Attention and the Economic Response to Immigration Enforcement

Greg C. Wright

University of California, Merced
gwright4@ucmerced.edu

July 8, 2026

Abstract

Does the community response to immigration enforcement depend on the operation itself or on the public attention it receives? Across 51 single-worksite ICE events in 26 states (2018–2026), a per-event difference-in-differences compares heavily Latino-immigrant neighborhoods with similar nearby ones. Events that receive more Spanish-language news coverage see larger foot-traffic declines, while arrest counts explain essentially none of the difference across events. Applying an instrumental-variables design that exploits competing news shocks from major foreign disasters, I estimate that one standard deviation of additional news coverage reduces foot traffic in these neighborhoods by about 2.6 percent. English-language news and search show the same pattern. Civic and institutional foot traffic falls more than commercial activity.

Keywords: public attention, immigration enforcement, chilling effects, foot traffic.

JEL Codes: J15, R23, K37.

Acknowledgments: I thank the UC Merced Community and Labor Center for support and Albert Kochaphum for helpful conversations. In preparing this manuscript I used OpenAI's ChatGPT and Codex to assist with coding, analysis checks, proofreading, and manuscript editing.

1 Introduction

Immigration enforcement can affect the lives of individuals beyond those who are detained. In part, this is because the community response may depend as much on the information environment surrounding an event as on the event itself. A workplace raid that goes largely unnoticed may have little effect beyond the worksite, whereas a smaller raid that becomes highly visible may lead residents to avoid public places. A long literature in economics and sociology documents these "chilling effects," which can lead to reduced use of healthcare, schooling, and public benefits due to fear of detention [Watson, 2014, Asad, 2023]. The intensification of U.S. Immigration and Customs Enforcement (ICE) activity under the second Trump administration provides a real-time test. For example, Lester et al. [2026] document an 8–10 percent foot-traffic decline and a 20–25 percent spending decline in the most heavily Latin American–born neighborhoods of Los Angeles–Orange County in the weeks following ICE's May 14, 2025 announcement of a major regional enforcement operation.

In this paper I ask whether the observed chilling effects are due to the operations themselves or to the public attention they receive. If the scale of an operation is central to its effect, arrest counts should help explain which events produce larger responses. On the other hand, if the broader *perception* of risk is what matters, the response will depend on whether news of the event reaches people who were not directly targeted. A raid affects both margins at once, detaining a limited number of individuals while also generating news coverage, social-media activity, rumor, and protest.

I exploit cross-event variation to separate these channels. The sample comprises 51 ICE enforcement events from April 2018 through January 2026 across 26 states, spanning a small cluster of high-profile 2018–2019 worksite actions and a much larger cluster of 2025–2026 operations. For each event I estimate a distributed-lag difference-in-differences (DiD) specification on Advan cellphone foot-traffic data, comparing Census Block Groups (CBGs) in the top decile of Latin-American foreign-born (FB) share to those in the bottom decile.[1] I then extract a summary estimate of the average response. Because an average can hide very different event-level responses, I first measure how much the true response varies. Roughly 70 percent of the cross-event variation reflects genuine differences rather than estimation noise, and the predicted foot-traffic response to a new event ranges from roughly a 5 percent decline to a 5 percent increase. I then ask which features of the event predict where its response falls in that range.

The strongest predictor of foot traffic is the public attention garnered by the event, not

[1] This is the specification Lester et al. [2026] use for their single LA event, which I implement across each event in the sample for comparison.

its operational scale. I exploit post-event spikes in the share of Spanish-language news stories mentioning "ICE," the agency name most salient to the population at risk. This measure is constructed from a collection of US Spanish-language outlets with daily coverage across the sample period. Events that receive more Spanish-language news coverage are associated with larger declines in foot traffic, and the same pattern appears in English-language news and Google search data. In a regression across events that weights each event by the precision of its estimate, the Spanish-language news coverage spike accounts for roughly 30 percent of the true cross-event variation in responses, while arrest counts account for essentially none, consistent with the attention mechanism.

Timing evidence and an instrumental-variables design support a causal reading. First, the effect of news coverage on foot traffic is absent in every pre-event week and emerges only after the event. The early foot-traffic response also does not predict later news coverage, suggesting the effect is not due to reverse causality. Second, I exploit news of major foreign disasters, which crowds out public attention to enforcement actions, as an instrument for the news coverage spike, with out-of-state U.S. disasters providing a complementary version. These news shocks strongly predict the amount of coverage given to enforcement events and occur independently of the operation's timing.[2] The instrumented estimates imply that one standard deviation of additional news coverage reduces foot traffic by about 2.6 percent, slightly larger than the OLS association.

The pattern is robust to several further checks. It holds when attention and arrests are included together, when noisier event estimates are down-weighted, when pre-event media attention is accounted for, and when I estimate within-event changes rather than levels. Balance checks show that coverage is unrelated to pre-event foot-traffic trends or placebo responses. The effect holds, and is if anything stronger, when the sample is restricted to the 2025–2026 enforcement wave alone, so it does not merely reflect differences between the two waves.

National Google Trends and national Spanish-language news both predict the local foot-traffic decline, while metro-level search and local Spanish-language news add little once national attention is accounted for. The response is broad, extending beyond the immediate worksite and the raided industry. I decompose the response by venue and find that foot traffic around schools, parks, and museums falls more than at commercial sites such as retail and food services, a pattern more consistent with a broad withdrawal from public life than a localized demand shock.[3] The economic footprint of visible enforcement is therefore not

[2]Since competing news shifts attention only for events whose coverage is sensitive to the news cycle, the instrument's estimate applies mainly to mid-visibility events, not the most visible operations.

[3]Religious organizations and hospitals are not present in the data. See Section 4.7.

limited to lost commercial activity but includes reduced participation in civic life.

