The Arabic AI-GPR Index: Measuring Geopolitical Risk using Artificial Intelligence and Arabic-Language News

Karim Iddouch[1]

June 2026

Abstract

This paper introduces the *Arabic AI-GPR Index*, the first measure of geopolitical risk constructed from Arabic-language news. Building on the Geopolitical Risk (GPR) framework of Caldara and Iacoviello (2022) and the large-language-model approach of Iacoviello and Tong (2026), I assemble a corpus of millions of Arabic-language articles from leading pan-Arab outlets and use artificial intelligence to identify and measure their geopolitical-risk content. Whereas existing indices read global tension through English-language, predominantly Western newspapers, the Arabic AI-GPR captures how geopolitical risk is perceived and reported within the Arab world itself. The resulting daily index spans 2015–2026 and rises sharply around major regional episodes (the Saudi-led intervention in Yemen, the 2017 Gulf diplomatic crisis, the killing of Qassem Soleimani, the 2023 Gaza war, and the 2025 Israel–Iran confrontation), while co-moving with, yet meaningfully diverging from, the original English-based GPR. The index offers a regional, high-frequency proxy for political tension well suited to studying oil markets, Gulf and MENA asset prices, trade, and macroeconomic fluctuations.

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

JEL Classification: C55, C82, D80, F51, G15.

[1]Email: karimiddouch@gmail.com.

1 Introduction

Geopolitical risk, the threat and the realization of adverse events such as wars, terrorism, and tensions between states, is a first-order driver of economic activity. Episodes of heightened geopolitical tension depress investment and hiring, raise oil and commodity prices, disrupt trade and supply chains, and move asset prices, risk premia, and cross-border capital flows. Because the risk itself is latent and cannot be observed directly, a central empirical challenge is to quantify it from observable data. Caldara and Iacoviello (2022) provide the standard solution with their Geopolitical Risk (GPR) index, which measures risk from the frequency of newspaper articles discussing adverse geopolitical events. The GPR has since become a benchmark across macroeconomics and finance, complementing related text-based measures of economic uncertainty such as the Economic Policy Uncertainty index of Baker, Bloom, and Davis (2016).

Yet virtually all existing measures of geopolitical risk read global tension through a narrow lens: English-language newspapers, overwhelmingly published in the United States and the United Kingdom. This Anglocentric vantage point is problematic precisely where geopolitical risk is most concentrated: the Middle East and North Africa (MENA). Events in the region reach these indices only after passing through the editorial priorities, framing, and publication rhythms of Western newsrooms, which devote attention to MENA selectively and often with a lag. Local and pan-Arab outlets, by contrast, cover regional escalations earlier, more continuously, and in far greater detail, and they capture tensions (intra-Gulf disputes, Red Sea and Bab-el-Mandeb shipping risk, Maghreb border frictions, Levant conflicts) that barely register in the international press. To date, no geopolitical risk index has been constructed from Arabic-language news.

This paper fills that gap. I introduce the Arabic AI-GPR Index, the first daily measure of geopolitical risk built from Arabic-language news. I assemble a large corpus of Arabic articles drawn from a curated set of leading pan-Arab outlets, accessed through the GDELT global news database, spanning February 2015 to mid-2026. Building directly on Iacoviello and Tong (2026), I use artificial intelligence to read and assess the geopolitical-risk content of news, moving beyond keyword counting toward semantic understanding: rather than tallying the appearance of fixed terms, a large language model evaluates whether an article is centrally about current geopolitical risk and how intense that risk is. I pair this with a transparent, fully reproducible theme-based baseline that requires no article-level inference, allowing the index to be both auditable and refreshed at low cost every month.

The resulting index displays strong face validity. It rises sharply around the defining episodes of the past decade in the region: the Saudi-led intervention in Yemen in 2015, the

2017 Gulf diplomatic crisis, the killing of Qassem Soleimani in early 2020, Russia's 2022 invasion of Ukraine, the Gaza war beginning in October 2023 (which produces the largest and most persistent elevation in the sample), and the June 2025 Israel–Iran confrontation, which marks the single highest reading in the series. The Arabic index co-moves with the original English-based GPR, confirming that the two capture a common underlying phenomenon, but it diverges in economically meaningful ways: it reacts more strongly and more promptly to events of direct regional consequence, reflecting a distinctly Arab vantage point on global instability rather than a Western one.

These differences are not merely descriptive; they matter for measurement and application. A MENA-focused, high-frequency index of geopolitical risk is a more faithful proxy for the tensions that shape outcomes in the region (oil and energy markets, Gulf equity and sovereign-debt markets, tourism and remittance flows, and bilateral trade) than an index assembled from distant newsrooms. By measuring risk closer to its source, the Arabic AI-GPR offers researchers and policymakers a sharper instrument for studying how geopolitical shocks transmit to MENA economies and to the global markets they influence.

This paper makes four contributions. First, it provides the first index of geopolitical risk built from Arabic-language news, extending a literature that has so far been almost exclusively English-based. Second, it adapts the large-language-model measurement approach end to end to a new language and a new media environment, and documents that an AI-based reading of the news captures geopolitical risk more precisely than keyword matching. Third, it lays the groundwork for a family of forthcoming sub-indices: by country, by event type (wars, terrorism, diplomatic tensions, coups, nuclear threats), by the distinction between anticipated threats and realized acts, and by oil-supply and bilateral dimensions. Fourth, it is released as a public, monthly-updated resource, with a planned extension of the series back across earlier decades toward the longest available record of Arabic geopolitical risk. The remainder of the paper proceeds as follows. Section 2 describes the data and the measurement methodology. Section 3 presents the index and compares it with the original GPR. Section 4 introduces the regional and categorical decompositions. Section 5 discusses applications, and Section 6 concludes.

2 Data and Measurement

This section describes how the index is built, from the underlying news corpus through the screening, scoring, and aggregation steps to the decompositions and the validation strategy. The guiding principle is that every quantity in the paper is derived from a single, transparent pipeline, so that the headline index and its sub-indices are mutually consistent and the whole

construction can be reproduced and refreshed each month.

The Arabic news corpus

The raw material is the Global Database of Events, Language, and Tone (GDELT), specifically its Global Knowledge Graph (GKG) version 2.0. GDELT monitors news worldwide in real time, machine-translates non-English content into English, and records for each article a rich set of structured annotations: the source outlet, the publication timestamp, thematic tags, geographic location tags, and the named persons and organizations the article mentions. Three features make it well suited to the present purpose: it covers Arabic-language outlets continuously and at daily frequency, it preserves the source language of each article (so Arabic reporting can be isolated), and its thematic and locational tagging allows the same corpus to be sliced in many internally consistent ways.

Three choices define the sample. First, rather than ingest the entire Arabic web, I restrict attention to a curated set of leading pan-Arab newspapers and broadcasters with regional reach and editorial consistency. Curation trades breadth for quality: it removes aggregators, low-signal sources, and spam, and it keeps the composition of the corpus stable over time, which matters when the object of interest is a change in coverage rather than a change in the set of sources. Second, I keep only articles whose original source language is Arabic, identified from GDELT's translation metadata. This is the defining restriction of the project: the index reflects how regional tension is reported within the Arab world, not how Western newsrooms report on the region. Third, the sample runs from February 2015, the earliest date of consistent GKG 2.0 coverage, to mid-2026, and the unit of observation is the article-day. Because the same story is often syndicated, I deduplicate by article URL within each day, so that a single report counts once. The complete list of outlets and the exact tag definitions are provided in the replication materials.

Measuring geopolitical risk: a two-track design

Geopolitical risk is latent: it cannot be read off any single observable, only inferred from how intensively the news discusses adverse geopolitical events. I infer it on two complementary tracks that measure the same underlying quantity in different ways.

The first track is a theme-based baseline. An article is classified as geopolitical if GDELT assigns it any theme in a fixed set covering armed conflict, military and security activity, terrorism and insurgency, sieges and blockades, assassinations, sanctions, nuclear and other unconventional weapons, ceasefires, and coups. The daily measure is the share of curated Arabic coverage that is geopolitical in this sense. Using a share rather than a raw count is

deliberate: it normalizes for the sheer volume of the news flow, which varies across days and grows over the sample, so that the index reflects the prominence of geopolitical risk rather than the size of the press. This track requires no article-level inference, is fully transparent, and can be recomputed cheaply every month, which is why it serves as the backbone of the published series.

The second track is a large-language-model index that follows Iacoviello and Tong (2026) and supplies what the baseline cannot, an intensity weighting. Here a large language model reads each article and judges, in a single zero-shot pass, whether the article is centrally about current geopolitical risk and, if so, how severe that risk is, returning a structured score. The model is queried at temperature zero for determinism, and the prompt is written in Arabic so that the model reasons over the language of the source rather than a translation. This semantic reading distinguishes a front-page account of an active war from an incidental mention of a past conflict, a distinction that pure keyword or theme matching cannot make. The LLM track is computationally heavier and is still being scored across the full corpus, so the results in this paper rest on the theme-based baseline, with the LLM index to follow in the next version. The two are designed to be combined: the baseline locates the days and topics that carry geopolitical attention, and the LLM weights them by intensity.

