

Article

A Geopolitical Shock, Oil Exposure, and Heterogeneous Equity-Market Responses: Firm-Level Evidence from Saudi Arabia

Ibraheem Alaskar ^{1,*} , Ibrahim N. Khatatbeh ^{2,3,*}  and Ahmad Bash ¹ 

¹ Department of Insurance and Banking, College of Business Studies, Public Authority for Applied Education and Training (PAAET), Kuwait City 12064, Kuwait; ay.bash@paaet.edu.kw

² Department of Banking and Financial Sciences, Business School, The Hashemite University, Zarqa 13133, Jordan

³ Department of Accounting and Finance, School of Business, American University of Ras Al Khaimah (AURAK), Ras Al Khaimah 72603, United Arab Emirates

* Correspondence: i.alaskar@paaet.edu.kw (I.A.); ibrahim.khatatbeh@aurak.ac.ae or ibrahim.khatatbeh@hu.edu.jo (I.N.K.)

Abstract

We study the firm-level equity-market response of 230 Tadawul Main Market firms to the outbreak of the U.S.–Israeli military action against Iran on 28 February 2026, using a market-model event study with the domestic TASI benchmark. Because a single geopolitical shock combines a market-wide risk-premium effect with an oil-price effect that accrues to oil-linked firms, the aggregate response is theoretically ambiguous, and we find that its sign in fact depends on how firms are weighted. Equal-weighted abnormal returns are negative around the event, whereas a value-weighted portfolio, dominated by the large oil-linked firms, rises significantly, reaching a cumulative abnormal return of +4.65% at the [0, +2] window (versus −2.84% equal-weighted). In cross-sectional regressions controlling for firm size, leverage, profitability, and sector fixed effects, oil-linked firms earn cumulative abnormal returns approximately 9.0 percentage points higher than non-oil firms over the [0, +5] window, a differential robust to the exclusion of Saudi Aramco, to median (robust) regression, to winsorization, and to leave-one-out estimation. We show that inference in this common-event setting is sensitive to cross-sectional dependence: under a Kolari–Pynnönen adjustment, only the immediate [0, +1] and [0, +2] aggregate windows remain statistically significant, so we treat the cross-sectional differential, rather than the aggregate response, as the paper’s most robust result. Overall, the evidence documents economically large cross-sectional heterogeneity that aggregate index responses conceal; it is consistent with, but does not separately identify, a common risk-premium effect and an oil-related cash-flow effect.

Keywords: event study; geopolitical risk; oil exposure; cross-sectional heterogeneity; emerging equity markets; Saudi Arabia; Tadawul Main Market



Received: 21 August 2026

Revised: 18 September 2026

Accepted: 21 September 2026

Published: 23 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

On the morning of 28 February 2026, coordinated airstrikes were launched against Iranian nuclear and military infrastructure, including the targeted killing of Supreme Leader Ali Khamenei ([BBC News 2026](https://www.bbc.com/news/world-middle-east-67890123)). The strikes triggered an immediate Iranian retaliation: the Islamic Revolutionary Guard Corps announced the closure of the Strait of Hormuz to vessels of “unfriendly nations,” attacks on the Al-Udeid airbase in Qatar, and missile strikes

on Qatari LNG infrastructure at Ras Laffan ([Middle East Council on Global Affairs 2026](#)). Within four weeks, Qatar declared force majeure on long-term LNG contracts ([Reuters 2026b](#)), Brent crude rose to a four-year high ([Reuters 2026a](#)), and the Gulf Cooperation Council countries found themselves at the geographic center of an escalating regional conflict. We refer to this sequence throughout as the 2026 geopolitical episode, and to its opening as the conflict outbreak.

For the equity market of Saudi Arabia, the largest crude exporter and the most important economy in the GCC, the episode raised an immediate empirical question: how does a regional conflict that simultaneously raises oil prices and increases geopolitical risk affect the valuation of equities in an oil-exporting economy? Standard event-study intuition would predict an unambiguous negative response in equity prices. Yet the same shock that elevates regional risk also delivers a positive cash-flow surprise to the country's most important sector, crude production and downstream petrochemicals. The net response of Saudi equities is therefore theoretically ambiguous and depends on the relative weights of these two transmission channels in the cross-section of listed firms. This ambiguity is not incidental to our framing: it is the theoretical prediction of the two-channel framework we set out in Section 2. The efficient market hypothesis ([Fama 1970](#)) supplies the null against which the event is tested, while the risk-premium channel ([Pastor and Veronesi 2013](#); [Caldara and Iacoviello 2022](#)) and the cash-flow channel ([Kilian 2009](#); [Kilian and Park 2009](#)) supply the two economically motivated forces whose relative weight determines the aggregate response and whose differential impact across firms determines the cross-sectional response. We treat this episode as an exploratory, single-event case study: our aim is to characterize how one major oil-exporting market priced one specific geopolitical shock at the firm level, not to estimate parameters that generalize across conflicts.

In this paper, we estimate the response of the Tadawul Main Market to the conflict outbreak using a standard market-model event-study framework, and exploit heterogeneity in firm-level exposure to examine how the aggregate response is distributed across firms. We construct a sample of 230 Main Market firms with sufficient pre-event trading history and apply the market model with the Tadawul All Share Index (TASI) as the benchmark. Cumulative average abnormal returns (CAARs) are computed over eleven event windows spanning the period from ten trading days before the event to twenty trading days after. Statistical inference proceeds through three procedures: the parametric cross-sectional *t*-test on the mean cumulative abnormal return (CAR); the non-parametric Wilcoxon signed-rank test, a test of a symmetric location shift around zero; and, because a single common event date induces cross-sectional correlation among abnormal returns, the [Kolari and Pynnönen \(2010\)](#) adjusted test, which corrects the inflation of significance that the first two tests would otherwise suffer. As we show, this correction is decisive: most of the aggregate significance does not survive it. The construction and interpretation of this adjusted test are set out in Section 3.2.2. To characterize the cross-sectional structure of the response, we then run cross-sectional regressions of firm-level CARs on an indicator for oil-linked sectors, firm size, leverage, profitability, and sector fixed effects.

Three features of the design should be stated at the outset. First, the event we study is not treated as a fully unanticipated shock. Regional tensions were elevated in the weeks preceding the outbreak, and the pre-event abnormal returns we document are consistent with partial anticipation of escalation in general terms; our identifying assumption is the weaker one that the specific timing and severity of the outbreak were not priced before it occurred, and Section 3.1 sets out the evidence for and against that assumption. Second, the paper documents the cross-sectional pattern of abnormal returns; it does not separately identify the economic mechanisms that generate it. We therefore distinguish throughout between what the results show and what they are consistent with. Third, oil-linked firms,

a category we define broadly to include crude producers, refiners, oilfield services, oil shipping, and petrochemicals, earned CARs approximately 9.0 percentage points higher than non-oil firms over the $[0, +5]$ window, controlling for firm size, leverage, profitability, and sector fixed effects. This effect is economically large, statistically significant at the 1% level using heteroskedasticity-consistent (HC1) standard errors, and robust to excluding Saudi Aramco, to median regression, to winsorization, and to leave-one-out estimation. This cross-sectional differential, rather than the fragile aggregate response, is the paper's central and most robust finding.

Our findings contribute to three strands of the literature. The event-study literature on geopolitical shocks and equity markets is large but has focused primarily on the responses of importing or politically neutral economies, and has been conducted overwhelmingly at the index level. Our contribution is to resolve the response of a major oil-exporting market to a regional conflict on its own doorstep at the level of the individual firm, and to show that the differential between oil-linked and non-oil firms is economically large at a horizon at which the aggregate index response is close to zero. We also contribute to the literature on the Saudi market specifically. That literature is not unexplored: [Ayadi and Ben Mbarek \(2025\)](#) estimate the dynamic response of TASI returns to Saudi-specific geopolitical risk and oil-price uncertainty using monthly local projections over 1998–2024, [Medhioub \(2025\)](#) examines investor herding in MENA markets under global geopolitical risk, and [Alqahtani and Chevallier \(2020\)](#) document oil-driven volatility spillovers into GCC equities. What distinguishes the present study is its unit of observation and its horizon: these studies characterize the aggregate market's average response over months or its second-moment linkages, whereas we measure the dispersion of firm-level responses over the trading days immediately surrounding a single identified event. The two approaches are complementary, and our results are informative about a margin that index-level and longer-horizon designs cannot observe.

The remainder of the paper is organized as follows: Section 2 reviews the literature, Section 3 describes the data and methodology, Section 4 reports the results, Section 5 discusses their interpretation and limitations, and Section 6 concludes.

2. Literature Review

2.1. Geopolitical Risk and Asset Pricing

The theoretical foundation for examining how markets price a sudden geopolitical event is the efficient market hypothesis ([Fama 1970](#)), which holds that asset prices fully and rapidly incorporate all publicly available information. Under the semi-strong form, an unanticipated shock should be reflected in prices during the first trading session at which the news can be acted upon, leaving no exploitable abnormal returns thereafter. Our empirical design follows the conventional market-model event-study framework of [Brown and Warner \(1985\)](#) and [MacKinlay \(1997\)](#).

[Caldara and Iacoviello \(2022\)](#) define geopolitical risk as the threat and realization of adverse events arising from wars and interstate tensions, and show that elevated readings of their news-based index are followed by weaker investment and greater downside risk. The first channel through which such risk reaches prices is the discount rate: when investors feel less safe, they demand a larger premium for holding equities, the discount rate rises, and prices fall. [Pastor and Veronesi \(2013\)](#) provide a formal theoretical framework for this mechanism focusing on political uncertainty. Recent evidence extends the political-risk channel to security-level effects on financial-market functioning: [Liu et al. \(2025\)](#) document that firm-level political risk impairs liquidity in the corporate bond market, showing that political-risk exposure has measurable financial-market consequences at the security level and not only at the aggregate level. The second channel runs through

the cash flows themselves, as a shock alters expected earnings through demand, supply disruption, sanctions, or commodity prices. A third mechanism operates through investor sentiment. [Baker et al. \(2016\)](#) show that economic policy uncertainty increased stock price volatility and reduced investment and employment, while [Baker and Wurgler \(2006\)](#) show that investor sentiment affects the cross-section of stock returns. These channels rarely operate uniformly: [Bossman and Gubareva \(2023\)](#) show that the effect of geopolitical risk on equities is both market-specific and asymmetric across calm, stressed, and bullish states, so that the same shock can depress some markets while leaving others resilient. [Medhioub \(2025\)](#) reports a related asymmetry in investor behavior across MENA markets, finding that herding is concentrated in down-market quantiles and that the measured effect of global geopolitical risk on return dispersion is not uniform across countries, with Saudi Arabia among the markets in which it is not detected.

2.2. Equity Market Reactions to Military Conflicts

The empirical literature on conflict and equity prices is large and far from unanimous. Three regularities recur. First, markets usually fall on impact. The 2022 Russian invasion of Ukraine generated negative cumulative abnormal returns across sectors and indices in ten countries close to the conflict ([Aguilar Rodrigues et al. 2025](#)), across the stock market indices of forty-seven countries ([Boubaker et al. 2022](#)), and across ninety-four countries worldwide ([Boungou and Yatie 2022](#)). Second, the reaction is anything but uniform, with more globally integrated economies experiencing the largest losses. [Federle et al. \(2026\)](#) sharpen this regularity into a “proximity penalty”: in a four-week window around the Russian invasion, equity returns fall the closer a country is to Ukraine, a gradient that holds even across firms within a country and that trade linkages explain only about two-thirds of it, the remainder tracing to military-disaster risk. [Hoffmann et al. \(2025\)](#) study seventeen invasions and confirm that emerging markets are considerably more sensitive than advanced ones. Third, the effect does not last. [Park and Newaz \(2018\)](#) find that terrorist attacks have barely any measurable effect on average, and [Markoulis et al. \(2020\)](#) argue that even significant reactions tend to disappear within days. The sign can even reverse depending on the attack outcome: [Bash et al. \(2021\)](#) find that a failed terror attack on the Pakistan Stock Exchange generated positive CAARs, suggesting that a thwarted threat signals market resilience rather than vulnerability.

Wars carry macroeconomic and supply consequences that a single bombing does not, and a pooled average is of limited informational value when the underlying reactions offset one another across countries and sectors. [Martins \(2024\)](#) finds that the October 2023 Middle East escalation led to an increase in commodity futures prices. [Pandey \(2025\)](#) documents falling equity markets and a flight to the dollar following the April 2024 regional escalation, with the magnitude of the losses varying according to market maturity and geographic proximity. [Pandey et al. \(2024\)](#) show that country-level characteristics shape the reaction to the October 2023 conflict, with markets in happier countries proving less sensitive to the shock. [Ijaz et al. \(2025\)](#) find that the October 2023 conflict produced significant negative abnormal returns concentrated in markets with direct economic ties to the belligerents, including the United Arab Emirates, Bahrain, and Kuwait, while oil was among the few assets to react strongly, underscoring that exposure, rather than geography alone, governs the cross-section of the response. This literature has, however, examined the equity market of a major oil exporter far less often, and almost never at the firm level.

2.3. Oil Prices as a Transmission Mechanism

The effect of an oil-price change on equities depends on what caused the change. [Kilian \(2009\)](#) established that not all oil shocks are alike, and [Kilian and Park \(2009\)](#) carried

the distinction into the stock market: a price rise driven by strong demand lifts equities while one driven by supply disruption or precautionary hoarding pushes them down. [Ready \(2018\)](#) arrives at a similar conclusion for supply-driven increases. A conflict that raises crude prices by threatening supply is bad news for the discount rate, through the geopolitical risk it creates, and good news for the cash flows of any firm whose revenue is denominated in barrels. [Smales \(2019\)](#) documents this mechanism, with geopolitical risk lifting oil returns while depressing equities.

At the national index level, the country's trade position helps determine the direction of the response. [Mokni \(2020\)](#) shows that oil-specific demand shocks help exporters and hurt importers. [Sokhanvar and Bouri \(2023\)](#) find that conflict-driven energy-price hikes appreciate the currencies of commodity exporters relative to those of importers, evidence that the same supply shock redistributes value toward exporting economies. [Sakaki \(2019\)](#) finds that real oil-price shocks depress every S&P 500 sector except energy and utilities. At the firm level, however, the effect is heterogeneous: producers and refiners gain, while firms that buy energy or sell to squeezed consumers may well lose, exporter or not.