A broader economics literature uses interior-enforcement intensity and federal programs, especially Secure Communities, to estimate effects of immigration enforcement on safety-net participation, family decisions, childcare, healthcare, and labor markets [Watson, 2014, Amuedo-Dorantes et al., 2018, 2020, Alsan and Yang, 2024, Ali et al., 2024, Herring and Barnow, 2025, East et al., 2023]. A related sociology literature describes the institutional-avoidance strategies of mixed-status families [Asad, 2023]. Recent work studies the current wave. Amuedo-Dorantes and Antman [2022] link removals and raid awareness to lower immigrant labor-market engagement, Ciancio and García-Jimeno [2026] study consumption responses to learned enforcement risk, and De Balanzó et al. [2026], Hernandez [2026], and Cox and East [2026] estimate economic effects of the 2025 wave using spending, mobility, and labor-market data. These papers show that enforcement moves economic activity when risk becomes visible or learned. Two features of these designs nevertheless leave the role of attention, separate from enforcement, unmeasured. Variation built on program rollouts or enforcement intensity bundles any change in attention together with the change in enforcement itself, and single-episode studies hold the event fixed but have no variation in event visibility. I instead hold the enforcement event fixed and ask why two raids of similar operational severity produce very different neighborhood responses, varying the event's public visibility rather than the event itself.

These results also add to evidence on how public attention, more than operational scale, shapes the response to risk. A related COVID-mobility literature [Goolsbee and Syverson, 2021, Baker et al., 2020, Chetty et al., 2024] shows that perceptions of risk changed individual behavior independently of the formal COVID mandates. The case here is similar. An immigration enforcement operation produces both a worksite disruption and public attention, but the community response tracks the attention. The same logic may extend to other salient but partly hidden risks whose economic footprint depends on how widely they are noticed. Where prior work instruments or exploits quasi-experimental variation in enforcement itself, I hold the event fixed and instrument its visibility instead, adapting the competing-news logic of Eisensee and Strömberg [2007] to isolate the attention component of a specific event. To my knowledge, this approach is new to the immigration literature. The event-by-event design also changes what can be learned from a null result. Because the 51 events are estimated separately and then modeled as a distribution, the analysis can distinguish a world in which enforcement rarely moves behavior from one in which it moves behavior strongly but only where it is widely noticed.

2 Data

The analysis combines the event inventory with outcome, exposure, and attention data. Outcome and exposure data include cellphone foot traffic, SafeGraph consumer spending, and American Community Survey (ACS) demographics. Attention and news-shock data include Google Trends, Mediacloud US Spanish-language and GDELT 2.0 English-language news volume, Global Disaster Alert and Coordination System (GDACS) disaster alerts, and Nielsen Designated Market Area (DMA) definitions.

2.1 Event inventory

I compile a list of single-worksite ICE enforcement events from news sources, agency press releases, and litigation records. The inventory covers two distinct enforcement waves from April 2018 through August 2019 and January 2025 through January 2026 (the second term's worksite operations).[4] It restricts attention to events that satisfy four criteria. First, the event date must fall within these two windows. Second, the event must correspond to a specific worksite or commercial address. Third, the worksite must be in one of 26 states for which CBG-level Latin-American foreign-born data from the ACS are available. Finally, the worksite latitude-longitude must be geocoded either to the named business address or, when that is not feasible, to the city or county centroid.[5] Two classes of candidate events fall outside these criteria. Multi-day or metro-wide operations, such as National Guard deployments and sustained urban surges, lack the single discrete enforcement date at a single worksite that the event-time design requires. I also remove actions that were not ICE-led, including one Border Patrol worksite operation and one FBI labor-trafficking investigation in which the workers were treated as victims rather than removal targets, so the sample is cleanly ICE worksite enforcement. These screens identify 55 single-worksite ICE-led events. Four have insufficient point-of-interest (POI) coverage within the event catchment (the 50 km radius defined in Section 2.2) to estimate the event-study specification, leaving a main analysis sample of 51 events spanning 26 states. The Online Appendix summarizes the event-level variables by enforcement wave. The 2018–2019 events report more arrests on average but generate much smaller news spikes than the 2025–2026 events, which motivates the wave-robustness check in Section 4.4.

[4]The first Trump term's worksite operations included Bean Station TN, Salem OH Fresh Mark, Mt. Pleasant IA Precast, Allen TX CVE Technology, Sumner TX Load Trail, and the August 2019 Mississippi poultry operation, which hit seven plants across five companies on a single day and which is included as a single Morton-MS-anchored compound event.

[5]Four events are explicitly coded as city-centroid locations, one additional event uses the earlier city-query fallback, and one event uses a county-level fallback.

A publicly compiled inventory could be biased toward visible raids. No complete administrative list of single-worksite ICE actions is publicly available for these windows, so the inventory draws on several sources rather than national news alone. These include agency releases, litigation records, immigration-policy trackers, state and local reporting, and prior compilations. The result includes a set of events that received little attention. Of the 55 events, 32 are coded medium or low publicity and 29 have an English-GDELT news spike at or below 0.01 (Section 2.6), and the four POI-coverage exclusions include three medium-publicity events and one high-publicity event. These facts do not establish a complete universe of obscure raids, but they make it unlikely that the effect of news coverage is mechanically generated by retaining only highly visible events. The Online Appendix reports more information about the site selection.

2.2 Foot traffic

I obtain weekly visitor counts at the point-of-interest level from Advan Research's *Foot Traffic/Weekly Patterns Plus* dataset [Advan Research, 2025], derived from anonymized opt-in mobile devices. Each POI is geocoded to a latitude and longitude, assigned a North American Industry Classification System (NAICS) code, and assigned to its CBG. For each event, all Advan POIs within 50 km of the worksite anchor are collected and assembled into a POI $\times$ week panel spanning ± 12 weeks of the event date.

The resulting multi-event panel contains approximately 69 million POI-week observations and 3.3 million distinct POI–event pairs across 1.5 million unique POIs. Within each event's catchment, POI foot traffic is aggregated to CBG-week totals and restricted to CBGs containing at least 10 POIs. A CBG's POI count is computed from POIs that report foot traffic in at least half of the event's sample weeks, so short-lived or intermittently reported POIs do not determine which CBGs enter. This restriction concentrates the analysis on retail-active CBGs.

2.3 Consumer spending

Spending is measured with SafeGraph's *SpendPatterns* dataset [SafeGraph, 2022], which reports monthly POI-level aggregates of consumer spending from a large opt-in financial-institution panel. Each row contains an array of daily spending values that are aggregated to weeks, producing a POI $\times$ week spending panel comparable to the foot-traffic data. POI-week values are then aggregated to CBG-week using the same catchment definitions as the foot-traffic panel.

2.4 ACS exposure measure

Treatment is defined as CBGs in the top decile of share of residents born in Latin America ("high exposure"), with bottom-decile CBGs as the comparison group. The 2018–2022 5-year ACS provides the Latin-American foreign-born share, computed at the tract level and allocated to all CBGs within each tract.

