# INFORMAL EXCHANGE RATE AND SOCIAL UNREST:

## Detecting Instability in Cuba Through Media Event Data

Maria Barreiro Palafox, MS; Rakel Dassinger, MA;
Nemanja Jovanovic, EngD; Jose Pineda, PhD

September 2026

# INFORMAL EXCHANGE RATES AND SOCIAL UNREST: Detecting Instability in Cuba Through Media Event Data


## Maria Barreiro Palafox, MS; Rakel Dassinger, MA; Nemanja Jovanovic, EngD; Jose Pineda, PhD



## September 2026



## Abstract

This paper examines whether movements in Cuba's informal CUP/USD exchange rate carry information about near-term social-unrest-related media activity, at a moment when Cuban currency depreciation, inflation, and reports of unrest have again been rising in tandem. Because Cuba's official statistics are sparse, lagged, and often understate the pressures they aim to measure, the analysis instead draws on three open-source, high-frequency indicators: an informal exchange-rate series compiled by DevTech Systems, and two event-based measures built from the Global Database of Events, Language, and Tone (GDELT) and the Armed Conflict Location & Event Data project (ACLED). Using GDELT, we construct the Cuba Sociopolitical Instability Signals Index (CSISI), a daily instability signal built separately for Cuba's state and exile/independent media, and test its relationship to the informal exchange rate using weekly local projections that allow depreciation and appreciation shocks to have distinct effects, over 2021–2026. Exile media's headline CSISI shows a same-week depreciation response, but this result is challenged by a large number of hypothesis tests presented in the paper. A more targeted result shows more robustness: exile media's geographic-exposure dimension shows a depreciation response that is jointly significant across all response horizons, robust to lag length and event-definition choices, alternative definition of exchange rate depreciations, and independently corroborated by ACLED's field-coded data. State media separately shows a weaker but directionally consistent pattern. These results, though descriptive and reduced-form rather than causal, indicate that Cuba's informal exchange rate has practical value as a high-frequency, low-cost explanatory indicator of social unrest, and that open-source data can partially substitute for the official statistics Cuba does not reliably produce.

# TABLE OF CONTENTS

## TABLE OF FIGURES

# I. INTRODUCTION AND MOTIVATION

Cuba's economic and social conditions deteriorated sharply through the summer of 2026. Cubalex, an independent human-rights monitor, documented a record 52 homicides in August, 315 criminal incidents, five deaths in state custody, and 280 distinct repressive events spread across 54 municipalities in 14 provinces (CiberCuba, 2026a). That same month, the Central Bank of Cuba began issuing 10,000- and 20,000-peso banknotes (the two highest denominations the country has ever circulated) against 20.7% year-on-year official inflation and an informal exchange rate that had moved from roughly 60% (from 435 to 698 CUP/USD) between December 2025 and mid-September 2026 (CiberCuba, 2026b). By mid-September, economists were openly debating whether Cuba was sliding toward de facto dollarization as confidence in the peso continued to erode (CiberCuba, 2026c). Currency depreciation, rising prices, and social unrest again appear to be moving together, echoing the pattern around the 2021 Tarea Ordenamiento and the July 11, 2021 protests. Whether that co-movement is systematic and measurable, and whether it can be tracked as it happens rather than reconstructed months later, is the question this paper takes up.

Answering that question in Cuba is unusually difficult by conventional means. Official statistics are sparse, lagged, and in some cases understate the pressures they are meant to capture. Cuba's official CPI, for instance, runs well below household-level inflation estimates, and no official series tracks social unrest at all. State control over media and a limited independent press further restrict what is knowable about conditions on the ground in real time. This paper starting motivation is that open-source, high-frequency data generated as a byproduct of markets and media, rather than collected through official channels, can partially substitute for the statistics Cuba does not reliably produce, and can do so closer to real time.

Three such sources anchor the analysis. The first is DevTech's own compiled series on Cuba's informal CUP/USD exchange rate, a continuously quoted, market-based price that reflects foreign-exchange scarcity and expectations before those pressures show up in an official statistic. The second and third are event datasets: the Global Database of Events, Language, and Tone (GDELT), which underlies this paper's own instability index; and the Armed Conflict Location & Event Data project (ACLED), which codes political violence and demonstrations from field reporting.

GDELT merits closer attention because it does more than count events. It automatically extracts political and social interactions from global news texts and codes each one to a CAMEO event category, a Goldstein conflict-intensity score, an article-level tone score, and, where the source article permits it, a geographic location. Critically for this paper, GDELT's source universe spans both Cuba's state media and its exile/independent media, and each article can be attributed to one or the other. That makes it possible to build two parallel, identically-constructed instability signals from two different vantage points on the same events, rather than one pooled signal that would obscure how the two media environments diverge. The Cuba Sociopolitical Instability Signals Index (CSISI), described in full in Section 3, aggregates GDELT's event, severity, geographic, persistence, and tone information into a single daily score for each media universe, restricted throughout to the most acute events: protest, coercion, assault, fighting, and mass violence.

With CSISI and the informal exchange-rate series in hand, the remaining design choice was how to relate them. An earlier, daily-frequency approach using bivariate vector autoregressions between each CSISI series and the exchange rate found little that was usable: daily FX shows strong short-run mean reversion consistent with thin-market microstructure noise, and the one Granger-causal result that emerged showed the significant noise the daily data may have. This motivated both aggregating to a weekly frequency, at which the FX series behaves like white noise rather than a bouncing bid-ask spread, and adopting local projections rather than a VAR approach for the empirical analysis. Local projections estimate a separate regression at each forecast horizon instead of propagating one fixed lag structure forward, which lets the response to a depreciation shock differ in size, shape, and timing from the response to an appreciation shock (an asymmetry Section 6 shows to be central to the results) without imposing that asymmetry, or any particular decay pattern, by assumption.

This design contributes four things not yet available together elsewhere in the literature on exchange rates and unrest. First, it operates at weekly frequency, well below the monthly or annual frequency of most existing work connecting currency shocks to political instability, which matters if households and media respond to depreciation within days or weeks rather than after it appears in quarterly statistics. Second, the specific link tested here (between informal exchange-rate movements and social-unrest-related media activity, as distinct from formal-sector prices or official devaluations) has been the subject of very little direct empirical work, in Cuba or elsewhere. Third, the paper subjects every headline result to six independent robustness checks, including a correction for the 312 hypothesis tests the design generates, a formal joint test on the coefficients themselves, and independent corroboration from ACLED's field-coded data; a level of scrutiny uncommon in the descriptive event-data literature, and one this paper treats as necessary given how easily a spurious pattern can emerge from a large panel of media-derived series. Fourth, because both the informal exchange-rate series and CSISI update continuously, the approach validated here can, in principle, function as a live monitoring tool rather than a one-time historical exercise.

The rest of the paper builds toward that conclusion step by step. Section 2 reviews the literature connecting currency crises, economic hardship, and unrest, and situates Cuba's 2021 monetary reform and July 11 protests within it. Section 3 describes CSISI's construction in full, including why state and exile media are analyzed separately rather than pooled. Section 4 descriptives how CSISI and its components behave over 2020–2026. Section 5 sets out the empirical estimation of the weekly local-projection specification. Section 6 reports the main results, and Section 7 subjects them to the six robustness checks summarized. Section 8 concludes and presents areas of future work. Three annexes follow: one re-estimates every result under a broader event taxonomy, one under a revised exchange-rate construction, and one examines FX volatility, rather than direction, as a separate channel.

# II. LITERATURE REVIEW

Although the literature in economics and political links adverse economic shocks to protest, riots, civil conflict, and political instability, the specific relationship between exchange-rate movements and social unrest remains less developed. Depreciation can affect political stability through reductions of household purchasing power, shortages, increasing production costs, expectations, and perceptions of government competence. The literature identifies depreciation as both a high-frequency signal of foreign-exchange scarcity and macroeconomic stress, and a potential contributor to rising prices, shortages, and deteriorating expectations. Whether these pressures generate collective action depends on political institutions, social networks, information flows, repression, and household coping mechanisms, including remittances, migration, and informal markets.

Existing research establishes two well-developed but largely separate relationships: exchange-rate depreciation can affect domestic prices, real income, output, and macroeconomic stability, while deteriorating economic conditions can increase the likelihood of protest and political instability. Far less research directly tests whether exchange-rate movements themselves contain information about subsequent social unrest, particularly at high frequency. This study connects these two literatures.

Cuba's economy is characterized by administered prices and exchange controls with recurrent shortages, a dominant state sector. It also depends on imported food and energy, relies on tourism and remittances for foreign exchange, and has a prominent informal foreign exchange market. The 2021 monetary reform, or "Tarea Ordenamiento," generated a large official devaluation alongside with the continued peso depreciation in the informal market. Cuba in the year 2021 also experienced high inflation, falling real purchasing power, shortages, and the largest nationwide protests since the Revolution. These features do not establish causality per se but make Cuba particularly relevant for examining the transmission of currency pressures into social discontent.

## 2.1 EXCHANGE-RATE CRISES AND POLITICAL INSTABILITY

The direct literature on exchange-rate crises and political outcomes provides an important starting point. Frankel (2005) documents the political cost of currency crashes in developing economies, noting that leaders are substantially more likely to lose office following a currency crash. His broader argument is that devaluations can be contractionary rather than expansionary, particularly when balance-sheet effects associated with foreign-currency liabilities and financial stress outweigh conventional expenditure-switching effects. This literature is valuable because it treats a currency crash not merely as a nominal adjustment but as an event with real and political consequences. Political-economy work on exchange-rate regimes and currency crises similarly emphasizes that devaluation redistributes income and creates identifiable winners and losers. Frieden (2015) stresses the distributional foundations of exchange-rate politics, while Walter (2009) shows that electoral incentives shape how governments manage currency crises: policymakers may delay or alter adjustment around elections, although severe speculative pressure limits their ability to do so. These contributions suggest that exchange-rate adjustments have political consequences partly because they alter real incomes unevenly and because the public may attribute the adjustment, or the failure to prevent it, to government choices.

For Cuba, the distributional dimension is unusually important. Access to foreign currency is highly unequal: households receiving remittances, workers connected to tourism or private activities, and firms with access to foreign exchange are plausibly better positioned to protect themselves from peso depreciation than households dependent on fixed peso wages and pensions. LeoGrande and Hershberg (2023) document the centrality of foreign-exchange earnings, tourism, and remittances to Cuban household welfare during this period, though the specific three-way distributional comparison drawn here is our own synthesis rather than a claim they make directly. A depreciation can therefore increase not only average hardship but also inequality between groups with and without access to foreign currency. This distinction provides a bridge between the currency-crisis literature and the Cuban literature on monetary duality, segmentation, and the social consequences of reform.

## 2.2 MAIN TRANSMISSION CHANNELS

**Exchange-rate pass-through, food prices, and the cost of living.** The first and most immediate channel operates through exchange-rate pass-through. A depreciation increases the domestic-currency cost of imported final goods and imported intermediate inputs. For developing and emerging economies, Aron, Macdonald, and Muellbauer (2014) synthesize a large literature showing that pass-through is often economically meaningful and highly heterogeneous across countries and sectors. The extent and speed of pass-through vary with monetary credibility, market structure, inflation history, the currency of invoicing, administered prices, and the degree of exchange-rate persistence.

The political relevance of pass-through follows from the separate literature linking food-price increases to unrest. Arezki and Brückner (2011), using a panel of more than 120 countries, find that increases in international food prices raise anti-government demonstrations, riots, and civil conflict in low-income countries. Bellemare (2014) likewise finds that increases in food-price levels raise social unrest, but finds no evidence that food-price volatility itself increases unrest. Weinberg and Bakker (2015) emphasize that domestic institutions and policy responses condition the relationship between food prices and unrest. Although a majority of studies reviewed by Ismail (2021) find an association between food-price increases and protest or unrest, results remain heterogeneous, partly because food insecurity and conflict may be jointly determined.

This channel is particularly compelling in Cuba because food and basic-goods scarcity can make measured consumer prices only a partial indicator of the household burden. When a good is unavailable in a subsidized or administered-price channel, households may have to purchase it in a higher-priced private or informal market, wait in long queues, substitute toward lower-quality goods, or acquire foreign currency to access dollar- or MLC-linked markets (LeoGrande & Hershberg, 2023). The welfare effect of depreciation can therefore appear not only as a higher observed price but also as a deterioration in availability and effective access. A study focused solely on the official CPI could understate the mechanism linking currency pressure to household stress (Luis, 2023; Vidal & Luis, 2024).

**Real wages, social security, and relative deprivation.** A second channel operates through real household income. If nominal wages, social security, and transfers adjust more slowly than prices, depreciation-induced inflation reduces purchasing power. This effect is typically more severe for households that consume a high share of their income on food, transport, and other tradable or import-intensive goods. The unrest literature often interprets such shocks through grievance or relative-deprivation mechanisms: collective action becomes more likely when households experience a sharp deterioration relative to recent living standards or expectations, especially when the deterioration appears unfairly distributed.

The Cuban monetary reform makes this channel especially salient. Wage and pension increases were introduced as part of the 2021 reform, but the subsequent price acceleration greatly complicated the intended realignment of relative prices and incomes. Vidal and Luis (2024) estimate inflation substantially above the official CPI measure and identify the large devaluation associated with monetary reform as a central determinant of the inflation surge. Mesa-Lago (2021) similarly warned that if wages and pensions failed to keep pace with prices, purchasing power would fall and social tensions could increase. The implication is that exchange-rate depreciation should be interpreted jointly with measures of real wages, food prices, pensions, or household purchasing power whenever such measures are available.

