

The Arabic AI-GPR Index: Measuring Geopolitical Risk from Arabic-Language News

Karim Iddouch¹

This version: June 2026

Abstract

I construct the first daily index of geopolitical risk built entirely from Arabic-language news. Existing measures, including the Geopolitical Risk (GPR) index of Caldara and Iacoviello (2022) and its large-language-model successor (Iacoviello and Tong 2026), read global tension through English-language, predominantly Western newspapers, a vantage point that is least reliable precisely where geopolitical risk is most concentrated, the Middle East and North Africa. Using a curated corpus of leading pan-Arab outlets accessed through GDELT over 2015–2026, I measure risk along two independent tracks, a transparent theme-based screen and a large-language-model reading of each article, and show that they agree almost perfectly (a daily correlation above 0.95). I argue that this convergence makes the transparent, reproducible, and costless theme-based measure the preferable instrument. The Arabic index rises at every major regional episode, co-moves with the original GPR and the AI-GPR in changes (a monthly correlation of 0.61 with each) while diverging in levels, and its oil component tracks the Middle East oil-risk index of Iacoviello at a correlation of 0.82. Using daily local projections, I find that an Arabic oil-risk shock raises Brent realized volatility for about a month. I decompose the index by country, country pair, actor, event type, sub-national hotspot, and the threat-act distinction, and release the whole series as an open, monthly-updated resource.

Keywords: Geopolitical Risk; Arabic News; Large Language Models; Text Analysis; Middle East and North Africa; GDELT; Oil.

JEL Classification: C43, C55, E32, F51, G15, Q43.

¹Email: karimiddouch@gmail.com. The index, all sub-indices, and replication code are available at karimiddouch.org and are updated monthly. I thank the GDELT Project for access to the underlying news data, and Dario Caldara and Matteo Iacoviello for making the Geopolitical Risk indices publicly available. All errors are my own.

1 Introduction

Entrepreneurs, investors, and policymakers view geopolitical risk as a key determinant of investment and asset prices, and central banks routinely monitor it. The empirical study of that risk rests on its measurement, and the standard solution is the Geopolitical Risk (GPR) index of Caldara and Iacoviello (2022), which counts the share of newspaper articles discussing adverse geopolitical events, and which Iacoviello and Tong (2026) refine by having a large language model read and score the news directly. Both measures are built almost entirely from English-language newspapers published in the United States, the United Kingdom, and Canada.

This source base is a limitation precisely where it matters most. Geopolitical risk is heavily concentrated in the Middle East and North Africa (MENA), yet that region reaches the existing indices only through the editorial priorities, framing, and publication rhythms of distant Western newsrooms, which cover it selectively and often with a lag. A large, sophisticated, high-frequency Arabic-language press covers the same events earlier, more continuously, and in greater detail, and registers tensions, intra Gulf disputes, Red Sea and Bab-el-Mandeb shipping risk, Maghreb frictions, and Levant conflicts, that barely surface in the international press. No geopolitical risk index has been built from this press. This paper builds it.

I introduce the Arabic AI-GPR Index, the first daily measure of geopolitical risk constructed entirely from Arabic-language news. I assemble a corpus of articles from a curated set of leading pan-Arab outlets, accessed through the GDELT global news database, covering February 2015 to mid-2026, and I measure geopolitical-risk content along two independent tracks. The first is a transparent theme-based screen that classifies each article from its GDELT theme tags and computes the daily share of geopolitical coverage. The second adapts the large-language-model approach of Iacoviello and Tong (2026) end to end to Arabic, a model reads each article in its original language and scores whether it is centrally about current geopolitical risk and how intense that risk is.