The deciles are computed within each event's catchment, so the threshold for "top decile" varies across events. In heavily-Latino metros, the top-decile cutoff is roughly 32 percent foreign-born from Latin America, while in less-Latino metros the threshold is much lower in absolute terms. This within-event normalization means the cross-event comparison is between "most-Latino" and "least-Latino" CBGs in each respective metro.

2.5 Google Trends attention

For each event, attention is measured using weekly Google Trends search interest for three keywords ("ICE raid", "redada", "deportación") at the U.S. national, event-state, and Nielsen Designated Market Area (DMA) levels. Because Trends rescales each query within a geography and timeframe, within-geography pre-vs-post changes are used. The baseline is the four-week average over $\tau \in [-8, -5]$ relative to the event Monday. Trends are restricted to the 2024–2026 window because cross-era rescaling is noisy for these low-baseline terms.

2.6 News-volume series

Mediacloud US Spanish-language news. The primary news-volume series measures attention in the Spanish-language news environment the affected population consumes, on a daily scale consistent across both enforcement waves. I query the Mediacloud Online News Archive [Media Cloud, 2025] for daily story counts in its curated US Spanish-language collection of 407 sources, including national, state, and local outlets. The headline query is the noun "ICE" restricted to this collection. Counts are normalized by the daily total of all stories in the collection to produce a share-of-coverage time series that is comparable across years. Complementary Spanish keywords are also pulled, and a state-local Spanish-news series is constructed by restricting the same query to sources tagged as published in the event's home state.

GDELT English-language US news. As a corroborating measure on a different platform, language, and methodology, I query the GDELT 2.0 Document API [Leetaru and Schrodt, 2013] for daily article counts matching "`ICE raid`" in US English-language sources.

This produces a national daily series, in normalized "Volume Intensity" units, spanning January 2018 through March 2026. GDELT closely tracks the Mediacloud series at the per-event spike level.

2.7 Competing-news instruments

The preferred news-shock instrument comes from the Global Disaster Alert and Coordination System (GDACS) event feed [Global Disaster Awareness and Coordination System, 2026]. The most severe ("red-alert") sudden-onset disasters (earthquakes, tropical cyclones, floods, volcanic eruptions, wildfires) outside the United States, Latin America, and the Caribbean are obtained. The Latin America and Caribbean exclusion is deliberate, since homeland shocks could affect Latino immigrant communities directly through family ties, remittances, or Spanish-language news demand rather than only through news-cycle competition. For each ICE event, the instrument counts days in the post-event window overlapping the first seven days of such an alert.

A complementary news-shock measure is also constructed using a hand-coded calendar of major U.S. natural disasters in states other than the ICE event state. The foreign- and U.S.-disaster instruments identify different margins of news-cycle access, so the instrumental-variables (IV) estimates are interpreted as local effects for events whose attention is moved by the corresponding source of competing news shocks.

3 Empirical Strategy

3.1 Per-event distributed-lag DiD

For each event e in the 51-event sample (single-worksite ICE enforcement events from April 2018 through January 2026, see Section 2.1), the same distributed-lag difference-in-differences model is estimated:

$$\log(1 + y^e_{c,t}) = \alpha^e_c + \lambda^e_t + \sum_{\tau \neq -1} \gamma^e_\tau \left(\text{HighFB}^e_c \times \mathbf{1}[t - T^e_0 = \tau] \right) + \varepsilon^e_{c,t}, \tag{1}$$

where $y^e_{c,t}$ is weekly visitor count (foot traffic) at CBG c in week t, restricted to the catchment of event e. T^e_0 is the Monday of the week containing event e's date. The dummy HighFB^e_c equals one for CBGs in the top decile of Latin-American foreign-born share within event e's catchment and zero for CBGs in the bottom decile (middle deciles are dropped). α^e_c and λ^e_t are CBG and calendar-week fixed effects, both event-specific because Equation (1)

is estimated separately for each event. Standard errors are clustered at the CBG level. Because the outcome is a count, all first-step estimates are also re-estimated by Poisson pseudo-maximum likelihood on the raw visit counts, which requires no log transformation (Table C3).

The event-time window is $\tau \in [-8, +8]$ weeks and all seventeen weeks enter the estimation. The omitted category is $\tau = -1$, the last full pre-event week, so each γ_τ^e measures the high-FB versus low-FB foot-traffic gap in week τ relative to that reference week rather than a comparison of the event week to all other weeks. The per-event summary estimate β_{late}^e is the average of γ_τ^e over the late window $\tau \in [4, 8]$. Its standard error is the standard error of that linear combination, computed from the CBG-clustered covariance matrix. Because most events do not fall on a Monday, the $\tau = 0$ week mixes pre- and post-event days.[6]

3.2 Identifying variation

The design has two steps. First, Equation (1) uses within-event variation between high-FB and low-FB CBGs in the same catchment. CBG fixed effects absorb time-invariant differences in population, retail density, and baseline foot traffic. Calendar-week fixed effects absorb shocks common to all CBGs in the catchment in a given week. The event estimate is therefore the differential foot-traffic response of high-FB versus low-FB CBGs in the same metro and calendar weeks. Second, the 51 event estimates test whether differences in the news coverage spike across events predict differences in the estimated response.

The design does not require raids to be unanticipated, and some initiatives are clearly partially anticipated. The initial question is whether the news coverage generated by each event explains why otherwise similar raids produce different responses. Section 4.3 implements an instrumental-variables design to isolate plausibly exogenous variation in coverage from the surrounding news cycle, and Section 4.2 adds pre-event coverage as a control.

3.3 News-coverage and event severity measures

To test whether the chilling response is driven by news coverage or by operational scale, two categories of cross-event measures are constructed. The headline news-coverage measure is the post-event spike in the Mediacloud US Spanish-language news-volume series for the query "ICE" (Section 2.6), measured as the share of stories in the 407-source US Spanish-language collection that mention the term. The spike is computed as the peak over $\tau \in [0, 8]$

[6] Re-anchoring $\tau = 0$ to the first Monday on or after the event date, or dropping the partial event week entirely, leaves the conclusion unchanged. The late-window estimate excludes the partial week by construction, and the effect of news coverage below moves from -0.018 to -0.016 log points per standard deviation under re-anchoring (Online Appendix).

minus the baseline over $\tau \in [-8, -5]$.