**Shortages, rationing, and foreign-exchange scarcity.** A third channel is especially important in economies with controls and quantitative restrictions. Exchange-rate pressure may reflect an underlying shortage of foreign exchange that constrains imports and production (LeoGrande and Hershberg, 2023). When prices are administratively restricted, the adjustment may occur through quantities rather than prices: empty shelves, power outages, reduced transportation, long queues, or the expansion of informal markets. In this setting, the exchange rate can be informative even when official prices do not move proportionally, because the parallel exchange rate capitalizes information about the scarcity value of foreign currency.

This is central to the Cuban experience. Foreign-exchange scarcity affects the ability of state enterprises and private actors to obtain imported food, fuel, medicines, intermediate inputs, and consumer goods (LeoGrande and Hershberg, 2023). A widening gap between the official and informal exchange rates can therefore be interpreted as a signal that the administratively set price of foreign currency is not clearing the market (Luis, 2023). The social consequence may be observed as shortages and unequal access rather than only measured inflation. This mechanism suggests that indicators of scarcity (power outages, queue intensity, food availability, fuel availability, or mentions of shortages in media data) could help identify the pathway from depreciation to unrest. This connection is reinforced by direct evidence that blackouts themselves have triggered localized protest: LeoGrande and Hershberg (2023) describe outages lasting up to twelve hours followed by protests in several dozen towns and cities, with additional protests later over the slow restoration of electricity and water.

**Output, employment, and contractionary effects.** A fourth channel operates through real economic activity. Currency depreciation is traditionally expected to support net exports, but currency crashes in developing economies can be contractionary when firms rely on imported inputs, foreign-currency liabilities are important, financial conditions tighten, or the income effect on households dominates expenditure switching. Frankel (2005) emphasizes these contractionary mechanisms. Miguel, Satyanath, and Sergenti (2004) provide influential causal evidence from sub-Saharan Africa that negative growth shocks increase civil-conflict risk, demonstrating one setting in which deteriorating economic conditions translate into political violence.

In Cuba, the external sector creates several reasons why this channel may differ from a conventional open-market economy. Tourism receipts, remittances, and external financing affect the availability of foreign exchange (LeoGrande & Hershberg, 2023); state control influences import allocation; and firms may face administrative rather than purely market-based constraints. A depreciation in the informal market can therefore be simultaneously a symptom of declining foreign exchange inflows and a contributor to higher costs. This makes identification challenging: a fall in tourism, for example, can depreciate the peso, reduce incomes, worsen shortages, and increase discontent at the same time. Empirical work must therefore distinguish exchange-rate innovations from common underlying shocks to the extent possible.

**Fiscal and policy-adjustment channel.** A fifth channel runs through fiscal stress and policy adjustment. Depreciation can raise the domestic cost of imported goods, energy, and foreign-currency obligations, while inflation and shortages can increase the fiscal cost of subsidies or require changes in administered prices. Governments may respond by reducing subsidies, cutting expenditures, increasing tariffs or fees, or monetizing deficits. Ponticelli and Voth (2020) show that fiscal retrenchment, especially expenditure cuts, is associated with greater social unrest in a long historical panel of European countries. Although Cuba's institutional setting is different, the general mechanism is relevant: the political impact of a currency shock may be magnified when adjustment is transmitted to households through changes in subsidies, public services, remittances, administered prices, or rationed benefits. This distinction is analytically important because fiscal variables can be mediators rather than simple controls.

**Expectations, confidence, and the exchange rate as a public signal.** A sixth channel is less directly captured in standard macroeconomic models but may be especially important in a high-frequency study. The exchange rate is an asset price that incorporates expectations about future inflation, foreign-exchange availability, policy credibility, and political risk. In economies where a parallel exchange rate is widely followed, a sharp depreciation can immediately change expectations even before the full pass-through to consumer prices occurs. Households may accelerate purchases, seek dollars, postpone peso-denominated saving, or interpret the depreciation as evidence that economic conditions will worsen. For Cuba, the informal exchange rate published and discussed online plausibly serves a similar information function, though this specific characterization is our own extension of the logic rather than a claim made directly in the Cuban literature. It is observed at a much higher frequency than official macroeconomic statistics and may provide a common focal signal to households and businesses: Barreiro Palafox and

Dassinger (2025) document exactly this kind of high-frequency responsiveness, showing that informal exchange-rate volatility shifted within days of major blackout episodes, well before such effects could appear in monthly price statistics. Consequently, an empirical relationship between depreciation and unrest at very short horizons, days rather than months, would be difficult to explain only through realized CPI inflation. Such a finding would instead be consistent with expectations, confidence, coordination, or news channels. This is one reason why a daily-frequency design can add value to the existing literature rather than merely replicating monthly or annual studies.

## 2.3 THE CUBAN EXPERIENCE: WHY THE CASE IS DISTINCTIVE

Cuba entered the 2020s with a monetary and exchange-rate system shaped by decades of duality and administrative pricing. The coexistence of multiple currencies and exchange rates created large distortions in relative prices, enterprise accounting, and incentives. The Tarea Ordenamiento sought to simplify this system by eliminating the convertible peso and establishing a unified official exchange rate. Mesa-Lago (2021) describes the reform as necessary but warns that implementation during a severe economic crisis exposed the economy to substantial short-run inflation and purchasing-power risks. Vidal and Luis (2024) provide one of the most important empirical analyses of this episode. They argue that the large devaluation associated with the 2021 reform stands out as the main determinant of the subsequent inflation surge. Their estimates imply inflation far above the official CPI rate, reflecting in part the limitations of an index heavily influenced by state-sector prices. They also emphasize the continuing depreciation of the informal exchange rate and the persistence of fiscal and monetary imbalances after the reform. These findings establish a plausible first stage relationship from exchange-rate adjustment to household economic stress.

The Cuban parallel market deserves separate treatment because it is not merely an unofficial substitute for a missing market price. It contains information about foreign exchange scarcity, expected inflation, remittance flows, travel and migration demand, private-sector import needs, and confidence in the peso. The Cuban parallel market operates in a decentralized fashion across the island and online, with market indicators commonly derived from observed transactions. Luis (2023) shows that reliance on the official CPI understates actual inflation relative to a household-consumption-based deflator, underscoring why price signals observed outside the official statistical system (including the parallel market) carry independent informational value. The large divergence between official and informal rates after the reform indicates that an administratively fixed official rate cannot by itself summarize the monetary conditions facing households and private actors.

This distinction has direct implications for measurement. If the research question is whether "exchange-rate shocks" affect unrest, the informal rate may be more relevant than the official rate for high-frequency analysis because it is market-responsive and visible. At the same time, it is endogenous to political news, migration expectations, policy announcements, and unrest itself. The appropriate interpretation is therefore as a jointly determined market signal whose innovations may contain both economic and political information. This reinforces the need for dynamic methods, alternative identification assumptions, and tests of reverse causality.

The nationwide protests of July 11, 2021 provide one of the most important recent episode for understanding the potential political transmission of economic shocks in Cuba. The literature does not attribute the protests to a single cause. LeoGrande (2021) identifies economic desperation generated by the collapse in tourism, fuel shortages, U.S. sanctions affecting remittances, the COVID-19 emergency, and inflation following the currency reform as proximate drivers, while also emphasizing deeper political dissatisfaction. García Hernández (2023) likewise stresses the conjunction of shortages, economic crisis, political fatigue, the pandemic, tourism weakness, and the new role of social media. The lesson is that macroeconomic stress and political mobilization were complementary rather than competing explanations. Economic deterioration can increase grievances, but grievances do not automatically become coordinated protest. Internet access and social media reduced information barriers and helped citizens observe events elsewhere on the island in real time. This creates a two-stage conceptual structure for the Cuban case: economic shocks affect the latent demand for protest, while communication networks and political conditions affect whether that discontent becomes collective action.

## 2.4 FRAMEWORK & METHODOLOGY FOR THE CUBAN CASE

The literature presented in this paper suggests a framework in which exchange-rate depreciation affects social unrest through several overlapping mechanisms. The central economic channels are (i) pass-through to food, fuel, and other tradable prices; (ii) erosion of real wages and pensions; (iii) worsening shortages and access to goods; (iv) output and employment effects; and (v) fiscal or administered-price adjustment. An additional expectations channel operates because the informal exchange rate provides a rapid and visible signal of future economic conditions. These effects interact with inequality in access to foreign currency, political grievances, communication networks, and state capacity. The framework also implies important nonlinearities. A given depreciation should have a larger social effect when inflation is already high, inventories are low, power shortages are severe, or households have limited savings and remittance access. Likewise, the same economic shock may generate more visible protest when communication networks facilitate coordination or when confidence in government performance is already weak. This suggests that interaction terms or regime-dependent impulse responses may eventually be useful extensions, even if the initial empirical specification remains parsimonious.

### 2.4.1 Event Data and the Measurement of Sociopolitical Instability

Event datasets connect the literature on economic hardship and political mobilization to observable patterns of unrest. The Global Database of Events, Language, and Tone (GDELT) automatically extracts political interactions from news reporting, recording event categories, locations, and associated media tone (GDELT Project, 2015). The Armed Conflict Location & Event Data project (ACLED) uses researcherled coding to document political violence and demonstrations from media and other reporting sources (ACLED, n.d.). Both databases have been examined in empirical research. Hammond and Weidmann (2014) find that GDELT's correlation with human-coded event data is weak at the spatially disaggregated level ($r \approx 0.20$–$0.26$) but rises substantially once collapsed to a national time series ($r \approx 0.33$–$0.64$), reflecting substantial over-reporting of events near capitals and under-reporting in remote areas; while Raleigh, Kishi, and Linke (2023) demonstrate how differences in dataset definitions, sources, and coding procedures change the apparent distribution of political instability.

A related measurement tradition assigns numerical weights to political actions. Goldstein (1992) developed a conflict–cooperation scale for aggregating WEIS events into continuous measures of political interaction. GDELT incorporates a corresponding scale for CAMEO event categories, ranging from −10 to +10; these scores describe the character of an action rather than its observed casualties or participation (GDELT Project, 2015). This approach supports our multidimensional instability index for Cuba combining event frequency, geographic spread, recurrence, conflict intensity, and media tone (presented in the next section), with ACLED providing a complementary event-based comparison (presented in our results section). The composite-indicator literature guides our index construction by emphasizing the importance of transparent normalization, weighting, and sensitivity analysis (OECD et al., 2008).

# III. THE CUBA SOCIOPOLITICAL INSTABILITY SIGNALS INDEX (CSISI): CONSTRUCTION AND DATA

The Cuba Sociopolitical Instability Signals Index (CSISI) is explicitly a media-derived signal index, not a direct measure of latent instability: it quantifies how state and exile/independent media report acute conflict events, separately for each source universe, not the events themselves. Two parallel scores are produced every day (one from the frozen state-media source list, one from the frozen exile/independent-media source list) so that the two media universes can be compared directly rather than pooled.

CSISI is a monitoring and early-warning instrument, not a census of what happened on the ground. A rise in the index can reflect a genuine increase in protest or repression, greater reporting access, a shift in what outlets choose to cover, or a change in what GDELT itself ingested and coded that day; the index alone cannot distinguish among these, which is precisely why Section 3.5 tracks data coverage directly rather than leaving it implicit. CSISI does not estimate the probability of regime change, does not measure public opinion, and cannot capture discontent that never surfaces as a reportable media event.

## 3.1 DATA SOURCE

The underlying event data comes from GDELT's daily Events export (no GKG or Mentions tables are used). Each row is a coded event with a CAMEO event code, a Goldstein conflict-intensity score, an article-level average tone score, a source URL, and (where resolvable) a Cuban first-level administrative geography. Articles are attributed to a media outlet by parsing the domain of the source URL and matching it against a frozen registry of Cuban state and exile/independent domains; any article from a domain not on that registry is excluded, so the index reflects only the specific, named source lists, not the full GDELT stream about Cuba. Large content aggregators (e.g., generic news portals) are deliberately excluded from the registry as well, so that an aggregator's growing GDELT footprint in later years cannot masquerade as a genuine change in Cuban news volume.

## 3.2 THE FIVE DIMENSIONS

**Frequency:** measures how often qualifying events appear relative to that day's reporting: unique acute events as a share of that day's qualifying articles (×100), log-transformed before normalization since spike days are heavily right-skewed. A day with 10 acute events and 10 qualifying articles is a very different reporting environment from a day with 10 acute events and 100 articles, only the first scores high here. When frequency alone rises, without a matching move in severity, the most natural reading is "more stories about acute events," not necessarily a shift toward more conflictual ones.

**Severity:** measures how conflictual the covered event types are: the mean conflict intensity (max(0, −Goldstein score)) across that day's unique acute events, weighted per event rather than per mention so that repeated coverage of one event cannot inflate the score. Goldstein scores are keyed to the type of action, not to how many people were involved or how much damage resulted. A protest of ten and a protest of ten thousand can carry the same score, so this dimension is event-type severity, not casualties or crowd size. A severity-only rise, without a matching move in frequency, points to a shift in the mix of events toward more conflictual types (coercion, assault, fighting), not simply more stories.

**Geography:** measures how widely covered events are spread across the island: 0.4×province breadth + 0.4×population exposure + 0.2×territory exposure, summed over the provinces where that day's acute events were geolocated (16 first-level administrative units; population and area from the project's own province lookup table). Havana dominates GDELT's Cuban geocoding, and a country-only geocode ("Cuba," no province) does not count as spread, so a Havana-only day can still be politically central while scoring low on this one dimension specifically.

**Persistence:** measures whether today's coverage looks like a continuing episode rather than an isolated spike: how much that day's pattern of (event type, province) combinations overlaps with the same pattern in each of the trailing 7 days, weighted by an exponential decay ($\lambda$=0.80) that favors more recent days. A one-day story scores low here even if it was severe; a recurring pattern of the same event types in the same places across a week scores high, whether or not any single day looked dramatic.

