows, a risk dashboard can treat Brent realized volatility as a persistent external stress covariate rather than a signal that only matters during the pandemic. A simple implementation would combine standardized oil volatility, absolute EGP movement, and recent EGX volatility into a monitoring panel, with elevated readings prompting scenario expansion rather than automatic directional trading. The distinction matters: the evidence supports a risk-management use case more strongly than a return-timing strategy.

For asset allocation, the results also imply that oil can affect diversification benefits even when expected return exposure is weak. During high-volatility months, a stronger common risk component between Brent and Egyptian equities can reduce the protection expected from holding the two assets together. This interpretation is consistent with quantile and portfolio evidence from MENA markets reported by Mensi et al. (2024), and it argues for evaluating correlations and hedge ratios conditionally rather than assuming that unconditional diversification properties persist through stress episodes.

Currency risk should be integrated into the same process. Absolute EGP movement is strongly related to realized equity volatility, and the local-versus-USD wealth comparison shows that currency denomination can dominate the investment interpretation. For international investors, equity and currency hedges should therefore be evaluated jointly rather than independently. For domestic institutions, periods of large FX adjustment may require higher risk limits even when local-currency equity prices are rising.

For policymakers and market regulators, the findings suggest that oil shocks should be monitored as part of a broader external-risk dashboard rather than as a deterministic driver of stock-market direction. Global risk sentiment, currency adjustment, and oil-market turbulence can arrive together. Policy communication that separates inflationary or supply-driven oil stress from demand-driven oil movements may reduce the tendency to interpret every oil-price increase as equivalent. At the market-infrastructure level, heightened oil and FX volatility can also serve as a trigger for liquidity surveillance and stress testing.

7. Limitations and Research Agenda

First, the study is observational and reduced-form. Brent returns are not structurally identified supply, aggregate-demand, or oil-specific demand shocks. A stronger causal design would augment the data with global oil production, inventories, global real activity, and an identification strategy such as a structural VAR or external instruments.

Second, 132 monthly observations limit the complexity that can be estimated reliably. The return-stability analysis uses 36-month rolling windows, while the expanded volatility-persistence analysis uses 48-month windows to balance local adaptation against estimation precision. These rolling regressions are descriptive of

parameter stability rather than formal regime estimates, and overlapping windows make adjacent coefficients mechanically dependent. High-dimensional state-space, threshold, or Markov-switching models could overfit this sample unless they are supported by a longer history or stronger regularization.

Third, EGX30 is an aggregate index. Oil-sensitive industries may respond differently from banks, real estate, telecommunications, consumer goods, or exporters. Sector-level or firm-level data could determine whether the weak aggregate mean effect masks offsetting heterogeneous exposures. A referee raises this point directly: a sector- or firm-level decomposition (energy, banking, real estate, telecommunications, consumer goods) would show whether the aggregate index result masks heterogeneous industry exposures, but it requires firm-level panel data outside the archived aggregate-index workbook used here, and we did not have access to a verified firm-level dataset for this revision. We flag it as the most direct next step for future work rather than attempt it with data not available to us.

Fourth, the exchange-rate control uses the month-end EGP/USD midpoint supplied in the workbook. Egypt experienced discrete exchange-rate reforms and parallel-market pressures that a single official monthly series may not fully capture. Future work could distinguish official and market rates, inflation-adjusted returns, and hedged versus unhedged investor positions.

Fifth, realized volatility is computed from daily closing returns and therefore does not exploit intraday information. Higher-frequency data would provide a more efficient realized-variance measure and permit jump decomposition. Even so, the daily-based result is robust to an alternative monthly standard-deviation measure and to exclusion of 2020, which reduces concern that it is purely a measurement artifact.

Sixth, the raw workbook did not retain row-level source URLs for every historical observation. The replication package preserves the exact supplied workbook, cleaned analysis data, date-reconstruction rule, and executable analysis code, but final publication should archive the original vendor or official-source links to strengthen provenance documentation.