The headline measure is complemented with three corroborating measures: the analogous GDELT national US English-language news-volume spike for "ICE raid," the national Google Trends spike for "ICE raid," and the DMA-level Google Trends spike for the local media market in which the event occurred. The two Trends measures are available on the 2024–2026 subsample and are defined as within-geography pre-vs-post changes, since Google Trends rescales each query within a geography and timeframe.

To capture the severity of the enforcement event, I use two measures of underlying operational scale. These are the count of arrests reported by ICE at the worksite, Arrests_e, and the same measure in log form, $\log(1 + \text{Arrests}_e)$, to accommodate the right-skewed distribution.[7] The arrests measure has coverage of 49 of 51 fitted events.[8]

3.4 Horse-race specification

To assess whether attention rather than severity predicts chilling, a cross-event ordinary least squares (OLS) regression relates the per-event late-window coefficient to a measure of attention and, in the severity specification, the log-arrests control:

$$\beta^e_{\text{late}} \;=\; \pi_0 \;+\; \pi_S \,\text{Spike}_e \;+\; \pi_A \,\log(1 + \text{Arrests}_e) \;+\; u_e, \tag{2}$$

where Spike_e is the post-event spike in one of the four news-coverage measures defined above. Heteroskedasticity-robust (HC1) standard errors are reported. Table 1 reports the headline per-standard-deviation estimates alongside the instrumented versions. Table C2 reports the full per-unit specification, including a standalone arrests specification and single-regressor specifications using English-GDELT and national Google Trends as corroborating measures. All specifications use the same event windows. The Trends specification uses the 2024–2026 subsample. Comparing the single-regressor estimates across sources shows whether the same effect of news coverage appears outside the headline Spanish-language measure.

3.5 News-shock instruments

The OLS horse race treats realized post-event attention as the explanatory variable. I use predetermined competing-news shocks as instruments for attention generated by the

[7] No sampled event has zero reported arrests (the minimum is 2), so the transformation is innocuous. The attention estimates are identical and the arrests coefficient is essentially unchanged using $\log(\text{Arrests}_e)$ instead.

[8] The two events with missing arrest counts in the inventory are Newark NJ (an unnamed business) and Baldwin County AL (an unnamed construction site).

surrounding news cycle. The preferred instrument uses major foreign disaster news shocks. It counts days in the post-event window that overlap with the first seven days of a GDACS red-alert sudden disaster outside the United States, Latin America, and the Caribbean. The exclusion of Latin America and the Caribbean is intended to remove shocks that could directly affect Latino immigrant communities through homeland ties rather than through news-cycle competition only. The identifying assumption is that these shocks move media coverage of the enforcement action without directly changing the relative foot-traffic response in high-FB versus low-FB neighborhoods. Section 4.3 presents the IV estimates; Appendix B and the Online Appendix report the corresponding balance and robustness checks.

Out-of-state U.S. disasters also serve as a complementary instrument. This measure captures domestic disaster news shocks in states other than the ICE event state. The foreign- and U.S.-disaster instruments identify different local margins of news-cycle access. The IV estimates should therefore be read as local effects for events whose public attention is shifted by a particular source of competing news shocks.

4 Results

The response is shown in event time first, pooling the 51 event panels so that the average path, its pre-trends, and the scaling of the response with news coverage are all visible. The cross-event analysis then presents the paper's central estimates, with OLS and instrumented effects of news coverage side by side, using competing news shocks from major foreign disasters to isolate variation in the attention an event receives. Timing and design-validity checks address reverse causality and design artifacts, and the remaining sections trace the geography, venues, and spending margins of the response.

4.1 The response in event time

Before estimating the cross-event effect of news coverage, Figure 1 shows when the response emerges, pooling the 51 event panels and tracing its path in event time. Panel A plots the level path γ_τ, the high-FB versus low-FB gap in weekly foot traffic for an event of average attention, and Panel B plots the interaction path δ_τ, the additional gap per standard deviation of the Spanish-language news coverage spike. Estimates are presented for the full sample as well as for the subsample of 23 events with no overlapping-catchment contamination and statistically flat pre-trends.

Three features stand out. First, the pooled level path shows little average response. The path γ_τ moves around zero after the event in both samples. Second, the pre-period path

for the 23-event subsample is flat, while the full sample has a downward-drifting pre-trend, suggesting the subsample estimates are more reliable.[9] Third, the interaction estimates in Panel B are flat and near zero through the entire pre-period in both samples and turn negative only after the event, reaching −0.02 log points per standard deviation of news coverage in the late window.

Figure 2 estimates the effect of news coverage separately at each event time, regressing the per-event coefficients on the news-coverage spike with each event weighted by the precision of its estimate. The effect of news coverage is indistinguishable from zero in every pre-event week and emerges only after the event, stabilizing near −0.02 log points per standard deviation over the late window.[10] The heterogeneity that the rest of the paper studies therefore appears after the enforcement event, and its alignment with news coverage has no counterpart before the event.

4.2 The effect of news coverage

The average neighborhood response to an enforcement event is close to zero, but this average masks substantial cross-event heterogeneity. After accounting for sampling error in the event-level estimates, about 70 percent of the observed cross-event variance remains as true variation in the underlying response, and the implied prediction interval for a new event runs from a roughly 5.5 percent decline to a 4.8 percent increase. Table C1 reports the full random-effects decomposition, and Figure C1 shows a forest plot of the 51 event estimates. The question is what explains where in that range an event falls.

Table 1 reports the paper's central estimates. Public attention strongly predicts the variation, and operational scale does not. Column (1) regresses each event estimate on the post-event Spanish-language ICE news spike, weighting each event by the precision of its first-step estimate (a random-effects meta-regression). The spike explains around 30 percent of the between-event variance. Columns (2) and (3) report the OLS counterpart, and columns (4) and (5) instrument the spike with foreign-disaster news shocks (Section 3.5). Across estimators, a one standard deviation larger news-coverage spike predicts a 1.8 to 2.7 percent larger foot-traffic decline, and the instrumented estimates are, if anything, slightly larger than OLS. Adding the log arrest count leaves the attention coefficient essentially unchanged, while the arrest coefficient is small and statistically insignificant in every specification.[11] The

[9] A formal sensitivity analysis of the pooled path to pre-trend violations appears in Appendix A.