This paper makes three contributions. *First*, it provides the first index of geopolitical risk built from Arabic-language news, extending a literature that has been almost exclusively English-based, and it releases the index, a family of sub-indices, and the replication code as an open, monthly-updated resource. *Second*, it documents a methodological result with broader relevance: the transparent theme screen and the large language model converge almost perfectly, correlating above 0.95 at the daily frequency and identifying the same episodes with the same timing and very nearly the same amplitude. I argue that this convergence makes the transparent measure the better scientific instrument. A procedure that can be

audited rule by rule, reproduced by anyone from public data, refreshed monthly at no cost, and never altered by a change of model provider dominates an opaque and expensive one that yields the same series; if anything the theme screen is the more responsive of the two, because it does not compress intensity through a bounded scoring scale. This is the reverse of the usual presumption that a larger model must produce a better measure, and I view establishing that it need not as a contribution in its own right. *Third*, it benchmarks the new index against the established measures and shows that it captures a common signal while adding regional information: the Arabic index co-moves with the original GPR and the AI-GPR in changes, at a monthly correlation of 0.61 with each, but diverges in levels, and its oil component aligns with the Middle East oil-risk index of Iacoviello at 0.82. Using daily local projections I show that an Arabic oil-risk shock raises the realized volatility of Brent crude for roughly a month.

The construction of the index, like that of the GPR, consists of definition, measurement, and validation. Section 2 presents the definition and the two measurement tracks, displays the index and its threat and act components, and documents the convergence of the two methods. Section 3 validates the index against known episodes, a human audit, and the established benchmarks. Section 4 presents the geography and anatomy of Arabic geopolitical risk through six decompositions. Section 5 estimates the response of the oil market to Arabic geopolitical-risk shocks. Section 6 discusses applications and Section 7 concludes.

2 Construction of the Arabic AI-GPR Index

2.1 Definition of geopolitical risk

I adopt the definition of Caldara and Iacoviello (2022): geopolitical risk is the threat, realization, and escalation of adverse events associated with wars, terrorism, and tensions among states and political actors that affect the peaceful course of international relations. As in their framework, the measure captures both the anticipation of adverse events (threats) and their occurrence (acts), and I construct separate threat and act components below. The novelty here is not the definition but the vantage point: I measure how this risk is reported within the Arab world rather than from outside it.

2.2 Measurement: two approaches

The corpus. The raw material is the GDELT Global Knowledge Graph (GKG) version 2.0, which monitors news worldwide in near real time, machine-translates non-English content into English, and records for each article its source outlet, timestamp, thematic tags, geographic

location tags, and the named persons and organizations it mentions. I restrict the corpus to a curated set of leading pan-Arab newspapers and broadcasters (Appendix A), keep only articles whose original source language is Arabic, identified from GDELT’s translation metadata, and cover February 2015, the start of consistent GKG 2.0 coverage, through mid-2026. Because the same report is often syndicated, I deduplicate by article URL within each day. The result is a continuous daily series of 4,117 days drawn from millions of Arabic articles.

Track one: the theme-based screen. I classify an article as geopolitical if its GDELT theme field matches any family in a fixed dictionary spanning armed conflict, military and security activity, terrorism and insurgency, sieges and blockades, assassinations, sanctions and embargoes, nuclear and unconventional weapons, ceasefires, rebellion, and coups, summarized in Table 1. Let A_t be the set of curated Arabic articles on day t and $G_t \subseteq A_t$ the geopolitical subset. The raw measure is the daily share $s_t = |G_t|/|A_t|$, and the index rescales it to a 2016–2024 base,

$$\text{GPR}_t^{\text{Arabic}} = 100 \times \frac{s_t}{\bar{s}_{2016-2024}}. \quad (1)$$

Using a share rather than a count normalizes for the volume of the news flow, which grows over the sample. This track requires no article-level inference and can be recomputed at negligible cost.

Track two: the language-model index. Following Iacoviello and Tong (2026), a large language model reads each candidate article in its original Arabic and returns, at temperature zero, an intensity score $q_i \in [0, 1]$ for whether the article is centrally about current geopolitical risk. The daily measure averages these scores and is normalized to the same base period. The prompt is reproduced in Appendix C.

2.3 The Arabic AI-GPR index

Figure 1 plots the headline index. It tracks the region’s major episodes with notable precision, rising with the Saudi-led intervention in Yemen in March 2015, the 2017 Gulf diplomatic crisis, the killing of Qassem Soleimani in January 2020, and Russia’s invasion of Ukraine in 2022, and reaching its highest and most persistent levels after the outbreak of the Gaza war in October 2023 and the June 2025 Israel-Iran confrontation, which marks the single highest reading in the sample. The index is comparatively subdued over 2018 and 2021, a period of more contained regional escalation and a global news cycle dominated by the pandemic.