2.4. Saudi Arabia and GCC Financial Markets

Saudi Arabia provides a natural setting in which to examine the interaction of these channels. It is the world's largest crude exporter and the largest and most liquid stock market in the Middle East, admitted to the MSCI Emerging Markets index in 2019. [Alqahtani et al. \(2020\)](#) find that crude oil prices predict GCC stock market returns, and [Abdou et al. \(2024\)](#) report that oil became the dominant driver of the Saudi market after 2006. [Alqahtani and Chevallier \(2020\)](#) show, using a dynamic conditional correlation framework, that GCC equity returns are most strongly correlated with oil-market shocks and that Saudi Arabia is among the most responsive markets in the region, with correlations strengthening during crisis periods; this establishes at the second moment the oil sensitivity whose first-moment cross-sectional distribution we examine here. The effect is sign-dependent across sectors: [Amin et al. \(2022\)](#) find that positive oil shocks lift petrochemical and energy valuations while weighing on banks and financial-services firms, the very split this paper examines. Consistent with this sectoral divide, [Bouri et al. \(2023\)](#) show that oil implied volatility and geopolitical risk move GCC stock sectors in a regime-dependent way, with the returns of Energy, Materials, Industrials, and Financials turning negative in bear markets and positive in bull markets.

The geopolitical evidence is contradictory. [Alqahtani et al. \(2021\)](#) find that Saudi Arabia barely reacts to local geopolitical risk over the long run, a resilience attributed to fiscal strength and strategic weight. [Elsayed et al. \(2021\)](#) cast Saudi Arabia as one of the principal transmitters of return spillovers across MENA markets during crises, and [Saâdaoui et al. \(2023\)](#) detect clear causality from geopolitical factors to the Saudi market precisely when a crisis is underway. [Ayadi and Ben Mbarek \(2025\)](#) bring local-projection methods to this question, estimating the dynamic response of TASI returns to Saudi-specific geopolitical risk, oil-price uncertainty, climate-policy uncertainty, and U.S. monetary-policy uncertainty on monthly data from 1998 to 2024. They report that the sign of the geopolitical-risk response reverses across sub-periods while oil-price uncertainty depresses returns throughout, evidence that the market's relationship to geopolitical risk is neither stable nor uniform. Their design identifies average dynamic responses at a monthly frequency for the market as a whole; ours identifies the dispersion of responses across firms at a daily frequency around one dated event. The two are complementary, and the instability they document is one reason to be cautious about generalizing from a single episode, as we note in Section 5. [Al Nafea \(2026\)](#) combines event-study and GARCH models to trace the Saudi market's reaction to recent regional escalation and finds that oil prices

act as a dual moderator, stabilizing energy-linked sectors while intensifying pressure on non-energy sectors; that analysis, however, remains at the level of TASI and its sectoral indices, whereas we resolve the response to the individual firm and examine variation in oil exposure within sectors.

The Saudi and GCC evidence establishes that oil and geopolitical risk both matter for this market and that their effects are regime- and period-dependent. What it has not established is how a single conflict-onset event is distributed across the firms of that market. The work is overwhelmingly index-level or sector-index-level; firm-level heterogeneity within an oil-exporting economy around a dated geopolitical event has scarcely been examined. That is the gap this paper addresses.

The framework we use is deliberately parsimonious. One shock reaches the market through two channels that work against each other: the risk-premium channel drags the whole market down, while the cash-flow channel, working through a supply-driven rise in crude, lifts firms whose earnings move with oil. The aggregate response is the time-varying balance of the two and cannot be signed in advance, whereas the cross-sectional response should line up with firms' exposure to oil. We emphasize that this framework motivates the empirical design and organizes the interpretation of the results; the design does not permit the two channels to be estimated separately, as Section 3.2.2 explains.

The preceding framework yields two testable hypotheses. Because a supply-driven, war-related oil-price shock raises the discount rate applied to all equities (the risk-premium channel) while simultaneously raising the expected cash flows of oil-linked firms (the cash-flow channel), and because these two channels act in opposite directions on the aggregate market, our first hypothesis is deliberately two-sided:

H1 (aggregate response). *The onset of the conflict generates a non-zero aggregate abnormal return on the Tadawul Main Market, but its sign is not signed a priori: it depends on the relative magnitude of the risk-premium and cash-flow channels and on the weighting scheme applied across firms.*

H2 (cross-sectional differential). *Firms with oil-linked business activities earn higher cumulative abnormal returns than non-oil-linked firms following the shock, because the cash-flow channel accrues disproportionately to them. This differential is directional (positive) and is the paper's primary focal test.*

3. Data and Methodology

3.1. Data

The event examined in this paper is the outbreak of military action against Iran on the morning of Saturday, 28 February 2026. Because the Tadawul exchange operates on a Sunday-through-Thursday trading week, with markets closed on Friday and Saturday, the first trading session at which Saudi equity prices could incorporate the news was Sunday, 1 March 2026. We accordingly set the event day ($t = 0$) to 1 March 2026.

Our sample universe is the Tadawul Main Market, excluding the parallel Nomu market, REITs, ETFs, and other non-equity instruments. Daily closing prices and the TASI benchmark are obtained from the Saudi Exchange Historical Reports portal (saudiexchange.sa). The closing prices are adjusted for corporate actions (stock splits, bonus and rights issues, and cash dividends), so that measured returns reflect price changes net of such actions rather than mechanical adjustments. The price sample spans 1 January 2025 through 11 May 2026; the observations beyond March ensure that the longest event window, $[0, +20]$, which ends on 5 April 2026 in event time, is fully covered by trading data. Firm-level fundamentals, including market capitalization, total assets, total debt, stockholders' equity, net income, return on equity (ROE), and shares outstanding, are obtained from the Sahmk

Developer API (Sahmk.sa). We use annual financial statement values denominated in Saudi riyals for the fiscal year ending 31 December 2025. Because these statements were not necessarily public on the event date, we audited each firm's disclosure date and, in the cross-sectional analysis, rely on a primary specification that uses only information available before the event; the audit and the resulting specification are described in Section 4.3.

To control for global equity and commodity influences in the supplementary analyses of Appendix B, we also obtain three daily control series over the sample period: the Brent crude oil spot price (U.S. Energy Information Administration series DCOILBRENTU, obtained via FRED), the MSCI World price index, and the MSCI GCC Countries price index (both in U.S. dollars). These series are used only for the continuous Brent-beta exposure measure and the augmented benchmark model reported in Appendix B; the main event-study results use the domestic TASI benchmark alone.

We adjust the Global Industry Classification Standard (GICS) by merging the small Energy sector (six firms) into Basic Materials. This merge is adopted for two reasons. First, the Energy cell is small (six firms). Second, the GICS Energy sector and the oil-linked set overlap heavily but do not coincide: four firms (Aramco, Rabigh, Arabian Drilling, ADES) are in both, while the Energy sector also contains two firms we do not code as oil-linked (National Gas, Luberef), and the oil-linked set extends well beyond Energy to the petrochemical complex and to SARCO and Bahri. Retaining a separate Energy sector dummy alongside the oil-sector indicator would therefore be largely redundant and would add little independent identifying variation, which is why we merge Energy into Basic Materials for the sector fixed effects. We do not claim perfect collinearity between the two variables, since the oil-sector indicator additionally equals one for eleven petrochemical firms outside Energy and is therefore not a linear function of the Energy dummy alone; but the near-redundancy inflates the variance of both coefficients and, with sector fixed effects included, an Energy dummy contributes no within-sector variation for identifying the oil-sector effect. Section 4.4 reports estimates under the narrow six-firm Energy-only definition of oil-linked status alongside the broad seventeen-firm definition.

The value-weighted composition of the merged sector is highly asymmetric. By pre-event market capitalization, Saudi Aramco alone accounts for approximately 88% of the sector's aggregate market value; the petrochemical complex accounts for a further 9%; and cement, metals, and non-oil materials firms account for the remaining 3%. In firm-count terms, however, oil-linked firms are only 17 of 44 (39%) of the sector. Because the sector-level CAAR statistics in reported later are equal-weighted across firms, they give greater weight to the numerically dominant non-oil-linked majority. The cross-sectional regression in Section 4.3, by contrast, controls directly for firm size via log(market capitalization) and provides a sharper test of the oil-linked differential. A firm-level breakdown of the merged sector, listing its principal constituents with oil-linked designations and pre-event market capitalizations, is provided in Table A2 in Appendix A.

We apply three standard filters to the initial universe of 251 listed Main Market firms. First, a listing-date filter requires the first valid price observation to occur on or before 14 February 2025. Second, a liquidity filter requires the firm to have a non-missing, non-zero daily return on at least 70% of trading days in the estimation window. Third, a suspension filter excludes firms with any missing returns in the $[-1, +20]$ event window. The application of these filters yields a final sample of 230 firms.

Figure 1 reports the distribution of the final 230-firm sample across the original eleven-sector GICS classification. The largest sectors are Financial Services (42 firms), Industrials (38), Basic Materials (38), Consumer Cyclical (34), and Consumer Defensive (28). Three sectors contain fewer than ten firms each (Energy, Technology, Utilities), and we interpret sector-level results for these small categories with caution.




Figure 1. Distribution of firms by GICS sector: eleven-sector classification.

Event Identification and the Identifying Assumption

The identifying assumption of an event study is that the abnormal returns measured in the event window are attributable to the information released at $t = 0$ rather than to information already impounded in prices. We do not claim that the 2026 episode was an unanticipated shock in the strong sense, and the pre-event evidence is not consistent with such a claim. The negative pre-event abnormal returns documented below (Section 4.1) call for an explicit chronology of the two weeks preceding the event. By mid-February 2026, elevated Iranian nuclear rhetoric and political signaling had raised regional risk perceptions but had not resolved into a specific event expectation; press coverage during this period referred to “heightened diplomatic tension” without predicting a specific date for military action. Public war-warning statements intensified during the week of 20–25 February; media commentary during this period discussed the possibility of coordinated strikes without specifying a date. This background is consistent with the negative pre-event drift over the $[-5, -1]$ window (a CAAR of -1.73%). It is also consistent with the observed positive abnormal return on day $t = -1$ (Thursday, 26 February 2026), which coincided with a scheduled third round of U.S.–Iran diplomatic negotiations in Geneva, publicly reported the prior day, at which Iran’s foreign minister expressed cautious optimism that a deal was within reach (News On Air 2026). We interpret the positive $t = -1$ return as a briefly held expectation of de-escalation associated with those talks, an expectation that was sharply reversed when coordinated military action began two days later, on 28 February.

The consequence is that the measured post-event abnormal returns are best understood as the price response to the residual, unanticipated component of the event. We do not characterize them as a lower bound on the total valuation effect, since pre-event movements may reflect anticipation, unrelated news, or temporary overreaction in unknown proportions, and the sign and size of any bias cannot be established from the data.

3.2. Methodology

Table 1 summarizes the variables used in the empirical analysis, including symbols, definitions, units, and sources.

Table 1. Variables and data sources.

Symbol	Definition	Unit	Source
$R_{i,t}$	Log return of firm i on trading day t	log-return	Saudi Exchange
$R_{m,t}$	Log return of TASI benchmark on day t	log-return	Saudi Exchange
α_i, β_i	Firm-specific market-model intercept and slope (OLS)	coefficient	Estimated (Equation (1))
$AR_{i,\tau}$	Firm i 's abnormal return on day t	log-return	Computed (Equation (2))
$CAR_{i,[\tau_1, \tau_2]}$	Firm i 's cumulative abnormal return over window w	log-return	Computed (Equation (3))
$AAR_\tau / CAAR_{[\tau_1, \tau_2]}$	Cross-sectional average AR/cumulative average	log-return	Computed (Equation (4))
Oil_i	Binary indicator = 1 if firm i is oil-linked (see Section 3.2.3)	0/1	Author-constructed
$\ln(MC_i)$	Natural log of pre-event equity market capitalization	log SAR	Sahmk.sa
D/E_i^*	Debt-to-equity ratio, winsorized at the 1st and 99th percentiles (FY2025; used in augmented specification only, see Section 4.3)	ratio	Sahmk.sa (winsorized)
ROE_i	Return on equity, FY ending 31 December 2025 (used in augmented specification only, see Section 4.3)	ratio	Sahmk.sa
$Sector_{i,s}$	Sector-fixed-effect indicator (10-sector classification)	0/1	GICS (adjusted)
β_i^{Brent}	Pre-event sensitivity of firm i 's return to Brent crude return (Appendix B.1)	coefficient	Estimated (Appendix B)
$Brent_t$	Daily Brent crude oil spot price return	log-return	FRED (DCOILBRETEU)
$World_t$	Daily MSCI World price index return (USD)	log-return	MSCI
GCC_t	Daily MSCI GCC Countries price index return (USD)	log-return	MSCI

We estimate abnormal returns using the conventional market model of MacKinlay (1997) and Brown and Warner (1985). For each firm i , daily log returns over the estimation window are regressed on the contemporaneous log return of the TASI benchmark:

$$R_{i,t} = \alpha_i + \beta_i R_{m,t} + \varepsilon_{i,t} \quad (1)$$

where $R_{i,t}$ is the log return of firm i on trading day t , $R_{m,t}$ is the log return of the TASI benchmark, α_i and β_i are firm-specific parameters estimated by ordinary least squares over the estimation window, and $\varepsilon_{i,t}$ is a mean-zero idiosyncratic disturbance. The main estimation window spans 250 trading sessions ending on 11 February 2026, corresponding to event time $t = -11$, so that the estimation window ends strictly before the first day of the pre-event window ($t = -10$, 12 February 2026) and does not overlap it. Robustness checks using 200-session and 100-session estimation windows are also reported, both ending on the same date.