From articles to a daily index

Both tracks aggregate to a daily series. For the theme-based baseline, the daily geopolitical share is normalized so that its average over the 2016–2024 window equals 100, a stable interior period that excludes the partial first year and the most recent, still-evolving episodes; the resulting level is therefore interpretable as a percentage of its own typical value. For readability the figures plot a 30-day moving average, while the underlying daily series is retained for any quantitative analysis. The same normalization convention is applied to the sector (oil and energy) index, so that levels are comparable across the family.

Building the sub-indices

A central advantage of building the index from structured news data is that the same daily screen can be conditioned on GDELT's annotations to produce internally consistent decompositions. The country and bilateral indices condition on the location tags attached to each article, counting coverage that mentions a given country, or that mentions a given pair of countries together. The city index conditions on the sub-national place names in the same tags. The event-type index conditions on the thematic tags, partitioning geopolitical coverage into wars, terrorism, military activity, unrest, sanctions, and the like. The oil and energy

index adds an energy-theme filter, retaining geopolitical articles that also concern oil or gas and tagging them by corridor. The actor index conditions on the person and organization tags, attributing coverage to specific armed groups and named leaders. Because all of these are derived from the one screen, each sub-index is a conditional slice of the headline series rather than a separate construction, and a single article may contribute to several slices.

Validation strategy

The index is validated along four lines, developed in the sections that follow. First, face validity: the series should rise around known episodes and subside otherwise, both in the aggregate and within each decomposition, where the test is sharpest, since a country's index should peak during its own crises and an actor's around its own salience. Second, external comparison: the Arabic index should co-move with the original English-language GPR of Caldara and Iacoviello (2022), confirming a shared signal, while diverging where the Arab and Western news agendas differ. Third, human coding: a sample of articles is hand-labelled to check that the automated screen captures what a reader would call geopolitical. Fourth, market validation: the oil and energy index is compared with the price of Brent crude, the variable it should track most directly. Together these provide complementary evidence, from narrative to statistical, that the index measures what it claims to.

3 The Arabic AI-GPR Index

Preliminary evidence

Figure 1 offers a first look at the index. It plots the daily theme-based baseline of the Arabic AI-GPR from February 2015 through mid-2026, normalized to a 2016–2024 mean of 100, together with a 30-day moving average. Even in this preliminary, keyword-light form, the series tracks the region's major geopolitical episodes with notable precision: it rises around the Saudi-led intervention in Yemen (2015), the 2017 Gulf diplomatic crisis, the killing of Qassem Soleimani (January 2020), and Russia's invasion of Ukraine (2022), and it reaches its highest and most persistent levels following the outbreak of the Gaza war in October 2023. The index is comparatively subdued over 2018–2021, consistent with a period of more contained regional military escalation and a global news cycle dominated by the COVID-19 pandemic. Taken together, these patterns provide initial and encouraging evidence that Arabic-language news carries a coherent, economically meaningful signal of geopolitical risk.

Two caveats frame this evidence. First, the figure reports the *theme-based* baseline, which classifies articles from their GDELT geopolitical theme tags; it is fast, transparent, and fully

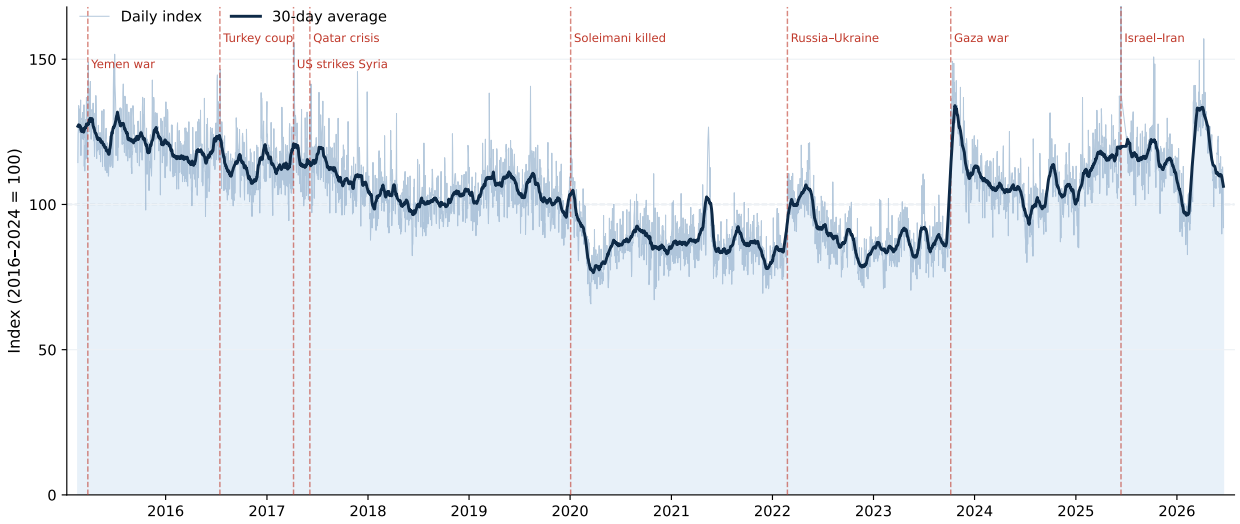

Figure 1: The Arabic AI-GPR Index (theme-based baseline), daily, 2015–2026. The series is the daily share of curated Arabic-language news classified as geopolitical (using GDELT theme tags), normalized to a 2016–2024 mean of 100; the solid line is a 30-day moving average. Dashed markers denote major events. Source: author's calculations using GDELT data.

reproducible, but it measures the *share* of coverage devoted to geopolitical risk rather than its intensity. Second, the full *LLM-based* index, in which a large language model reads and scores the intensity of each article following Iacoviello and Tong (2026), is under construction and will be released shortly. Preliminary comparisons indicate that it yields a sharper, intensity-weighted signal that separates incidental mentions from front-page crises; results based on the LLM index will be reported in the next version of this paper.

4 Decompositions of the Index

The geography of risk: a country-level decomposition

A central advantage of building the index from the news is that it can be decomposed by the countries each article concerns. Using the location tags that GDELT attaches to every article, I compute, for each day, the share of Arabic geopolitical coverage that mentions each country. Figure 2 plots this country-level decomposition for the most prominent states, and Figure 3 shows the same series as individual panels on a common scale.

The geography is revealing. Israel's share is moderate through 2015–2023 and then rises dramatically after October 2023, when the Gaza war makes it the single most-covered country in the Arabic press, a position it holds through 2026, with a further spike during the June 2025 Israel–Iran confrontation. Syria dominates the early sample (2015–2018), reflecting its

civil war and the territorial peak of ISIS, and re-emerges around 2024–2025 with the collapse of the Assad regime. Yemen peaks in 2015 at the onset of the Saudi-led intervention, while Iran rises sharply in 2025. Egypt, by contrast, maintains a steady baseline presence rather than sharp spikes, consistent with its demographic weight and centrality in pan-Arab media rather than with acute conflict.

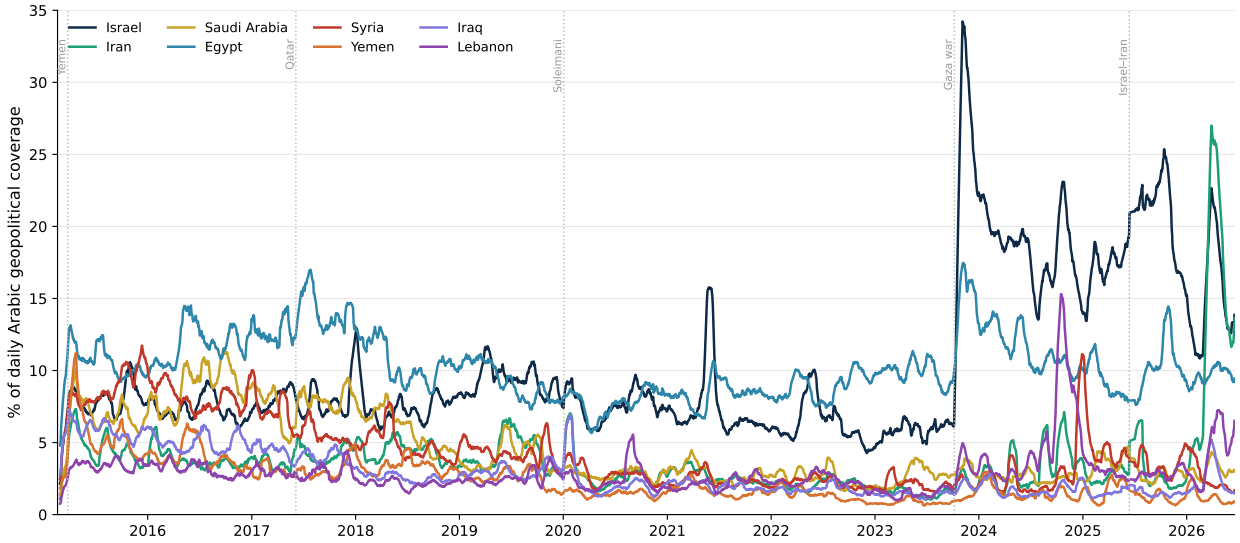

Figure 2: Country-level Arabic GPR (theme-based), 2015–2026. Each line is the 30-day moving average of the share of daily Arabic geopolitical coverage that mentions the country. Dashed markers denote major events. Source: author's calculations using GDELT location tags.