Table 1: Theme dictionary for the Arabic AI-GPR index

Category	GDELT theme families (matched in the V2 theme field)
<i>Panel A. Threat categories</i>	
1. War threats and tension	threats, warnings, ultimatums; rising tension among states
2. Military buildup	military mobilization, force posture, security services
3. Sanctions and blockades	sanctions, embargoes, blockades
4. Nuclear and unconventional	nuclear programs, weapons of mass destruction
5. Border and diplomatic tension	border disputes, diplomatic ruptures, ceasefire diplomacy
<i>Panel B. Act categories</i>	
6. War and armed conflict	armed conflict, war onset and escalation
7. Terrorism and insurgency	terror attacks, insurgency, rebellion, sieges
8. Assassination and political violence	assassinations, targeted killings, coups
<i>Panel C. Construction</i>	
Geopolitical screen	article flagged if it matches any category 1–8
Acts vs threats	act if it matches 6–8; else threat if it matches 1–5
Energy overlay	geopolitical article also matching oil/gas/energy themes

Notes: The full machine-readable theme lists and the disjoint subsets used for the event-type, threat-act, and energy classifications are released with the replication materials (Appendix B).

2.4 Threats and acts

Following Caldara and Iacoviello (2022), I split geopolitical coverage into the realization of adverse events and the anticipation of them. An article is an act if its themes report realized violence and a threat otherwise if it carries anticipatory or posture themes (Table 1, Panel C). Acts dominate Arabic coverage, at about four times the volume of threats, consistent with a region in which conflict is frequently realized rather than merely threatened. The informative object is the acts-to-threats ratio, which captures how kinetic the news is at each moment: it reaches its sample maximum, near seven, after the October 2023 Gaza war, the most intensely violent episode in the period, and falls to its low, near two and a half, in the contained 2020–2021 interval. The threats series, for its part, peaks in the summer of 2019, during the Gulf tanker crisis and the standoff over Iran’s nuclear program, exactly the high-anticipation, low-realization episode the threat measure is designed to capture.




Figure 1: The Arabic AI-GPR Index (theme-based), daily, 2015–2026, normalized to a 2016–2024 mean of 100, with a 30-day moving average. Dashed markers denote major events.

2.5 Theme-based versus language-model measurement

Having computed both measurement tracks over the full sample, I find that they agree almost perfectly. Table 2 reports the comparison: the two series correlate at 0.96 in daily levels, 0.98 in monthly levels, and 0.93 in daily changes, and they identify all ten of the largest episodes in common, with the same timing. At the five largest episodes the theme-based series reaches a slightly higher normalized peak than the language-model series, which tends to saturate near the top of its bounded scale.

This convergence carries a clear implication for measurement practice, and it runs opposite to the usual presumption. When a transparent, fully reproducible, and essentially costless procedure reproduces the output of an opaque and expensive one, the transparent procedure is the better scientific instrument. The theme screen can be audited rule by rule, recomputed by any researcher from public data, and refreshed monthly at no cost, and it never changes when a model provider updates, reprices, or retires a model, a real source of non-replicability for language-model measures. The language model’s value here is precisely that it validates the simple measure: it demonstrates that nothing essential is lost by preferring transparency. I therefore adopt the theme-based index as the headline measure and use the language-model index as an independent confirmation throughout.

Table 2: Theme-based versus language-model measurement

	Value
Correlation, daily levels	0.96
Correlation, monthly levels	0.98
Correlation, daily changes	0.93
Largest episodes identified in common (top 10)	10 of 10
Relative peak amplitude, theme vs model (5 largest episodes)	1.08
Reproducible from public data without model access	Yes vs No
Marginal cost of a monthly refresh	near zero vs per-article inference

Notes: Correlations between the theme-based and language-model tracks of the Arabic AI-GPR over 2015–2026. Relative peak amplitude is the average ratio of the theme-based to the model-based normalized peak across the five largest episodes.