[10] The dispersion evidence has the same timing. If the post-event heterogeneity were simply a continuation of pre-existing differences across event catchments, dispersion would look similar at matched distances before and after the reference week. Instead, dispersion is larger after the event at each matched distance (Appendix A).

[11] Arrests are a coarse proxy for operational scale. As a result, a small coefficient does not imply that scale

Figure 1: The pooled response in event time

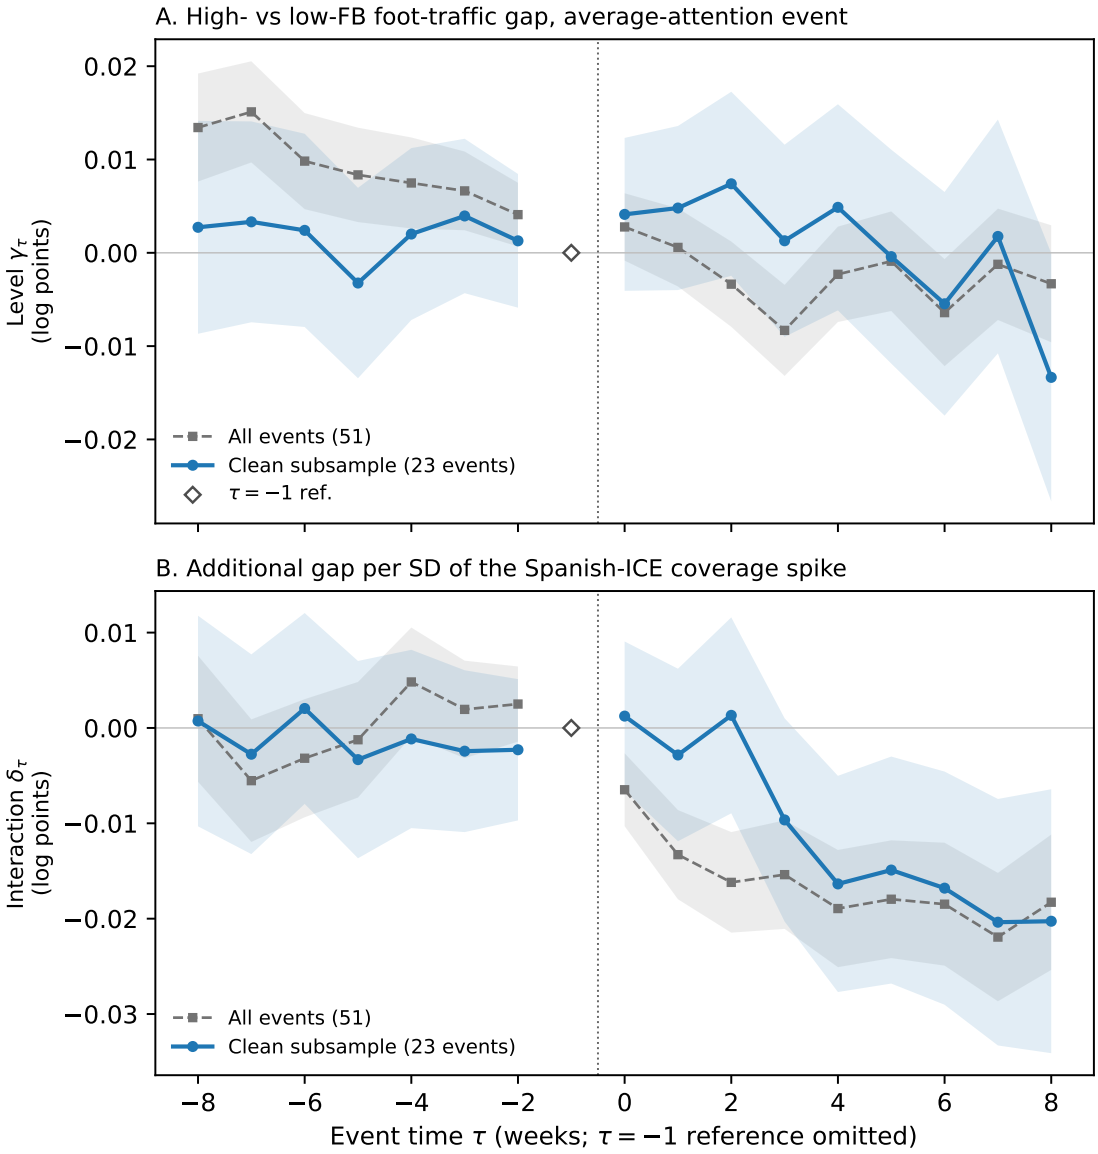

Notes. Pooled distributed-lag DiD with event × CBG and event × calendar-week fixed effects, standard errors clustered at the CBG level. High-FB (low-FB) CBGs are in the top (bottom) decile of Latin-American foreign-born share within the event's catchment. Panel A plots the high-FB versus low-FB foot-traffic gap γ_τ for an average-attention event. Panel B plots the additional gap δ_τ per standard deviation of the Spanish-language news-coverage spike, standardized at the event level on the full sample. The clean subsample is the 23 events with no overlapping-catchment contamination and flat pre-trends ($p > 0.10$). Shaded bands are 95 percent confidence intervals, and the reference week $\tau = -1$ is omitted and marked at zero.

sign of the event estimate is ordered by attention. Negative estimates concentrate among high-attention events, while positive estimates concentrate among low-attention events. The

does not matter, only that the dimension captured by arrest counts does not predict the response.