Two features stand out. First, each country's peaks align closely with its known episodes, a country-level analogue of the face validity documented for the aggregate index and reassuring evidence that the measure localizes risk correctly. Second, the decomposition is highly concentrated: a small set of Levantine and Gulf-adjacent states accounts for most geopolitical attention, while the Maghreb registers far lower and steadier coverage. The country decomposition thus turns the aggregate index into a time-varying map of which actors drive regional tension, and it provides the natural building block for the directed, country-pair (bilateral) measures developed below. As above, these series use the theme-based baseline and GDELT location tags; because a single article can mention several countries, the country shares need not sum to one, and the corresponding LLM-scored country indices are under construction.

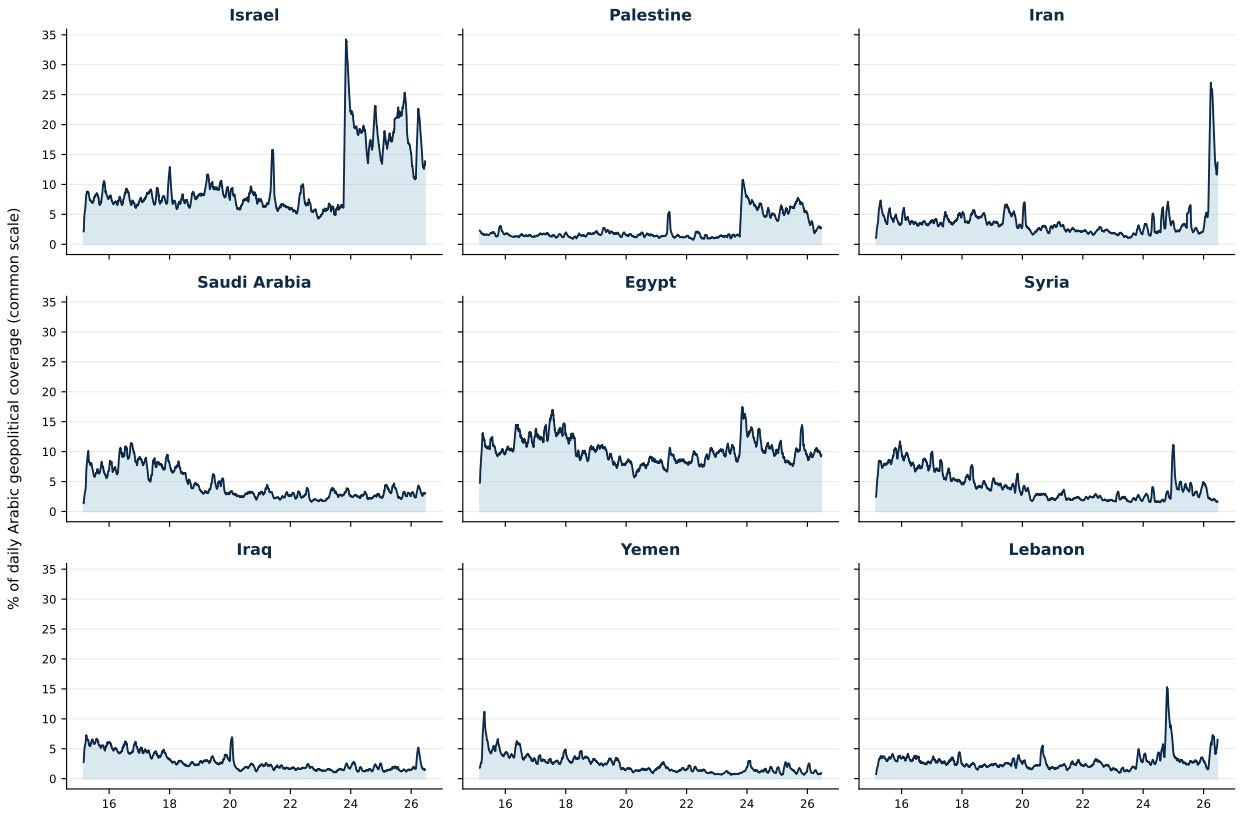

Figure 3: Country-level Arabic GPR by country, individual panels on a common vertical scale, 2015–2026. Each panel plots the 30-day moving average of the share of daily Arabic geopolitical coverage that mentions the country. Source: author's calculations using GDELT location tags.

The anatomy of risk: a decomposition by event type

A second decomposition asks not where risk is concentrated but what kind of risk it is. I assign each geopolitical article to one or more categories using its GDELT theme tags, wars and armed conflict, military and security activity, terrorism, protests and civil unrest, sanctions and blockades, nuclear and other weapons of mass destruction, coups, and assassinations, and compute each category's share of daily coverage. Figure 4 plots the leading categories over time, and Figure 5 shows how their relative composition shifts.

Wars and armed conflict are the dominant category throughout, averaging roughly a quarter of all geopolitical coverage, followed by military and security activity (about a fifth) and terrorism (about a seventh). The time profile is informative. Terrorism is relatively more prominent in the early sample (2015–2018), the period of the Islamic State's territorial peak, and recedes thereafter. War and armed conflict, by contrast, surges from late 2023 with the Gaza war and again in 2025 with the Israel–Iran confrontation, when it accounts for the

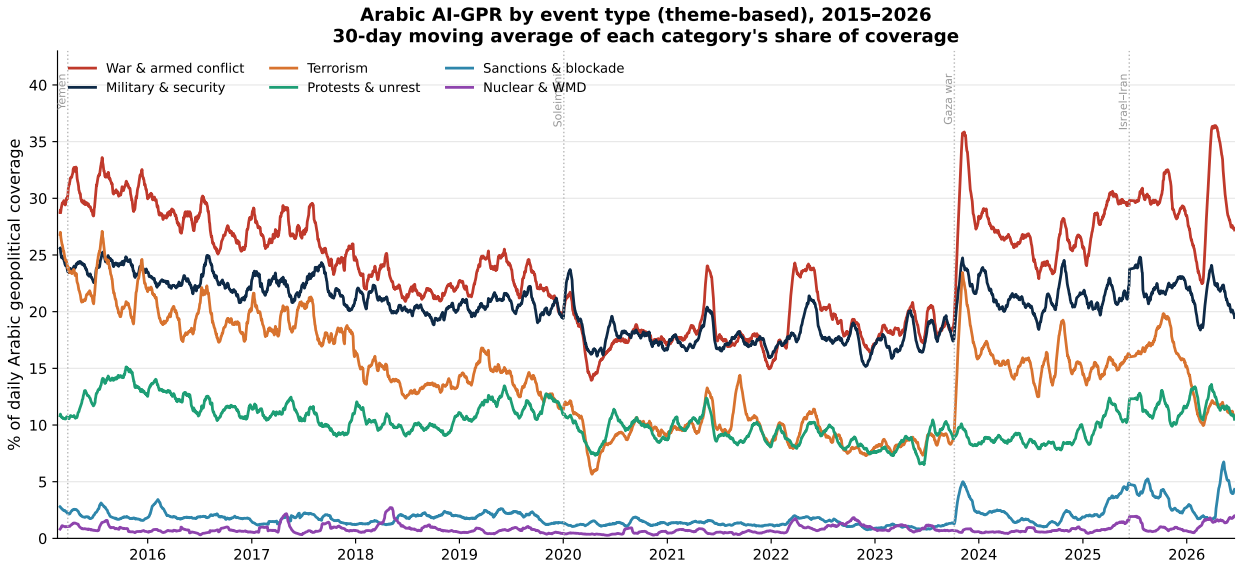

Figure 4: Arabic AI-GPR by event type (theme-based), 2015–2026. Each line is the 30-day moving average of the share of daily Arabic geopolitical coverage carrying that event-type theme. Dashed markers denote major events. Source: author's calculations using GDELT theme tags.

single largest share of coverage on record. Protests and civil unrest hold a steady baseline of around one article in ten, while sanctions and nuclear coverage remain small but spike around specific episodes such as the killing of Qassem Soleimani and recurrent disputes over Iran's nuclear program.

The composition view in Figure 5 makes the structural shift visible: the Arabic geopolitical conversation moves from a mix in which terrorism and unrest figure prominently toward one increasingly dominated by interstate war. This is the kind of qualitative change that a single aggregate index cannot reveal, and it is economically meaningful, since wars, sanctions, and terrorism transmit to oil prices, trade, and asset markets through different channels. As before, the categories are based on theme tags and are not mutually exclusive, so an article may contribute to more than one; the composition shares are normalized to sum to one, and the corresponding LLM-scored event-type indices are under construction.