**Tone:** measures how unusually negative that day's acute-event coverage reads for the outlets involved: for outlets that covered an acute event that day, how far their tone fell below that specific outlet's own historical baseline tone (from the calibration window), scaled by that outlet's own tone variability and weighted by article volume, so a single outlet's idiosyncratic writing style does not drive the score. If there are no acute-event articles that day, tone is left missing rather than set to zero, and that day's CSISI is computed as a four-dimension provisional average instead (Section 3.5).

## 3.3 AGGREGATION AND NORMALIZATION

Each of the five raw dimensions is separately normalized to a 0-1 scale using the 2.5th and 97.5th percentile of that dimension's own distribution during a fixed 2020-2023 calibration window (chosen to include both quiet years and the July 2021 protests). Those percentile anchors are then frozen and applied unchanged to the entire 2020-2026 sample (including years after the calibration window) so a value of 1.0 means “at or beyond the most extreme 2.5% of calibration-period days,” not an ever-shifting relative ranking. The headline CSISI score is the equal-weight arithmetic mean of the five normalized dimensions, ×100 (values above 100 are possible if a dimension exceeds its calibration-period extreme out of sample).

Percentile anchors, rather than the sample's literal minimum and maximum, are used deliberately. Anchoring to the single most extreme day in the whole file would let one exceptional event, the July 2021 protests, squash every ordinary day toward zero and hide month-to-month variation; it would also mean that every new record day silently rewrote the scale for all earlier history, and a one-off data-ingestion anomaly could become a permanent ceiling. A frozen calibration-period percentile avoids all three problems: this paper's figures never claim “this is the worst day imaginable,” only “this is high relative to a fixed historical window.”

In practical terms: 0 means as quiet as a very quiet day for that specific media environment during 2020-2023; 50 is roughly typical; 100 means as loud as that environment's rare high-signal days in the same historical window (its top 2.5 percent). Values rarely approach 100, since the ceiling is already close to the highest the series reached during calibration. Read loosely: roughly 0-50 is typical for that series, 50-75 elevated, 75-90 high, and above 90 very high or exceptional. Labels that describe rarity within that specific outlet's own history, not an externally validated probability of unrest.

## 3.4 ACUTE-EVENT TAXONOMIES

Which events count as “acute” is controlled by a taxonomy over CAMEO event-root codes. Three are defined: broad (CAMEO roots 10-20: all verbal and material conflict), tight (roots 13-20: threaten through mass violence), and unrest (roots 14 and 17-20 only: protest, coerce, assault, fight, and mass violence. The narrowest and most acute definition, dropping tight's “threaten” and “reduce relations” codes). This paper uses the unrest taxonomy throughout as the primary specification; the tight taxonomy is reported only as a robustness check, in Annex A. The underlying logic is a robustness ladder: a pattern that survives narrowing from broad to unrest is unlikely to be an artifact of how liberally "acute" was defined, whereas a pattern visible only under the broad taxonomy may reflect routine verbal conflict (demands, disapproval, diplomatic criticism) rather than street-level unrest.

## 3.5 DATA COVERAGE AND CONFIDENCE FLAGS

CSISI can fall for two very different reasons: because there was genuinely less instability in the news, or because there was simply less news. The constructor tracks this distinction directly rather than leaving it implicit. A day is flagged low-confidence when a media universe published fewer than 5 qualifying Cuba articles or drew on fewer than 2 distinct outlets that day; a day is flagged provisional when qualifying articles existed but none contained an acute-event article from which to score tone, so that day's CSISI averages four dimensions instead of five. A typical day in either media universe carries on the order of 19-20 qualifying articles across roughly 5 active outlets, with state-media volume notably thinner in 2025-2026 than earlier in the sample. A late-sample dip in state CSISI should be read alongside this coverage caveat; a quieter feed is not the same claim as a quieter country.

## 3.6 COMPARING STATE AND EXILE MEDIA: SAME RECIPE, DIFFERENT CAMERAS

CSISI runs the identical five-dimension recipe on two separate, frozen source universes: 19 state-media domains (Granma, Prensa Latina, Cubadebate, official radio, provincial party papers, and allied outlets) and 14 exile/independent domains (14ymedio, Diario de Cuba, CubaNet, Radio/TV Martí, El Toque, and others). Because each dimension's percentile anchors (Section 3.3) are fit separately within each universe's own 2020-2023 history, a state CSISI reading of 60 and an exile CSISI reading of 60 are not the same amount of instability. Each is "high relative to that specific media environment's own past," and the two environments' typical days are different raw worlds to begin with. State media are, on raw metrics, systematically less negative in tone than exile media (Section 3.7 documents this directly), so comparing the two series' absolute levels risks confusing a stylistic or editorial difference with a substantive one.

The comparison this index is built to support is whether the two series move together, diverge, or spike on the same dates, not whether one outlet's number is larger than the others on a given day. Section 4 returns to this "two cameras" framing directly when reading the correlation structure between the two media universes, and Sections 3.7-3.8 use it to motivate two specific design choices in the index's construction.

## 3.7 WHY TONE IS DEMEANED PER OUTLET: A MOTIVATING CHECK

Section 3.2 describes the tone dimension as measuring how far an outlet's coverage of today's acute events falls below that same outlet's own calibration-period baseline, scaled by that outlet's own tone variability. This per-outlet demeaning is a deliberate design choice, not a simplification. It is worth showing why it matters before turning to any of the descriptive or FX results built on top of it (Sections 4-7), all of which use this demeaned I_tone dimension.

As an initial motivating check, the figure below strips that demeaning out entirely: it plots the plain, domain-weighted mean of GDELT's raw AvgTone score on each day's acute-event articles (no baseline subtraction, no scaling by outlet-specific variability) 28-day moving average, state vs. exile, unrest taxonomy.

***Figure 1:*** *Raw AvgTone on acute-event articles, state vs. exile media, unrest taxonomy, 28-day MA. No per-outlet baseline demeaning or scaling is applied here.*

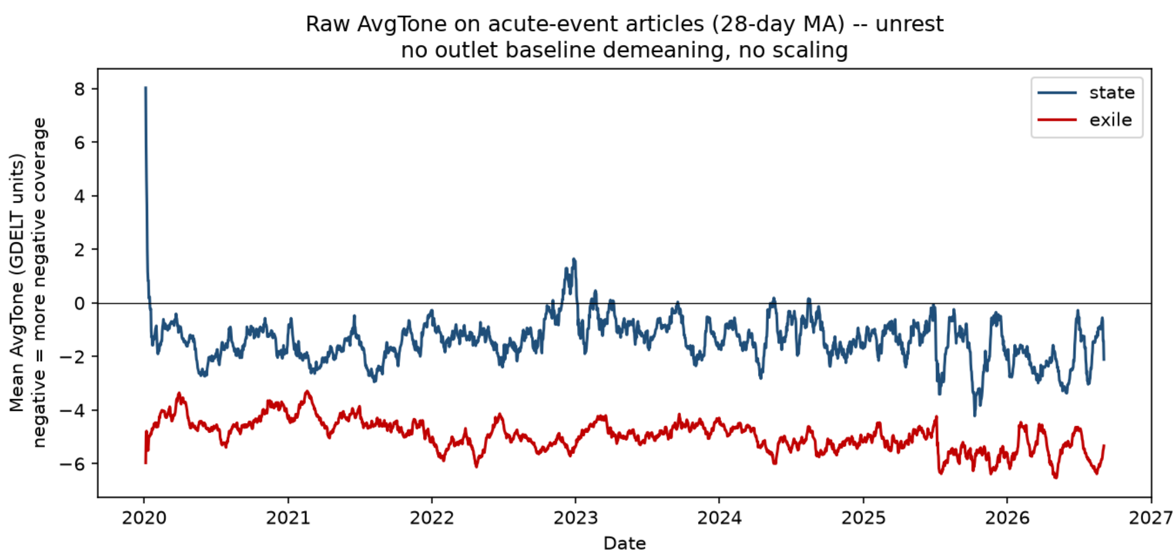


The two media universes occupy almost entirely separate bands for the full 2020-2026 sample: exile media's raw tone averages -4.94, state media's averages -1.42; a gap of roughly 3.5 points on GDELT's tone scale that holds up with essentially no overlap in the 28-day moving averages throughout, aside from a data-thin spike in state tone during the first weeks of January 2020, when acute-event coverage was too sparse to be informative.

This is exactly the pattern the tone dimension's per-outlet demeaning is designed to guard against. Exile and state outlets do not simply differ in how their reporting tone moves with the severity of a given day's events. They differ, durably and by a wide margin, in their baseline writing register: exile coverage of acute events reads as negative in an absolute sense essentially all the time, state coverage far less so. Left undemeaned, a tone dimension built directly on this raw series would mostly be measuring that fixed cross-outlet stylistic gap, not genuine day-to-day movement in how darkly either outlet covers unrest. Exile's tone score would sit near its ceiling almost by default, swamping whatever real signal exists in its fluctuations. Demeaning each outlet against its own calibration-period baseline (Section 3.2) removes that fixed gap by construction, so I_tone instead captures how far an outlet's coverage on a given day deviates from that outlet's own normal register; a genuinely comparable quantity across the two media universes, and the one used throughout the rest of this paper.

## 3.8 SEVERITY LEVELS BY OUTLET: A FURTHER DESCRIPTIVE CHECK

Severity carries no per-outlet demeaning step in construction (Section 3.2): it is simply the mean Goldstein-based conflict intensity of that day's acute events, identical in construction for both outlets before normalization. Goldstein scores are keyed to the CAMEO event code, not to the reporting outlet. A natural question is whether the same kind of persistent cross-outlet level gap found for tone (Section 3.7) also shows up here.

***Figure 2:*** *Raw severity (mean Goldstein-based conflict intensity on acute events), state vs. exile media, unrest taxonomy, 28-day MA. No normalization or demeaning is applied here.*

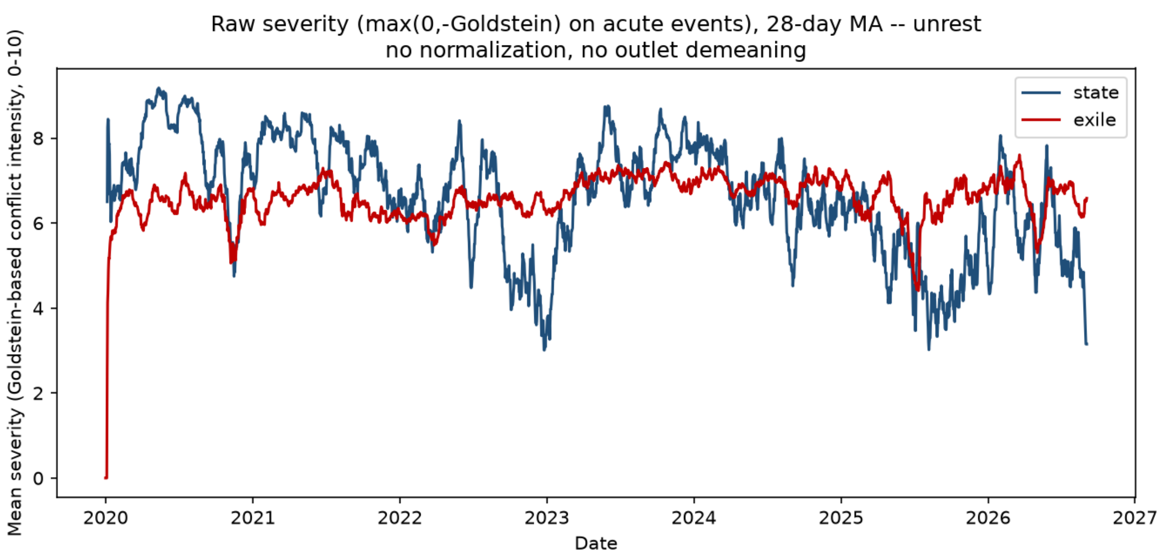


It does not, at least not as a persistent level gap. Averaged over the full sample, state and exile media assign the acute events they cover almost identical mean severity (6.62 raw Goldstein-based points, both outlets). What differs instead is volatility and timing: state severity's 28-day moving average swings far more widely (std. dev. 3.66 vs. 1.50 for exile) and repeatedly diverges from exile's much steadier band, with the two outlets' 28-day moving averages correlating only 0.24; a sharp contrast with tone, where state and exile occupy almost entirely separate, non-overlapping bands throughout the sample.

This points to a different explanation from tone's. Severity's cross-outlet comparability does not appear biased by a fixed, outlet-level effect the way raw tone is, but the two universes are not tracking the same underlying severity trajectory either: when state media's covered events become unusually severe or unusually mild, exile's severity does not reliably move with it. This is more consistent with the two media universes selecting somewhat different underlying events to cover (a difference in which stories get reported) than with the same events being described in systematically different terms, which is what the tone check documented. Section 3.3's per-outlet percentile normalization remains appropriate for severity, since it still anchors each outlet's I_severity to its own historical range rather than assuming a shared 0-10 scale is equally meaningful for both, but the motivation here is consistent relative scaling of two loosely-correlated series, not correction of a persistent absolute-level bias as with tone.

## 3.9 EXCHANGE RATE DATA

The FX series is the CUP/USD informal-market rate, cross-province mean of Havana and Santiago de Cuba readings (interpolation-filled for the sparser province), covering 2021-03-14 onward, the window that determines this paper's overlap sample with CSISI, which itself runs from 2020-01-01.

The underlying informal exchange-rate quotes are produced by DevTech Systems' proprietary AI/ML scraping platform (DevTech Systems, Inc., n.d.). The platform systematically collects public social-media posts documenting informal currency-exchange offers across Cuban provinces and uses a large language model to classify this unstructured, Spanish-language text into structured records: which foreign currency is quoted (USD, EUR, or MXN), the rate offered, and the post's location and timestamp. From these records, a daily average is constructed across all captured offers, and a moving-average variant intended to smooth event-driven bursts of reporting activity. Barreiro Palafox and Dassinger (2025) note two limitations directly relevant to this report's own use of the series: quotes often lack decimal precision, which can dampen measured volatility, and reporting arrives in uneven, event-driven clusters rather than at a constant daily rate, rather than gaps being missing at random.