Finally, the VIX-dependent interaction reported in Appendix A is exploratory. It suggests that the oil-return coefficient becomes more positive in high-VIX states, but the interaction disappears when 2020 is excluded. This is precisely why the paper does not build its main conclusion around a crisis threshold. Future work with a longer sample could test state dependence more rigorously using pre-specified regimes or time-varying parameter methods. The new structural-break tests in Section 3.4 (Table 4) reinforce this limitation rather than resolve it: all three candidate currency-regime dates, and an independently detected break around April 2023, show statistically significant parameter instability, so the reduced-form return coefficient should continue to be interpreted regime-by-regime rather than as a single stable population parameter. Relatedly, the non-overlapping-window evidence in Section 4.8 (Table 13) has only two and five independent windows at the 48- and 24-month block lengths respectively, which is too few to characterize the full distribution of the oil-return coefficient across regimes with any precision; a longer sample would let this be examined with proper out-of-sample cross-validation rather than a small number of manually chosen blocks. Finally, the formal GARCH-family volatility modeling and full structural-VAR oil-shock identification discussed in Sections 2.1 and 3.5 remain unaddressed in this revision for the data-availability reasons given there, and we repeat them here as concrete, scoped items for future work rather than leave them as an open-ended aspiration.

A further, more technical item concerns the distributional assumptions used for inference. The residual diagnostics reported in Section 4.8 (Table 12) show that the return-model residuals are mildly skewed and leptokurtic (skewness -0.39 , kurtosis 3.88 , Jarque–Bera $p = 0.022$); the Newey–West HAC standard errors used throughout this paper are robust to this non-normality for inference on the mean coefficients, but they do not themselves provide a full distributional model of return or volatility tail risk. A growing statistical literature develops flexible parametric families – built from fractional and conformable-calculus generalizations, transmuted distributions, and generalized-order-statistics characterizations – designed specifically to capture asymmetry and heavy tails more flexibly than the normal distribution or standard exponential-family alternatives (Gad et al., 2026a; Gad and Hamouda, 2026; Ali et al., 2025; Eldeeb et al., 2025; Abdul-Moniem et al., 2024a; Abdul-Rahman et al., 2024; Gad et al., 2023; Abdul-Rahman et al., 2023; Hassan et al., 2020). Fitting one of these flexible families to the conditional distribution of EGX30 returns or realized volatility – for instance, replacing the normal innovation assumption in the GARCH-family extension already flagged above (Section 3.5) with a fractional or transmuted

alternative – is a concrete, citable direction for future work on tail-risk characterization in this setting; we note it explicitly as an item on the research agenda rather than attempt it within the scope of the present revision, which remains focused on the mean- and second-moment oil-equity-currency relationships that motivated the paper.

8. Conclusion

This paper re-examines the 2015-2025 relationship between Brent crude oil and the Egyptian EGX30 using a reproducible design that separates mean returns from market risk. The original positive OLS oil coefficient can be replicated, but the underlying levels model contains a nonstationary exchange-rate regressor and its explanatory power was overstated in the original report. Once exchange-rate changes and HAC inference are used, the oil-return coefficient is positive but statistically weak. More importantly, it is highly sensitive to March 2020, disappears when the acute COVID shock is removed, is not robust in a Huber regression, does not generate two-lag Granger predictability, and is not significant for USD-adjusted EGX30 returns.

The main result lies in volatility. Monthly Brent realized volatility is strongly associated with EGX30 realized volatility after controlling for VIX, currency movement, and lagged equity volatility. The estimated full-sample elasticity is approximately 0.37 and remains statistically significant when March-April 2020 or all of 2020 is excluded and when volatility is measured differently. Most importantly, the new time-varying analysis finds a positive oil-volatility coefficient in all 84 rolling 48-month windows, with 81 windows significant at 10% and 67 at 5%. Egyptian equity exposure to oil therefore appears more stable in the second moment than in the conditional mean, and that conclusion is not driven by a single crisis episode.

The practical implication is a change in emphasis. Investors should be cautious about using oil-price direction as a stable forecast signal for EGX30 returns, but oil-market turbulence is informative for Egyptian equity risk across a wide range of rolling samples. Exchange-rate movement and global uncertainty remain essential parts of the same macro-financial system, while the local-versus-USD comparison shows that nominal equity gains can be materially altered by currency repricing. Treating oil volatility, currency stress, and equity-market persistence jointly provides a more credible account of how international energy shocks reach an emerging equity market and offers a stronger basis for stress testing than a single unconditional oil-return beta. This revision strengthens that conclusion with six additional checks requested in review: the exchange-rate level and the EGX30 are not cointegrated, so the stationarity-aware specification is the appropriate primary one rather than an ad hoc alternative to the levels regression; inference on both the return and volatility oil coefficients is stable across HAC lag truncations of 0–8 months; the return relationship exhibits statistically significant structural breaks around each of Egypt's three currency-regime changes; multicollinearity does not explain the volatility model's insignificant VIX coefficient; the oil-volatility channel's out-of-sample forecast improvement over a simple autoregressive benchmark, while directionally favorable, is not statistically significant with the sample available; and the null return result is more properly read as inconclusive given the study's limited power than as evidence of no effect. None of these checks overturn the paper's central distinction between a fragile return channel and a persistent volatility channel; if anything, they show that distinction survives scrutiny specifically designed to break it.