The relational structure of risk: bilateral country pairs

A third decomposition turns from single countries to the pairs of countries that appear together. For each geopolitical article I extract the set of countries it mentions and record every unordered pair, so that a single report naming Israel, Iran, and the United States contributes to three pairs. Counting how often two countries are named together in the same

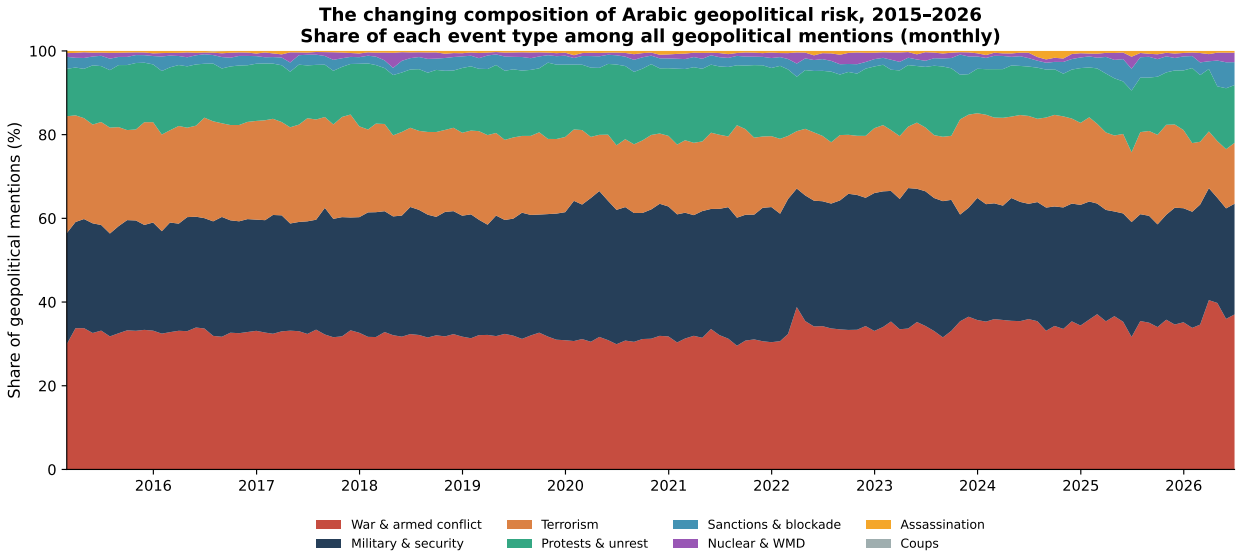

Figure 5: The changing composition of Arabic geopolitical risk, 2015–2026. Monthly share of each event type among all geopolitical theme mentions. Source: author's calculations using GDELT theme tags.

geopolitical article yields a relational map of the Arabic geopolitical conversation: a high co-mention count signals that two countries are, in the eyes of Arabic news, entangled in the same tensions. Normalizing each pair's daily count by total geopolitical coverage gives a bilateral tension share that is comparable over time.

Figure 6 traces seven salient pairs, and each tracks its conflict closely. The Saudi Arabia–Yemen series dominates the early sample, peaking with the 2015 Saudi-led intervention and fading as the war becomes protracted. Israel–Lebanon and Israel–Syria carry steady weight throughout, the former spiking sharply in late 2023 and 2024. The defining feature of the late sample is the Iran–Israel pair, which is comparatively quiet for most of the decade and then rises to the largest bilateral share on record in mid-2025, the period of direct Israel–Iran confrontation. The Russia–Ukraine series, by contrast, registers a clear jump in early 2022 and recedes, confirming that the Arabic press tracks extra-regional conflict but anchors its attention regionally.

Figure 7 summarizes the full relational structure. The United States and Israel are the most connected nodes, named alongside almost every regional actor, which reflects their role as recurring reference points in coverage of regional tension. Dense clusters appear among the Levant states (Syria, Iraq, Lebanon, Israel) and around the Gulf (Saudi Arabia, Yemen, Iran, Egypt), recovering the region's principal fault lines from the text alone. This pairwise view complements the country and event-type decompositions: it identifies not only how much risk attaches to each country and what kind it is, but which relationships drive it,

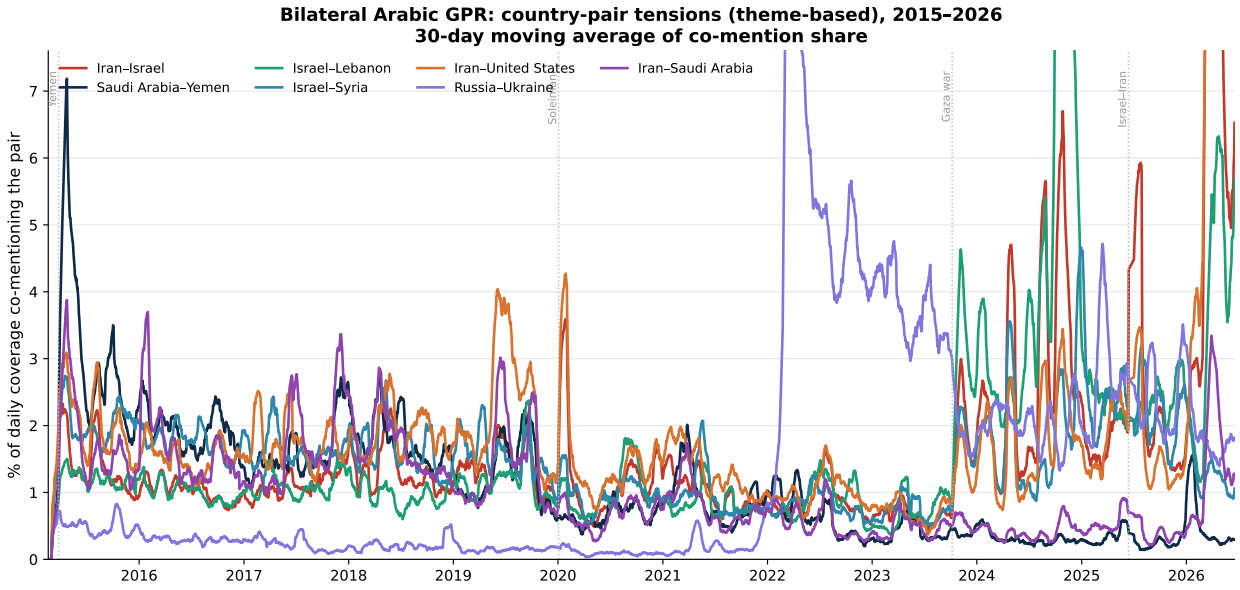

Figure 6: Bilateral Arabic AI-GPR: country-pair tensions (theme-based), 2015–2026. Each line is the 30-day moving average of the share of daily Arabic geopolitical coverage that co-mentions both members of the pair. Dashed markers denote major events. Source: author's calculations using GDELT location tags.

information directly relevant to studying contagion across oil markets, defense spending, and MENA asset prices. As with the other decompositions, pairs rest on co-mention rather than directed causal attribution, and the corresponding LLM-scored bilateral indices are under construction.

Risk by actor: armed groups and named leaders

A fourth decomposition attributes risk to the specific actors that drive it. Using the organization and person tags that GDELT extracts from each article, I identify mentions of the principal armed groups and political and military leaders in the region and compute, for each, the share of daily geopolitical coverage in which it appears. Figure 8 plots the leading armed groups and Figure 9 the most prominent named leaders.

The actor series are perhaps the sharpest validation of the index, because each peaks precisely at the event that made its subject salient. Among armed groups, Hamas jumps from near silence to the single most-covered actor with the October 2023 Gaza war; Hezbollah spikes in the autumn of 2024 during its open confrontation with Israel; the Houthis peak first in 2015 at the onset of the Yemen war 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 leaders, the killing of Qassem Soleimani in January 2020 appears as a clean