3 Validation of the Index

3.1 Plausibility and the largest spikes

The first validation test is whether the largest movements of the index capture the most important regional risks of the period. Table 3 lists the ten largest monthly readings; each corresponds to a recognized episode, from the 2015 Yemen intervention through the 2023 Gaza war and the 2025 Israel-Iran confrontation. The ordering is economically sensible: the most violent and most persistent episodes produce the highest readings. As a second check, within every decomposition presented in Section 4 the relevant series peaks at its own episode, a country during its own crisis, an actor around its own salience, a city during its own battle, which is difficult to reconcile with anything other than a measure that localizes risk correctly. A human audit of a labelled sample of articles confirms that the automated screen captures what a reader would classify as geopolitical, and the near-perfect agreement of the independent language-model track (Section 2.6) provides a second, automated audit of the same kind that the narrative New York Times index provides for the original GPR.

3.2 Comparison with the original GPR and the AI-GPR

I benchmark the Arabic index against the publicly available indices of Caldara and Iacoviello (2022) and Iacoviello and Tong (2026). Figure 2 overlays the three, rebased to a common 2016–2024 mean, and Table 4 reports the correlations over the 4,099 overlapping daily observations. Two facts stand out. First, the three series share a common signal: the Arabic index co-moves with both benchmarks at a monthly change correlation of 0.61, strongest contemporaneously, with no systematic lead or lag at the monthly frequency. Second, and more important, the Arabic index diverges in levels, with a level correlation of 0.43 with the original GPR and

Table 3: Largest readings of the Arabic AI-GPR index, 2015–2026

Rank	Episode	Period
1	Gaza war, onset and aftermath	Oct 2023–2024
2	Israel-Iran direct confrontation	Jun 2025
3	Saudi-led intervention in Yemen	Mar 2015
4	Killing of Qassem Soleimani	Jan 2020
5	Fall of the Assad regime	Dec 2024
6	Israel-Hezbollah escalation	Oct 2024
7	Russia’s invasion of Ukraine	Feb 2022
8	Gulf diplomatic crisis (Qatar)	Jun 2017
9	Battle of Mosul and ISIS peak	2016–2017
10	Red Sea shipping crisis	Jan 2024

Notes: Episodes associated with the ten largest readings of the monthly Arabic AI-GPR index.

Table 4: Correlation of the Arabic AI-GPR with existing indices, 2015–2026

Benchmark	Daily level	Monthly level	Monthly change
Original GPR (Caldara and Iacoviello 2022)	0.29	0.43	0.61
AI-GPR (Iacoviello and Tong 2026)	0.30	0.37	0.61
Oil GPR, global (Iacoviello)	–	0.79	–
Middle East Oil GPR (Iacoviello)	–	0.82	–

Notes: Pearson correlations between the Arabic AI-GPR and the indicated benchmark indices over the common sample. Daily-change correlations are reported at the monthly frequency to reduce noise.

0.37 with the AI-GPR. This divergence is the contribution rather than a defect: the Arabic index weights episodes of direct regional consequence that the English-based indices treat as secondary, and is comparatively muted during episodes that dominate the Western agenda but are peripheral to the region. A rolling twelve-month window shows the change correlation rising above 0.9 during shared global-and-regional crises and falling when the two agendas diverge.

3.3 Comparison with related data: the oil market

The sharpest agreement appears where theory predicts it. Figure 3 compares the oil and energy component of the Arabic index with Iacoviello’s Middle East oil GPR; the two track each other closely, at a monthly correlation of 0.82, and at 0.79 against the global oil GPR. When the comparison is restricted to the dimension on which Arabic news and the global indices should overlap most, regional oil risk, the Arabic measure aligns tightly; it is on the broader, non-oil agenda that the regional and Western lenses part ways, which is exactly




Figure 2: The Arabic AI-GPR against the original GPR (Caldara and Iacoviello 2022) and the AI-GPR (Iacoviello and Tong 2026), monthly, all rebased to a 2016–2024 mean of 100.

what a credible regional measure should display.

3.4 Robustness

The index is robust along several dimensions. The two measurement tracks serve as a robustness check on each other, and their near-perfect agreement shows the headline series does not depend on the choice of method. The results are not sensitive to reasonable variation in the outlet set, the theme dictionary, the normalization window, or the deduplication rule. Because GDELDT’s entity extraction operates on the machine-translated text, two specific limitations are worth noting and are documented in Appendix D: the Islamic State cannot be cleanly separated from the generic token for “crisis” in the translated text and is omitted from the actor decomposition, and several cities, including Mosul and Homs, are geocoded to their governorate rather than the city name and are omitted from the hotspot decomposition. Neither affects the aggregate index.