The firm-level abnormal return on event-window trading day t is defined as:

$$AR_{i,\tau} = R_{i,\tau} - (\hat{\alpha}_i + \hat{\beta}_i R_{m,\tau}) \quad (2)$$

where $\hat{\alpha}_i$ and $\hat{\beta}_i$ denote the ordinary least squares (OLS) estimates of the market-model parameters from Equation (1). Firm-level cumulative abnormal returns over an event window $[\tau_1, \tau_2]$ are then computed by summing daily abnormal returns:

$$CAR_{i,[\tau_1,\tau_2]} = \sum_{\tau=\tau_1}^{\tau_2} AR_{i,\tau} \quad (3)$$

Daily average abnormal returns (AAR) and cumulative average abnormal returns (CAAR) are computed by averaging across firms and cumulating across the window:

$$AAR_{\tau} = \frac{1}{N} \sum_{i=1}^N AR_{i,\tau}, \quad CAAR_{[\tau_1,\tau_2]} = \frac{1}{N} \sum_{i=1}^N CAR_{i,[\tau_1,\tau_2]} \quad (4)$$

We report eleven event windows spanning pre-event periods ($[-10, -1]$ and $[-5, -1]$), the event day ($[0, 0]$), the short symmetric window ($[-1, +1]$), short post-event windows ($[0, +1]$, $[0, +2]$, $[0, +3]$, $[0, +5]$), windows isolating the post-event recovery segment ($[-1, +5]$ and $[+1, +5]$), and a longer post-event window ($[0, +20]$). The cross-sectional *t*-test compares the sample mean CAR to its sampling distribution. The Wilcoxon signed-rank test uses both the signs and the ranks of the absolute CAR values and is robust to departures from normality. More precisely, it is a test of a symmetric location shift of the CAR distribution around zero. Like the cross-sectional *t*-test, however, it assumes that abnormal returns are independent across firms; because a single common event date violates that assumption, the Wilcoxon test is subject to the same over-rejection as the *t*-test, which is why we rely on the Kolari–Pynnönen adjusted test of Section 3.2.2 for aggregate inference. Because a single common event date induces cross-sectional correlation among abnormal returns, both of these tests can overstate significance; we therefore also report the Kolari and Pynnönen (2010) adjusted test, described in Section 3.2.2, and treat it as the primary basis for aggregate inference.

All computations were performed in Python 3.12. Returns, abnormal returns, cumulative abnormal returns, and the event-study test statistics (the cross-sectional *t*-test, the Wilcoxon signed-rank test, and the Kolari–Pynnönen adjusted test) were implemented using NumPy 2.4 and SciPy 1.17, with data handled in pandas 3.0. Market-model parameters were estimated by ordinary least squares. The cross-sectional regressions and their heteroskedasticity-consistent (HC1) standard errors, together with the median (least-absolute-deviation) and multiplicity-adjusted (Holm) procedures, were estimated using statsmodels 0.15. The Kolari–Pynnönen adjustment factor was computed from the average pairwise correlation of the estimation-window residuals across the 230 sample firms. Estimation and event windows were aligned to the Saudi Exchange trading calendar (Sunday–Thursday); event-time indexing and all window definitions follow Section 3.1. Replication code is available as described in the Data Availability Statement.

3.2.1. Benchmark Model and Systematic Factors

The market model in Equation (1) uses the domestic market index as the sole benchmark. This choice has a specific implication for what the abnormal returns measure. Because TASI is itself an oil-sensitive index, the market-model residual nets out the component of the shock that is common to the Saudi market as a whole, including the market-wide component of the contemporaneous oil-price move and of any global equity-market movement to the extent that these are transmitted through the domestic index. The abnormal returns we report are therefore firm-level returns relative to the domestic market, not returns relative to a global or multi-factor benchmark. This is the standard interpretation of a single-index event study (Brown and Warner 1985; MacKinlay 1997), and it is well suited to the cross-sectional question that is the paper’s focus, since it differences out precisely

the common component that the cross-sectional comparison does not rely on. A further, practical reason reinforces this choice: the Saudi Exchange trades Sunday through Thursday while global equity and oil-price series trade Monday through Friday, so the event day itself (Sunday, 1 March 2026) has no contemporaneous global-market observation. A global or multi-factor benchmark therefore cannot price the single most important day in the study, whereas the domestic index can.

It also implies two limitations that should be stated plainly. First, the aggregate CAARs in Section 4.1 are not estimates of the total value effect of the episode on Saudi equities; they measure only the extent to which the average sample firm departed from what its historical beta to TASI would predict. A shock that moved the entire market proportionally to beta would produce CAARs near zero by construction. Second, Saudi equities during this window were plausibly affected simultaneously by global equity-market movements, further regional developments, oil-price movements, and other systematic factors, none of which enter the primary market model of Equation (1) as a separate regressor (the augmented model of Appendix B.2 adds global-equity and oil factors as a robustness check). The domestic single-index model is our primary specification, for the reasons given above. To assess whether omitted global and commodity factors materially affect the results, we also estimate an augmented benchmark model that adds the MSCI World price-index return and the Brent crude return to the TASI benchmark, reported in Appendix B.2. The oil-linked differential remains positive under this augmented model. That analysis is nonetheless constrained by the calendar mismatch noted above: the augmented model cannot price the Sunday event day, which is the single most important day in the study. We therefore do not claim that the abnormal returns are fully purged of omitted systematic factors beyond what the domestic index absorbs, and we treat this as a limitation on the interpretation of the aggregate abnormal returns, to which we return in Section 5. The cross-sectional specification in Equation (5) is less exposed to this concern than the aggregate CAARs, because any systematic factor that affects all sample firms identically is absorbed by the intercept and does not contaminate the estimated oil-linked differential; it would contaminate that differential only to the extent that its loading differs systematically between oil-linked and non-oil firms, which, in the case of the oil price itself, is precisely the channel under study.

3.2.2. Cross-Sectional Dependence and Event-Study Inference

All 230 sample firms share the same event date. Abnormal returns are therefore very likely to be correlated across firms within the event window, both because the event is common and because the market model removes only the linear projection on TASI and not any residual common factor. This has a direct consequence for inference: the conventional cross-sectional *t*-test applied to aggregate CAARs assumes cross-sectionally independent abnormal returns, and when that assumption fails, the estimated standard error understates the true sampling variability, so that *t*-statistics are inflated, and *p*-values are too small. The same concern applies to the Wilcoxon signed-rank test, which relaxes the distributional assumption but not the independence assumption.

We implement the cross-sectionally adjusted test of Kolari and Pynnönen (2010), which extends the standardized cross-sectional test of Boehmer et al. (1991) to allow for a common non-zero correlation among abnormal returns. The average pairwise correlation among the 230 firms' estimation-window residuals is $\bar{r} = 0.098$, implying an adjustment factor of $\sqrt{(1 - \bar{r}) / [1 + (N - 1)\bar{r}]} = 0.196$. $\sqrt{[(1 - \bar{r}) / (1 + (N - 1)\bar{r})]} = 0.196$; that is, the effective number of independent observations is roughly one-twenty-fifth of the nominal sample size. Table 2 reports, for every event window, the naive cross-sectional *t*-statistic alongside the Kolari–Pynnönen adjusted *Z*-statistic and its *p*-value.

The adjustment changes the statistical picture materially. Under the naive test, nine of the eleven windows are significant at 5% or better; under the Kolari–Pynnönen adjustment, only the [0, +2] window survives at 5% and the [0, +1] window at 10%. We therefore treat the immediate two- to three-session reaction as the only robust aggregate result and read the event-day, pre-event, and long-horizon significance levels as not established. It would be tempting to assume that t-statistics as large as these would survive any plausible correction, but they do not: the event-day [0, 0] statistic, for instance, falls from -5.36 to an insignificant -0.90 once the adjustment is applied. The cross-sectional regression of firm-level CARs on the oil-linked indicator and firm controls (specified in Section 3.2.3) is far less exposed to this problem, because its coefficient of interest is a within-date difference between groups of firms; a common shock that moves all firms together shifts the intercept rather than the oil-linked coefficient. We regard that cross-sectional differential, rather than the aggregate CAAR significance levels, as the paper’s most robust contribution.

Table 2. Cross-sectionally adjusted inference (Kolari and Pynnönen 2010), equal-weighted.

Window	Mean CAR	Naive t (p)	KP-Adjusted Z (p)	Significant (KP)
[−10, −1]	−1.32%	−3.75 (<0.001)	−0.68 (0.494)	No
[−5, −1]	−1.73%	−6.57 (<0.001)	−1.26 (0.206)	No
[0, 0]	−0.74%	−5.36 (<0.001)	−0.90 (0.368)	No
[−1, +1]	−0.44%	−1.89 (0.060)	−0.39 (0.700)	No
[0, +1]	−1.62%	−9.30 (<0.001)	−1.78 (0.075)	10%
[0, +2]	−2.84%	−11.15 (<0.001)	−2.09 (0.036)	5%
[0, +3]	−1.03%	−4.15 (<0.001)	−0.73 (0.464)	No
[0, +5]	−0.11%	−0.36 (0.716)	−0.06 (0.954)	No
[−1, +5]	+1.07%	+3.20 (0.002)	+0.63 (0.527)	No
[+1, +5]	+0.63%	+2.13 (0.034)	+0.37 (0.709)	No
[0, +20]	+1.57%	+2.30 (0.022)	+0.50 (0.619)	No

Notes: $\bar{r} = 0.098$; adjustment factor = 0.196. Under the adjusted test only the [0, +1] and [0, +2] windows remain significant. All other windows, including the event day, the pre-event windows, and the [0, +20] window, are statistically indistinguishable from zero once cross-sectional dependence is accounted for.

3.2.3. Classification of Oil-Linked Firms

Our cross-sectional analysis hinges on a binary oil-sector indicator. This indicator is best understood as a proxy for membership in an oil-linked business category rather than as a continuous measure of a firm’s sensitivity to oil-price movements; it captures a discrete re-rating of oil-linked firms as a group, not the degree of each firm’s price exposure. (A continuous Brent-beta measure is examined separately in Appendix B.1). The GICS Energy sector contains six firms: Saudi Aramco (2222), Rabigh Refining and Petrochemicals (2380), Arabian Drilling (2381), ADES Holding (2382), National Gas and Industrialization (2080), and Saudi Arabian Refineries for Lubricants/Luberef (2223). This exchange-level Energy classification does not, however, correspond to economic oil exposure: it excludes the entire Saudi petrochemicals complex (classified by GICS as Basic Materials but economically oil-linked through feedstock costs, output prices, and corporate linkages), and it includes two firms: National Gas and Luberef, whose activities we do not classify as crude-oil-linked in our sense. This narrow definition excludes the Saudi petrochemicals complex, which is classified by GICS as Basic Materials but is economically oil-linked through feedstock costs, output prices, and corporate linkages.

We therefore define the oil-sector indicator on business-activity grounds rather than by GICS sector. It equals one for seventeen firms: the crude producer Saudi Aramco (2222); the refiners Rabigh Refining and Petrochemicals (2380) and Saudi Arabian Refineries/SARCO (2030); the oilfield-services firms Arabian Drilling (2381) and ADES Holding (2382); the

oil-shipping firm National Shipping/Bahri (4030); and eleven petrochemical and chemical-feedstock producers: SABIC (2010), SAFCO (2020), Tasnee (2060), Chemanol (2001), Alujain (2170), Nama Chemicals (2210), SIIG (2250), YANSAB (2290), SIPCHEM (2310), APPC (2330), and Saudi Kayan (2350). Note that SARCO and Bahri are classified by GICS under Industrials, not Energy, and that the two GICS-Energy firms whose activities are not crude-oil-linked in our sense—National Gas and Industrialization (2080) and Luberef (2223)—are coded non-oil-linked. The classification is applied on the basis of a firm's principal line of business and is held fixed throughout; it does not use any information from the event window. The regression code uses exactly this seventeen-firm list, which is the same list reported in Appendix A Table A2 and used in the descriptive statistics, all cross-sectional regressions, the leave-one-out exercise, the narrow-versus-broad comparison, and the Brent-beta validation.

The rationale for the broad definition is that the narrow exchange classification would misclassify the petrochemical complex, whose feedstock costs and output prices are contractually and economically tied to crude and its derivatives, as unexposed. Widening the definition also has the practical consequence that the indicator varies within the Basic Materials sector, which is what allows the coefficient in Equation (5) to be identified from within-sector variation once sector fixed effects are included.

Two properties of the measure should nonetheless be recognized. First, it is binary, and economic exposure to oil prices is continuous: a crude producer, a refiner whose margin depends on the crack spread, an oilfield-services firm whose revenue depends on drilling activity, and a petrochemical producer for whom crude is an input cost do not share the same sign or magnitude of exposure to a given oil-price move, yet all take the value one. Firms with material but secondary oil exposure, for example, transport or utility firms with fuel-cost sensitivity, take the value zero. The indicator therefore proxies for membership of an oil-linked business category rather than measuring the degree of a firm's economic exposure to oil-price changes, and it should not be interpreted as a firm-level oil beta. Second, the classification is not orthogonal to other firm characteristics, a point we take up in Section 4.3. A continuous measure, a revenue share derived from crude-related products, or an estimated historical sensitivity of returns to Brent crude, is the natural refinement. We construct the second of these, a pre-event Brent-crude sensitivity for each firm, and report it in Appendix B.1. Oil-linked firms have significantly higher Brent betas than non-oil firms, confirming that the classified firms are indeed more oil-sensitive; the continuous beta does not, however, predict event-window CARs, whereas the binary indicator does. We read this as evidence that the oil-linked differential in event returns operates as a discrete re-rating of the oil-linked business category rather than as a smooth function of each firm's historical price sensitivity. A revenue-share measure, which would require firm-by-firm hand-collection of product-level disclosures across the Tadawul universe, remains beyond the scope of this paper, and we identify the development of a fuller continuous exposure measure as the principal measurement refinement for future work.