Data Availability and Reproducibility

The submission archive contains the original source workbook supplied by the authors, a cleaned 2015-2025 monthly analysis file, daily EGX30 and Brent data used to construct realized volatility, leave-one-month-out and rolling-coefficient files, a structured results file, and executable Python code for the reported analysis and figures. The data-processing note documents the date-localization correction. The authors should verify and archive the final raw-source URLs or vendor citations before submission.

Declarations

Funding: No funding information was supplied in the project files; authors should confirm the final declaration at submission.

Conflict of interest: The final declaration must be confirmed by all authors before submission.

Ethics: The study uses publicly available market and macro-financial time-series data and contains no human-participant or personally identifiable information.

A. Integrated Robustness and Data-Audit Evidence

Appendix A is part of this manuscript and is not a separate supplementary file. It reports detailed diagnostics that support the interpretation in the main text.

A.1. Correlations and influential observations

Table A.I. Pearson correlations in the 2015-2025 monthly sample. Correlations are descriptive unconditional associations.

Variable	EGX ret	Oil ret	VIX	FX change	EGX RV	Oil RV
EGX ret	1.00	0.27	-0.28	0.36	-0.13	-0.18
Oil ret	0.27	1.00	-0.31	0.06	-0.38	-0.31
VIX	-0.28	-0.31	1.00	-0.05	0.24	0.56
FX change	0.36	0.06	-0.05	1.00	0.38	0.07
EGX RV	-0.13	-0.38	0.24	0.38	1.00	0.50
Oil RV	-0.18	-0.31	0.56	0.07	0.50	1.00

Table A.II. Most influential observations in the stationarity-aware return regression according to Cook's distance.

Month	Cook D	EGX return	Oil return	VIX	FX change
2020-03	1.875	-30.451	-79.824	53.540	0.870
2024-03	1.082	-7.456	4.513	13.010	42.417
2016-11	0.871	31.171	4.395	13.330	70.591
2020-04	0.060	9.538	10.549	34.150	-0.002
2016-07	0.046	13.967	-15.704	11.870	0.002
2020-05	0.032	-3.215	33.511	27.510	0.883

A.1.1. Additional visual diagnostics for volatility transmission The following figures provide two complementary diagnostics for the second-moment result. The cross-market scatter shows that the positive association between log Brent and log EGX30 realized volatility remains visible after 2020 is removed. The observed-versus-fitted plot shows how closely the parsimonious dynamic volatility model tracks broad movements in monthly EGX30 risk without claiming precise point forecasting.

A.2. Quantile-return evidence

The quantile results reinforce the mean-regression conclusion. Oil returns are not significant at the 0.10, 0.25, 0.50, 0.75, or 0.90 conditional quantiles. VIX is most strongly negative in the lower-to-middle part of the return distribution, while EGP depreciation is positively associated with the median and upper return quantiles.




Figure A.II. Cross-market volatility comovement. Points plot monthly log Brent realized volatility against log EGX30 realized volatility; color indexes calendar year and marker size increases with absolute currency movement. Solid and dashed fitted lines use all months and the sample excluding 2020, respectively.




Figure A.III. Observed and model-implied EGX30 realized volatility. The fitted series is generated by the main dynamic volatility regression and is shown as a descriptive fit diagnostic rather than an out-of-sample forecast.

Table A.III. Quantile regressions for monthly EGX30 returns. Oil returns are not statistically significant across the reported conditional quantiles

Quantile	Oil coef.	Oil p	VIX coef.	VIX p	FX coef.	FX p
0.1	0.143	0.382	-0.325	0.085	-0.053	0.788
0.25	0.115	0.188	-0.382	0.002	0.287	0.003
0.5	0.037	0.503	-0.225	0.016	0.411	<0.001
0.75	0.013	0.808	-0.015	0.878	0.375	<0.001
0.9	-0.077	0.364	-0.091	0.527	0.358	<0.001

These estimates are exploratory because quantile-regression standard errors do not address every form of serial dependence in a short monthly sample.

A.3. Exploratory VIX-dependent oil exposure

Table A.IV. Exploratory state-dependence test. The full-sample interaction is strong but disappears when the 2020 crisis is excluded, so it is not treated as a stable threshold effect.