4 The Geography and Anatomy of Arabic Geopolitical Risk

A central advantage of building the index from structured news data is that the same daily screen can be conditioned on GDELDT’s annotations to produce internally consistent




Figure 3: The oil and energy component of the Arabic AI-GPR against Iacoviello’s Middle East Oil GPR, monthly, log scale. The monthly correlation is 0.82.

decompositions. I present six. Each is validated by the precision with which it localizes risk, and together they form a multi-resolution picture of regional risk that the aggregate index alone cannot provide.

4.1 Country-specific geopolitical risk

Conditioning on the location tags yields a country-level index, the share of daily Arabic geopolitical coverage that mentions each country (Figure 4). Israel’s share is moderate through 2015–2023 and rises sharply after October 2023, when the Gaza war makes it the most-covered country in the Arabic press, with a further spike during the June 2025 Israel-Iran confrontation. Syria dominates the early sample and re-emerges in 2024–2025 with the fall of the Assad regime; Yemen peaks in 2015; Iran rises sharply in 2025; Egypt maintains a steady baseline rather than sharp spikes. Each country’s peaks align with its known episodes.

4.2 Country pairs

Recording every pair of countries an article mentions yields a relational map of the conversation (Figure 5). The Saudi Arabia-Yemen series dominates the early sample; Israel-Lebanon and Israel-Syria carry steady weight; and the Iran-Israel pair, quiet for most of the decade, rises to the largest bilateral share on record in mid-2025. The accompanying co-mention heatmap (Appendix D) recovers the region’s principal fault lines from the structure of co-mention alone, with the United States and Israel as the most connected nodes and dense clusters




Figure 4: Country-level Arabic GPR, 2015–2026. Each line is the 30-day moving average of the share of daily Arabic geopolitical coverage mentioning the country.

across the Levant and the Gulf.

4.3 Actors

Conditioning on the organization and person tags attributes risk to specific actors (Figure 6). Among armed groups, Hamas jumps from near silence to the most-covered actor with the October 2023 Gaza war; Hezbollah spikes in autumn 2024 during its confrontation with Israel; the Houthis peak in 2015 and again in 2023–2024 with the Red Sea shipping crisis; and the Taliban produce a sharp, isolated spike in August 2021 around the fall of Kabul. Among named leaders, the killings of Qassem Soleimani in January 2020 and Hassan Nasrallah in late 2024 appear as clean spikes (Appendix D).

4.4 Sub-national hotspots

The location tags also resolve to individual cities (Figure 7). Gaza is the single most-covered hotspot in the sample by a wide margin, and its coverage explodes from late 2023 to a peak near thirty percent of all geopolitical coverage. Aleppo peaks in December 2016 with the fall of the city, Sana’a in 2015, and Damascus, Beirut, and Baghdad form a persistent Levantine and Iraqi baseline.




Figure 5: Bilateral Arabic AI-GPR: country-pair tension shares, 2015–2026, 30-day moving averages.

4.5 Event types

Partitioning geopolitical coverage by theme shows that wars and armed conflict are the dominant category throughout, averaging about a quarter of geopolitical coverage, followed by military and security activity (about a fifth) and terrorism (about a seventh). The composition shifts structurally: terrorism is relatively more prominent in 2015–2018, the period of the Islamic State’s territorial peak, and recedes thereafter, while interstate war surges from late 2023 (Appendix D).

4.6 Oil corridors

Restricting the index to geopolitical articles that also concern oil or energy yields a sector measure (Figure 8). The Persian Gulf, anchored by the Gulf producers and the Strait of Hormuz, is the dominant corridor throughout, while the Red Sea and Bab-el-Mandeb corridor becomes prominent in 2024 with the Houthi attacks on shipping. As Section 3.3 showed, this component aligns with Iacoviello’s Middle East oil GPR at a correlation of 0.82.