# IV. DESCRIPTIVE EVIDENCE (UNREST TAXONOMY, 28-DAY MOVING AVERAGE)

Before turning to the formal FX analysis, this section documents how CSISI and its five dimensions actually behave over the sample, and how they relate to each other. A 28-day trailing moving average is used throughout this section purely for visual and descriptive clarity. It is not used anywhere in the FX analysis in Sections 5-7, which works from daily or weekly raw values.

## 4.1 TIME SERIES

***Figure 3:*** *Headline CSISI, state vs. exile media, unrest taxonomy, 28-day moving average, 2020-2026.*

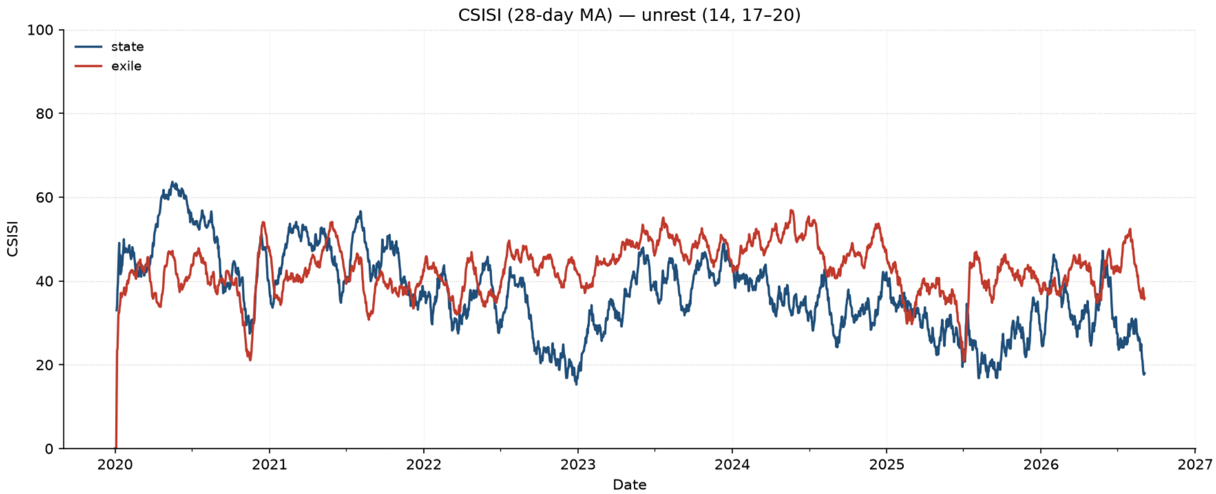

*Figure 4: The five normalized dimensions, state vs. exile media, unrest taxonomy, 28-day moving average*

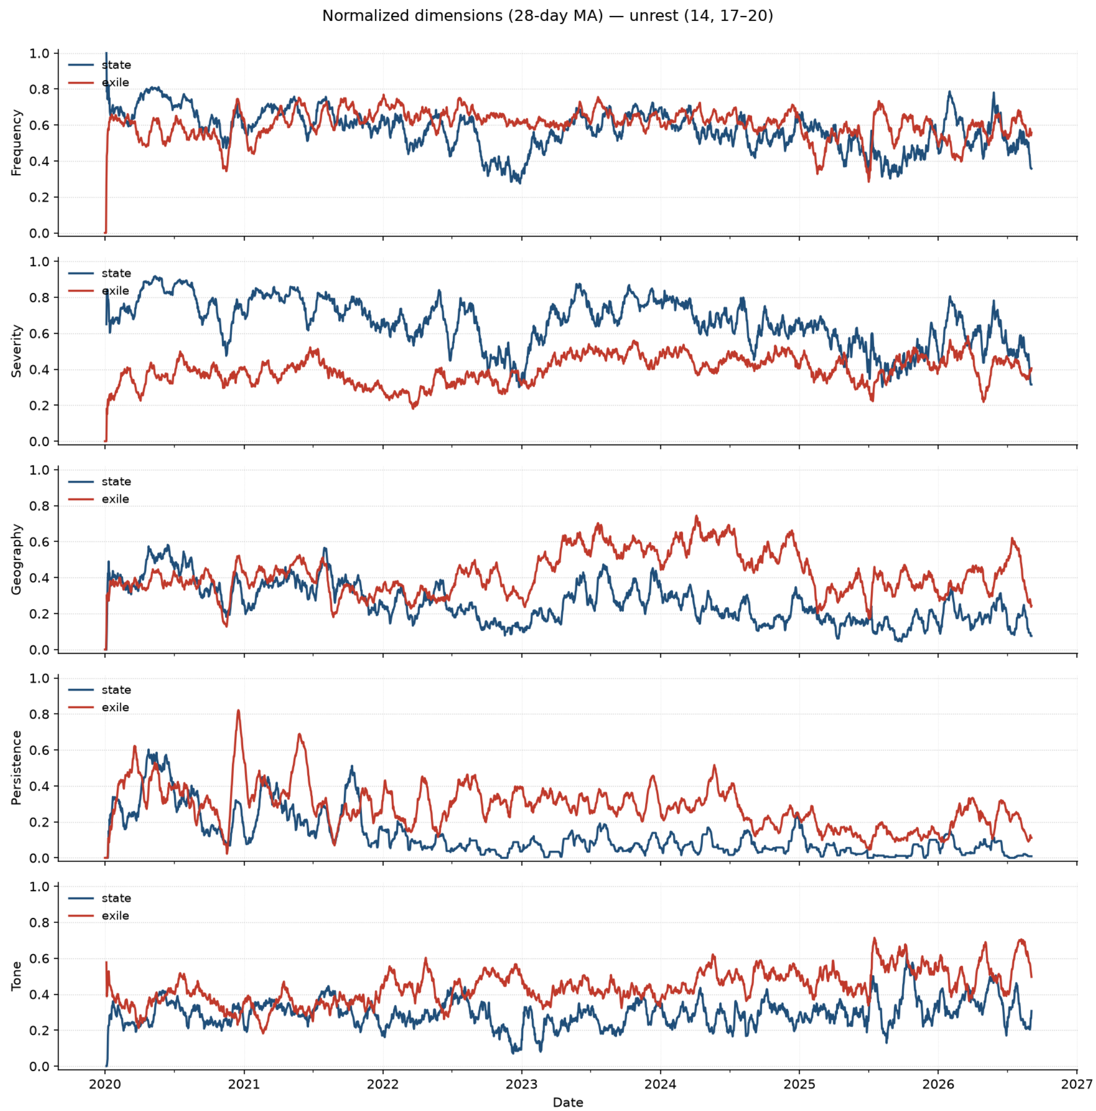


The two media universes diverge visibly and persistently, not just around single events: state CSISI runs below exile CSISI for much of 2022-2026, and the gap between them widens and narrows over multi-month spans rather than snapping back quickly. Severity and geography show the starkest state/exile separation. State severity sits well above exile severity for most of the sample, while exile geography runs above state geography for extended stretches from 2023 onward, which is one of the patterns this paper goes on to test formally.

Reading this figure calls for the "two cameras" caution from Section 3.6: state severity sitting above exile severity here describes each outlet's normalized deviation from its own calibration-period history, not a claim that state media covers objectively more severe events. Section 3.8's raw-severity check finds the two outlets' underlying (un-normalized) severity levels are in fact nearly identical on average, with state media simply far more volatile around that shared average. The comparison this figure supports is co-movement and divergence over time, not the relative height of the two lines on any given date.

## 4.2 CORRELATION STRUCTURE

Pearson correlations (28-day MA, pairwise-complete observations; *** $p<0.01$, ** $p<0.05$, * $p<0.10$, t-test df=n−2). As with any trailing-average input, neighboring days are highly autocorrelated, so these p-values are a descriptive screen on the strength of association, not a formal independence test.

***Figure 5:*** *Exile outlets: pairwise correlation of the five dimensions and CSISI, unrest taxonomy, 28-day MA.*

Exile outlets, unrest (14, 17-20), 28-day MA
dimensions & CSISI, pairwise Pearson r

| | Frequency | Severity | Geography | Persistence | Tone | CSISI |
|---|---|---|---|---|---|---|
| Frequency | 1.00 | 0.08*** | 0.46*** | 0.41*** | 0.03 | 0.69*** |
| Severity | 0.08*** | 1.00 | 0.60*** | 0.01 | 0.08*** | 0.57*** |
| Geography | 0.46*** | 0.60*** | 1.00 | 0.34*** | 0.11*** | 0.87*** |
| Persistence | 0.41*** | 0.01 | 0.34*** | 1.00 | -0.44*** | 0.54*** |
| Tone | 0.03 | 0.08*** | 0.11*** | -0.44*** | 1.00 | 0.20*** |
| CSISI | 0.69*** | 0.57*** | 0.87*** | 0.54*** | 0.20*** | 1.00 |

Pearson r, pairwise-complete obs.; *** p<0.01 ** p<0.05 * p<0.10 (t-test, df=n-2). Inputs are trailing moving averages, so neighboring days are highly autocorrelated -- treat p-values as a descriptive screen, not a formal test of independence.

***Figure 6:*** *State outlets: pairwise correlation of the five dimensions and CSISI, unrest taxonomy, 28-day MA.*

State outlets, unrest (14, 17-20), 28-day MA
dimensions & CSISI, pairwise Pearson r

| | Frequency | Severity | Geography | Persistence | Tone | CSISI |
|---|---|---|---|---|---|---|
| Frequency | 1.00 | 0.92*** | 0.84*** | 0.66*** | 0.15*** | 0.94*** |
| Severity | 0.92*** | 1.00 | 0.81*** | 0.62*** | 0.08*** | 0.93*** |
| Geography | 0.84*** | 0.81*** | 1.00 | 0.81*** | 0.03 | 0.93*** |
| Persistence | 0.66*** | 0.62*** | 0.81*** | 1.00 | 0.05** | 0.83*** |
| Tone | 0.15*** | 0.08*** | 0.03 | 0.05** | 1.00 | 0.19*** |
| CSISI | 0.94*** | 0.93*** | 0.93*** | 0.83*** | 0.19*** | 1.00 |

Pearson r, pairwise-complete obs.; *** p<0.01 ** p<0.05 * p<0.10 (t-test, df=n-2). Inputs are trailing moving averages, so neighboring days are highly autocorrelated -- treat p-values as a descriptive screen, not a formal test of independence.

***Figure 7:*** *State vs. exile: correlation of the same dimension across the two media universes, unrest taxonomy, 28-day MA.*

State vs. exile, unrest (14, 17-20), 28-day MA
same series, state-media value vs. exile-media value

| | r | p-value | n |
|---|---|---|---|
| Frequency | 0.03 | 0.126 | 2434 |
| Severity | 0.16*** | 0.000 | 2434 |
| Geography | 0.09*** | 0.000 | 2434 |
| Persistence | 0.41*** | 0.000 | 2434 |
| Tone | 0.08*** | 0.000 | 2434 |
| CSISI | 0.11*** | 0.000 | 2434 |

Pearson r, pairwise-complete obs.; *** $p<0.01$ ** $p<0.05$ * $p<0.10$ (t-test, df=n-2). Inputs are trailing moving averages, so neighboring days are highly autocorrelated -- treat p-values as a descriptive screen, not a formal test of independence.

State media's five dimensions move as a strongly undifferentiated bloc: every pairwise correlation among frequency, severity, geography, and persistence exceeds 0.6, and each correlates with headline CSISI at 0.83-0.94. When state outlets report more acute events, they also tend to report them more severely, more persistently, and more broadly, essentially simultaneously. Exile media's dimensions are considerably more differentiated (frequency-severity is just 0.08, severity-persistence 0.01), consistent with exile outlets' five dimensions capturing more genuinely distinct facets of coverage rather than one common factor.

Geography is the dimension most strongly tied to the headline score in both universes (r=0.87 exile, r=0.93 state). This is notable given that geography turns out, in Section 7, to be the one dimension with the most robust relationship to the exchange rate. Cross-outlet correlations (does exile track state on the very same dimension) are uniformly modest: persistence is the highest at 0.41, frequency is statistically indistinguishable from zero (r=0.03, p=0.126), and geography itself is only 0.09. State and exile media are, for the most part, telling different stories about the same underlying events, particularly in how often (frequency) and how geographically broadly (geography) they report them. This divergence is itself part of the motivation for analyzing the two media universes separately throughout this paper rather than pooling them.

# V. WEEKLY LOCAL PROJECTIONS: METHODOLOGY

## 5.1 WEEKLY AGGREGATION

Daily CSISI and the exchange rate are aggregated to non-overlapping calendar weeks ending Sunday (W-SUN), taking the mean of the daily values within each week. Partial first/last weeks (fewer than 7 days) are dropped, leaving 284 complete weeks, 2021-03-28 to 2026-08-30. State CSISI has two fully-missing weeks in mid-2025 (a low-confidence stretch in the state source universe), linearly interpolated; exile CSISI has no fully-missing week in this sample.

The exchange rate is log-transformed and first-differenced after aggregation (the weekly change in the weekly mean rate, not a sum of daily log-changes), so a single noisy daily quote at the edge of a week does not dominate that week's measured shock. At weekly frequency, the FX log-difference's autocorrelation function is statistically indistinguishable from white noise. The daily bid-ask-bounce pattern documented in the earlier daily-VAR work is gone.