To characterize the cross-sectional structure of the response, we estimate the following cross-sectional regression separately for each of three event windows $\omega \in \{[0, +1], [0, +5], [0, +20]\}$:

$$CAR_{i,\omega} = \gamma_0 + \gamma_1 Oil_i + \gamma_2 \ln(MC_i) + \gamma_3 D/E_i^* + \gamma_4 ROE_i + \sum_s \delta_s Sector_{i,s} + u_i \quad (5)$$

where the dependent variable is firm i 's cumulative abnormal return over event window ω . The variables of interest are: Oil_i , the binary oil-sector indicator; $\ln(MC_i)$, the natural logarithm of pre-event equity market capitalization in Saudi riyals; D/E_i^* , the debt-to-equity ratio winsorized at the 1st and the 99th percentiles¹; and ROE_i , the return on

equity. The summation indexes a set of sector fixed effects based on the regression-specific ten-sector classification, with Basic Materials serving as the omitted reference category. Heteroskedasticity-robust (HC1) standard errors are reported. The coefficient of primary interest is γ_1 , which captures the within-sector differential response of oil-linked versus non-oil-linked firms.

4. Results

4.1. Aggregate Market Response

Table 3 reports the mean cumulative abnormal return, the cross-sectional *t*-statistic, the median CAR, and the two-sided Wilcoxon signed-rank *p*-value across eleven event windows. The results indicate a sharp negative response concentrated in the first three trading days after the event, followed by a reversal: by [0, +5] the mean CAAR is no longer distinguishable from zero, and by [0, +20] it is positive and significant at the 5% level under the conventional cross-sectional test, though, as Section 3.2.2 shows, this [0, +20] significance does not survive the cross-sectional-dependence adjustment. The significance levels in this table should be read subject to the cross-sectional dependence discussed in Section 3.2.2.

Table 3. Mean and median CARs and statistical tests across event windows.

Window	N	Mean CAR	t-Stat	Median CAR	Wilcoxon <i>p</i>
[−10, −1]	230	−1.32%	−3.75 ***	−1.24 ***	0.0001
[−5, −1]	230	−1.73%	−6.57 ***	−1.71 ***	<0.001
[0, 0]	230	−0.74%	−5.36 ***	−0.56 ***	<0.001
[−1, +1]	230	−0.44%	−1.89 *	−0.55 **	0.0124
[0, +1]	230	−1.62%	−9.30 ***	−1.64 ***	<0.001
[0, +2]	230	−2.84%	−11.15 ***	−2.76 ***	<0.001
[0, +3]	230	−1.03%	−4.15 ***	−1.15 ***	<0.001
[0, +5]	230	−0.11%	−0.36	−0.72	0.1104
[−1, +5]	230	+1.07%	+3.20 ***	+0.52 ***	0.0040
[+1, +5]	230	+0.63%	+2.13 **	−0.14	0.4230
[0, +20]	230	+1.57%	+2.30 **	+1.00 *	0.0935

Note: Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Stars on the *t*-statistic are derived from the cross-sectional *t*-test *p*-value; stars on the median CAR are derived from the two-sided Wilcoxon signed-rank *p*-value. The estimation window comprises 250 trading sessions ending on 11 February 2026 (event time $t = -11$), strictly before the pre-event window. Reported significance levels assume cross-sectional independence of abnormal returns; see Section 3.2.2.

The sharpest immediate response occurs in the [0, +1] window, where the mean CAR of −1.62% and the median CAR of −1.64% are both significant at the 1% level. The effect deepens through the third trading day: the [0, +2] window records a mean CAR of −2.84%, the largest negative CAR in the table. Inspection of daily abnormal returns (in Appendix A, Table A1) shows that the three-day decline is concentrated in event days 0, +1, and +2 (daily AARs of −0.74, −0.88, and −1.22%, respectively).

The negative response also has a pre-event component. The [−5, −1] window records a mean CAR of −1.73%, and the longer [−10, −1] window shows −1.31%, both significant under the conventional *t*-test; we note, however, that neither survives the cross-sectional-dependence adjustment of Section 3.2.2, so the pre-event drift should be read as suggestive rather than statistically established. The pre-event drift is not monotonic: a sharp positive AAR appears on day $t = -1$ (+1.18%), which we interpret as markets briefly pricing in de-escalation associated with the scheduled Geneva talks (see Section 3.1) before that expectation was sharply reversed by the strikes two days later. As set out in Section 3.1, this pre-event pattern implies that part of the value effect of the escalation may have been

impounded before the event day; we do not, however, characterize the post-event abnormal returns as a lower bound on the total effect, since the direction and size of any such bias cannot be established from the data.

The longer post-event horizons tell a substantially different story. The $[0, +5]$ window has an equal-weighted mean CAR of essentially zero (-0.11% , insignificant even before any dependence adjustment), indicating that the average sample firm's cumulative abnormal return over the first six trading sessions is not statistically distinguishable from zero. The $[+1, +5]$ window records a positive equal-weighted mean CAR of $+0.63\%$. The $[0, +20]$ window shows an equal-weighted mean CAR of $+1.57\%$.

We therefore regard the short windows as the cleaner measure of the market's response to the identified event, and the $[0, +20]$ window as descriptive of the subsequent evolution of returns. The transitory character of the equal-weighted response, a sharp initial decline followed by recovery of the point estimate, echoes the transitory nature of geopolitical shocks documented by Markoulis et al. (2020) and the possibility of reversal reported by Bash et al. (2021); we emphasize, however, that under the dependence-adjusted test only the immediate $[0, +1]$ and $[0, +2]$ windows are statistically significant, so this pattern is a feature of the point estimates rather than of established significance. Figures 2 and 3 visualize the daily and cumulative structure of the response.




Figure 2. Daily average abnormal returns (AAR) around the event, 250-session estimation window ending at $t = -11$. Blue bars denote positive AAR, red bars negative. The dashed vertical line marks the event day ($t = 0$, 1 March 2026); the horizontal line at zero separates positive from negative abnormal returns.

Equal-Weighted Versus Value-Weighted Response

The equal-weighted mean CAR treats a 6-trillion-riyal firm and a 300-million-riyal firm identically. Because the oil-linked complex is dominated by a few very large firms, Saudi Aramco alone accounts for roughly 88% of the merged oil-linked sector by market capitalization, the equal- and value-weighted responses need not agree, and in fact they diverge sharply. We construct equal-weighted and value-weighted (by pre-event market capitalization) portfolios of the 230 sample firms and test their cumulative abnormal performance against the distribution of the portfolio's daily abnormal return over the estimation period. Table 4 reports the two portfolios side by side.

Because these tests assess the cumulative performance of a single portfolio series against its own pre-event time-series distribution, they are the calendar-time portfolio approach of Jaffe (1974) and are, by construction, robust to the cross-sectional correlation of firm-level abnormal returns that motivates the adjustment in Section 3.2.2. The equal-

weighted decline at $[0, +2]$ (-2.84% , $p = 0.002$) therefore survives dependence-robust inference here, consistent with its survival under the Kolari–Pynnönen test; the value-weighted appreciation is likewise robust.




Figure 3. Cumulative average abnormal return (CAAR) cumulated from $t = 0$, 250-session estimation window, with 95% confidence interval. The horizontal line at zero marks no cumulative abnormal return; the dashed vertical line marks the event day ($t = 0$). The series is reset to zero immediately before the event day; it is not cumulated from $t = -10$. The 95% confidence band is based on conventional cross-sectional standard errors and is therefore not adjusted for the cross-sectional dependence discussed in Section 3.2.2; it should not be interpreted as dependence-robust. The dependence-adjusted inference is reported in Table 2.

Table 4. Equal-weighted versus value-weighted portfolio abnormal returns.

Window	Equal-Weighted CAR (p)	Value-Weighted CAR (p)
$[0, 0]$	-0.74% (0.163)	$+2.87\%$ (<0.001)
$[-1, +1]$	-0.44% (0.633)	$+2.42\%$ (0.004)
$[0, +1]$	-1.62% (0.031)	$+3.81\%$ (<0.001)
$[0, +2]$	-2.84% (0.002)	$+4.65\%$ (<0.001)
$[0, +3]$	-1.03% (0.331)	$+2.78\%$ (0.004)
$[0, +5]$	-0.11% (0.931)	$+3.58\%$ (0.003)
$[0, +20]$	$+1.57\%$ (0.519)	$+3.73\%$ (0.096)

Notes: p -values from the pre-event time-series distribution of each portfolio's daily abnormal return (estimation-period standard deviation). The value-weighted portfolio loads heavily on Saudi Aramco and the large petrochemical firms.

The contrast is the empirical heart of the paper. The equal-weighted portfolio, representing the typical listed firm, fell, driven by the war-risk premium common to all firms. The value-weighted portfolio, representing the market by capitalization, rose sharply and significantly, because it is dominated by the large oil-linked firms that benefited from the war-driven increase in crude prices. At $[0, +2]$, the equal-weighted portfolio is -2.84% while the value-weighted portfolio is $+4.65\%$, both significant under the portfolio time-series test. The sign of “the market’s reaction” is therefore not a property of the event but of the weighting scheme, and any single-number summary of the aggregate response is potentially misleading. This divergence is the most direct portfolio-level evidence in the paper for the oil-linked differential, and is consistent with the oil-price cash-flow channel: the weighting scheme that places more weight on oil-linked firms produces a more positive response, exactly as the cash-flow channel predicts. It also reframes the aggregate results of Section 4.1: the equal-weighted decline and the value-weighted appreciation are not in

tension but are two projections of the same underlying cross-sectional heterogeneity, which the regression analysis of Section 4.3 quantifies directly.

4.2. Sector-Level Decomposition

Tables 5 and 6 report CAARs and median CARs using the adjusted ten-sector GICS for the [0, +5] and [0, +1] event windows, respectively, both using the merged ten-sector classification described in Section 3.1, in which Energy is combined with Basic Materials. The two tables tell substantially different stories and together illuminate the dynamics of the cross-sectional response.

Table 5. Sector-level CAARs and median CARs for the [0, +5] event window, ordered from most negative to most positive.

Sector	N	Mean CAR	t-Stat	Median CAR	Wilcoxon <i>p</i>
Technology	5	−4.10%	−1.24	−2.86	0.3125
Real Estate	11	−1.48%	−1.16	−1.97	0.2402
Financial Services	42	−1.19%	−1.73 *	−1.51 **	0.0137
Consumer Cyclical	34	−1.06%	−1.78 *	−1.05	0.1087
Industrials	38	−0.58%	−0.97	−0.94	0.3965
Communication Services	9	−0.17%	−0.12	−1.33	0.6523
Utilities	4	+0.08%	+0.06	+0.56	1.0000
Healthcare	15	+1.05%	+0.96	−0.10	0.6387
Consumer Defensive	28	+1.39%	+2.10 **	+1.45 *	0.0628
Basic Materials	44	+1.48%	+1.53	−0.23	0.6901

Note: The Basic Materials sector includes seventeen firms classified as oil-linked under our broad definition, including Aramco. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6. Sector-level CAARs and median CARs for the [0, +1] event window.

Sector	N	Mean CAR	t-Stat	Median CAR	Wilcoxon <i>p</i>
Technology	5	−4.98%	−3.53 **	−3.30 *	0.0625
Real Estate	11	−2.74%	−2.73 **	−1.36 **	0.0137
Communication Services	9	−2.24%	−1.50	−2.37	0.1641
Healthcare	15	−2.15%	−3.67 ***	−2.45 ***	0.0067
Industrials	38	−1.86%	−4.43 ***	−1.90 ***	0.0001
Financial Services	42	−1.52%	−3.14 ***	−1.49 ***	<0.001
Consumer Cyclical	34	−1.47%	−4.53 ***	−1.77 ***	<0.001
Basic Materials	44	−1.29%	−4.33 ***	−1.19 ***	<0.001
Consumer Defensive	28	−0.82%	−1.77 *	−1.33 *	0.0815
Utilities	4	−0.43%	−0.42	+0.21	1.0000

Note: The Basic Materials sector includes seventeen firms classified as oil-linked under our broad definition, including Aramco. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In the [0, +1] window in Table 6, every one of the ten sectors records a negative mean CAR, and eight of the ten are statistically significant. The immediate response is broad-based: in the first two trading sessions, the negative impact is essentially universal across sectors, a pattern consistent with a common shock to the whole market. Basic Materials, despite containing the entire oil-linked complex, shows a mean CAR of −1.29% at this short horizon; any differential associated with oil exposure has not yet emerged.

The [0, +5] window in Table 5 looks different. The dispersion is wider, ranging from −4.11% for Technology to +1.47% for Basic Materials, and the sign of the mean CAR varies across sectors. Basic Materials swings from negative to positive between [0, +1] and [0, +5], a movement that coincides with the heavy weight of oil-linked firms within it (Aramco, SABIC, YANSAB, SIPCHEM, Saudi Kayan, and the rest of the petrochemical complex). Consumer Defensive's positive response is the pattern conventionally associated

with a risk-off rotation toward food and staple producers. These sector-level patterns are consistent with [Amin et al. \(2022\)](#), who document positive oil-shock benefits for Saudi petrochemicals and negative effects on financial-services firms, and align with [Sakaki \(2019\)](#), whose S&P 500 sector-level findings show that oil-price benefits accrue narrowly to producers. We note that the sector-level statistics in Tables 5 and 6 are equal-weighted and are not adjusted for the firm characteristics that differ across sectors; the regression evidence in Section 4.3 is the more informative test.

The contrast between Tables 5 and 6 carries a substantive empirical lesson. The immediate response was uniformly negative; the cross-sectional heterogeneity that this paper documents does not emerge until the second through fifth trading days after the event. Studies that examine only one-day or two-day windows around comparable geopolitical shocks therefore risk missing the slower-emerging differential response of fundamentally distinct exposure groups, a methodological caution that echoes the guidance of [Brown and Warner \(1985\)](#) and [MacKinlay \(1997\)](#) on window-length choice. Two cautions apply to the sector-level significance in Tables 5 and 6. First, the reported *t*-tests and Wilcoxon tests assume cross-sectional independence, and within a sector firms are, if anything, more strongly correlated than in the market as a whole; the sector-level significance levels are therefore subject to at least the same over-rejection as the aggregate results discussed in Section 3.2.2 and should be read as indicative. Second, several sectors contain very few firms, Technology (5), Utilities (4), Communication Services (9), and Healthcare (15), so their point estimates and significance are fragile and should not be over-interpreted; we retain them for completeness rather than as a basis for inference.