Specification	Oil x VIX interaction	HAC SE	p-value
Full sample interaction	0.0099	0.0022	<0.001
Exclude Mar-Apr 2020	0.0069	0.0072	0.340
Exclude all 2020	0.0037	0.0106	0.726




Figure A.I. Exploratory marginal oil-return coefficient across VIX levels in the full sample. The band is an approximate 95% HAC confidence interval. Because the interaction disappears when 2020 is excluded, this figure is descriptive of crisis concentration rather than evidence of a structural threshold.

Table A.V. Reproducible data-construction and date-localization audit.

Audit item	Series/range	Rule
Raw market sheets	EGX30, Brent, VIX	Daily historical prices in supplied workbook
Raw currency sheet	USD/EGP buy and sell	Daily/available observations in supplied workbook
Market-date repair	Numeric serial dates in EGX/Brent/VIX	Stored month/day interchanged to restore chronological trading sequence; text dates parsed mm/dd/yyyy
Exchange-rate dates	Excel serial dates	Used without month/day interchange
Monthly alignment	2015-01 to 2025-12	Last market observation each month; exchange rate matched as-of last available date
Return calculation	EGX, Brent, FX	100 x log price ratio
Volatility calculation	EGX and Brent daily returns	Square root of monthly sum of squared daily percentage log returns
Primary return N	132 months	No missing aligned monthly observations after processing
Primary volatility N	131 dynamic observations	One observation lost to lagged EGX realized volatility

A.4. Data-construction audit

The date repair is not an econometric adjustment and does not select observations according to outcomes. It corrects a locale-format inconsistency visible in the workbook itself. Applying the documented rule restores a continuous trading-day sequence and reproduces the original project's monthly return descriptive statistics. All processed analysis rows are retained in the replication package.

REFERENCES

1. Apergis, N., Miller, S.M. (2009), Do structural oil-market shocks affect stock prices? *Energy Economics*, 31(4), 569–575.
2. Basher, S.A., Haug, A.A., Sadorsky, P. (2012), Oil prices, exchange rates and emerging stock markets. *Energy Economics*, 34(1), 227–240. <https://doi.org/10.1016/j.eneco.2011.10.005>
3. Bastianin, A., Manera, M. (2018), How does stock market volatility react to oil price shocks? *Macroeconomic Dynamics*, 22(3), 666–682. <https://doi.org/10.1017/S1365100516000353>
4. Bataineh, K. (2024), Crude oil prices and the Egyptian economy: evidence from the stock market. *International Journal of Energy Economics and Policy*, 14(1), 383–392. <https://doi.org/10.32479/ijee.15198>
5. Degiannakis, S., Filis, G., Arora, V. (2018), Oil prices and stock markets: a review of the theory and empirical evidence. *The Energy Journal*, 39(5), 85–130. <https://doi.org/10.5547/01956574.39.5.sdeg>
6. Elgharib, W.A., Elmarzouky, M., Shohaieb, D. (2025), How do stock returns respond to a currency devaluation announcement? *Journal of Risk and Financial Management*, 18(12), 1–17, Article 663. <https://doi.org/10.3390/jrfm18120663>
7. Eissa, M.A., Al Refai, H., Chortareas, G. (2025a), Stock-market responses, oil-price dynamics, and geopolitical risk in the MEA region. *Applied Economics*. <https://doi.org/10.1080/00036846.2025.2545019>
8. Eissa, M.A., Al Refai, H., Chortareas, G. (2025b), Oil dependency, political instability and the stock market: a perspective from the Middle East and Africa. *Journal of Policy Modeling*, 47(6), 1201–1221. <https://doi.org/10.1016/j.jpolmod.2025.09.009>
9. Elsayed, A.H., Sohag, K., Sousa, R.M. (2024), Oil shocks and financial stability in MENA countries. *Resources Policy*, 89, 104653. <https://doi.org/10.1016/j.resourpol.2024.104653>
10. Filis, G., Degiannakis, S., Floros, C. (2011), Dynamic correlation between stock market and oil prices: the case of oil-importing and oil-exporting countries. *International Review of Financial Analysis*, 20(3), 152–164. <https://doi.org/10.1016/j.irfa.2011.02.014>
11. Heshmati, A., Khayat, N., Ors, S. (2020), The impact of the COVID-19 pandemic on the Egyptian stock market. *Emerging Markets Finance and Trade*, 56(10), 2262–2275.
12. Kilian, L. (2009), Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3), 1053–1069. <https://doi.org/10.1257/aer.99.3.1053>
13. Kilian, L., Park, C. (2009), The impact of oil price shocks on the U.S. stock market. *International Economic Review*, 50(4), 1267–1287. <https://doi.org/10.1111/j.1468-2354.2009.00568.x>
14. Le, T.H., Chang, Y. (2015), Effects of oil price shocks on the stock market performance: do nature of shocks and economies matter? *Energy Economics*, 51, 261–274.
15. Mahfouz, H.C., Youssef, E., Marzouk, W. (2021), Oil prices and stock market returns: evidence from Egypt. *Journal of Economic Studies*, 48(4), 1045–1066.
16. Mensi, W., Mishra, T., Ko, H.-U., Vo, X.V., Kang, S.H. (2024), Quantile dependence and portfolio management between oil, gold, silver, and MENA stock markets. *Research in International Business and Finance*, 70, 102296.