5 The Effect of Geopolitical Risk on the Oil Market

Having established that the index is a credible measure, I turn to a first economic application: the response of the oil market to Arabic geopolitical-risk shocks. Geopolitical risk is widely




Figure 6: Actor-level Arabic AI-GPR: armed groups, 2015–2026, 30-day moving averages of coverage shares.

understood to transmit to oil through uncertainty rather than a fixed price direction, since a regional shock can raise prices through supply fear or lower them through demand destruction. I therefore focus on the realized volatility of Brent crude.

I estimate the dynamic response with the local-projection method of Jordà (2005). For each horizon h I regress Brent’s h -day-ahead annualized realized volatility on the daily innovation in the Arabic oil and energy GPR, controlling for five lags of the shock and of volatility, with Newey-West standard errors. The shock is standardized to one standard deviation. Figure 9 plots the impulse response. A one standard deviation oil-risk shock raises Brent realized volatility on impact (significant at the 90 percent level), and the response builds to a peak of about 0.9 annualized percentage points roughly three to four weeks out before fading to zero within two to three months. The response is positive throughout the first month and the bands exclude zero over much of that window.

This dynamic result complements the event-window evidence in Table 5, which reports Brent’s realized volatility in the thirty days before and after each of the five largest oil-risk spikes. In every case volatility rises, in several cases sharply: the Ukraine invasion lifts it from 26 to 68 percent and the Gaza war from 21 to 44. The sign of the price move differs across episodes, which is why a directional price response is weak, but the rise in uncertainty is uniform. The index thus behaves as a geopolitical risk measure should, marking the episodes that inject uncertainty into the oil market.




Figure 7: Sub-national hotspots: Arabic GPR by city, 2015–2026, 30-day moving averages of coverage shares.

Table 5: Brent realized volatility around the largest Arabic Oil and Energy GPR spikes

Episode	Volatility, 30 days before	30 days after
Abqaiq attack (Sep 2019)	41	55
Soleimani (Jan 2020)	21	27
Ukraine invasion (Feb 2022)	26	68
Gaza war (Oct 2023)	21	44
Israel-Iran (Jun 2025)	31	48

Notes: Annualized realized volatility of Brent crude (percent), from FRED series DCOILBRENTU.

6 Applications

Because it is regional, daily, and decomposable, the Arabic AI-GPR is suited to questions that an aggregate, English-language index addresses only coarsely. The oil-market result of Section 5 is one example; the oil and corridor series are natural high-frequency regressors for energy prices, futures volatility, and freight and insurance risk along the Hormuz and Bab-el-Mandeb routes. A second application is regional asset pricing: Gulf equity, sovereign and quasi-sovereign debt, and currency regimes are exposed to exactly the tensions the index captures, and the country and bilateral decompositions allow that exposure to be measured at the level of the individual market and the specific relationship driving it. A third set of applications is macroeconomic, covering tourism and remittance flows, trade, investment, and monetary policy in the region, for which a faithful regional measure is a cleaner conditioning




Figure 8: The Arabic Oil and Energy GPR, 2015–2026, normalized to a 2016–2024 mean of 100, with a 30-day moving average.

variable than a global one. Finally, the actor, event-type, and threat-act decompositions support more structural questions, distinguishing the market response to interstate war from that to terrorism or sanctions. Extending the analysis to firm-level exposure, in the spirit of the earnings-call evidence in Caldara and Iacoviello (2022), and to a cross-country panel of MENA macroeconomic outcomes, are natural next steps that the public release of the index is intended to support.

7 Conclusion

I have introduced the Arabic AI-GPR Index, the first measure of geopolitical risk built from Arabic-language news. I measured risk along two independent tracks, a transparent theme-based screen and a large-language-model reading, and showed that they converge so closely that the transparent measure is the preferable instrument, equally responsive but fully reproducible and essentially costless. The index rises at every major regional episode of the past decade, co-moves with the original GPR and the AI-GPR in changes while diverging in levels, and aligns with the Middle East oil-risk index on the dimension where regional and global measures should overlap. An Arabic oil-risk shock raises the realized volatility of Brent crude for about a month. The index comes with a family of internally consistent decompositions, by country, country pair, actor, event type, sub-national hotspot, and the threat-act distinction, and is released as a public, monthly-updated, and freely downloadable