Figure 2: The effect of news coverage over event time

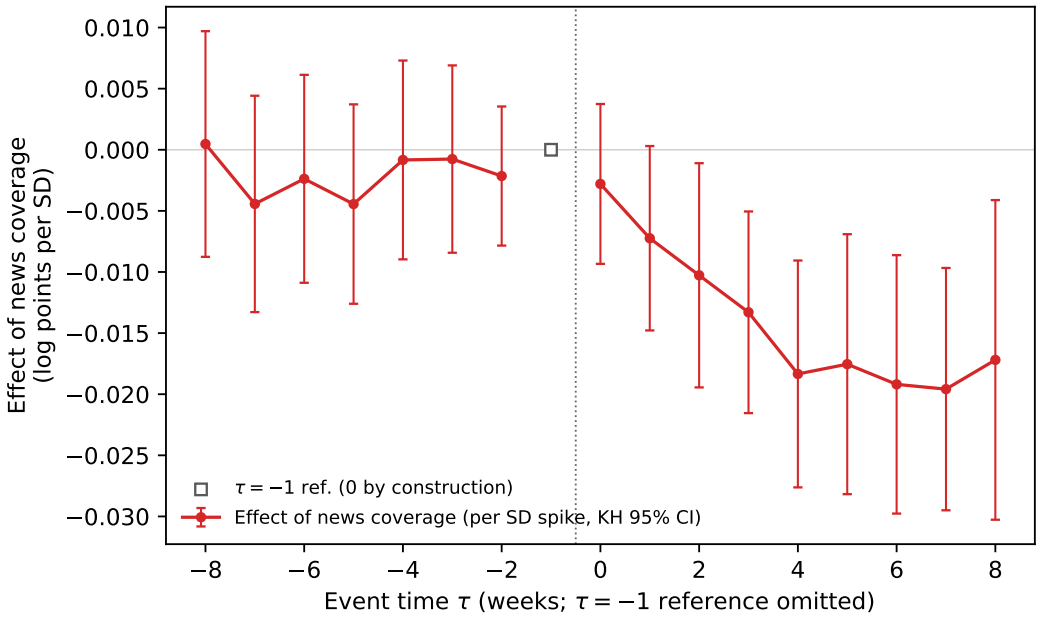

Notes. Each point is the coefficient from a precision-weighted regression of the 51 per-event event-study coefficients γ_τ^e on the standardized Spanish-language news-coverage spike, estimated separately at each event time τ, with Knapp–Hartung 95 percent confidence intervals (a small-sample-robust adjustment). Each event is weighted by the inverse of its coefficient's CBG-clustered sampling variance. The reference week $\tau = -1$ is zero for every event by construction and is marked separately. Estimates are pointwise across τ. The companion dispersion analysis appears in Appendix A and the underlying numbers in the Online Appendix.

positive estimates therefore arise where the evidence indicates little chilling, not a systematic increase in local activity after enforcement.

Table 1: The effect of news coverage on neighborhood foot traffic

	Precision-weighted	OLS		2SLS (news-shock IV)	
	(1)	(2)	(3)	(4)	(5)
Spanish-ICE spike (per SD)	-0.0184^{***}	-0.0222^{***}	-0.0241^{***}	-0.0256^{***}	-0.0267^{***}
	(0.0051)	(0.0062)	(0.0064)	(0.0082)	(0.0081)
log(1 + Arrests)			−0.0053		−0.0055
			(0.0042)		(0.0039)
First-stage F				52.7	51.2
True variance explained	30%				
Events	51	51	49	51	49

Notes: The dependent variable is the per-event late-window (weeks 4–8) foot-traffic estimate from Equation (1), in log points. The news coverage spike is the post-event peak in Mediacloud US Spanish-language ICE coverage, standardized to the 51-event sample (one SD = 0.024 in coverage-share units), so all columns are per standard deviation of coverage. Column (1) is a precision-weighted regression weighting each event by the inverse sampling variance of its estimate, with Knapp–Hartung standard errors. True variance explained is the share of between-event variance in true effects that the spike explains. Columns (2)–(3) are OLS with HC1 standard errors. Columns (4)–(5) instrument the spike with foreign disaster news pressure, the number of post-event days overlapping the first week of a GDACS red-alert disaster outside the United States, Latin America, and the Caribbean (Section 3.5), with robust standard errors. The arrests specifications have $N = 49$ because two events lack reported arrest counts. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The pattern is not specific to the headline measure, and it survives standard checks. The English-language GDELT news spike has the same sign and statistical significance and explains about 19 percent of the variance, and the US national Google Trends spike gives the same answer on the 2024–2026 subsample (Table C2). The effect survives controlling for operation type and sector and is robust to corrections for measurement error in arrests, making it unlikely that the effect of news coverage simply proxies for operational scale. Controlling for pre-event news attention in the four weeks before the raid does not explain away the post-event effect. Precision-weighted estimates, event-level bootstrap intervals, pre-event coefficients, and placebo estimates give the same conclusion. Estimating the first step by Poisson pseudo-maximum likelihood rather than in logs gives precision-weighted and OLS effects near −0.016 per standard deviation, with an instrumented estimate of the same magnitude but less precision (Table C3). The Online Appendix collects the remaining robustness tables, including an early-window table showing that the effect of news coverage builds over the first eight post-event weeks.

4.3 Instrumenting attention with news shocks

Realized news coverage is not randomly assigned. Features of an operation that make it salient to editors could also make it salient to residents, so the OLS association could overstate the causal effect of news coverage. This concern is addressed with an instrumental-variables strategy following the competing-news logic of Eisensee and Strömberg [2007]. In short, the idea is that predetermined national news-cycle conditions can shift attention to an enforcement event without altering the operation itself.

The preferred instrument uses major foreign disaster news shocks. It counts the number of days in the post-event window $\tau \in [0, 8]$ that fall within the first seven days of a GDACS red-alert sudden disaster outside the United States, Latin America, and the Caribbean. These disasters are predetermined with respect to the enforcement event and, to varying degrees, occupy national news attention. The outcome is a within-catchment differential between heavily Latino-immigrant neighborhoods and otherwise similar neighborhoods. A national news shock that moves attention equally in both groups nets out of this differential. The residual threat is therefore a home-region disaster that differentially affects the treated population. Latin America and the Caribbean are therefore excluded, since disasters there could directly affect Latino immigrant communities through family ties, remittances, or Spanish-language news demand. Out-of-state U.S. disasters are also reported as a complementary margin of domestic disaster news shocks. These instruments need not identify the same local average treatment effect. Each isolates events whose public attention is shifted by a particular source of competing news shocks.

Columns (4) and (5) of Table 1 report the second-stage estimates, and Table 2 reports the first stage. Foreign-disaster pressure has the expected crowd-out sign, with a first-stage F near 53. The second-stage estimates are similar to, and if anything slightly larger in magnitude than, the OLS estimates, consistent with attenuation from measurement error or with a larger effect for events whose coverage is movable by competing news. The overidentified specification that adds out-of-state U.S. disasters gives a similar result, although its second-stage estimate is imprecise when the U.S.-disaster instrument is used alone (Online Appendix). Reported arrests do not predict the foot-traffic response in any specification, and the IV estimate is unchanged when arrests are included as a control. These results make an arrest channel less likely, although arrest counts remain an imperfect measure of operational scale. The IV identifies the effect of news coverage for mid-visibility events at the margin of the news cycle. These are not the largest, most visible operations. Appendix B reports the instrument balance and exclusion-restriction checks, and the Online Appendix the full per-unit OLS-versus-IV comparison and further robustness.