<https://doi.org/10.1016/j.ribaf.2024.102296>

17. Newey, W.K., West, K.D. (1987), A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
18. Qutb, R., Abouelfarag, H.A. (2024), How oil price shocks affect the Egyptian stock market performance in the context of recent challenges. *Egyptian Journal of Development and Planning*, 32(2). <https://doi.org/10.21608/inp.2024.240890.1064>
19. Sadorsky, P. (1999), Oil price shocks and stock market activity. *Energy Economics*, 21(5), 449–469.
20. Shaoubi, A.A., El-Shawarby, S.A. (2017), Crude oil price shock effects on stock market return and volatility: evidence from Egypt. *OPEC Energy Review*, 41(4), 401–427.
21. Srouf, H.M., Youssef, S., Darwish, N.M., Abouelenein, M.F., Eissa, O.M.S. (2026), Investor sentiment and stock market volatility in the Egyptian Exchange: a GARCH and EGARCH analysis with persistence and asymmetry checks (2009–2025). *Statistics, Optimization and Information Computing*, 16(2), 1311–1333. <https://doi.org/10.19139/soic-2310-5070-4151>
22. Abdullah, N.A., Algamal, Z. (2026), Enhancing interval forecasting accuracy of Iraqi stock market prices based on ν -support vector regression. *Statistics, Optimization and Information Computing*, Online First. <https://doi.org/10.19139/soic-2310-5070-3316>
23. Darmanto, D., Efendi, A.H. (2026), Estimating value at risk of the Indonesian Sharia stock index using asymmetric GARCH-MIDAS. *Statistics, Optimization and Information Computing*, 16(3), 2200–2214. <https://doi.org/10.19139/soic-2310-5070-3899>
24. Gad, K. A. E., Atia, F. A., Moussa, M., Hamouda, H. M., & Abouelenein, M. F. (2026). The Fractional Exponential Distribution: A Gamma Subfamily from Conformable Calculus. *Results in Engineering*, 111224.
25. Gad, K. A. E., & Hamouda, H. M. (2026). A New Flexible Transmuted Distribution: Theory and Application. *Statistics, Optimization and Information Computing*, 4464, 22.
26. Ali, A. A., Ali, M., Gad, K. A., & Assar, S. M. (2025). Some Characterizations Based on Generalized Order Statistics from Power Inverted Topp–Leone Distribution. *Advances in Nonlinear Variational Inequalities*, 28(6s), 349–365.
27. Eldeeb, A. S., Abdul-Moniem, I. B., & Gad, K. A. E. (2025). A New Discrete Power Function Distribution with Its Mathematical Properties. *Journal of Information Systems Engineering & Management*, 10(44s).
28. Abdul-Moniem, I. B., Gad, K. A. E., & Assar, S. M. (2024). Characterization of Phani Distribution Based on Generalized Order Statistics. *Sohag Journal of Sciences*, 9(1), 32–39.
29. Abdul-Rahman, A. A., Abdul-Moniem, I. B., Gad, K. A. E., & Assar, S. M. (2024). Some Characterizations Based on Generalized Order Statistics from Weibull-Family of Life Distributions. *56th Annual International Conference on Data Sciences*.
30. Gad, K. A. E., Abdul-Moniem, I. B., & Assar, S. M. (2023). Some Characterizations Based on Generalized Order Statistics from Weibull-Weibull Distribution. *The Egyptian Statistical Journal*, 67(2), 1–16.
31. Abdul-Rahman, A. A., Abdul-Moniem, I. B., Gad, K. A. E., & Assar, S. M. (2023). Some Characterizations Based on GOS from Weibull-Weibull Distribution (Part II). *Egyptian Statistical Journal*, 67(2-2023).
32. Hassan, A. S., Abdul-Moniem, I. B., & Gad, K. A. E. (2020). A Generalized Transmuted Moment Exponential Distribution: Properties and Application. *Academic Journal of Applied Mathematical Sciences*, 6(5), 41-52.