4.3. Cross-Sectional Regression: The Oil-Linked Differential

The sector means in Tables 5 and 6 suggest that oil exposure may be a key driver of cross-sectional heterogeneity at the [0, +5] horizon, but they confound oil exposure with other features that vary across sectors. To isolate the oil-linked differential, we estimate the cross-sectional regression specified in Equation (5). Before reporting the regression estimates, Table 7 summarizes the regression sample and the pairwise correlations among the variables. ROE is expressed as a ratio (net income divided by book equity; e.g., 0.081 = 8.1%); one firm with near-zero book equity produces an extreme negative ROE, so we report the median as the more representative central value and winsorize debt-to-equity as described in Section 3.2.

Table 7. Descriptive statistics and pairwise correlations.

(A)—Descriptive statistics:					
Variable	Mean	SD	Min	Median	Max
CAR [0, +1] (%)	−1.62	2.67	−10.17	−1.64	11.83
CAR [0, +5] (%)	−0.08	4.7	−15.65	−0.70	20.52
ln(market cap)	21.76	1.72	18.54	21.52	29.43
Debt-to-equity (winsorized)	0.72	1.13	0	0.38	7.89
ROE (ratio)	−0.38	5.08	−72.31	0.08	1.59
Oil-linked (0/1)	0.072	0.259	0	0	1
(B)—Correlation matrix:					
	CAR [0, +5]	ln(mktcap)	D/E	ROE	Oil
CAR [0, +5]	1				
ln(market cap)	−0.03	1			
Debt-to-equity	0.05	−0.05	1		
ROE	0.02	0.11	−0.54	1	
Oil-linked	0.43	0.22	0.04	0.02	1

Note: (A) N = 222. The extreme ROE minimum (−72.31) reflects one firm with near-zero book equity; the median ROE of 0.08 is more representative. (B) N = 222. The strongest correlate of CAR [0, +5] is oil-linked status (0.43). The D/E–ROE correlation (−0.54) reflects that highly leveraged firms tend to have lower ROE.

No pair of explanatory variables is strongly correlated (the largest, between debt-to-equity and ROE, is -0.54), so multicollinearity among the regressors is not a concern. The 0.43 correlation between oil-linked status and CAR $[0, +5]$ previews the regression result. Table 8 then reports the cross-sectional regression results for the three event windows.

Table 8. Cross-sectional regressions of firm-level CARs on the oil-sector indicator, firm controls, and sector fixed effects (Model A).

Variable	CAR $[0, +1]$	CAR $[0, +5]$	CAR $[0, +20]$
Oil sector	0.0044 (0.0066)	0.0900 (0.0188) ***	0.0949 (0.0460) **
ln(market cap)	0.0005 (0.0013)	-0.0042 (0.0017) **	-0.0090 (0.0040) **
Debt-to-equity (wins.)	-0.0010 (0.0016)	0.0015 (0.0036)	-0.0010 (0.0079)
Return on equity	-0.0004 (0.0002) *	0.0003 (0.0005)	0.0009 (0.0010)
Sector fixed effects	Yes	Yes	Yes
N	222	222	222
R ²	0.079	0.262	0.120

Note: Sector fixed effects use the merged ten-sector classification (Energy combined with Basic Materials). Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. HC1-robust standard errors are reported in parentheses. HC1 standard errors do not correct for cross-sectional dependence in abnormal returns; see Section 3.2.2.

The headline result is the coefficient on the oil-sector indicator in the $[0, +5]$ window: 0.0900 (HC1 standard error 0.0188 , $p < 0.001$). The results indicate that oil-linked firms outperformed non-oil firms by approximately 9.0 percentage points of cumulative abnormal return over the six trading sessions following the conflict outbreak, after controlling for firm size, leverage, profitability, and sector fixed effects.² In raw terms, the mean CAR $[0, +5]$ of oil-linked firms is $+6.81\%$ (median $+5.85\%$) versus -0.66% (median -0.87%) for non-oil firms.

A concern with any cross-sectional regression using annual accounting controls is look-ahead bias: the fiscal-year-2025 statements from which debt-to-equity and ROE are drawn may not have been public on the event date. We audited the public disclosure date of every firm's FY2025 annual results and found that only 60 of the 230 sample firms had released them before the pre-event cut-off (the market close of 26 February 2026), while 164 released them on or after the event day and 6 had non-calendar or unavailable fiscal years. To ensure the result does not rely on information unavailable to the market, our primary specification therefore uses only pre-event-available regressors: the oil-sector indicator, log pre-event market capitalization (computed from the 26 February 2026 closing price and publicly known shares outstanding), and sector fixed effects. In this look-ahead-clean specification, the oil-linked differential is $+0.0864$ (HC1 standard error 0.0180 , $p < 0.001$, $N = 230$). Adding the FY2025 accounting controls (debt-to-equity and ROE) leaves the coefficient essentially unchanged at $+0.0900$, and dropping them entirely, or replacing the full specification with the raw oil-linked group difference ($+0.0787$, $p < 0.001$), moves it only within the range $+0.079$ to $+0.090$. The differential is therefore invariant to the potentially contaminated controls, and the headline result does not depend on any post-event accounting information. The specification with FY2025 controls is estimated on 222 firms rather than 230, because eight firms lack a complete set of accounting controls; sixteen of the seventeen oil-linked firms remain in this 222-firm sample, the exception being Saudi Industrial Investment Group.

The magnitude and timing of this differential are consistent with the oil-price cash-flow channel articulated in Kilian and Park (2009) and with subsequent GCC evidence in Alqahtani et al. (2020) and Abdou et al. (2024). The result should not, however, be read as an identification of that channel, for two distinct reasons. First, as set out in Section 3.2.3, the binary indicator measures membership of an oil-linked business category

rather than the degree of a firm's economic exposure to oil prices. Second, the binary oil-sector indicator may capture a portion of variation attributable to features that are themselves correlated with oil-linked status: (a) government ownership stakes, which enter through the Public Investment Fund's holdings in Aramco, SABIC, and other petrochemical firms; (b) foreign investor participation, correlated with MSCI Emerging Markets inclusion and with oil-linked status; and (c) strategic and Vision-2030 importance, likewise correlated with oil-linked status. The 9.0-percentage-point coefficient is therefore best interpreted as reflecting the joint contribution of the oil-price cash-flow channel and these correlated firm characteristics, and the evidence does not allow us to conclusively distinguish between them. We view a continuous measure of oil exposure (for example, revenue share from crude-derivative products, or an estimated historical beta to Brent crude returns) as an important direction for future work; developing such a measure for the Tadawul universe is beyond the scope of this paper, as we discuss in the limitations of Section 5.

The pattern across windows is informative. In the very short $[0, +1]$ window, the oil-sector coefficient is small ($+0.0044$) and statistically insignificant, consistent with the broad-based, uniformly negative sectoral response documented in Table 6. By the $[0, +5]$ window, the estimated differential is large and precisely estimated. Over the longer $[0, +20]$ window, the corrected coefficient is $+0.095$ and remains statistically significant at the 5% level ($p = 0.039$). As noted in Section 4.1, the $[0, +20]$ result at both the aggregate and cross-sectional level characterizes the opening phase of the episode as a whole, since developments over the subsequent trading weeks accumulate within the window; it is not an estimate of the delayed response to the initial event.

To assess whether this differential reflects the event rather than a pre-existing divergence between oil-linked and non-oil firms, and to address the full event-time profile, Table 9 reports the daily oil-linked-minus-non-oil abnormal-return differential from $t = -10$ through $t = +20$. Two features support an event-consistent interpretation. First, over the pre-event window the differential is small and shows no systematic trend: the mean daily differential over $[-10, -1]$ is $+0.10\%$ and is not statistically distinguishable from zero ($t = 0.48$, $p = 0.64$), so the two groups were not already diverging before the event. Second, the differential is essentially zero on the event day itself (-0.07% , $p = 0.89$) and then emerges over the following sessions, reaching $+0.74\%$ at $t = +1$, $+2.23\%$ at $t = +2$ ($p = 0.04$), and $+4.69\%$ at $t = +5$ ($p < 0.001$), for a cumulative oil-linked outperformance of $+7.5\%$ by $t = +5$. The profile is not monotonic thereafter: the differential reverses sharply at $t = +7$ (-4.28% , $p < 0.001$) before drifting upward again, so we do not characterize the post-event path as a smooth, permanent divergence. What the profile does establish is that the differential is absent before and on the event day and emerges only afterward, which is the pattern that a differential emerging after the event, rather than a pre-existing divergence, would produce.

The control variables show modest patterns. Log market capitalization has a statistically significant negative coefficient of -0.0042 in the $[0, +5]$ window, suggesting that larger firms were slightly more adversely affected. Return on equity has a marginally significant negative coefficient of -0.0004 in the $[0, +1]$ window. The winsorized debt-to-equity ratio enters with small and statistically insignificant coefficients across all windows. That the oil-sector coefficient remains large and statistically significant with these controls and sector fixed effects included indicates that the differential is not accounted for by size, leverage, profitability, or sector membership; it does not establish that the differential reflects the oil-price channel specifically, for the reasons given above.

Table 9. Daily oil-linked-minus-non-oil abnormal-return differential (selected event-time sessions).

t	Oil AAR	Non-Oil AAR	Differential	p
−10	+1.52%	−0.11%	+1.63%	0.014
−5	−0.96%	−0.64%	−0.33%	0.264
−1	+1.12%	+1.19%	−0.06%	0.910
0	−0.81%	−0.74%	−0.07%	0.890
+1	−0.19%	−0.93%	+0.74%	0.174
+2	+0.85%	−1.39%	+2.23%	0.036
+3	+1.77%	+1.81%	−0.04%	0.950
+5	+5.03%	+0.34%	+4.69%	<0.001
+7	−4.06%	+0.22%	−4.28%	<0.001
+20	+1.43%	+0.40%	+1.03%	0.158

Notes: *p*-values from Welch two-sample *t*-tests comparing daily abnormal returns of the 17 oil-linked and 213 non-oil firms. Full path for all sessions *t* = −10 to +20 available on request; pre-event mean daily differential +0.10% (*t* = 0.48, *p* = 0.64).

4.4. Robustness Checks

Influence and specification robustness of the oil differential. Because the 9.0-percentage-point oil-linked differential is economically large, we verify that it is not an artifact of a few influential firms or of the treatment of extreme observations. Table 10 collects the checks. Excluding Saudi Aramco leaves the coefficient essentially unchanged at +0.089 (*p* < 0.001), so the result is not an Aramco effect. A median (least-absolute-deviation) regression yields +0.097, and winsorizing the dependent variable at the 1st and 99th percentiles yields +0.085, so the result is not driven by outliers. Leave-one-out estimation, dropping each oil-linked firm in turn, produces coefficients ranging only from +0.081 to +0.097. Adding a separate negative-book-equity indicator (for the small number of firms whose winsorized debt-to-equity ratio is affected by negative equity) leaves the coefficient at +0.090. The differential is therefore robust across every specification we examine.

Table 10. Robustness of the oil-linked differential, CAR [0, +5].

Specification	Oil Coefficient	<i>p</i> -Value
Main (17-firm binary, HC1, sector FE)	+0.0900	<0.001
Excluding Saudi Aramco	+0.0893	<0.001
Median (LAD) regression	+0.0972	<0.001
Winsorized CAR (1st/99th pct)	+0.0853	<0.001
With negative-equity indicator	+0.0900	<0.001
Leave-one-out (range across oil firms)	+0.081 to +0.097	—

Notes: raw group means of CAR [0, +5] are +6.81% (median +5.85%) for the 17 oil-linked firms and −0.66% (median −0.87%) for the 213 non-oil firms. The continuous Brent-beta and augmented-factor-model analyses are reported in Appendix B.

We also examine whether the results depend on the combination of controls included in the main specification. Model A comprises four firm-level regressors: the oil-sector indicator, ln(market capitalization), winsorized debt-to-equity, and return on equity, together with ten-sector fixed effects. This is the specification reported in Table 8 and replicated in Tables 11 and 12. We estimate two alternative specifications. In Model B, we remove the sector fixed effects while retaining all four firm-level regressors, allowing the oil-sector coefficient to be identified from both between-sector and within-sector variation. In Model C, we remove the oil-sector indicator while retaining the sector fixed effects and the other three firm-level regressors: ln(market capitalization), winsorized debt-to-equity, and return on equity. Model C provides a benchmark for assessing how much cross-sectional variation

in CARs is attributable specifically to oil exposure rather than to sector membership or other firm characteristics.

Table 11. Three-model cross-sectional regression results for the [0, +5] event window.

Variable	Model A	Model B	Model C
Intercept	0.0777 ** (0.0364)	0.0674 * (0.0342)	0.0448 (0.0431)
Oil sector	0.0900 *** (0.0188)	0.0831 *** (0.0186)	—
ln(market cap)	−0.0042 ** (0.0017)	−0.0035 ** (0.0016)	−0.0014 (0.0020)
Debt-to-equity (wins.)	0.0014 (0.0036)	0.0019 (0.0032)	0.0038 (0.0031)
Return on equity	0.0003 (0.0005)	0.0005 (0.0005)	0.0005 (0.0005)
Sector: Communication Services	0.0217 (0.0153)	—	−0.0113 (0.0177)
Sector: Consumer Cyclical	−0.0003 (0.0083)	—	−0.0280 ** (0.0117)
Sector: Consumer Defensive	0.0242 *** (0.0091)	—	−0.0032 (0.0122)
Sector: Financial Services	0.0008 (0.0091)	—	−0.0271 ** (0.0127)
Sector: Healthcare	0.0267 ** (0.0126)	—	−0.0038 (0.0151)
Sector: Industrials	0.0028 (0.0082)	—	−0.0201 * (0.0117)
Sector: Real Estate	−0.0008 (0.0131)	—	−0.0299 * (0.0160)
Sector: Technology	−0.0288 (0.0289)	—	−0.0585 * (0.0308)
Sector: Utilities	0.0212 (0.0171)	—	−0.0146 (0.0179)
N	222	222	222
R ²	0.262	0.205	0.085
Sector fixed effects	Yes	No	Yes
Oil-sector indicator	0.0900 *** (0.0188)	0.0831 *** (0.0186)	—

Note: HC1-robust standard errors in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12. Three-model cross-sectional regression results for the [0, +1] event window.