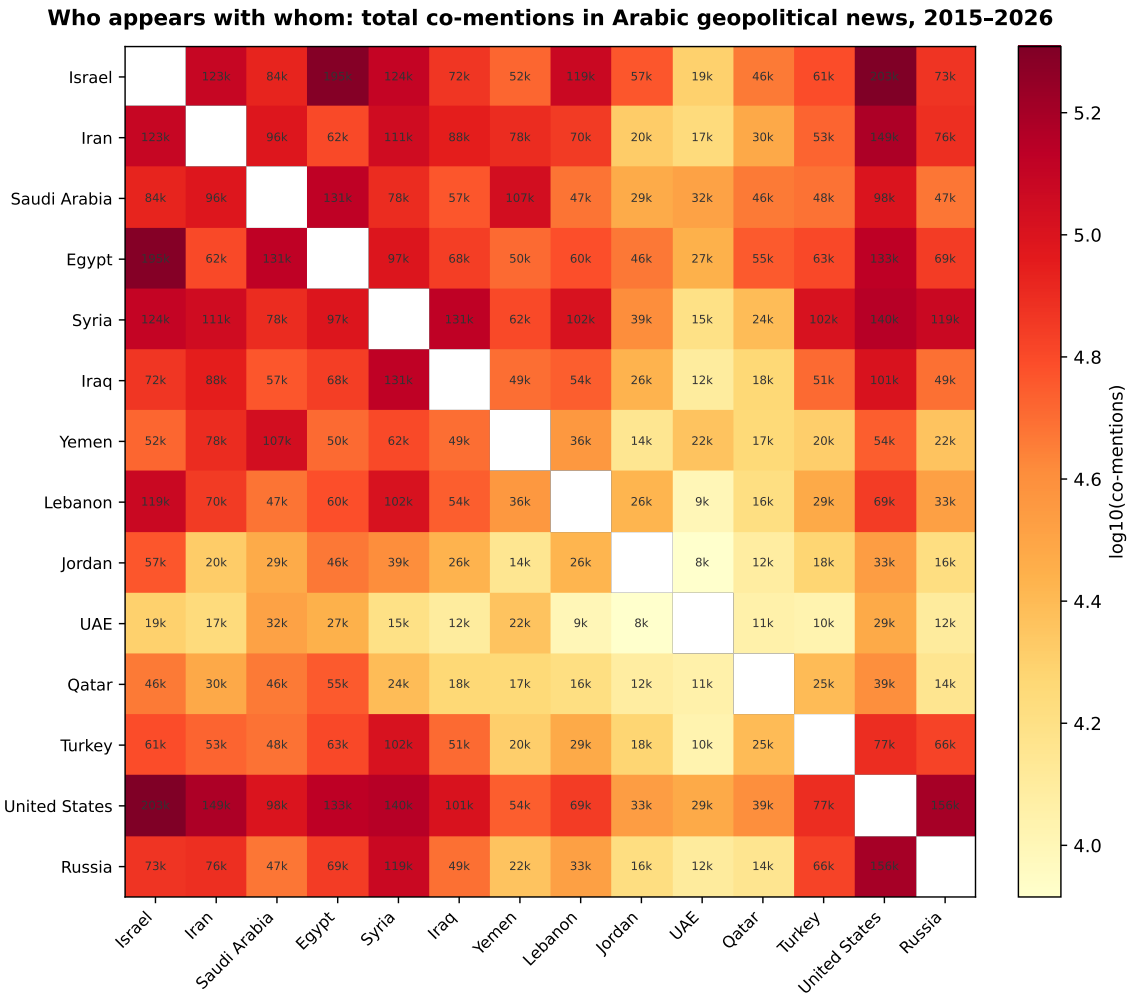

Figure 7: Who appears with whom: total co-mentions of each country pair in Arabic geopolitical news, 2015–2026 (log scale). Darker cells indicate pairs that are named together more often. Source: author's calculations using GDELT location tags.

spike, as does the death of Hassan Nasrallah in late 2024, while Benjamin Netanyahu rises to sustained prominence from late 2023 onward. That coverage of an actor surges exactly when that actor is central, and subsides otherwise, is direct evidence that the index tracks real events rather than noise.

This actor view is the natural complement to the country, event-type, and bilateral decompositions: it identifies the protagonists of regional risk and offers a foundation for actor-specific risk measures, for instance the share of tension attributable to non-state armed groups as against interstate actors. As with the other decompositions, the actor tags are drawn from GDELT entity extraction run on the machine-translated text, a single actor can appear in many articles, and coverage of the Islamic State is being incorporated in a forthcoming revision.

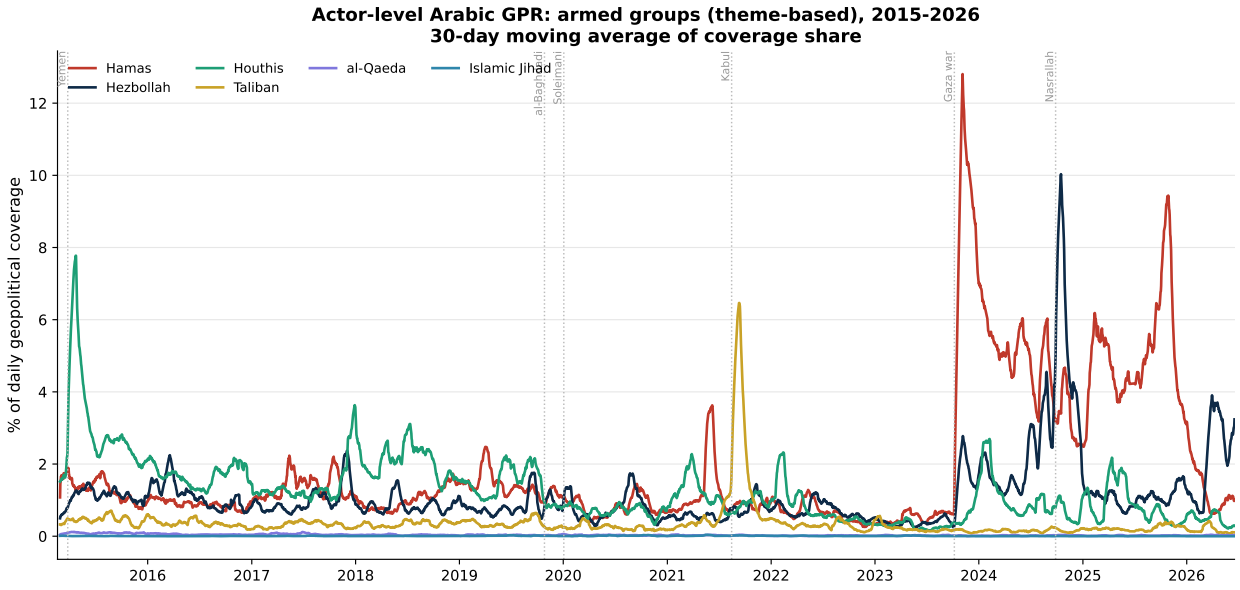

Figure 8: Actor-level Arabic AI-GPR: armed groups (theme-based), 2015–2026. Each line is the 30-day moving average of the share of daily Arabic geopolitical coverage that mentions the group. Dashed markers denote major events. Source: author's calculations using GDELT organization tags.

Oil and energy: a sector-specific risk index

Geopolitical risk matters for the world economy in large part through energy markets, and the Arab region sits at the centre of global oil and gas supply. I therefore construct an Arabic Oil and Energy GPR by restricting the index to geopolitical articles that also concern oil or energy, identified through GDELT energy theme tags. Figure 10 plots the resulting daily series, normalized to a 2016–2024 mean of 100, and Figure 11 decomposes it by the principal oil corridors of the region.

The series behaves exactly as an oil-risk measure should. It is quiescent through the mid-2010s, lifts with the September 2019 attack on the Abqaiq processing facility in Saudi Arabia, rises again with the killing of Qassem Soleimani in January 2020 and Russia's 2022 invasion of Ukraine, and reaches its all-time peak, more than five times its baseline, during the 2025 Israel–Iran confrontation and the associated threats to Gulf energy infrastructure and the Strait of Hormuz. The corridor decomposition in Figure 11 locates this risk geographically: 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 and the rerouting of traffic away from the Suez Canal.

This oil-energy index is the regional analogue of the oil-specific GPR in Caldara and

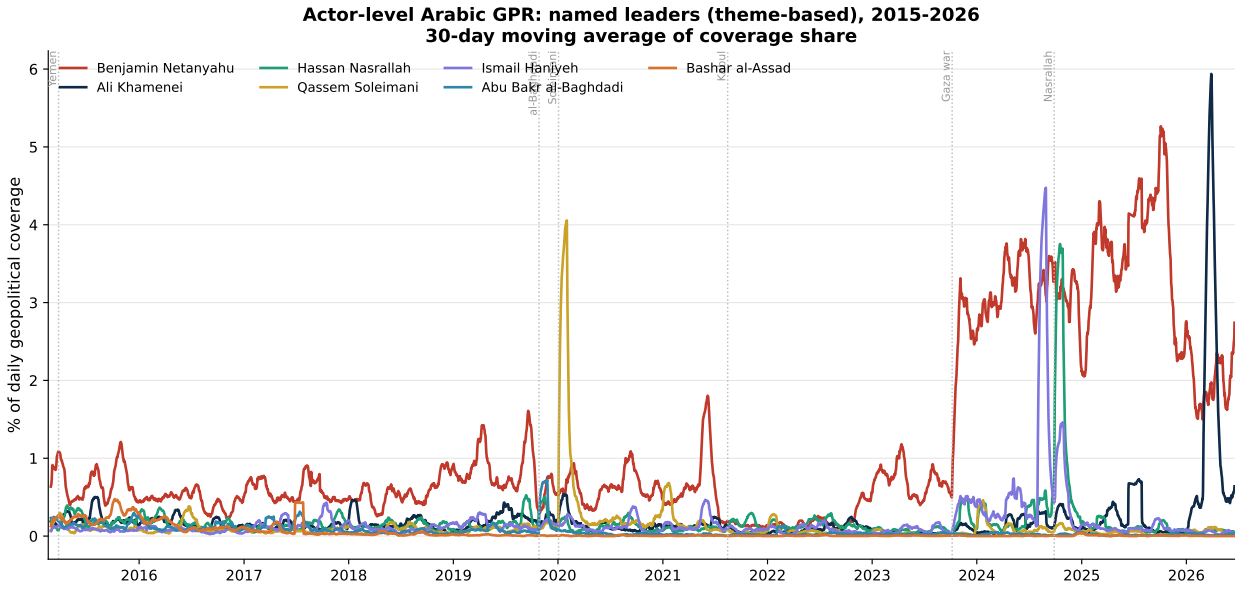

Figure 9: Actor-level Arabic AI-GPR: named leaders (theme-based), 2015–2026. Each line is the 30-day moving average of the share of daily Arabic geopolitical coverage that mentions the person. Dashed markers denote major events. Source: author's calculations using GDELT person tags.