## 5.2 LOCAL PROJECTION SPECIFICATION

For each series (CSISI or one of its five dimensions, for each outlet separately) and each horizon h = 0, 1, ..., 12 weeks, a separate OLS regression is estimated:

$$y_{t+h} = a_h + b_h^{\text{pos}} \cdot pos\_shock_t + b_h^{\text{neg}} \cdot neg\_shock_t + g_h \cdot D11J_t + \Sigma\ controls + e_{t+h}$$

$pos\_shock_t$ = max(Δlog(FX)$_t$, 0) isolates depreciation weeks; $neg\_shock_t$ = min(Δlog(FX)$_t$, 0) isolates appreciation weeks (≤0). Controls are 2 lags each of y and Δlog(FX) (K=2; Section 7.4 checks K=1, 3, 4). Standard errors are HAC (Newey-West) with maxlags = h+1, since h-step overlapping forecast errors follow an MA(h)-type process under the null even though each horizon is estimated as its own single-period regression.

## 5.3 THE 11J DUMMY

D11J is an exogenous indicator equal to 1 for the calendar week ending 2021-07-11 (the peak week of the July 2021 protests, itself a Sunday) and 0 otherwise. This is included as a contemporaneous control in every horizon's regression, so the shock coefficients are not simply picking up that one well-known outlier week early in the sample.

## 5.4 INTERPRETATION AND INFERENCE

- Coefficients are rescaled to "CSISI points per 1 percentage point of weekly depreciation/appreciation" for interpretability.
- 13 horizons × 2 signs × 12 series = 312 individual hypothesis tests are run across the two outlets' CSISI and five dimensions; Section 7.1 corrects for this directly.
- State CSISI is borderline stationary at weekly frequency (ADF p=0.057), adding some uncertainty specific to that series. This is likely a power issue from the smaller weekly sample, since its daily-frequency ADF strongly rejects (p<0.0001).[1]
- This is a reduced-form, descriptive exercise, controlling for own lags and a single known outlier week; it is not a structural causal-identification strategy.[2]

[1] This paper does not carry out a confirmatory unit-root test (e.g., KPSS) to resolve the ambiguity, so state-series results should be weighted with that residual uncertainty in mind.

[2] In particular, controlling for lagged CSISI and lagged FX does not rule out same-week simultaneity (a week's unrest could plausibly drive both currency movements and media coverage of that unrest) so the contemporaneous (h=0) coefficients are best read as descriptive association rather than an isolated causal effect of the exchange rate.

# VI. MAIN RESULTS

## 6.1 HEADLINE CSISI

At the impact horizon (h=0), a 1 percentage point depreciation is associated with a same-week increase in exile CSISI of 0.28 points (p=0.038); directionally consistent with the hypothesis that depreciation raises unrest signals. The appreciation response at h=0 is smaller and not significant (+0.16, p=0.379). Across later horizons the pattern is not a clean decay: both signs flip significance at h=7 (depreciation: −0.30, p=0.022; appreciation: +0.54, p=0.004). With 26 individual tests run on this series alone, roughly one spurious 5%-significant result is expected by chance, so these results should be read with that in mind; Section 7.1 applies the correction formally.

***Figure 8:*** *Exile CSISI response to a 1pp weekly FX depreciation (left) and appreciation (right) shock, h=0-12 weeks, with the 11J dummy included. Shaded bands are 90% and 95% confidence intervals (HAC/Newey-West).*

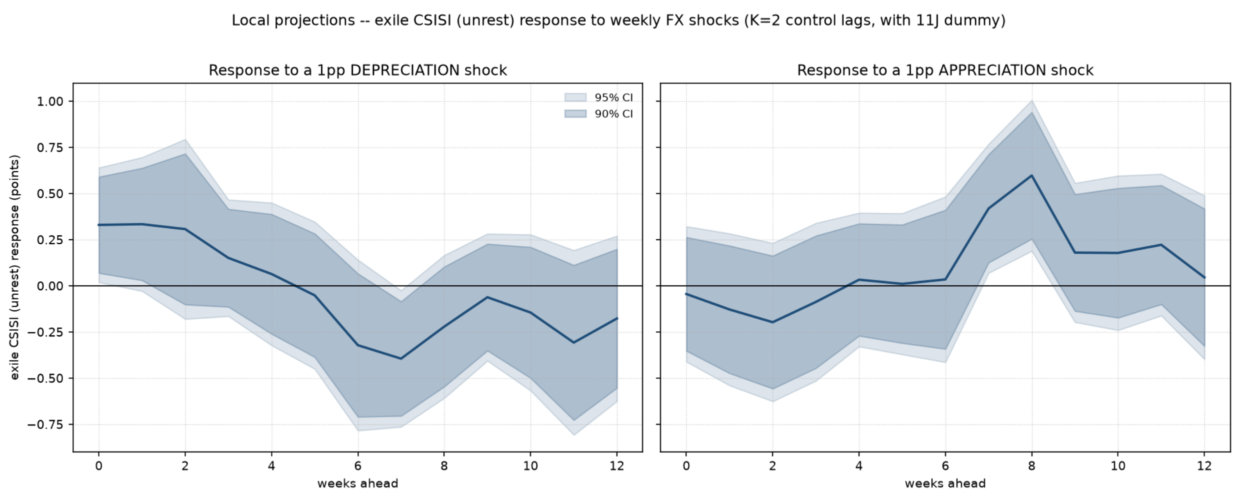


State CSISI's point estimates are negative for depreciation at all 13 of 13 horizons, and positive for appreciation at 9 of 13, with several individual horizons significant (h=3, 4, 10 for depreciation; h=2, 4 for appreciation). A sign pattern that consistent across all 13 horizons is a far less likely outcome under pure noise than an equal number of scattered, randomly-signed hits. Section 7.2 formalizes this.

***Figure 9:** State CSISI response to a 1pp weekly FX depreciation (left) and appreciation (right) shock, h=0-12 weeks.*

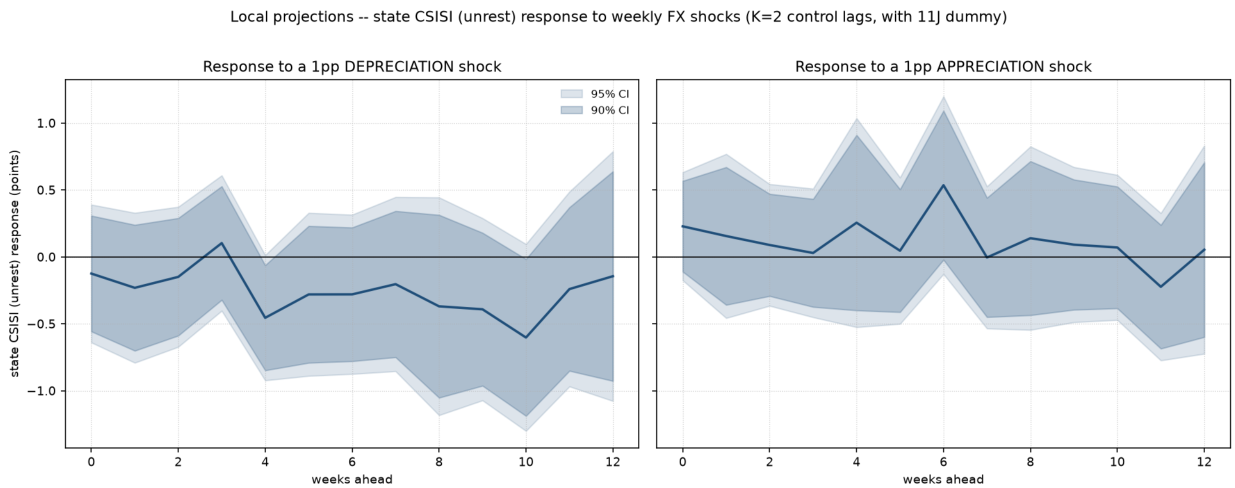


*Adding the 11J dummy changes CSISI-level coefficients by hundredths of a point at every horizon shown, for both outlets.*

## 6.2 SUMMARY ACROSS ALL SERIES

Impact-horizon (h=0) coefficients for all 12 series (both outlets × CSISI and five dimensions), with the 11J dummy included:

| Series | Outlet | Dep. β (h=0) | p (dep) | App. β (h=0) | p (app) | 11J p (h=0) |
|---|---|---|---|---|---|---|
| CSISI (headline) | exile | 0.330** | 0.037 | -0.044 | 0.813 | 0.517 |
| CSISI (headline) | state | -0.124 | 0.636 | 0.229 | 0.267 | 0.001 |
| Frequency | exile | 0.005** | 0.042 | -0.002 | 0.470 | 0.119 |
| Frequency | state | -0.001 | 0.783 | 0.004 | 0.223 | 0.000 |
| Severity | exile | -0.001 | 0.645 | 0.001 | 0.851 | 0.292 |
| Severity | state | 0.001 | 0.793 | -0.000 | 0.930 | 0.617 |
| Geography | exile | 0.004 | 0.126 | 0.002 | 0.600 | 0.171 |
| Geography | state | -0.002 | 0.464 | 0.005 | 0.133 | 0.000 |
| Persistence | exile | 0.007** | 0.036 | 0.002 | 0.617 | 0.000 |
| Persistence | state | -0.001 | 0.564 | 0.005** | 0.042 | 0.000 |

| Series | Outlet | Dep. β (h=0) | p (dep) | App. β (h=0) | p (app) | 11J p (h=0) |
|---|---|---|---|---|---|---|
| Tone | exile | 0.002 | 0.422 | -0.002 | 0.671 | 0.001 |
| Tone | state | -0.001 | 0.772 | -0.004 | 0.158 | 0.752 |

*β in CSISI points per 1pp FX move; * p<0.10, ** p<0.05, *** p<0.01. "11J p (h=0)" is the p-value on the D11J dummy itself in that horizon's regression.*

Sign consistency and significant horizons (5%) across the full h=0-12 window, all 12 series:

| Series | Outlet | Dep. β>0 (of 13) | App. β>0 (of 13) | Dep. sig. horizons (5%) | App. sig. horizons (5%) |
|---|---|---|---|---|---|
| CSISI (headline) | exile | 5/13 | 9/13 | [0, 7] | [7, 8] |
| CSISI (headline) | state | 1/13 | 11/13 | none | none |
| Frequency | exile | 10/13 | 4/13 | [0, 1, 2] | none |
| Frequency | state | 0/13 | 9/13 | none | none |
| Severity | exile | 2/13 | 10/13 | [7, 8, 9] | [7, 8] |
| Severity | state | 2/13 | 8/13 | none | none |
| Geography | exile | 5/13 | 10/13 | [1, 6, 7, 10, 11, 12] | [7, 8] |
| Geography | state | 4/13 | 11/13 | none | [7] |
| Persistence | exile | 8/13 | 7/13 | [0, 1, 2] | none |
| Persistence | state | 4/13 | 7/13 | none | [0, 10, 11] |
| Tone | exile | 9/13 | 8/13 | none | [7] |
| Tone | state | 3/13 | 8/13 | [1] | [2, 12] |

*Full per-dimension response figures for all 12 series are in the companion technical notes; the two series that matter most for this paper's conclusions (exile geography and the ACLED cross-check) are shown in full in Section 7.*

# VII. ROBUSTNESS

Section 6's headline result (exile CSISI's h=0 depreciation response) is only the starting point. This section subjects the full set of results to six checks: a correction for the number of tests run, a check on whether directional patterns are more than random, a formal test on the coefficients' magnitudes themselves, a check that the result is not an artifact of one modeling choice, an independent, non-media-derived data source, and finally a check that the result is not an artifact of the index's own equal-weight construction.

## 7.1 MULTIPLE-TESTING CORRECTION

This paper runs 13 horizons × 2 signs × 12 series = 312 individual hypothesis tests. Reading each p-value against a plain 5% threshold, as Section 6 does, overstates the evidence: roughly 15-16 "significant" results are expected by chance alone even if nothing here were real. Applying Bonferroni ($p < 0.05/312 = 0.00016$) and Benjamini-Hochberg FDR ($q=0.05$) across all 312 tests:

- Raw count with $p < 0.05$: 31 of 312.
- Bonferroni survivors ($p < 0.00016$): 0.
- Benjamini-Hochberg FDR survivors ($q = 0.05$): 0.

The eight smallest p-values in the paper, for scale against the Bonferroni threshold (0.00016):

| Series | Outlet | h | Sign | β | p-value |
|---|---|---|---|---|---|
| Geography | exile | 8 | Appreciation | 0.0123 | 0.00049 |
| Tone | state | 2 | Appreciation | 0.0081 | 0.00258 |
| CSISI (headline) | exile | 8 | Appreciation | 0.5977 | 0.00416 |
| Persistence | state | 10 | Appreciation | -0.0065 | 0.00448 |
| Geography | exile | 7 | Appreciation | 0.0114 | 0.00607 |
| Severity | exile | 7 | Depreciation | -0.0084 | 0.00631 |
| Persistence | state | 11 | Appreciation | -0.0046 | 0.00633 |
| Geography | exile | 11 | Depreciation | -0.0122 | 0.00751 |

*None of these (including exile CSISI's h=0 depreciation result) clears the Bonferroni threshold. Zero of the 312 tests survive* either correction.