Variable	Model A	Model B	Model C
Intercept	−0.0256 (0.0270)	−0.0180 (0.0261)	−0.0272 (0.0278)
Oil sector	0.0044 (0.0066)	0.0065 (0.0063)	—
ln(market cap)	0.0005 (0.0013)	0.0001 (0.0012)	0.0007 (0.0013)
Debt-to-equity (wins.)	−0.0010 (0.0016)	−0.0008 (0.0015)	−0.0009 (0.0016)
Return on equity	−0.0004 * (0.0002)	−0.0004 * (0.0002)	−0.0004 * (0.0002)
Sector: Communication Services	−0.0084 (0.0172)	—	−0.0100 (0.0171)
Sector: Consumer Cyclical	0.0002 (0.0047)	—	−0.0012 (0.0045)
Sector: Consumer Defensive	0.0090 (0.0058)	—	0.0076 (0.0057)
Sector: Financial Services	−0.0019 (0.0061)	—	−0.0033 (0.0060)
Sector: Healthcare	−0.0070 (0.0067)	—	−0.0085 (0.0066)
Sector: Industrials	−0.0033 (0.0052)	—	−0.0044 (0.0052)
Sector: Real Estate	−0.0126 (0.0103)	—	−0.0141 (0.0103)
Sector: Technology	−0.0346 ** (0.0138)	—	−0.0361 *** (0.0138)
Sector: Utilities	0.0102 (0.0098)	—	0.0084 (0.0098)
N	222	222	222
R ²	0.079	0.008	0.078
Sector fixed effects	Yes	No	Yes
Oil-sector indicator	Yes	Yes	No

Note: HC1-robust standard errors in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11 presents the three-model comparison for the main [0, +5] window. First, the oil-sector coefficient is stable across Model A and Model B. The difference is well within one standard error, indicating that the oil-channel finding is not driven by the within-sector identification design of Model A; it would emerge essentially unchanged from a between-sector comparison. Second, when we remove the oil-sector indicator, the explanatory power of the model falls: R² declines from 0.262 to 0.085; we quantify the oil-sector indicator's incremental contribution formally below. Third, the sector fixed effects

shift in sign and magnitude between Model A and Model C. In Model A, with the oil-sector indicator controlling for oil exposure within Basic Materials, the Consumer Defensive and Healthcare dummies become significantly positive (the textbook defensive sectors), while in Model C, without the oil control, the Financial Services, Consumer Cyclical, Industrials, Real Estate, and Technology dummies become significantly negative relative to the Basic Materials reference.

To formalize the contribution of the oil-linked indicator beyond the informal comparison of model R^2 values, we conduct a nested-model test. Adding the oil-linked indicator to the specification containing firm controls and sector fixed effects raises the R^2 from 0.085 to 0.262, an incremental R^2 of 0.177. The corresponding partial R^2 is 0.193: the oil-linked indicator explains 19.3% of the variation in CAR $[0, +5]$ that is left unexplained by firm size, leverage, profitability, and sector membership. A nested-model F-test rejects the exclusion of the oil-linked indicator at any conventional level ($F(1, 208) = 49.76, p < 0.001$). The oil-linked differential is therefore not merely a large point estimate but a statistically decisive addition to the model's explanatory power, conditional on the full set of controls and sector fixed effects.

Table 12 reports the same three-model comparison for the short $[0, +1]$ window and tells a different story. The oil-sector coefficient is small and insignificant across both Model A and Model B, consistent with the uniform-decline pattern of Table 6 in which all sectors experienced negative CARs over the first two trading days. The R^2 values are lower than at $[0, +5]$ for Model A and for Model B. Model B at $[0, +1]$ has near-zero explanatory power, indicating that without sector fixed effects, almost none of the variation in firm-level CARs over the immediate post-event window can be explained by the four firm-level variables. The ROE carries a significant negative coefficient, indicating higher-profitability firms experienced slightly more negative immediate abnormal returns, a pattern that may reflect the concentration of high-ROE firms among large, foreign-investor-owned names, although we do not test this. The Technology sector dummy is significantly negative at the 5% level in both Model A and Model C, but this finding is based on only five firms.

Tables 11 and 12 clarify the timing of the oil-linked differential. Table 11 shows that the 9.0-percentage-point outperformance of oil-linked firms over the $[0, +5]$ window is robust to the inclusion or exclusion of sector fixed effects and that removing the oil-sector indicator in Model C eliminates most of the cross-sectional explanatory power at this horizon. Table 12 shows that this differential is not present in the shorter $[0, +1]$ window. This pattern is consistent with the interpretation that the oil-linked differential emerges over several trading sessions rather than on the event day itself.

We re-estimate the market model with 200-session and 100-session estimation windows. The aggregate CAR results are quantitatively similar to the main 250-session results: the $[0, +1]$ mean CARs are -1.74% and -1.65% respectively, both significant at the 1% level under both the t -test and the Wilcoxon signed-rank test. The $[0, +20]$ CAARs are more strongly positive in both robustness specifications ($+1.82\%$ and $+2.62\%$ respectively, significant at 1% under the naive cross-sectional test but, like the primary-specification $[0, +20]$ result, not surviving the cross-sectional-dependence adjustment of Section 3.2.2). The short-window results are therefore stable across benchmark specifications while the long-window result is not, which reinforces the cautious reading of the $[0, +20]$ horizon developed in Section 4.1.

Multiple-window inference and choice of primary horizon. Because we report eleven event windows, family-wise error is a concern. We designate the $[0, +5]$ window as the primary focal horizon, designated on theoretical grounds on the reasoning set out in Section 2: a supply-driven oil-price shock is expected to be incorporated into oil-linked equity valuations over several sessions rather than on the event day, so a short multi-session

post-event window is the theoretically appropriate horizon for the oil-linked differential; the remaining ten windows are secondary and exploratory. For the family of eleven aggregate CAAR window tests, we additionally report multiplicity-adjusted p -values. Applying the Holm step-down correction to the already dependence-adjusted (Kolari–Pynnönen) p -values, none of the eleven aggregate windows remains significant at the 10% level (for example, the $[0, +2]$ window moves from a dependence-adjusted p of 0.036 to a Holm-adjusted p of 0.40). We read this as reinforcing, not overturning, the paper’s position: the aggregate CAAR path is statistically fragile under stringent correction, which is precisely why we treat the cross-sectional oil-linked differential, a single, pre-specified comparison at the $[0, +5]$ horizon, not one of a family of window tests, as the robust core of the analysis. That differential remains significant at the 1% level under heteroskedasticity-robust (HC1) inference and across the full battery of robustness checks in this section.

Narrow versus broad oil-linked definition. To address the concern that the binary indicator pools distinct business models, we re-estimate the $[0, +5]$ cross-sectional regression under the narrow six-firm Energy-only definition of oil-linked status (the six firms in the GICS Energy sector—Saudi Aramco, Rabigh Refining and Petrochemicals, Arabian Drilling, ADES Holding, National Gas and Industrialization, and Luberef). Under this narrow definition, the oil-linked coefficient is +0.010 (HC1 standard error 0.029, $p = 0.72$), statistically indistinguishable from zero, compared with +0.090 ($p < 0.001$) under the broad seventeen-firm definition. The differential is therefore driven by the petrochemical complex rather than by the crude-and-services core alone, which both justifies the broad classification and reinforces the interpretation of the effect as an oil-linked business-category differential operating through downstream as well as upstream firms. We note that the narrow definition cannot be identified with a separate Energy sector fixed effect, because the six Energy firms coincide exactly with the Energy sector, a concrete illustration of the near-redundancy discussed in Section 3.1.

Figures 4 and 5 demonstrate that the estimated CAAR path is robust to alternative estimation lengths.




Figure 4. Robustness CAAR: 200-session and 100-session estimation windows. The shaded band is the 95% confidence interval [based on conventional cross-sectional standard errors]. The horizontal line at zero marks the level of no cumulative abnormal return. The dashed vertical line marks the event day ($t = 0$).




Figure 5. CAAR robustness comparison across estimation-window lengths. The horizontal line at zero marks no cumulative abnormal return; the dashed vertical line marks the event day ($t = 0$).

5. Discussion

Three implications stand out. First, the sign of the Tadawul Main Market's aggregate response depends on how firms are weighted, and its statistical robustness is limited. The equal-weighted portfolio declined (a mean CAR of -2.84% at $[0, +2]$), while the value-weighted portfolio, dominated by the large oil-linked firms, rose ($+4.65\%$ at $[0, +2]$). The equal-weighted point estimate subsequently recovers to $+1.57\%$ at $[0, +20]$, but under the cross-sectional-dependence adjustment only the immediate $[0, +1]$ and $[0, +2]$ windows are statistically significant; the apparent longer-horizon recovery is not statistically established. One interpretation consistent with this timing pattern is that the initial response was dominated by an increase in the required return on equity common to all firms, while the longer-horizon response increasingly reflected the positive cash-flow implications of higher crude prices for a major exporting economy. We emphasize that this is an interpretation rather than a finding. The design estimates abnormal returns; it does not decompose them into discount-rate and cash-flow components, and it does not separately identify a war-risk premium from an oil-driven cash-flow effect. Alternative readings, a common liquidity or sentiment effect that reversed, a shift in the composition of marginal investors, or the accumulation of unrelated favorable news within the twenty-day window, are not excluded by the evidence presented here. What the results establish is the shape and magnitude of the return path, not the mechanism that generated it.

Second, the aggregate response, whatever its sign and fragile significance, conceals a sharp and consistent cross-sectional structure. Over the $[0, +1]$ window, the decline is common to all sectors. By $[0, +5]$, oil exposure differentiates firms, with the cross-sectional regression yielding a significant oil-sector coefficient of approximately $+0.0900$. At $[0, +20]$, the corrected cross-sectional differential is $+0.095$ and remains statistically significant at the 5% level ($p = 0.039$); we do not, however, emphasize this longer-horizon estimate, since the $[0, +20]$ window aggregates substantial post-event information (Section 4.1) and its aggregate CAAR does not survive the cross-sectional-dependence adjustment. Oil-linked firms, broadly defined to include the full petrochemical complex, earned strongly positive abnormal returns at the $[0, +5]$ horizon, while risk-sensitive sectors such as Real Estate, Financial Services, and Consumer Cyclical continued to bear the brunt of the negative response. This is the paper's most robust finding, and it is a statement about dispersion rather than about mechanism: analyses focused only on aggregate index returns or short-

window CARs will miss economically large cross-sectional variation that emerges in the days after the immediate impact. It is also the finding least affected by the cross-sectional dependence discussed in Section 3.2.2, since it rests on a comparison between groups of firms observed on the same dates.

Third, the timing pattern of the oil-linked differential, negligible at $[0, +1]$, emerging strongly by $[0, +5]$, and remaining positive in cumulative terms through $[0, +20]$ despite a non-monotonic daily path (Table 9), has methodological implications for event studies of commodity-price shocks more broadly. Event-study windows that are too short (one or two days) may understate cross-sectional heterogeneity because the slower channels of price formation have not yet fully operated. Windows that are too long may capture developments unrelated to the event. The $[0, +5]$ window appears, at least in this setting, to be a useful focal event window for identifying the differential impact of a commodity-price shock on listed firms. The contrast between the uniformly negative sectoral pattern in Table 6 and the sign-dispersed pattern in Table 5, both estimated on the same shock and same sample but at two different horizons, is the clearest illustration in our results of this view.

5.1. Positioning Relative to the Existing Saudi Evidence

Our results neither confirm nor contradict the longer-horizon Saudi evidence, because they answer a different question. Ayadi and Ben Mbarek (2025) estimate how TASI returns respond on average to innovations in geopolitical risk and oil-price uncertainty at a monthly frequency and find that the sign of the geopolitical-risk response is not stable across sub-periods. Medhioub (2025) finds that global geopolitical risk does not significantly affect return dispersion in the Saudi market, in contrast to several other MENA markets. Alqahtani and Chevallier (2020) find that GCC equity returns co-move most strongly with oil-market shocks and that this co-movement intensifies in crises. Read together, these results suggest that the Saudi market's aggregate relationship to geopolitical risk is weak or unstable while its relationship to oil is strong and state-dependent. Our firm-level evidence is consistent with that picture and adds a specific margin to it: within a single episode, the aggregate response is close to zero over the $[0, +5]$ window precisely because it averages over firms whose responses differ by nearly nine percentage points. Aggregate stability, on this reading, is not the same as an absence of effect. We do not claim that the Saudi geopolitical-risk literature is unexplored; we claim that the firm-level dispersion documented here is not observable in the designs that literature employs.

5.2. Limitations

Seven limitations qualify the results. First, inference under cross-sectional dependence. All firms share a single event date, so abnormal returns are correlated across firms, and the cross-sectional t -test and Wilcoxon test both assume an independence that does not hold. We address this directly by reporting the Kolari and Pynnönen (2010) adjusted test (Section 3.2.2) and a portfolio time-series test (Section 4.1), under which only the immediate $[0, +1]$ and $[0, +2]$ aggregate windows remain significant while the pre-event, event-day, and $[0, +20]$ windows do not. A residual limitation remains: the Kolari–Pynnönen adjustment assumes a single common correlation across all firms, whereas dependence is plausibly stronger within sectors than across the market as a whole, so the sector-level significance levels in particular should be read as indicative rather than exact.

Second, the benchmark model. Our primary abnormal returns are measured relative to a single domestic index (TASI). We assess sensitivity to omitted global and commodity factors through an augmented benchmark model that adds the MSCI World and Brent crude returns (Appendix B.2), under which the oil-linked differential remains positive. A residual

limitation is that the Saudi and global markets trade on different calendars, so the Sunday event day (1 March 2026) has no contemporaneous global-market observation, and the augmented model cannot price the single most important day; this is itself the reason the primary specification uses the domestic index. The aggregate CAARs should accordingly be read as returns relative to the Saudi market rather than as total valuation effects.