Iacoviello (2022) and Iacoviello and Tong (2026), and it is the sub-index most directly suited to empirical work on energy markets. A natural validation, pursued in ongoing work, compares the series with the price of Brent crude: the episodes that move the Arabic Oil and Energy GPR coincide with the major oil-price spikes of the sample. As with the other decompositions, the index rests on the theme-based baseline and GDELT tags, and the corresponding LLM-scored oil index is under construction.

Validation against the oil price

To check the oil and energy index against market data, I compare it with the price of Brent crude (FRED series DCOILBRENTEU). Figure 12 overlays the two. They co-move clearly around supply-threatening episodes, most visibly the 2022 invasion of Ukraine and the 2025 Israel–Iran confrontation, but the continuous correlation between the two levels is modest (about 0.2). This is expected: the oil price level is governed largely by supply, demand, and macroeconomic conditions (the 2020 demand collapse being the clearest case), whereas the index measures the prominence of oil-related geopolitical tension.

The more revealing link is with oil-market volatility, the channel through which geopolitical risk is usually thought to transmit. Table 1 reports Brent's annualized realized volatility in the thirty days before and after each of the five largest Oil and Energy GPR spikes. In

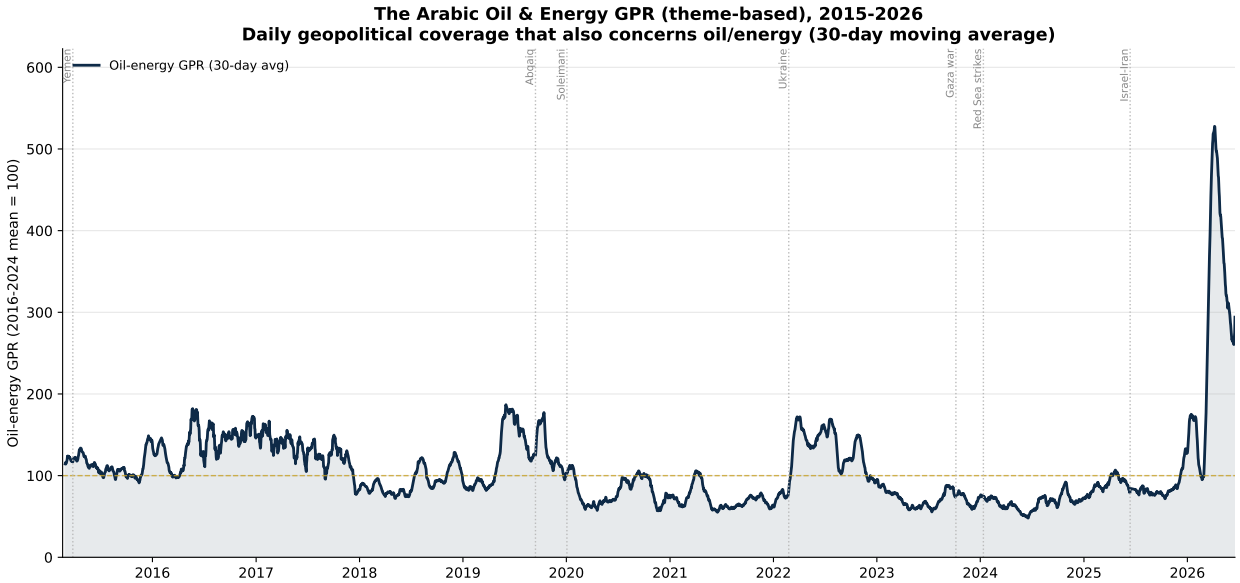

Figure 10: The Arabic Oil and Energy GPR (theme-based), daily, 2015–2026. The series is the share of daily Arabic geopolitical coverage that also concerns oil or energy (using GDELT energy theme tags), normalized to a 2016–2024 mean of 100; the line is a 30-day moving average. Dashed markers denote major events. Source: author's calculations using GDELT data.

every case volatility rises, in several cases sharply: the Ukraine invasion lifts it from 26 to 68 percent, the Gaza war from 21 to 44. The direction of the price move differs across episodes, rising when supply fear dominates and falling when demand destruction or a broader risk-off move dominates, which is precisely why a directional price correlation is weak; the increase in uncertainty, by contrast, is uniform. The index thus behaves as a geopolitical risk measure should, marking the episodes that inject uncertainty into the oil market rather than predicting the sign of the price change.

Episode	Brent 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

Table 1: Brent crude annualized realized volatility (percent) in the 30 calendar days before and after each major Oil and Energy GPR spike. Source: author's calculations using FRED series DCOILBRENTEU.

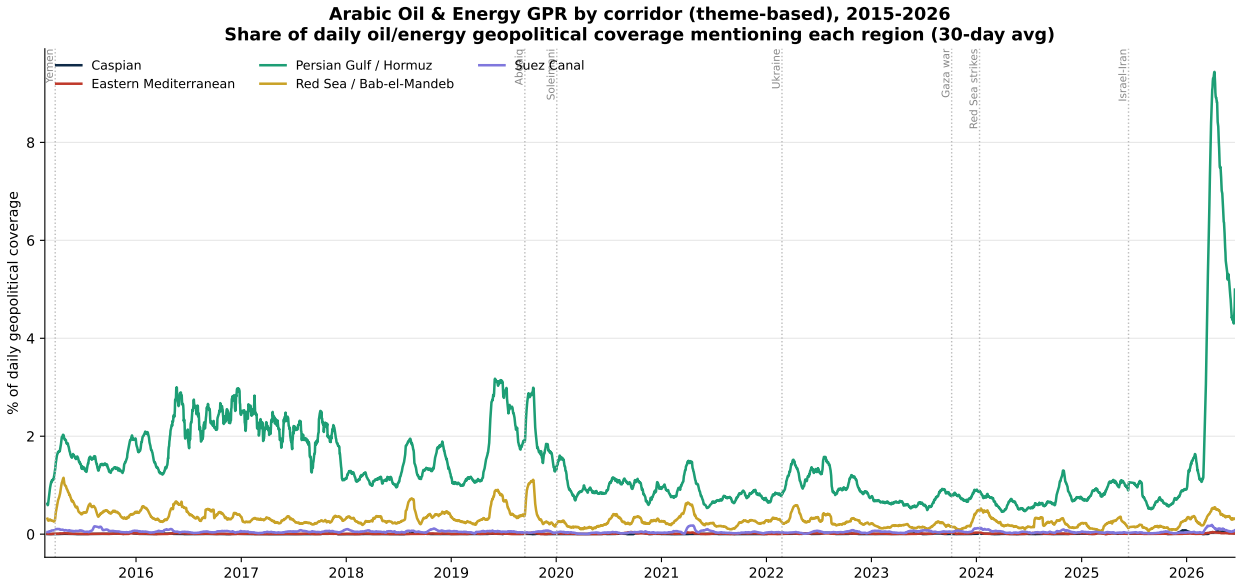

Figure 11: The Arabic Oil and Energy GPR by corridor (theme-based), 2015–2026. Each line is the 30-day moving average of the share of daily oil/energy geopolitical coverage that mentions the corridor, where the Persian Gulf corridor is keyed to the Gulf producers and the Strait of Hormuz. Source: author's calculations using GDELT location tags.

Threats and acts: anticipation versus realization

A further distinction, central to the original GPR framework of Caldara and Iacoviello (2022), separates geopolitical *threats*, the anticipation of adverse events, from geopolitical *acts*, their realization. I implement it at the article level: a geopolitical article counts as an act if its themes report realized violence (armed conflict, terrorism, sieges, assassinations, killings, active rebellion or war) and otherwise as a threat if it carries anticipatory or posture themes (explicit threats, military buildups, nuclear and other unconventional weapons, sanctions, blockades, ceasefire diplomacy, border tension). Acts take priority, so an article reporting actual fighting is classified as an act even when it also mentions a threat. Figure 13 plots the two as shares of daily coverage, and Figure 14 shows their ratio.