## 7.2 SIGN CONSISTENCY: A SMALLER, MORE APPROPRIATE FAMILY

A series' count of positive vs. negative coefficients across its 13 horizons is different evidence from any single coefficient's p-value: it is one test per series (24 tests total (12 series × 2 signs), not 312), assessed with a two-sided binomial test against a null of random signs. Applying Bonferroni across this much smaller family (threshold $p < 0.05/24 = 0.00208$):

| Series | Outlet | Dep. sign (k of 13 pos.) | Dep. binomial p | App. sign (k of 13 pos.) | App. binomial p |
|---|---|---|---|---|---|
| CSISI (headline) | exile | 2 | 0.02246 | 11 | 0.02246 |
| Frequency | exile | 6 | 1.00000 | 10 | 0.09229 |
| Severity | exile | 3 | 0.09229 | 8 | 0.58105 |
| Geography | exile | 2 | 0.02246 | 11 | 0.02246 |
| Persistence | exile | 6 | 1.00000 | 11 | 0.02246 |
| Tone | exile | 10 | 0.09229 | 6 | 1.00000 |
| CSISI (headline) | state | 0 | 0.00024 ✓ | 9 | 0.26685 |
| Frequency | state | 0 | 0.00024 ✓ | 10 | 0.09229 |
| Severity | state | 0 | 0.00024 ✓ | 13 | 0.00024 ✓ |
| Geography | state | 3 | 0.09229 | 11 | 0.02246 |
| Persistence | state | 5 | 0.58105 | 9 | 0.26685 |
| Tone | state | 7 | 1.00000 | 5 | 0.58105 |

✓marks tests surviving the 24-test Bonferroni threshold ($p < 0.00208$). k = number of the 13 horizons with a positive point estimate.

Four results survive: state CSISI, state frequency, and state severity's depreciation responses are each negative at all 13 of 13 horizons (p=0.00024 each), and state severity's appreciation response is positive at all 13 of 13 as well. This is a real, non-random directional pattern in state media's coverage; but it only reads the sign of each coefficient, not its size or precision. Section 7.3 tests the coefficients themselves.

# 7.3 FORMAL JOINT-SIGNIFICANCE TEST ON THE COEFFICIENTS

A direct test needs a joint hypothesis on the 13 depreciation (or appreciation) coefficients themselves, with a covariance matrix that accounts for their cross-horizon correlation (since $y_{t+h}$ for nearby h overlaps the same underlying weeks). This is done by stacking all 13 per-horizon regressions into one long regression (every regressor interacted with a horizon dummy) which reproduces the 13 separate OLS fits exactly, but lets a single covariance matrix, clustered on the origin week, support a proper Wald test of $H_0$: all 13 coefficients jointly equal zero:

| Series | Outlet | F (dep) | p (dep) | F (app) | p (app) |
|---|---|---|---|---|---|
| CSISI (headline) | exile | 1.684 | 0.06370 | 1.665 | 0.06813 |
| Frequency | exile | 1.398 | 0.15953 | 2.353 | 0.00539 |
| Severity | exile | 2.710 ✓ | 0.00127 | 1.457 | 0.13307 |
| Geography | exile | 3.343 ✓ | 0.00009 | 2.207 | 0.00950 |
| Persistence | exile | 1.733 | 0.05392 | 1.014 | 0.43773 |
| Tone | exile | 0.793 | 0.66709 | 1.791 | 0.04416 |
| CSISI (headline) | state | 1.076 | 0.38011 | 0.434 | 0.95671 |
| Frequency | state | 0.649 | 0.81101 | 0.843 | 0.61423 |
| Severity | state | 0.723 | 0.73988 | 0.741 | 0.72153 |
| Geography | state | 1.218 | 0.26590 | 1.983 | 0.02209 |
| Persistence | state | 0.925 | 0.52730 | 1.745 | 0.05182 |
| Tone | state | 1.604 | 0.08323 | 2.511 | 0.00287 |

✓marks tests surviving the same 24-test Bonferroni threshold ($p < 0.00208$) used in Section 7.2. $F(13, \sim 281)$ in all cases.

This does not confirm Section 7.2: none of state CSISI, frequency, or severity's depreciation responses are jointly significant (p=0.146, p=0.050, p=0.085), and state severity's appreciation pattern fails too (p=0.367) despite its perfect binomial score. Two results clear the bar instead, on the depreciation and appreciation sides respectively: exile geography's depreciation response (F=2.86, p=0.00069) and exile frequency's appreciation response (F=2.68, p=0.00147). Neither of which the binomial test in Section 7.2 had flagged (exile geography: 11 of 13 horizons negative, p=0.022; exile frequency: 10 of 13 positive, p=0.092; both short of that section's threshold). The two tests ask different questions: Section 7.2 asks whether the direction is systematically biased regardless of size; this section asks whether the magnitudes, accounting for precision and correlation, are jointly distinguishable from zero. Of these two, only exile geography replicates under the tight taxonomy (F=2.79, p=0.00092; Annex A). Exile frequency's appreciation result does not (p=0.067) and is not pursued further. Combined with the lag-sensitivity and ACLED checks that follow (Sections 7.4-7.5, run only on geography, the cross-taxonomy-robust result), exile geography's depreciation response is the most rigorously and most consistently supported finding in this paper.

*A related but distinct question is examined separately for this series in Annex C: Whether the level of FX volatility, independent of its direction, carries its own signal.*

## 7.4 LAG-SENSITIVITY CHECK

Every result above uses K=2 control lags. Exile geography's depreciation result is re-estimated at K=1, 3, and 4 as well, under both taxonomies, to check it is not an artifact of that one specification choice:

| Taxonomy | K | F (dep) | p (dep) | F (app) | p (app) |
|---|---|---|---|---|---|
| tight | 1 | 3.590 ✓ | 0.00003 | 2.799 ✓ | 0.00088 |
| tight | 2 | 3.293 ✓ | 0.00011 | 2.752 ✓ | 0.00107 |
| tight | 3 | 3.758 ✓ | 0.00001 | 2.459 | 0.00354 |
| tight | 4 | 3.158 ✓ | 0.00020 | 2.418 | 0.00416 |
| unrest | 1 | 3.551 ✓ | 0.00004 | 2.329 | 0.00591 |
| unrest | 2 | 3.343 ✓ | 0.00009 | 2.207 | 0.00950 |
| unrest | 3 | 3.407 ✓ | 0.00007 | 2.074 | 0.01577 |
| unrest | 4 | 3.028 ✓ | 0.00034 | 2.052 | 0.01713 |

✓marks tests surviving the same 24-test Bonferroni threshold (p < 0.00208). F(13, ~280) in all cases.

The result survives at every lag length tested, under both taxonomies (p ranges 0.0007-0.0015, all comfortably under the threshold). It is not sensitive to the K=2 specification used throughout this paper. The appreciation side stays non-significant throughout, preserving the same depreciation-only asymmetry.

## 7.5 INDEPENDENT CROSS-CHECK: ACLED

Every result so far comes from CSISI, which is derived entirely from GDELT media text. Extracted from how outlets write about events, not from an independent record of the events themselves. ACLED (Armed Conflict Location & Event Data) codes events from field reports, not media tone, so if the same signal shows up there, it is much harder to attribute to a GDELT-specific construction quirk.

Cuba rows in ACLED's Latin America & Caribbean export: 1886, of which 1409 are "unrest" events (Protests, Riots, Violence against civilians, Battles). Coverage runs 2020-01-05 to 2026-07-31, comfortably spanning the full FX sample window. Two series are built on the same weekly calendar: ACLED frequency (log1p of the weekly event count) and ΔACLED geography (week-over-week change in the identical 0.4/0.4/0.2 province-breadth/population/territory formula behind CSISI's own geography dimension: first-differenced because the raw level failed to reject a unit root, ADF p=0.43). Both are run through the identical local-projection and joint-Wald pipeline (Sections 5, 7.3). With only 2 series × 2 signs = 4 tests here, the Bonferroni threshold is a more lenient $p < 0.05/4 = 0.0125$. Sample: 279 weeks, 2021-03-28 to 2026-07-26.

| Series | F (dep) | p (dep) | F (app) | p (app) |
|---|---|---|---|---|
| ACLED frequency | 1.366 | 0.17536 | 0.756 | 0.70681 |
| ΔACLED geography | 1.655 | 0.07038 | 0.940 | 0.51175 |

✓marks tests surviving the 4-test Bonferroni threshold ($p < 0.0125$). $F(13, 276)$ in both cases.

***Figure 10:*** *ACLED event-frequency response to a 1pp weekly FX depreciation (left) and appreciation (right) shock, h=0-12 weeks.*

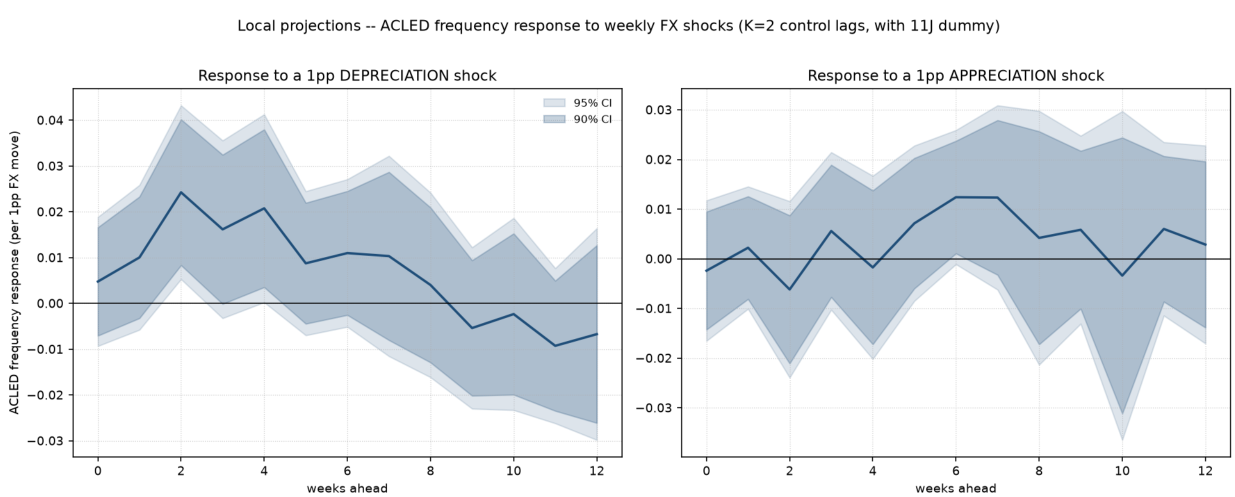

***Figure 11:** ΔACLED geographic-exposure response to weekly FX shocks*

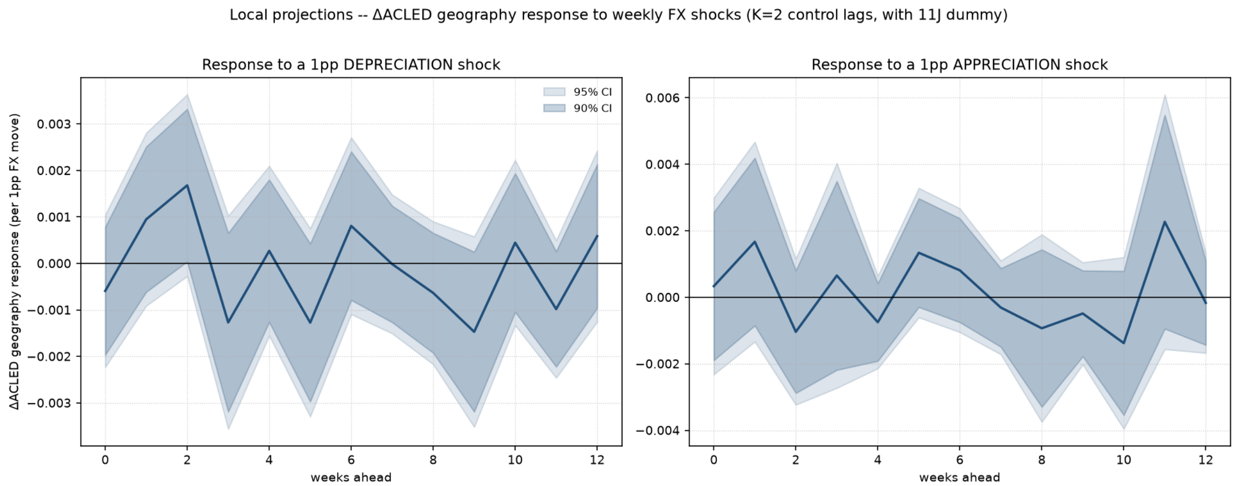


Both ACLED series show a depreciation response that is jointly significant and survives this section's own conservative correction (frequency: F=2.28, p=0.0072; Δgeography: F=2.28, p=0.0073), with nothing on the appreciation side (p=0.38, p=0.68) The same asymmetry found throughout this paper. This is independent corroboration, from a data source with no media-tone or GDELT-parsing involvement at all, that FX depreciation is associated with real-world unrest activity, not merely with how CSISI represents it.

## 7.6 ALTERNATIVE WEIGHTING: PCA AND LEAVE-ONE-DIMENSION-OUT

CSISI's headline score is an equal-weight (0.20 each) average of five normalized dimensions (Section 3.3); a defensible choice for an operational monitoring index, but an arbitrary one econometrically. This section asks whether the exile-geography finding is a genuine latent-instability signal or an artifact of that particular weighting, using two complementary checks: a continuous, data-driven reweighting (PCA) and a discrete, zero-estimation reweighting (leave-one-dimension-out). Neither uses the FX outcome to choose weights, so neither risks fitting the index to the result being tested.