Third, measurement of oil exposure. Our binary indicator classifies firms by business category rather than measuring the degree of exposure, and it cannot separately identify the oil-price cash-flow channel from correlated firm characteristics such as government ownership, foreign investor participation, or Vision-2030 strategic status. We construct a continuous measure, a pre-event sensitivity of firm returns to Brent crude, and report it in Appendix B.1; oil-linked firms have significantly higher Brent betas, but the continuous beta does not predict event-window CARs, whereas the binary indicator does. Two residual limitations follow: a revenue-share measure of oil exposure would require firm-by-firm hand-collection of product-level disclosures that we did not undertake, and the finding that the discrete indicator, rather than the continuous beta, drives the event-window differential means that the precise economic content of the “oil-linked” category is not fully resolved.

Fourth, the long-window results. The $[0, +20]$ window necessarily captures developments beyond the initial event and is sensitive to the estimation-window choice; conclusions about the immediate market reaction should rest on the shorter windows.

Fifth, we focus on a single event; comparative analysis across multiple geopolitical shocks in the same market would help distinguish shock-specific from more general responses, a caution reinforced by the sub-period instability reported by Ayadi and Ben Mbarek (2025). An intraday analysis of the event day could additionally provide sharper evidence on the speed of price discovery. Sixth, we do not screen individual firms for confounding corporate announcements (earnings releases, dividends, mergers, or capital actions) falling within the event windows. Such firm-specific news could contribute to individual abnormal returns; while it is unlikely to bias the oil-linked-versus-non-oil differential systematically, since there is no reason for confounding announcements to cluster differentially by oil-linked status, we cannot rule it out, and we identify a clean-sample screen for material corporate announcements as a useful refinement for future work. Seventh, our event-study design captures the market’s reaction to the information released at the onset of the conflict but does not model how the return-generating process itself may have shifted over the surrounding period; a regime-switching or structural-break analysis of the fuller episode, distinguishing pre-conflict, acute-conflict, and post-conflict regimes, would be a valuable complementary study, and we identify it as a direction for future research.

6. Conclusions

This paper documents three principal findings. First, the sign of the Tadawul Main Market’s aggregate response to the conflict outbreak depends on the weighting scheme: the equal-weighted cross-section declined (a -2.84% mean CAR at $[0, +2]$, the one aggregate result that survives a cross-sectional-dependence adjustment), while the value-weighted portfolio, dominated by the large oil-linked firms, rose ($+4.65\%$ at $[0, +2]$). The aggregate index response therefore conceals rather than summarizes the firm-level effects. Second, this aggregate response conceals a sharp and consistent cross-sectional structure: the immediate decline is common to all sectors, but by the $[0, +5]$ window oil-linked firms outperform non-oil firms by approximately nine percentage points of cumulative abnormal return, a differential robust to excluding Saudi Aramco, to median regression, to winsorization, and to leave-one-out estimation. Third, this differential is absent before and on the event day and emerges only over the following sessions, a timing pattern suggesting that the equity-

market incorporation of a commodity-price shock is gradual rather than instantaneous and implying that very short event windows may understate the cross-sectional heterogeneity such shocks generate.

The contribution of the paper is the documentation of this firm-level dispersion around a dated geopolitical event in a major oil-exporting market, together with the observation that it is largest at a horizon at which the aggregate response is close to zero. The paper does not establish the mechanism generating the dispersion. The evidence is consistent with an oil-related cash-flow channel operating alongside a market-wide increase in required returns, but the binary exposure classification, the single-index benchmark, and the event-study design do not permit these channels to be separately identified, and the results should not be read as doing so.

Because these findings derive from a single, idiosyncratic conflict episode, the following implications are offered as exploratory and case-specific rather than as generalizable policy prescriptions. For sovereign-wealth managers and domestic institutional investors, the cross-sectional heterogeneity documented here implies that index-level exposure provides only an imperfect proxy for firm-level exposure to regional geopolitical risk. Firm-level exposure to the oil-linked complex was associated with materially better performance over the six sessions of the $[0, +5]$ window following the event, a role consistent with the strategic weight of oil-linked assets in Gulf sovereign portfolios. To the extent this single-episode pattern generalizes, portfolio construction that treats “the Saudi market” as a single risk factor may miss the firm-level heterogeneity documented here; we frame this as a possible implication rather than a demonstrated result.

For financial regulators and market-integrity authorities, the finding that the cross-section of returns dispersed materially while the aggregate index response was muted suggests that surveillance systems focused on index-level circuit breakers may miss economically significant sector-level information flows during geopolitical crises. This suggests, as a possible rather than demonstrated implication, that enhanced sector-level monitoring during crises, particularly of the transmission of shocks across the financial-services and real-estate sectors, could help regulators detect anomalous trading; we do not test surveillance outcomes directly.

For macroprudential authorities and fiscal policymakers in oil-exporting economies, the immediate risk-off response hit the financial and real-estate sectors materially harder than the aggregate index and, in point-estimate terms, recovered only partially at longer horizons, though the longer-horizon pattern is not statistically established. The transmission of geopolitical shocks to the domestic banking system and property market may therefore warrant attention as a distinct concern, independent of any oil-price hedge implicit in the sovereign’s cash-flow position. Macroprudential monitoring of bank and real-estate stress during regional geopolitical events might therefore complement the fiscal insurance provided by higher oil revenues; we offer this as a conjecture consistent with the return patterns rather than as a tested policy conclusion, since we do not examine banking-system or property-market stress directly. These implications follow from the documented pattern of returns and should be weighed against the limitations set out in Section 5.2, in particular the inference and measurement caveats.

Author Contributions: Conceptualization, I.A., I.N.K. and A.B.; methodology, I.A., I.N.K. and A.B.; software, I.A., I.N.K. and A.B.; validation, I.A., I.N.K. and A.B.; formal analysis, I.A., I.N.K. and A.B.; investigation, I.A., I.N.K. and A.B.; data curation, I.A., I.N.K. and A.B.; writing—original draft preparation, I.A., I.N.K. and A.B.; writing—review and editing, I.A., I.N.K. and A.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The analytical code and author-generated datasets used in this study, including abnormal-return series, event-window statistics, and regression inputs, are available from the corresponding author upon reasonable request. The underlying raw data were obtained from third-party sources. Daily adjusted closing prices and the TASI benchmark index were collected from the Saudi Exchange Historical Reports portal (saudiexchange.sa), while firm-level financial fundamentals were obtained from the Sahmk Developer API (Sahmk.sa). Public disclosure dates of FY2025 annual results were obtained from Saudi Exchange issuer announcements. Additional control variables, including Brent crude oil prices and MSCI World and MSCI GCC Countries indices, were obtained from publicly available sources. Some third-party datasets are subject to provider licensing restrictions and cannot be redistributed.

Acknowledgments: During the preparation of this manuscript, the authors used a generative AI assistant for the purposes of language editing and improving the presentation of the text. The authors have reviewed and edited the output and take full responsibility for the content of this publication. The tool was not used to determine the study design, collect data, conduct the statistical analysis, or interpret the results.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Additional Tables

This appendix collects three Additional Tables referenced in the main text: Table A1 reports the daily average abnormal returns and their statistical tests for every event-time session from $t = -10$ through $t = +20$; Table A2 lists the full composition of the merged Basic Materials plus Energy sector, with each firm's oil-linked designation and pre-event market capitalization; and Table A3 maps event time to Saudi Exchange calendar dates. The continuous-exposure and augmented-benchmark analyses are reported separately in Appendix B.

Table A1. AAR% and median AAR% with statistical tests.

Day	N	AAR%	Median AR%	<i>t</i> -Test <i>p</i> -Value	Wilcoxon <i>p</i> -Value
−10	230	0.008	−0.177	0.9382	0.1336
−9	230	0.935	0.659	<0.001	<0.001
−8	230	−0.129	−0.150	0.2089	0.0268
−7	230	−0.257	−0.379	0.0117	<0.001
−6	230	−0.145	−0.205	0.1836	0.0140
−5	230	−0.662	−0.479	<0.001	<0.001
−4	230	−1.476	−1.311	<0.001	<0.001
−3	230	−0.775	−0.641	<0.001	<0.001
−2	230	−0.000	0.042	0.9987	0.9043
−1	230	1.182	1.109	<0.001	<0.001
0	230	−0.742	−0.565	<0.001	<0.001
+1	230	−0.880	−0.929	<0.001	<0.001
+2	230	−1.220	−1.091	<0.001	<0.001
+3	230	1.808	1.624	<0.001	<0.001
+4	230	0.232	0.089	0.0384	0.1263
+5	230	0.690	0.240	<0.001	0.0012
+6	230	−0.038	−0.146	0.8036	0.2274
+7	230	−0.093	−0.157	0.5134	0.4843
+8	230	−0.998	−0.839	<0.001	<0.001
+9	230	0.023	−0.106	0.8500	0.6614
+10	230	−0.010	−0.115	0.9386	0.4398
+11	230	0.708	0.503	<0.001	<0.001
+12	230	−0.183	−0.087	0.2478	0.1396
+13	230	0.666	0.515	<0.001	<0.001
+14	230	0.342	0.136	0.0023	0.0563

Table A1. Cont.

Day	N	AAR%	Median AR%	<i>t</i> -Test <i>p</i> -Value	Wilcoxon <i>p</i> -Value
+15	230	0.722	0.428	<0.001	<0.001
+16	230	−0.542	−0.555	<0.001	<0.001
+17	230	0.024	−0.079	0.8695	0.4563
+18	230	0.445	0.374	0.0022	<0.001
+19	230	0.144	−0.053	0.2965	0.3722
+20	230	0.475	0.196	<0.001	0.0032

Note: The *t*-test is the one-sample cross-sectional *t*-test. The Wilcoxon *p*-value is from the two-sided Wilcoxon signed-rank test. Reported significance levels assume cross-sectional independence of abnormal returns; see Section 3.2.2.

Table A2. Composition of the merged Basic Materials + Energy sector.

Ticker	Company	GICS Sub-Group	Oil-Linked	MCap (B SAR)
2222	Saudi Aramco	Energy—Producer	1	6040.3
2010	SABIC	Materials—Petrochemicals	1	162.6
2020	SAFCO	Materials—Petrochemicals	1	58.9
2382	ADES Holding	Energy—Services	1	20.2
2290	YANSAB	Materials—Petrochemicals	1	14.2
2380	Rabigh Refining & Petrochem.	Energy—Refiner	1	12.2
2310	SIPCHEM	Materials—Petrochemicals	1	10.2
2250	SIIG	Materials—Petrochemicals	1	8.8
2381	Arabian Drilling	Energy—Services	1	8.0
2350	Saudi Kayan	Materials—Petrochemicals	1	7.2
2330	Advanced Petrochemical	Materials—Petrochemicals	1	6.2
2060	Tasnee	Materials—Petrochemicals	1	5.6
2170	Alujain Holding	Materials—Petrochemicals	1	1.5
2210	Nama Chemicals	Materials—Petrochemicals	1	0.5
2001	Chemanol	Materials—Petrochemicals	1	0.4
1211	Ma'aden	Materials—Mining	0	273.0
2223	Luberef	Energy—Refiner	0	15.4
1322	AMAK	Materials—Mining	0	8.8
2080	National Gas & Ind.	Energy—Utilities	0	5.9
3030	Saudi Cement	Materials—Cement	0	5.0
3020	Yamama Cement	Materials—Cement	0	5.0
3040	Qassim Cement	Materials—Cement	0	4.7
1321	East Pipes	Materials—Metals	0	4.4
3050	Southern Province Cement	Materials—Cement	0	3.0
3092	Riyadh Cement	Materials—Cement	0	2.8
3060	Yanbu Cement	Materials—Cement	0	2.3
3010	Arabian Cement	Materials—Cement	0	2.2
3080	Eastern Province Cement	Materials—Cement	0	2.0
1320	Saudi Steel Pipes	Materials—Metals	0	2.0
1304	Al Yamamah Steel	Materials—Metals	0	1.9
2300	Saudi Paper Mfg.	Materials—Diversified	0	1.9
3003	City Cement	Materials—Cement	0	1.6
1202	Middle East Paper (MEPCO)	Materials—Diversified	0	1.5
3004	Northern Region Cement	Materials—Cement	0	1.3
4143	TALCO	Materials—Metals	0	1.2
3002	Najran Cement	Materials—Cement	0	1.1
2200	Arabian Pipes	Materials—Metals	0	0.9
3090	TCC	Materials—Cement	0	0.7
3005	Umm Al Qura Cement	Materials—Cement	0	0.7
1210	Basic Chemical Industries	Materials—Diversified	0	0.7
3091	Al Jouf Cement	Materials—Cement	0	0.6

Table A2. *Cont.*

Ticker	Company	GICS Sub-Group	Oil-Linked	MCap (B SAR)
3008	Al Kathiri Holding	Materials—Diversified	0	0.4
2090	National Gypsum	Materials—Diversified	0	0.4
2220	National Metal Mfg.	Materials—Metals	0	0.4

Note: The merged sector contains 44 firms. The table lists ticker, company name, GICS sub-group, oil-linked designation (1/0), and pre-event market capitalization (billion SAR). All 44 constituents are listed individually: the seventeen oil-linked firms (oil-linked = 1) followed by the twenty-seven non-oil-linked constituents of the cement, metals, and diversified-materials sub-groups (oil-linked = 0).

Table A3. Event-time to Saudi Exchange calendar-date mapping.

Event Time	Calendar Date	Weekday
t = −11	11 February 2026	Wednesday (last estimation day)
t = −10	12 February 2026	Thursday
t = −5	19 February 2026	Thursday
t = −1	26 February 2026	Thursday
t = 0	1 March 2026	Sunday (event day)
t = +1	2 March 2026	Monday
t = +2	3 March 2026	Tuesday
t = +3	4 March 2026	Wednesday
t = +5	8 March 2026	Sunday
t = +10	15 March 2026	Sunday
t = +11	16 March 2026	Monday
t = +12	24 March 2026	Tuesday (post-holiday gap)
t = +20	5 April 2026	Sunday

Notes: the Saudi Exchange trades Sunday–Thursday. The gap between t = +11 (16 March) and t = +12 (24 March) reflects an exchange closure. The [0, +5] window thus spans six trading sessions and the [0, +20] window twenty-one sessions.