Table 2: First-stage relationship between news shocks and attention

	Mediacloud Spanish-ICE		English-GDELT	
	(1) Foreign	(2) US disaster	(3) Foreign	(4) US disaster
Foreign disaster days (std.)	-0.0168^{***}		-0.0105^{***}	
	(0.0023)		(0.0014)	
Cross-state disaster days (std.)		-0.0078^{***}		-0.0031^{***}
		(0.0015)		(0.0009)
First-stage F	52.66	28.44	53.12	12.85
R^2	0.482	0.103	0.520	0.046
N	51	51	51	51

Notes: The dependent variable is the indicated post-event news coverage spike. Each column reports a single-instrument first stage with the instrument standardized. The foreign-disaster measure counts days in the ICE event's post-event window that fall within the first seven days of a GDACS red-alert sudden disaster outside the United States, Latin America, and the Caribbean. The US-disaster measure counts major natural disasters in a different US state from the ICE event. Heteroskedasticity-robust standard errors in parentheses. Stars: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

4.4 Timing and design validity

Because news coverage and foot traffic are measured after the same event, the estimated effect could in principle reflect coverage responding to the neighborhood response rather than the reverse. Table 3 reports three tests. First, the effect is re-estimated using only news coverage from the first two post-event weeks, which predates the late-window response by at least three weeks and so cannot be caused by it. This early-window, temporally ordered estimate is about half the full-window effect and is essentially unchanged by a control for arrests. Second, the reverse direction is tested directly. The early foot-traffic response does not predict late-window news coverage conditional on early news coverage. Third, a weekly cross-lag panel with event and event-time fixed effects shows the same asymmetry. Lagged news coverage predicts the next week's response differential, while the lagged response does not predict the next week's news coverage. These tests show no evidence of feedback from the observed response to measured news coverage, though they do not rule out every common cause. The arrests control addresses operational severity, the leading candidate common cause, and the news-shock instrument addresses common causes in the news environment.

Table 3: Temporal ordering of attention and the foot-traffic response

	Attention → response	Response → attention
Panel A: Cross-event ordering ($N = 51$ events)		
Early vs. late window	-0.0103^{**}	-0.0392
	(0.0051)	(0.1256)
+ arrests control	-0.0098^{*}	
	(0.0055)	
Panel B: Weekly cross-lag, event and event-time FE ($N = 459$ event-weeks, 51 events)		
One-week lag	-0.0586^{*}	0.0728
	(0.0295)	(0.0848)

Notes: Columns report the two temporal directions: "Attention → response" regresses the later response on earlier attention, and "Response → attention" regresses later attention on the earlier response. Attention is the weekly US Spanish-language Mediacloud ICE news-coverage ratio aligned to each event's Monday anchor, as a deviation from the event's pre-period baseline (mean over $\tau \in [-8, -5]$); $\gamma_{e,\tau}$ is the per-event distributed-lag DiD coefficient (high- vs. low-foreign-born CBGs, reference $\tau = -1$). All regressors and outcomes are standardized, so coefficients are per SD. In Panel A the forward cells are precision-weighted regressions weighted by the CBG-clustered sampling variance of the late-window mean, with Knapp–Hartung inference, and the reverse cell is OLS with HC1 standard errors; controlling additionally for the early response ($\tau \in [0, 2]$) attenuates the forward coefficient to -0.0035 (HC1 SE 0.0035), as expected when a noisily estimated response enters as a regressor. Panel B pools post-event weeks ($\tau \in [0, 8]$) with event and event-time fixed effects and standard errors two-way clustered by event and ISO calendar week; each specification also controls for the once-lagged dependent variable. Temporal precedence rules out backward-in-time causation, not common causes; the arrests control addresses the leading common-cause candidate (operational severity). The Panel B lagged-dependent-variable specifications with unit fixed effects in a short panel are subject to Nickell bias. Stars: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

One potential concern is that the heterogeneity I exploit across event estimates comes from differences in the "quality" of the estimates themselves. To address this, I measure design quality by the flatness of pre-event trends, the number of block groups in the catchment, the precision of the estimate, and exposure to overlapping nearby events. I can then ask whether the event quality is itself correlated with the news coverage spike. It is not, and the effect of news coverage is unchanged when these design characteristics are added as controls (Table C4). Further restricting the sample to the 34 events with no overlapping-catchment contamination, and then to the 23 events with both no contamination and statistically flat pre-trends, leaves the share of variation attributable to true cross-event differences near 60 percent and preserves the effect of news coverage, although precision falls in the smallest subsample.

The effect is also a within-wave phenomenon rather than an artifact of the difference between the 2018–2019 and 2025–2026 enforcement waves. Estimated on the 2025–2026

wave alone, which contains 44 of the 51 events, it is if anything stronger (Online Appendix).

4.5 National vs. local attention

Holding the search platform and search terms fixed and varying only geography, national Google Trends predicts larger foot-traffic declines than metro-level (DMA) Trends, and DMA Trends adds little once national attention is included. National attention therefore provides the more predictive variation, and reported arrests remain null. Table 4 reports the national-versus-DMA Google Trends horse race for the 2024–2026 subsample.

Table 4: Local vs. National Coverage of Enforcement Events

	(1) Nat.	(2) + arr.	(3) + DMA	(4) DMA	(5) DMA + arr.
National Trends spike $\text{Spike}_e^{\text{Nat}}$	-0.000525^{***}	-0.000564^{***}	-0.000621^{***}		
	(0.000158)	(0.000167)	(0.000228)		
DMA Trends spike $\text{Spike}_e^{\text{DMA}}$			$+0.000074$	-0.000326^{*}	-0.000330^{*}
			(0.000215)	(0.000174)	(0.000184)
$\log(1 + \text{Arrests}_e)$		-0.001275			$+0.000589$
		(0.005263)			(0.006522)
Constant	$+0.010989$	$+0.015012$	$+0.012604$	$+0.006794$	$+0.005279$
	(0.008648)	(0.017686)	(0.008953)	(0.008799)	(0.021078)
N	44	42	41	41	39
R^2	0.215	0.232	0.239	0.085	0.079