Acts dominate Arabic geopolitical coverage throughout, averaging about four times the volume of threats, which is consistent with a region in which conflict is frequently realized rather than merely threatened. The more informative object is the ratio in Figure 14, which captures how kinetic the news is at each moment. It reaches its sample maximum (about seven) in the weeks after the October 2023 Gaza war, the most intensely violent episode in the period, and remains elevated through the 2024–2025 escalations. It falls to its lows (about two and a half) in 2020–2021, a comparatively contained interval dominated by the pandemic, when anticipation and diplomacy occupied a larger share of the conversation. The

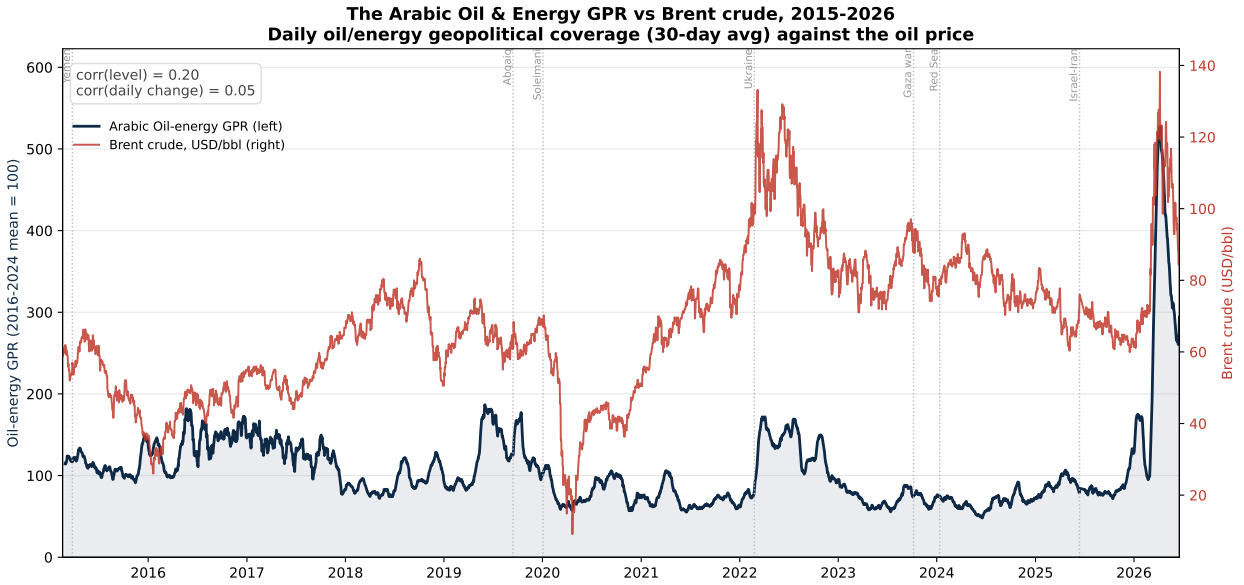

Figure 12: The Arabic Oil and Energy GPR (left axis, 2016–2024 mean = 100) against the price of Brent crude (right axis, USD per barrel), 2015–2026. Source: author's calculations and FRED series DCOILBRENTEU.

threats series, for its part, peaks in the summer of 2019, during the Gulf tanker crisis and the standoff over Iran's nuclear program, precisely the kind of high-anticipation, low-realization episode the threat measure is designed to capture. The decomposition thus separates two economically distinct phases of risk: anticipation, which raises uncertainty and risk premia ahead of events, and realization, which inflicts the damage. As elsewhere, the split rests on the theme-based baseline, an article may carry several themes, and the corresponding LLM-scored threat and act indices are under construction.

Zooming in: sub-national conflict hotspots

The decompositions so far operate at the level of countries, actors, and sectors. A final cut goes finer, to the individual cities and towns where conflict is actually fought. Using the sub-national place names that GDELT records in its location tags, I track the share of daily geopolitical coverage that names each of the principal conflict hotspots of the region. Figure 15 plots the leading hotspots over time and Figure 16 ranks them by total coverage.

The city-level view is the most granular validation of the index and the most vivid. Gaza is the single most-covered hotspot in the entire sample by a wide margin, and its coverage explodes from late 2023 to a peak of nearly thirty percent of all geopolitical coverage, the highest reading any single place attains. Aleppo peaks in December 2016 with the fall of the city, Sana'a in 2015 at the onset of the Yemen war, and Idlib through the late-stage Syrian

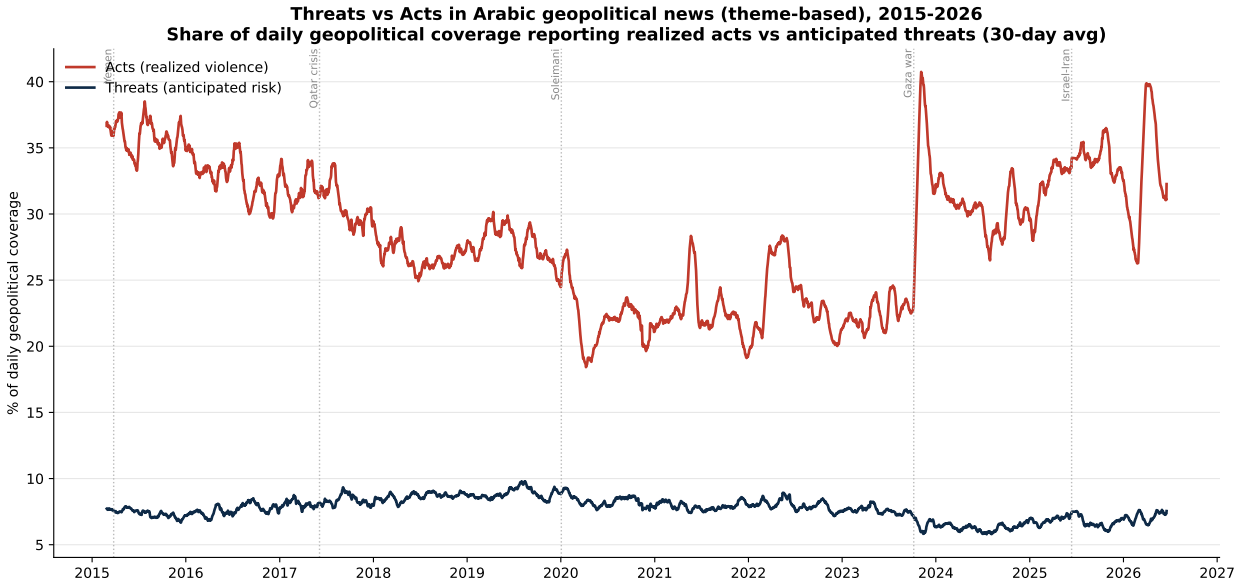

Figure 13: Threats versus acts in Arabic geopolitical news (theme-based), 2015–2026. Each line is the 30-day moving average of the share of daily geopolitical coverage classified as a realized act or an anticipated threat. Source: author's calculations using GDELT theme tags.

campaigns, while Damascus, Beirut, and Baghdad form a persistent Levantine and Iraqi baseline punctuated by their own crises. Each city's profile maps onto its known history with striking fidelity, which is reassuring evidence that the index localizes risk correctly even at the level of individual towns.

Two caveats apply. First, as elsewhere, these are theme-based shares from GDELT tags, and a single article can name several places. Second, sub-national detection is only as good as GDELT's geocoding: a number of important cities (Mosul, Homs, and Benghazi among them) are systematically resolved to their governorate or country rather than to the city name in this corpus, and so cannot be isolated reliably; they are therefore omitted here. The hotspots that do geocode cleanly nonetheless cover the central theatres of regional conflict and, together with the country, event-type, bilateral, actor, and oil decompositions, complete a multi-resolution picture of Arabic geopolitical risk, from the region as a whole down to the individual city.

5 Applications

Because it is regional, daily, and decomposable, the Arabic AI-GPR is suited to a range of empirical questions that an aggregate, English-language index addresses only coarsely. The most immediate is energy. The oil and energy sub-index, together with the Persian Gulf and

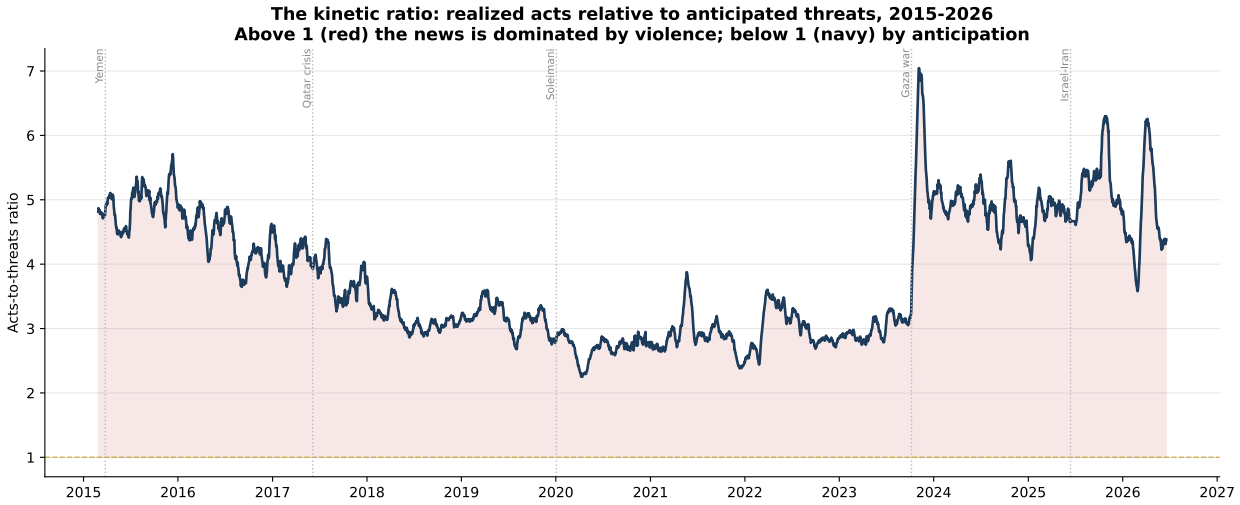

Figure 14: The kinetic ratio: realized acts relative to anticipated threats, 2015–2026 (30-day moving average). Higher values indicate coverage dominated by violence rather than anticipation. Source: author's calculations using GDELT theme tags.