Appendix B. Continuous Oil Exposure and Augmented Benchmark Model

We report here (i) a continuous measure of oil exposure in place of, or alongside, the binary indicator, and (ii) an expected-return model augmented with global equity and oil factors. Each has a material limitation that we state plainly, which is why these analyses supplement rather than replace the main results.

Appendix B.1. Continuous Pre-Event Brent Beta

For each firm, we estimate a pre-event sensitivity to crude oil by regressing its daily return on the TASI return and the daily Brent crude return (FRED series DCOILBRETEU) over the 250-session estimation window, and take the Brent loading as a continuous oil-exposure measure. Oil-linked firms have significantly higher Brent betas than non-oil firms (mean +0.029 versus −0.037; correlation between the binary indicator and the continuous beta is +0.31), which directly confirms that the firms we classify as oil-linked are in fact more sensitive to crude-price movements. The difference is statistically significant: a Welch two-sample *t*-test gives $t = 4.59$ ($p = 0.0002$), and a non-parametric Mann–Whitney test confirms the result ($p < 0.001$). However, when the standardized Brent beta is entered into the cross-sectional CAR [0, +5] regression, it is not itself a significant predictor (coefficient −0.006, $p = 0.10$), while the binary oil-sector indicator remains large and significant (+0.097, $p < 0.001$) in the same specification. We read this as evidence that the oil differential in event-window CARs operates as a discrete business-category effect; the oil-linked complex was re-rated as a group on the war-driven supply shock, rather than as a smooth function of each firm’s historical price sensitivity. This is an informative finding in its own right, but

it means the continuous measure does not supersede the binary classification as the more informative regressor here.

Appendix B.2. Augmented Benchmark Model (TASI + MSCI World + Brent)

We re-estimate the expected-return model adding the MSCI World price index return (in USD) and the Brent crude return to the TASI benchmark, and recompute the oil-linked differential. This analysis is constrained by a calendar mismatch that is itself informative: the Saudi Exchange trades Sunday–Thursday, while the global equity and oil series trade Monday–Friday, so the two calendars overlap only on Monday through Thursday. Critically, the event day itself (Sunday, 1 March 2026) has no global-market observation, as do several other event-window sessions. On the restricted overlapping subsample: four of the six sessions in the $[0, +5]$ window, excluding the two Sunday observations at $t = 0$ and $t = +5$, the oil-linked differential is $+0.027$ ($p = 0.089$). Because the excluded Sundays include the event day itself and $t = +5$ (which Table 9 shows carries the largest oil-linked abnormal return, $+5.03\%$), this subsample omits two economically important sessions, and we therefore do not treat it as directly comparable to the full-window estimate. The impossibility of pricing the Sunday event day with Monday–Friday global data is, in our view, a substantive reason to prefer the domestic TASI benchmark as the primary specification: it is the only benchmark that trades on the same calendar as the event.

Notes

- ¹ Winsorization reduces the influence of two firms with negative stockholders' equity (Saudi Cable Company and Alhokair Fashion Retail), whose raw debt-to-equity ratios are mathematically valid but economically anomalous, and of firms with extreme leverage outliers that would otherwise dominate the regression slope.
- ² Because abnormal returns are measured in log points, the reported cumulative abnormal returns are log-return magnitudes; a coefficient of 0.090 corresponds to a simple return of $\exp(0.090) - 1 \approx 9.4\%$. For readability we refer to these as percentage points throughout, with this log-point interpretation understood.

References

- Abdou, Hussein A., Ahmed A. Elamer, Mohammad Zoynul Abedin, and Bassam A. Ibrahim. 2024. The impact of oil and global markets on Saudi stock market predictability: A machine learning approach. *Energy Economics* 132: 107416. [\[CrossRef\]](#)
- Aguiar Rodrigues, Inês Feveteiro, Luís Miguel Pereira Gomes, and Cláudia Maria Ferreira Pereira. 2025. How do neighboring stock markets and sensitive sectors discount the Russia–Ukraine military conflict? *Cogent Economics & Finance* 13: 2555411. [\[CrossRef\]](#)
- Al Nafea, Fahad Sulaiman Mohammad. 2026. Geopolitical escalation, oil price moderation, and asymmetric market responses: Evidence from the Saudi financial market using event study and GARCH models. *International Journal of Accounting and Economics Studies* 13: 557–65. [\[CrossRef\]](#)
- Alqahtani, Abdullah, and Julien Chevallier. 2020. Dynamic spillovers between Gulf Cooperation Council's stocks, VIX, oil and gold volatility indices. *Journal of Risk and Financial Management* 13: 69. [\[CrossRef\]](#)
- Alqahtani, Abdullah, Elie Bouri, and Xuan Vinh Vo. 2020. Predictability of GCC stock returns: The role of geopolitical risk and crude oil returns. *Economic Analysis and Policy* 68: 239–49. [\[CrossRef\]](#)
- Alqahtani, Abdullah, Thies Klein, and Anas Khalid. 2021. Oil price changes, uncertainty, and geopolitical risks: On the resilience of GCC countries to global tensions. *Energy* 236: 121541. [\[CrossRef\]](#)
- Amin, Mohammad Ferdous Bhuiyan, Mohammad Mafizur Rahman, and Mustafizur Rahman. 2022. Asymmetric linkages of oil prices, money supply, and TASI on sectoral stock prices in Saudi Arabia: A non-linear ARDL approach. *SAGE Open* 12: 1–17. [\[CrossRef\]](#)
- Ayadi, Elounis, and Nizar Ben Mbarek. 2025. The dynamic effects of economic uncertainties and geopolitical risks on Saudi stock market returns: Evidence from local projections. *Journal of Risk and Financial Management* 18: 264. [\[CrossRef\]](#)
- Baker, Malcolm, and Jeffrey Wurgler. 2006. Investor sentiment and the cross-section of stock returns. *The Journal of Finance* 61: 1645–80. [\[CrossRef\]](#)
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. Measuring economic policy uncertainty. *The Quarterly Journal of Economics* 131: 1593–636. [\[CrossRef\]](#)
- Bash, Ali, Khalid Alsaifi, and Abdullah M. Al-Awadhi. 2021. Failed terror attack and stock market returns: The case of Pakistan Stock Exchange. *Applied Economics Letters* 28: 1698–702. [\[CrossRef\]](#)

- BBC News. 2026. Iran's Supreme Leader killed in joint US–Israeli strikes, Trump says. *BBC News*, March 1. Available online: <https://www.bbc.com/news/articles/c70n9wlkx3lo> (accessed on 1 March 2026).
- Boehmer, Ekkehart, Jim Musumeci, and Annette B. Poulsen. 1991. Event-study methodology under conditions of event-induced variance. *Journal of Financial Economics* 30: 253–72. [\[CrossRef\]](#)
- Bossman, Ahmed, and Mariya Gubareva. 2023. Asymmetric impacts of geopolitical risk on stock markets: A comparative analysis of the E7 and G7 equities during the Russian–Ukrainian conflict. *Heliyon* 9: e13626. [\[CrossRef\]](#)
- Boubaker, Sabri, John W. Goodell, Devendra Kumar Pandey, and Vineeta Kumari. 2022. Heterogeneous impacts of wars on global equity markets: Evidence from the invasion of Ukraine. *Finance Research Letters* 48: 102934. [\[CrossRef\]](#)
- Bougou, Whelsy, and Alhonita Yatie. 2022. The impact of the Ukraine–Russia war on world stock market returns. *Economics Letters* 215: 110516. [\[CrossRef\]](#)
- Bouri, Elie, Rima Hammoud, and Chaker Abou Kassm. 2023. The effect of oil implied volatility and geopolitical risk on GCC stock sectors under various market conditions. *Energy Economics* 120: 106617. [\[CrossRef\]](#)
- Brown, Stephen J., and Jerold B. Warner. 1985. Using daily stock returns: The case of event studies. *Journal of Financial Economics* 14: 3–31.
- Caldara, Dario, and Matteo Iacoviello. 2022. Measuring geopolitical risk. *American Economic Review* 112: 1194–225. [\[CrossRef\]](#)
- Elsayed, Ahmed H., Nader Naifar, Shazia Nasreen, and Aviral Kumar Tiwari. 2021. Volatility transmission and spillover dynamics across financial markets: The role of geopolitical risk. *Annals of Operations Research* 305: 1–22. [\[CrossRef\]](#)
- Fama, Eugene F. 1970. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance* 25: 383–417. [\[CrossRef\]](#)
- Federle, Jonathan, Alexander Meier, Gernot J. Müller, and Victor Sehn. 2026. Proximity to war: The stock market response to the Russian invasion of Ukraine. *Journal of Money, Credit and Banking* 58: 681–703. [\[CrossRef\]](#)
- Hoffmann, Valentin, Dieter Neuenkirch, and Bo Sheng. 2025. What drives abnormal returns of stock markets in wartime? Evidence from 17 invasions. *European Journal of Political Economy* 86: 102624. [\[CrossRef\]](#)
- Ijaz, Muhammad S., Shaker Ali, Alan M. Du, and Muhammad Khurram. 2025. Analyzing financial market reactions to the Palestine–Israel conflict: An event study perspective. *International Review of Economics & Finance* 98: 103864. [\[CrossRef\]](#)
- Jaffe, Jeffrey F. 1974. Special information and insider trading. *The Journal of Business* 47: 410–28. [\[CrossRef\]](#)
- Kilian, Lutz. 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99: 1053–69. [\[CrossRef\]](#)
- Kilian, Lutz, and Cheolbeom Park. 2009. The impact of oil price shocks on the U.S. stock market. *International Economic Review* 50: 1267–87. [\[CrossRef\]](#)
- Kolari, James W., and Seppo Pynnönen. 2010. Event study testing with cross-sectional correlation of abnormal returns. *The Review of Financial Studies* 23: 3996–4025. [\[CrossRef\]](#)
- Liu, Yang, Jian Yang, and Thuan P. Pham. 2025. The hidden cost of firm-level political risk: Impairing liquidity in corporate bond markets. *Finance Research Letters* 85: 108058. [\[CrossRef\]](#)
- MacKinlay, A. Craig. 1997. Event studies in economics and finance. *Journal of Economic Literature* 35: 13–39.
- Markoulis, Stelios, Savvas Katsikides, and Nicos Achiotis. 2020. The effect of terrorism on stock markets: Evidence from the 21st century. *Terrorism and Political Violence* 32: 988–1010. [\[CrossRef\]](#)
- Martins, António Miguel. 2024. Global equity, commodities and bond market response to the Israel–Hamas war. *Finance Research Letters* 60: 104871.
- Medhioub, Imed. 2025. Impact of geopolitical risks on herding behavior in some MENA stock markets. *Journal of Risk and Financial Management* 18: 85. [\[CrossRef\]](#)
- Middle East Council on Global Affairs. 2026. Qatar's Strategic Balancing amid Escalation and Disruption. Available online: <https://mecouncil.org/publication/qatars-strategic-balancing-amid-escalation-and-disruption/> (accessed on 15 April 2026).
- Mokni, Khaled. 2020. Time-varying effect of oil price shocks on the stock market returns: Evidence from oil-importing and oil-exporting countries. *Energy Reports* 6: 605–19. [\[CrossRef\]](#)
- News On Air. 2026. Third round of diplomatic negotiations between US and Iran to convene in Geneva tomorrow. *All India Radio (Akashvani)*, February 25. Available online: <https://www.newsonair.gov.in/third-round-of-diplomatic-negotiations-between-us-and-iran-to-convene-in-geneva-tomorrow/> (accessed on 8 September 2026).
- Pandey, Devendra Kumar. 2025. Effects of the Israel–Iran conflict: Insights on global stock indices and currencies. *Journal of Economic Studies* 52: 762–83. [\[CrossRef\]](#)
- Pandey, Devendra Kumar, Vineeta Kumari, Alessandro Palma, and John W. Goodell. 2024. Are markets in happier countries less affected by tragic events? Evidence from market reaction to the Israel–Hamas conflict. *Finance Research Letters* 60: 104276. [\[CrossRef\]](#)
- Park, Juwon S., and Mohammad K. Newaz. 2018. Do terrorist attacks harm financial markets? A meta-analysis of event studies and the determinants of adverse impact. *Global Finance Journal* 37: 227–47. [\[CrossRef\]](#)
- Pastor, Lubos, and Pietro Veronesi. 2013. Political uncertainty and risk premia. *Journal of Financial Economics* 110: 520–45. [\[CrossRef\]](#)

- Ready, Robert C. 2018. Oil prices and the stock market. *Review of Finance* 22: 155–76. [CrossRef]
- Reuters. 2026a. Oil prices extend gains as U.S.–Iran war deadlock keeps supply off market. *Reuters*, April 30. Available online: <https://www.reuters.com/business/energy/oil-prices-extend-gains-us-iran-war-deadlock-keeps-supply-off-market-2026-04-30/> (accessed on 1 May 2026).
- Reuters. 2026b. QatarEnergy declares force majeure on LNG contracts. *Reuters*, March 24. Available online: <https://www.reuters.com/business/energy/qatarenergy-declares-force-majeure-lng-contracts-2026-03-24/> (accessed on 25 March 2026).
- Saâdaoui, Foued, Samir Ben Jabeur, and John W. Goodell. 2023. Geopolitical risk and the Saudi stock market: Evidence from a new wavelet packet multiresolution cross-causality. *Finance Research Letters* 53: 103654. [CrossRef]
- Sakaki, Hiroyuki. 2019. Oil price shocks and the equity market: Evidence for the S&P 500 sectoral indices. *Research in International Business and Finance* 49: 137–55. [CrossRef]
- Smales, Lee A. 2019. Geopolitical risk and volatility spillovers in oil and stock markets. *The Quarterly Review of Economics and Finance* 80: 358–66. [CrossRef]
- Sokhanvar, Amir, and Elie Bouri. 2023. Commodity price shocks related to the war in Ukraine and exchange rates of commodity exporters and importers. *Borsa Istanbul Review* 23: 44–54. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.