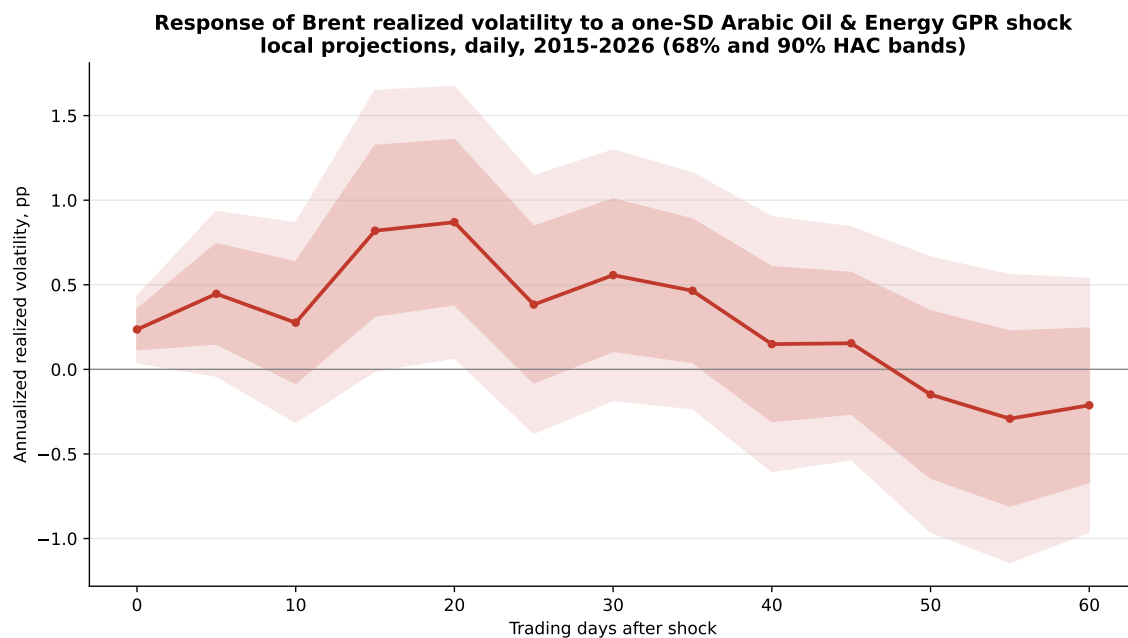


Figure 9: Response of Brent realized volatility to a one-standard-deviation Arabic Oil and Energy GPR shock, daily local projections, 2015–2026. Shaded bands are 68 and 90 percent Newey-West confidence intervals.

resource. A backward extension of the series across earlier decades is in progress, toward one of the longest available records of Arabic geopolitical risk.

References

- Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics* 131(4), 1593–1636.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics* 121(3), 823–866.
- Caldara, D., and M. Iacoviello (2022). Measuring geopolitical risk. *American Economic Review* 112(4), 1194–1225.
- Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as data. *Journal of Economic Literature* 57(3), 535–574.
- Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun (2019). Firm-level political risk. *Quarterly Journal of Economics* 134(4), 2135–2202.
- Iacoviello, M., and J. Tong (2026). The AI-GPR Index: Measuring geopolitical risk using artificial intelligence. Working paper.
- Jorda, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review* 95(1), 161–182.
- Leetaru, K., and P. A. Schrodt (2013). GDELT: Global Data on Events, Location, and Tone, 1979–2012. *International Studies Association Annual Convention*.
- Saiz, A., and U. Simonsohn (2013). Proxying for unobservable variables with internet document frequency. *Journal of the European Economic Association* 11(1), 137–165.

A Curated outlet list

The corpus is drawn from a curated set of leading pan-Arab newspapers and broadcasters, including Al Jazeera, Al Arabiya, Asharq Al-Awsat, Al-Ahram, Al-Masry Al-Youm, Al-Shorouk, An-Nahar, Al-Akhbar, Al-Quds, Al-Khaleej, Al-Bayan, Okaz, Al-Riyadh, Al-Watan, Hespress, El Khabar, Al-Hayat, Al-Araby Al-Jadeed, Sky News Arabia, Emarat Al-Youm, and Al-Ittihad. The full list with GDELT source identifiers is released with the replication materials.