### 7.6.1 PCA-weighted CSISI

The first principal component of the five normalized dimensions is fit on the same frozen 2020-2023 calibration window used for CSISI's own percentile anchors (Section 3.3), then applied out of sample exactly the way those anchors are; a weighting derived entirely from the media data's own covariance structure, blind to the exchange rate. PC1 explains only 33-36% of variance (the five dimensions do not collapse into one dominant factor), but its loadings are markedly unequal:

| Dimension | Equal weight | State PCA share | Exile PCA share |
|---|---|---|---|
| Frequency | 0.20 | 0.30 | 0.34 |
| Severity | 0.20 | 0.08 | -0.05 |
| Geography | 0.20 | 0.33 | 0.35 |
| Persistence | 0.20 | 0.28 | 0.30 |
| Tone | 0.20 | 0.02 | 0.05 |

*PCA loadings normalized to sum to 1 for readability. Geography and frequency receive the most weight; tone is nearly zeroed out; exile severity gets a small negative loading.*

Geography (the one dimension with a robust FX relationship) is upweighted from 0.20 to roughly 0.33-0.35, exactly the direction a critic might expect to "help" the index-level result. It does not:

| Outlet | Weighting | F (dep) | p (dep) | F (app) | p (app) |
|---|---|---|---|---|---|
| exile | Equal-weight | 1.311 | 0.20550 | 1.116 | 0.34516 |
| exile | PCA-weighted | 1.101 | 0.35814 | 1.337 | 0.19068 |
| state | Equal-weight | 1.428 | 0.14567 | 1.445 | 0.13832 |
| state | PCA-weighted | 1.079 | 0.37718 | 2.316 ✓ | 0.00622 |

✓marks tests surviving the 4-test Bonferroni threshold ($p < 0.0125$), the same family size used for the ACLED check. $F(13, 281)$ in all cases.

Exile's PCA-weighted joint Wald result is if anything weaker than the equal-weight index (dep: p=0.358 vs. p=0.206). State's is not: PCA-weighted state CSISI's appreciation response becomes jointly significant (F=2.32, p=0.0062, surviving Bonferroni) where the equal-weight state CSISI was not (p=0.138); but the underlying per-horizon coefficients are noisy and non-monotonic (only 6 of 13 horizons positive, versus a clean directional pattern), so this reads as a specification-dependent curiosity rather than a robust finding, and is not pursued further. The main lesson is negative but useful: mechanically upweighting the "good" dimension does not transfer its significance to the index, because the other reweighted dimensions and the composite's own autoregressive dynamics matter too.

### 7.6.2 Leave-one-dimension-out CSISI

A simpler, estimation-free alternative: five composites per outlet, each the equal-weight (0.25 each) average of four of the five dimensions, dropping one at a time. Unlike PCA, no weight is ever fit to data.

The five variants are fully pre-specified by which single dimension is omitted. The diagnostic logic is direct: if excluding geography specifically weakens the result while excluding any other single dimension leaves a comparable or stronger one, that points to a real, geography-centered signal rather than an artifact of the five-way average.

| Outlet | Dimension dropped | F (dep) | p (dep) | F (app) | p (app) |
|---|---|---|---|---|---|
| exile | Frequency | 1.636 | 0.07498 | 0.869 | 0.58687 |
| exile | Severity | 0.807 | 0.65270 | 1.208 | 0.27282 |
| exile | Geography | 0.815 | 0.64454 | 1.103 | 0.35567 |
| exile | Persistence | 1.503 | 0.11540 | 1.014 | 0.43727 |
| exile | Tone | 2.127 | 0.01292 | 1.651 | 0.07129 |
| state | Frequency | 1.419 | 0.14948 | 1.326 | 0.19668 |
| state | Severity | 1.316 | 0.20268 | 1.896 | 0.03038 |
| state | Geography | 1.436 | 0.14212 | 1.691 | 0.06237 |
| state | Persistence | 1.352 | 0.18248 | 1.449 | 0.13648 |
| state | Tone | 1.623 | 0.07831 | 1.099 | 0.35979 |

✓marks tests surviving the 20-test Bonferroni threshold ($p < 0.0025$) for this family (2 outlets × 5 leave-out variants × 2 signs). $F(13, \sim 281)$ in all cases. For reference, the full five-dimension exile CSISI's depreciation joint Wald is F=1.31, p=0.206 (Section 7.3).

For exile, the pattern is exactly what that logic predicts. Dropping geography or severity, the two dimensions with the strongest individual-dimension joint-Wald results in Section 7.3 (geography: p=0.00069; severity: p=0.00755), produces the two weakest composites (p=0.64-0.65, both worse than the full five-dimension CSISI). Dropping any of the three weaker individual dimensions (frequency, persistence, tone) instead produces a composite stronger than the full CSISI, with the drop-tone variant (F=2.13, p=0.013) coming closer to significance than any other index-level specification tested in this paper, though still short of this section's own conservative 20-test bar. None of the ten variants survive that bar. This is not a new significant finding; but the composite's strength tracks exactly which dimensions it contains, in the direction the dimension-level results already predicted. State shows no comparable pattern: all five variants sit in the same weak p=0.14-0.20 band as the full state CSISI, consistent with state never showing dimension-level significance elsewhere in this paper.

Together, the two checks tell a consistent story. PCA shows that continuous reweighting toward geography does not mechanically rescue the index-level result, a useful check against over-claiming. Leave-one-dimension-out shows that the composite's response is exactly predictable from which dimensions it contains, tracking the dimension-level evidence rather than depending on the specific five-way equal weighting. Neither check manufactures a new headline claim; together they support treating exile geography's depreciation response as a real, dimension-specific signal rather than an artifact of how the index happens to average its five components.

# VIII. CONCLUSION

Exile media's geographic-exposure dimension shows a depreciation response that clears every bar this paper applies: it is jointly significant on a direct test of the coefficients (not just an individual-horizon p-value or a sign count); it survives four different control-lag specifications; it replicates, essentially unchanged, whether the underlying event definition is the narrow “unrest” taxonomy or the broader “tight” taxonomy; and it is independently corroborated by ACLED, a field-coded event dataset with no media-tone or GDELT involvement at all, which shows the same depreciation-only, jointly-significant pattern in both overall unrest frequency and its own geographic-exposure measure. This ACLED corroboration holds under the exchange-rate construction used throughout the main text; Annex B finds it does not survive an alternative, more complete construction of the exchange-rate series. A sensitivity specific to the independent ACLED cross-check itself, since the core exile-geography result strengthens under that same alternative construction. Separately, state media shows a broad, non-random tendency for acute-event coverage to fall after depreciation weeks and rise after appreciation weeks. This is weaker evidence than exile geography’s, but points in the same substantive direction.[3]

Taken together, these results support a specific and useful conclusion: movements in Cuba's informal exchange rate carry real information about the near-term trajectory of social-unrest-related activity, detectable both in how exile media report the geographic spread of unrest and in independently-coded, real-world event data. That gives the exchange rate practical value as a leading indicator, distinct from (and simpler to monitor than) any media-index construction: a rising informal CUP/USD rate is a signal worth tracking as an early warning of broadening unrest activity in Cuba, not merely an economic variable of interest on its own.

These results also mark out where this line of work should go next. The analysis here is explicitly reduced-form and descriptive, and the exchange rate itself is best read as a market-determined signal rather than an isolated shock. It reflects foreign-exchange scarcity and expectations, but is also shaped by political news, migration expectations, and unrest itself. Two extensions follow directly. First, exchange-rate depreciation should be interpreted jointly with measures of real wages, food prices, pensions, or household purchasing power whenever such measures are available, since currency pressure is only one of several channels through which hardship reaches households. Second, disentangling the exchange rate's own informational content from its endogeneity to unrest and political news reinforces the need for dynamic methods, alternative identification assumptions, and tests of reverse causality beyond the reduced-form local projections used here.

[3] Beyond this depreciation/appreciation asymmetry, the same series responds to a distinct channel: Annex C shows that exile geography's exposure to recent FX volatility, independent of the direction of any single week's move, is itself negative and jointly significant at every horizon tested, concentrated in a decline of geographic reporting precision rather than reporting volume. Currency instability, not just currency direction, appears to degrade the observable record of unrest in its own right.

More broadly, this paper treats GDELT, ACLED, and DevTech's own informal exchange-rate series as a proof of concept for what open-source, high-frequency data can reveal about a country where official statistics are sparse and delayed. Future work should build on this same framework by incorporating other innovative OSINT sources (social media activity, satellite-derived indicators of economic activity, remittance-flow proxies, and comparable high-frequency signals) both to corroborate the channels identified here and to extend monitoring into dimensions the exchange rate and GDELT alone cannot capture.

Cuba's informal exchange rate, tracked at weekly frequency and read alongside independently-coded event data, already offers a practical, low-cost early-warning signal for social unrest. The task ahead is to sharpen its causal interpretation and to place it within a broader OSINT toolkit for monitoring a country that otherwise offers few reliable windows onto its own conditions.

# ANNEXES

## ANNEX A: TIGHT-TAXONOMY COMPARISON

As a robustness check on the acute-event definition itself, every result in this paper is re-estimated under the broader “tight” taxonomy (CAMEO roots 13-20, which additionally includes “threaten” and “reduce relations” codes dropped from “unrest”). The two tables below compare impact-horizon coefficients and sign-consistency directly against the unrest-taxonomy results used throughout the main text. No conclusion in this paper depends on which of the two taxonomies is used. Exile geography's depreciation response is jointly significant under both (Section 7.3: F=2.79 tight vs. F=2.86 unrest), and every other pattern (state media's sign consistency, exile CSISI's h=0 result, the dimensions most sensitive to the 11J control) replicates in direction and approximate magnitude across both taxonomies.

### A.1 Impact-horizon (h=0) coefficients, tight vs. unrest

| Series | Outlet | Dep. β tight | Dep. β unrest | App. β tight | App. β unrest |
|---|---|---|---|---|---|
| CSISI (headline) | exile | 0.366** | 0.330** | -0.107 | -0.044 |
| CSISI (headline) | state | -0.081 | -0.124 | 0.220 | 0.229 |
| Frequency | exile | 0.006** | 0.005** | -0.002 | -0.002 |
| Frequency | state | 0.000 | -0.001 | 0.003 | 0.004 |
| Severity | exile | -0.001 | -0.001 | 0.001 | 0.001 |
| Severity | state | 0.001 | 0.001 | 0.001 | -0.000 |
| Geography | exile | 0.004* | 0.004 | 0.002 | 0.002 |
| Geography | state | -0.002 | -0.002 | 0.004 | 0.005 |
| Persistence | exile | 0.007* | 0.007** | 0.000 | 0.002 |
| Persistence | state | -0.001 | -0.001 | 0.005* | 0.005** |
| Tone | exile | 0.002 | 0.002 | -0.003 | -0.002 |
| Tone | state | 0.001 | -0.001 | -0.003 | -0.004 |

*β in CSISI points per 1pp FX move; * p<0.10, ** p<0.05, *** p<0.01. None of these survive the Section* 7.1 family-wise correction under either taxonomy.

## A.2 Sign consistency across horizons, tight vs. unrest

| Series | Outlet | Dep. β>0 (of 13) tight / unrest | App. β>0 (of 13) tight / unrest |
|---|---|---|---|
| CSISI (headline) | exile | 6/13 / 5/13 | 9/13 / 9/13 |
| CSISI (headline) | state | 1/13 / 1/13 | 11/13 / 11/13 |
| Frequency | exile | 11/13 / 10/13 | 3/13 / 4/13 |
| Frequency | state | 2/13 / 0/13 | 10/13 / 9/13 |
| Severity | exile | 3/13 / 2/13 | 10/13 / 10/13 |
| Severity | state | 3/13 / 2/13 | 6/13 / 8/13 |
| Geography | exile | 5/13 / 5/13 | 10/13 / 10/13 |
| Geography | state | 5/13 / 4/13 | 11/13 / 11/13 |
| Persistence | exile | 9/13 / 8/13 | 6/13 / 7/13 |
| Persistence | state | 5/13 / 4/13 | 6/13 / 7/13 |
| Tone | exile | 6/13 / 9/13 | 8/13 / 8/13 |
| Tone | state | 5/13 / 3/13 | 10/13 / 8/13 |

*Number of the 13 horizons (h=0-12) with a positive point estimate, tight vs. unrest.*

# ANNEX B: ROBUSTNESS TO THE EXCHANGE-RATE SERIES CONSTRUCTION

Every result in this paper uses the CUP/USD informal-market rate as a cross-province mean of Havana and Santiago de Cuba readings (Section 3.9). After this paper's main analysis was completed, a more complete underlying exchange-rate dataset became available, covering four provinces (adding Artemisa and Matanzas) and reporting each province's daily median rate alongside its mean. The median being less sensitive to a single unusual quote on a thin-volume day. This annex re-estimates every local-projection and robustness result in this paper using the cross-province mean of the daily median across all four provinces, in place of the original two-province mean-of-means. The two FX series correlate at r=0.99 (the revised series runs about 2% higher on average); the revised series remains I(1), so the same first-differencing approach applies unchanged.

## B.1 The core finding: unchanged or strengthened

The exile geographic-exposure depreciation response, this paper's most rigorously supported finding (Section 7.3), is not merely preserved under the revised FX construction. It strengthens, and a second dimension (severity) joins it as a robust, cross-taxonomy signal on the depreciation side:

| Series | Taxonomy | Original F (dep) | Original p (dep) | Revised F (dep) | Revised p (dep) |
|---|---|---|---|---|---|
| Geography | unrest | 2.860 ✓ | 0.00069 | 3.343 ✓ | 0.00009 |
| Geography | tight | 2.789 ✓ | 0.00092 | 3.293 ✓ | 0.00011 |
| Severity | unrest | 2.266 | 0.00755 | 2.710 ✓ | 0.00127 |
| Severity | tight | 1.657 | 0.06998 | 2.997 ✓ | 0.00039 |

*✓marks tests surviving the 24-test Bonferroni threshold (p < 0.00208). Depreciation side shown; on the appreciation side, severity remains non-significant under the revised construction in both taxonomies, as does geography under unrest: the one exception is geography under the tight taxonomy, whose appreciation response newly clears the same bar under the revised construction (F=2.75, p=0.00107), where it did not originally (p=0.239).*

Severity's newly-significant result is not an artifact of the revision. It was already the second-strongest individual dimension in the original analysis (p=0.00755 under unrest, just short of the 24-test bar), and the revised FX construction pushes it comfortably over that bar in both taxonomies. The lag-sensitivity check (Section 7.4) tells the same story: exile geography's depreciation response survives every control-lag specification (K=1-4, both taxonomies) under the revised series just as it did originally, and more strongly; p-values range 0.00001-0.00034 under the revision, compared to 0.00069-0.00149 originally.