Red Sea corridor series, is a natural high-frequency regressor for studying oil and gas prices, the volatility of energy futures, and the risk premium embedded in freight and insurance rates along the Hormuz and Bab-el-Mandeb routes. Because the index is measured at the source of the disruptions, it should lead or sharpen the explanatory power of distance-based measures built from Western newspapers.

A second application is regional asset pricing. Gulf equity markets, sovereign and quasi-sovereign debt, and currency regimes are exposed to exactly the tensions this index captures, yet they are studied with global risk proxies that miss regional episodes or time them imprecisely. The country and bilateral decompositions allow this exposure to be measured at the level of the individual market and the specific relationship driving it, which is useful for event studies around episodes such as the 2017 Gulf diplomatic crisis, the 2019 Abqaiq attack, or the 2023–2025 escalations. A third set of applications is macroeconomic: tourism and remittance flows, trade and supply chains, investment, and the conduct of monetary policy in the region are all sensitive to geopolitical tension, and a faithful regional measure provides a cleaner conditioning variable than a global one. Finally, the actor, event-type, and city decompositions support more structural questions, distinguishing for example the market response to interstate war from that to terrorism or sanctions, or tracing how risk attributable to a particular armed group is priced. In each case the contribution of the index is the same: it brings the measurement of geopolitical risk closer, in language, geography, and time, to where that risk originates.

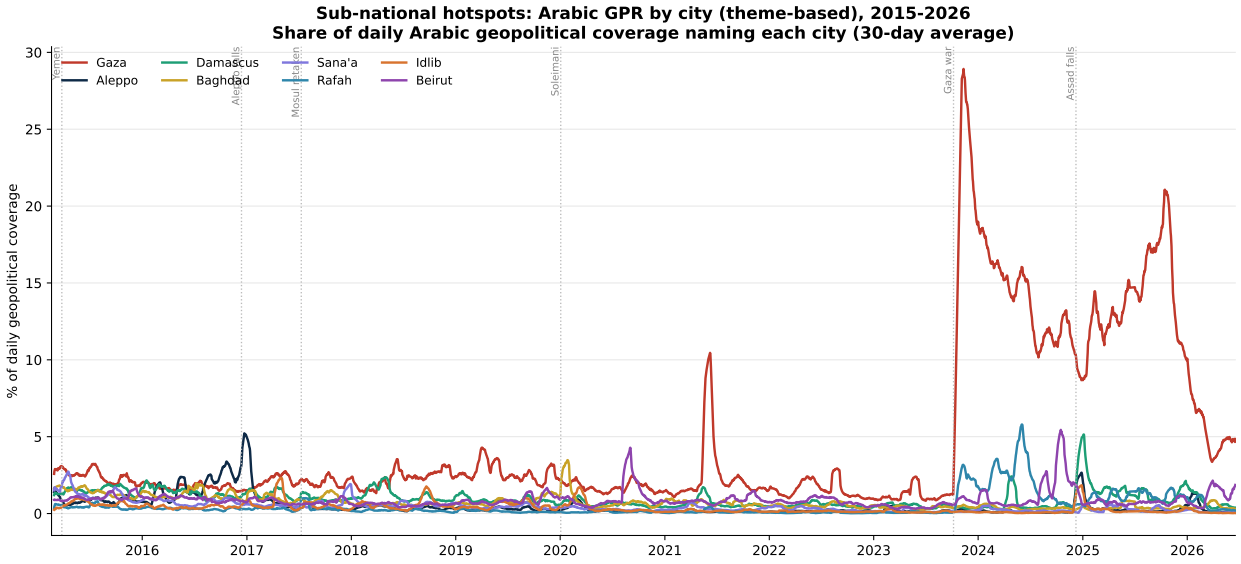

Figure 15: Sub-national hotspots: Arabic GPR by city (theme-based), 2015–2026. Each line is the 30-day moving average of the share of daily Arabic geopolitical coverage that names the city. Dashed markers denote major events. Source: author's calculations using GDELT sub-national location tags.

6 Conclusion

This paper has introduced the Arabic AI-GPR Index, the first measure of geopolitical risk built from Arabic-language news. Following the framework of Caldara and Iacoviello (2022) and the large-language-model approach of Iacoviello and Tong (2026), it assembles a curated corpus of Arabic reporting through GDELT and measures, at daily frequency from 2015 to 2026, how intensively that reporting concerns adverse geopolitical events. The resulting series displays strong face validity, rising around every major regional episode of the period, and it comes with a family of internally consistent decompositions, by country, by country pair, by event type, by actor, by oil corridor, and by city, that together render Arabic geopolitical risk at multiple resolutions, from the region as a whole down to the individual conflict town.

Two extensions are already underway. The first is the full large-language-model index, which will replace the theme-based share with an intensity-weighted reading of each article and is expected to sharpen the separation of front-page crises from incidental mentions. The second is a backward extension of the series across earlier decades, toward one of the longest available records of Arabic geopolitical risk. Alongside these, the index is released as a public resource, updated monthly and accompanied by an interactive interface to its sub-indices, so that researchers, analysts, and policymakers can study regional risk through the lens of the region's own news. By measuring geopolitical risk where it is reported first and in greatest

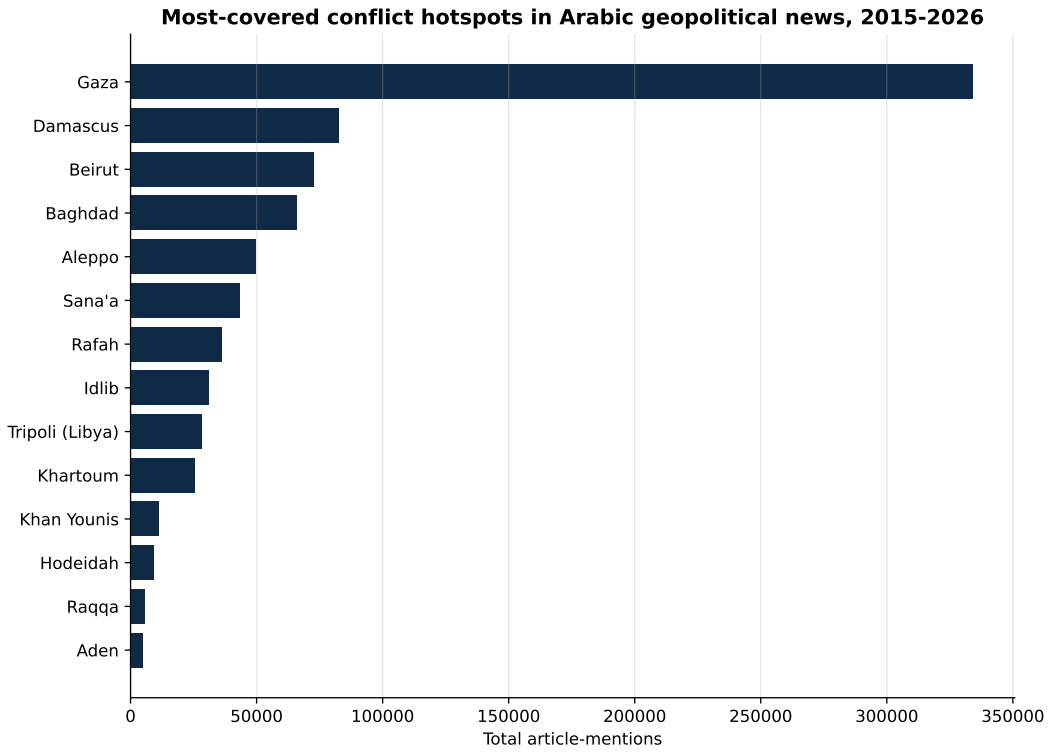

Figure 16: Most-covered conflict hotspots in Arabic geopolitical news, 2015–2026, by total article mentions. Source: author's calculations using GDELT sub-national location tags.

detail, the Arabic AI-GPR aims to be both a faithful regional barometer and a building block for the empirical study of how geopolitical shocks transmit to the economies of the Middle East and the wider world.

References

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics* 131(4), 1593–1636.

Caldara, D., and M. Iacoviello (2022). Measuring geopolitical risk. *American Economic Review* 112(4), 1194–1225.

Iacoviello, M., and J. Tong (2026). The AI-GPR Index: Measuring geopolitical risk using artificial intelligence. Working paper.

Leetaru, K., and P. A. Schrodt (2013). GDELT: Global Data on Events, Location, and Tone, 1979–2012. *International Studies Association Annual Convention*.