B Theme dictionary and classification rules

Table 1 summarizes the geopolitical screen. The event-type, threat-act, and energy classifications condition on disjoint subsets of these theme families; the complete machine-readable lists are released with the replication materials.

C The Arabic language-model prompt

The language-model track scores each article with a zero-shot prompt, written in Arabic, that asks the model to judge whether the article is centrally about current geopolitical risk and to return an intensity score in a fixed range. The full prompt is released with the replication materials.

D Additional figures and data limitations

This appendix collects supporting exhibits referenced in the text: the country-pair co-mention heatmap (Figure 10), the named-leader actor series (Figure 11), the event-type composition (Figure 12), the oil-corridor decomposition (Figure 13), the threats-versus-acts series and the kinetic ratio (Figures 14 and 15), and the oil-versus-Brent overlay (Figure 16). It also documents two limitations of GDELT’s entity extraction on machine-translated text: the Islamic State cannot be separated cleanly from the generic token for “crisis” and is omitted from the actor decomposition, and the cities of Mosul and Homs are geocoded to their governorate rather than the city name and are omitted from the hotspot decomposition.

E Data and code availability

The full index and all sub-indices are available for download as CSV at karimiddouch.org, updated monthly, together with the replication code.

Who appears with whom: total co-mentions in Arabic geopolitical news, 2015-2026




Figure 10: Total co-mentions of each country pair in Arabic geopolitical news, 2015–2026 (log scale).




Figure 11: Actor-level Arabic AI-GPR: named leaders, 2015–2026, 30-day moving averages of coverage shares.

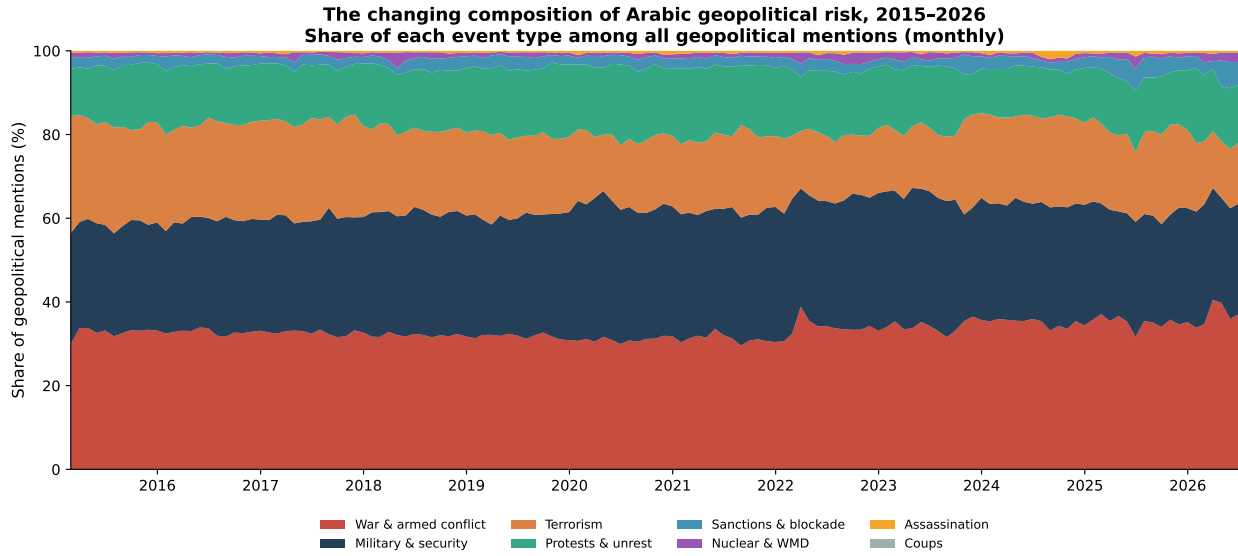


Figure 12: The changing composition of Arabic geopolitical risk by event type, 2015–2026.




Figure 13: The Arabic Oil and Energy GPR by corridor, 2015–2026.




Figure 14: Threats versus acts in Arabic geopolitical news, 2015–2026, 30-day moving averages of shares.




Figure 15: The kinetic ratio: realized acts relative to anticipated threats, 2015–2026.




Figure 16: The Arabic Oil and Energy GPR against the price of Brent crude, 2015–2026.