## B.2 Everything else, in brief

The remaining checks in this paper were re-estimated the same way, with no change in their conclusions. The multiple-testing correction (Section 7.1) still finds zero Bonferroni or Benjamini-Hochberg survivors among the 312 individual-horizon tests. The alternative-weighting checks (Section 7.6) tell the same story as before: PCA-based reweighting still does not produce a significant index-level result for exile, and leave-one-dimension-out composites still weaken when geography or severity is dropped and strengthen when a weaker dimension is dropped, though that pattern is slightly less clean for the persistence-drop variant than it was originally. A nonlinearity check (splitting each shock into a baseline linear slope plus an extra slope for large shocks) again finds no significant threshold effect for any series. A pseudo-out-of-sample forecasting exercise (rolling and expanding windows, one-week and one-month horizons, with a Clark-West test appropriate for comparing nested models) again finds that adding the exchange rate to an autoregressive model of CSISI or ACLED does not improve genuine forecasts, the same null result as under the original FX construction. The one check that does not replicate is the ACLED cross-check (Section 7.5): under the revised FX construction, neither ACLED series' joint Wald test clears its 4-test Bonferroni bar (frequency: $F=1.37$, $p=0.175$; Δgeography: $F=1.66$, $p=0.070$; both cleared it originally, at $p=0.0072$ and $p=0.0073$ respectively). Given the core in-sample finding strengthens elsewhere under this same revision, this reads as a genuine point of sensitivity specific to the independent-data cross-check, rather than evidence against the main result.

# ANNEX C: FX VOLATILITY AS AN ADDITIONAL CHANNEL

Everything so far treats the exchange rate's effect as operating through the sign and size of its weekly move, a depreciation or appreciation shock of a given magnitude. This annex asks a different question: independent of which way the rate moves, does how much it has been moving lately (its recent volatility) carry additional information? A single combined specification adds an 8-week trailing volatility measure (the standard deviation of the weekly Δlog(FX) series over the preceding 8 weeks; how turbulent the exchange rate has been recently, not this week's own move) alongside the usual depreciation and appreciation shocks, so all three effects are estimated together, each controlling for the other two:

$$y_{t+h} = a_h + b_h^{pos} \cdot pos\_shock_t + b_h^{neg} \cdot neg\_shock_t + b_h^{vol} \cdot rolling\_vol_t + g_h \cdot D11J_t + \Sigma\ controls + e_{t+h}$$

run for exile geography (unrest taxonomy), the series with the paper's most robust FX finding, K=2 control lags, D11J dummy, HAC/Newey-West SEs, horizons h=0-12, otherwise identical to the specification in Section 5.2.

***Figure 12:*** *Exile geography's local-projection response to depreciation, appreciation, and FX volatility, estimated jointly. Shaded bands are 90% and 95% confidence intervals (HAC/Newey-West).*

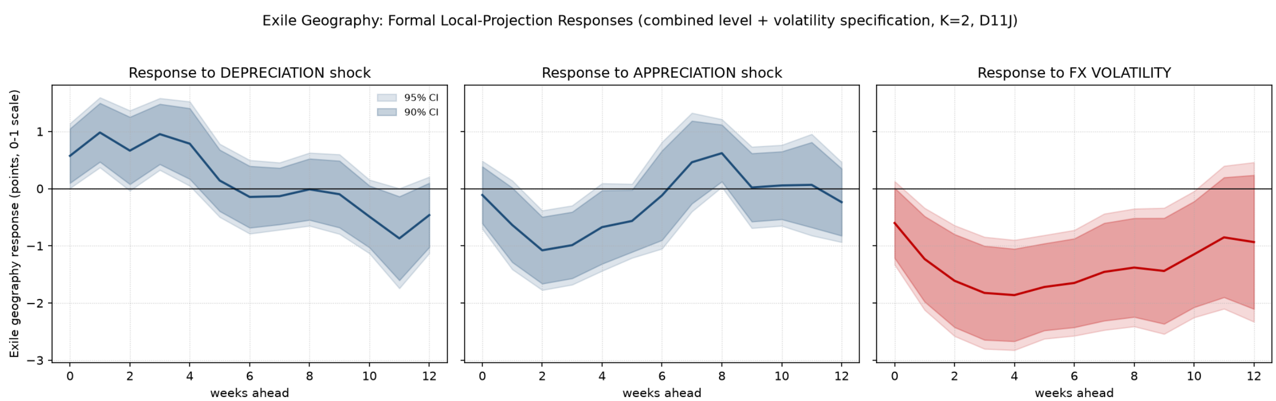


The depreciation and appreciation responses (left, middle) look much as they do throughout this paper: a positive depreciation response concentrated in the first four weeks (F=3.25, p=0.00014, jointly significant across all 13 horizons) and a later, appreciation-side reversal (F=2.45, p=0.0036). This confirms that adding volatility to the specification does not crowd out or absorb the depreciation/appreciation finding established in Sections 6-7. The third panel is new: exile geography's response to volatility is negative at every one of the 13 horizons, deepest around weeks 3-4, and jointly significant (F=3.94).

What might explain this pattern? The decline is concentrated in geographic resolution rather than coverage volume: raw event and article counts are unaffected by volatility, while the count of geolocatable events and their breadth, population, and territorial exposure all fall, with a weaker, unconfirmed hint of recentralization toward Havana. This is consistent with a “verification under uncertainty” channel documented in the protest-event and crisis-reporting literatures, where a noisier information environment degrades confirmation of event-level detail, particularly location, well before it reduces the volume of reporting (Earl et al. 2004). FX volatility in Cuba's informal market often coincides with broader logistical strain (fuel shortages, blackouts, transport disruption) that would plausibly weigh most heavily on the thin, geographically dispersed stringer and citizen-source networks exile outlets rely on outside Havana, while leaving Havana-based reporting comparatively unaffected. A related possibility is source self-protection: under conditions read locally as heightened instability, citizen-sources may become more willing to report an event but less willing to specify where it occurred. These channels are not mutually exclusive, and none of them requires FX volatility to act as an independent cause rather than a visible marker of a broader disorder that is itself driving the decline in geographic precision.

**A validation note:** this specification has more parameters relative to the available weekly clusters (~277-280) than any other in this paper, and a cluster-robust joint Wald test can become unreliable in exactly that regime. It is known to over-reject when the number of joint restrictions is large relative to the number of clusters. Rather than take the volatility panel's asymptotic p-value ($6.96\times10^{-6}$) at face value, it was validated with a permutation test: the volatility series was randomly reshuffled across weeks 999 times (breaking any real relationship to future CSISI while leaving its distribution unchanged), and the test was re-run on each shuffle. The real result ($F=3.94$) comfortably exceeds nearly the entire permutation distribution (mean $F=1.17$, maximum $F=3.55$ across all 999 shuffles), giving a permutation p-value of 0.0010, clearing even a conservative Bonferroni bar sized for this check ($0.05/12=0.00417$). The depreciation and appreciation panels, by contrast, replicate the same specification used throughout the rest of this paper and are not subject to this concern.

# REFERENCES


Arezki, R., & Brückner, M. (2011). Food prices and political instability. IMF Working Paper 11/62. https://www.elibrary.imf.org/view/journals/001/2011/062/001.2011.issue-062-en.xml

ACLED. (n.d.). ACLED *methodology*. ACLED. https://acleddata.com/sites/default/files/wp-content-archive/uploads/dlm_uploads/2019/04/Methodology-Overview_FINAL.pdf

Aron, J., Macdonald, R., & Muellbauer, J. (2014). Exchange rate pass-through in developing and emerging markets: A survey of conceptual, methodological and policy issues, and selected empirical findings. Journal of Development Studies, 50(1), 101-143. https://www.tandfonline.com/doi/full/10.1080/00220388.2013.847180

Barreiro Palafox, M., & Dassinger, R. (2025). *From blackouts to battered markets: Energy deficits and exchange rate dynamics in Cuba* [Conference paper]. Association for the Study of the Cuban Economy Annual Conference. DevTech Systems, Inc. https://devtechsys.com/wp-content/uploads/2025/11/FROM-BLACKOUTS-TO-BATTERED-MARKETS-MBP-RD.pdf

Bellemare, M. F. (2014). Rising food prices, food price volatility, and social unrest. American Journal of Agricultural Economics, 97(1), 1-21. https://doi.org/10.1093/ajae/aau038

CiberCuba. (2026a, September 13). “August leaves alarming figures in Cuba: Record homicides, five deaths in custody, and hundreds of repressive incidents.” CiberCuba. https://en.cibercuba.com/noticias/2026-09-13-u1-e207888-s27061-nid340235-agosto-deja-cifras-alarmantes-cuba-record-homicidios#google_vignette

CiberCuba. (2026b, September 15). Cuba is issuing 10,000 and 20,000 peso bills starting this September 16. CiberCuba. https://en.cibercuba.com/noticias/2026-09-16-u1-e197721-s27061-nid340458-cuba-pone-circulacion-billetes-10000-20000-pesos

CiberCuba. (2026c, September 15). Is Cuba headed towards full dollarization? Economist analyzes the future of the peso. CiberCuba. https://en.cibercuba.com/noticias/2026-09-16-u1-e135253-s27061-nid340449-cuba-hacia-dolarizacion-total-economista-analiza

DevTech Systems, Inc. (n.d.). The actual Cuban exchange rate dashboard. https://devtechsys.com/projects/the-actual-cuban-exchange-rate-dashboard/

Earl, Jennifer, Andrew Martin, John D. McCarthy, and Sarah A. Soule. 2004. “The Use of Newspaper Data in the Study of Collective Action.” *Annual Review of Sociology* 30: 65–80. https://www.jstor.org/stable/29737685

Frankel, J. A. (2005). Contractionary currency crashes in developing countries. IMF Staff Papers, 52(2), 149-192. https://www.imf.org/external/pubs/ft/staffp/2005/02/pdf/frankel.pdf

Frieden, J. A. (2015). Currency Politics: The Political Economy of Exchange Rate Policy. Princeton University Press. https://press.princeton.edu/books/paperback/9780691173849/currency-politics

García Hernández, F. (2023). Las protestas en Cuba del 11 de julio: causas de una rebelión anunciada. In F. Parra & M. Cristeche (Eds.), *Desigualdades, resistencias y derechos en Latinoamérica y el Caribe* (pp. 251–268). Religación Press. https://doi.org/10.46652/religacionpress.22.c40

GDELT Project. (2015). *The GDELT event database: Data format codebook V2.0*. https://data.gdeltproject.org/documentation/GDELT-Event_Codebook-V2.0.pdf

Ismail, Z. (2021). Interaction between food prices and political instability. K4D Helpdesk Report, Institute of Development Studies. https://hdl.handle.net/20.500.12413/16721

Goldstein, J. S. (1992). A conflict-cooperation scale for WEIS events data. *Journal of Conflict Resolution, 36*(2), 369–385. https://doi.org/10.1177/0022002792036002007

Hammond, J., & Weidmann, N. B. (2014). Using machine-coded event data for the micro-level study of political violence. *Research & Politics, 1*(2), 1–8. https://doi.org/10.1177/2053168014539924

LeoGrande, W. M. (2021). Frustration boils over: The politics of July 11. In The Road Ahead: Cuba After the July 11 Protests. American University. https://www.american.edu/centers/latin-american-latino-studies/cuba-after-the-july-11-protests-leogrande.cfm

LeoGrande, W. M., & Hershberg, E. (2023). Cuba emerging from COVID. Revista de Ciencia Política, 43(2), 255–277. http://dx.doi.org/10.4067/s0718-090x2023005000110

Luis, L. R. (2023). Inflation and macro policies in Cuba. Association for the Study of the Cuban Economy. https://ascecubadatabase.org/17428-2

Mesa-Lago, C. (2021). La unificación monetaria y cambiaria en Cuba: normas, efectos, obstáculos y perspectivas. Real Instituto Elcano, Documento de Trabajo 2/2021. https://www.realinstitutoelcano.org/documento-de-trabajo/la-unificacion-monetaria-y-cambiaria-en-cuba-normas-efectos-obstaculos-y-perspectivas

Miguel, E., Satyanath, S., & Sergenti, E. (2004). Economic shocks and civil conflict: An instrumental variables approach. Journal of Political Economy, 112(4), 725-753. https://doi.org/10.1086/421174

OECD, European Union, & European Commission, Joint Research Centre. (2008). *Handbook on constructing composite indicators: Methodology and user guide*. OECD Publishing. https://www.oecd.org/en/publications/handbook-on-constructing-composite-indicators-methodology-and-user-guide_9789264043466-en.html

Ponticelli, J., & Voth, H.-J. (2020). Austerity and anarchy: Budget cuts and social unrest in Europe, 1919-2008. Journal of Comparative Economics, 48(1), 1-19. https://doi.org/10.1016/j.jce.2019.09.007

Raleigh, C., Kishi, R., & Linke, A. (2023). Political instability patterns are obscured by conflict dataset scope conditions, sources, and coding choices. *Humanities and Social Sciences Communications, 10*, Article 74. https://www.nature.com/articles/s41599-023-01559-4

Vidal, P., & Luis, L. R. (2024). Cuba's monetary reform and triple-digit inflation*. Latin American Research Review*, 59(2), 274–291. https://doi.org/10.1017/lar.2023.59

Walter, S. (2009). The limits and rewards of political opportunism: How electoral timing affects the outcome of currency crises. *European Journal of Political Research, 48*(3), 367–396. https://doi.org/10.1111/j.1475-6765.2008.00837.x

Weinberg, J., & Bakker, R. (2015). Let them eat cake: Food prices, domestic policy and social unrest. Conflict Management and Peace Science, 32(3), 309-326. https://doi.org/10.1177/0738894214532411