Notes: OLS regressions of the per-event late-window estimate on national and DMA-level Google Trends spikes for "ICE raid" and, in columns (2) and (5), $\log(1 + \text{Arrests}_e)$. Both Trends measures are post-event spikes, defined as peak search interest over $\tau \in [0, 8]$ minus the baseline average over $\tau \in [-8, -5]$. The Google Trends sample is restricted to the 2024–2026 events for which weekly search data are internally consistent. The DMA measure is missing for a small set of events whose media markets have insufficient search volume. Heteroskedasticity-robust (HC1) standard errors in parentheses. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4.6 Mechanisms

The evidence so far shows that media attention predicts the size of the neighborhood response. I next ask what kind of response it produces. If direct disruption is the main channel, the effect should be strongest close to the worksite and in the raided sector. If perceived risk is the main channel, it should extend across geography and across sectors. When the event study is estimated separately by distance band and by sector, the negative response is not confined to the immediate worksite area, appearing in the 5–15 km and 15–30 km bands and remaining negative though smaller in the 30–50 km band. It is also not concentrated in the raided sector. The same-sector estimate is not negative, while other-sector POIs and both

consumer- and nonconsumer-facing POIs decline. The response therefore does not look like a same-sector demand shock or a purely retail channel. Table 5 reports the distance- and sector-split estimates.

Table 5: Mechanisms

Sample	Events	Mean estimate	Coef. on Spanish-ICE spike
Panel A. Geographic reach, mutually exclusive distance rings			
0–5 km from worksite	12	+0.031	−0.022
			(0.017)
5–15 km from worksite	29	−0.013	−0.026***
			(0.007)
15–30 km from worksite	36	+0.002	−0.023*
			(0.012)
30–50 km from worksite	40	−0.009	−0.009
			(0.006)
Panel B. Same-sector and broad-sector activity			
Same broad sector as raided worksite	20	+0.003	+0.012
			(0.017)
Other sectors	49	−0.011	−0.022***
			(0.007)
Consumer-facing POIs	51	−0.007	−0.018***
			(0.006)
Nonconsumer POIs	48	−0.003	−0.029***
			(0.008)

Notes. Each row reports a cross-event OLS regression of a subgroup-specific late-window estimate on the standardized Mediacloud Spanish-ICE news coverage spike. Standard errors are heteroskedasticity-robust and shown in parentheses. Panel A defines high-FB and low-FB CBGs once within the full 50 km event catchment and then estimates the event-study separately by distance ring. Panel B estimates the event-study separately by POI bucket. Same-sector POIs have the same broad NAICS-2 code as the raided worksite. Consumer-facing POIs are NAICS-2 codes 44, 45, 72, and 81. Stars denote $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Tests of the timing of coverage point the same way. Weekly news peaks, national search, and an eight-week English-news area-under-the-curve measure predict the response, while one-day peaks and days-above-threshold measures are weaker (Online Appendix). The mechanism evidence is most consistent with a broad perceived-risk response among heavily Latino-immigrant neighborhoods.

4.7 Civic vs. commercial response

I next identify the broad activity types most affected by the enforcement actions. Equation (1) is re-estimated separately for two POI bins. The "civic/community" bin includes NAICS-3 codes 712 and 611, which cover parks, museums, and educational services. The "commercial" bin includes NAICS-2 codes 44, 45, and 72, which cover retail and food services. NAICS-3 code 813 (religious organizations and civic associations) is omitted from the civic bin because the category is heavily contaminated by community-bank branches rather than civic organizations.

Twelve events have both a civic and a commercial estimate. Across these events, civic foot traffic in high-FB CBGs falls by an average of 0.120 log points, while the commercial decline is smaller and not distinguishable from zero. The within-event civic-minus-commercial differential averages -0.111 log points and is statistically significant. The activity that falls most sharply is therefore institutional contact at schools, parks, and museums, rather than purely commercial activity. The civic sample is small, and I cannot determine whether the gap is itself moderated by attention. The point estimates run opposite to that hypothesis, but the attention subgroups contain only six events each and are not statistically distinguishable. The civic result is therefore reported as an average contrast rather than an attention-moderated effect. The Online Appendix reports the inference and venue-coverage limits.

4.8 Spending

Equation (1) is re-estimated using $\log(1 + \text{weekly spending})$ from the exact-geocoded SafeGraph SpendPatterns panel. Thirty events have enough high- and low-FB CBG coverage to estimate a spending response. Individual-event estimates are much noisier than for foot traffic, and the aggregate is statistically indistinguishable from zero. Despite the noise, spending appears to fall more in events with larger Spanish-ICE and English-GDELT news spikes. Above-median-attention events show a differential decline of about 3 percent in heavily Latino-immigrant neighborhoods over the following eight weeks, while below-median-attention events show a small increase of about 4 percent, though the cross-event gradient is not statistically significant given the noise (Online Appendix). Because the panel records only card spending at covered merchants, it does not capture the full spending response.

5 Conclusion

This paper asks why ICE enforcement events generate large neighborhood responses in some cases and little response in others. Across 51 single-worksite events in 26 states, the average response is near zero, but the effect varies far more across events than sampling noise alone would produce. About 70 percent of the cross-event variation reflects true differences in the response rather than estimation noise. The amount of public attention an event receives explains part of that variation. Events that become visible in national Spanish-language news, English-language news, and web search are associated with larger declines in foot traffic in heavily Latino-immigrant neighborhoods, and attention explains roughly 30 percent of the cross-event variation while arrest counts explain essentially none.

Enforcement operates through an information environment as well as through direct detention. A raid disrupts a worksite, but it also sends a public signal about enforcement risk, and whether that signal reaches people outside the worksite depends on news coverage, search, protest, rumor, and the broader attention cycle. The foreign-disaster instrumental-variables design supports this reading for events where competing coverage could plausibly have crowded out the story, using out-of-state U.S. disasters as a second, complementary instrument. It does not imply that all enforcement news shocks are identical or that the largest events behave the same way, but predetermined disaster shocks shift news coverage, and that shift predicts the neighborhood response.

The largest response appears in civic and institutional venues that include schools, parks, museums, and related public-facing institutions. The cost of enforcement visibility therefore includes not only lost transactions but also reduced participation in public and institutional life.

Ultimately, the consequences of enforcement depend not only on the operation but also on the information environment it creates. Media coverage can turn a localized action into a broader signal of risk, and communities respond by reducing activity well beyond the worksite, the raided sector, and ordinary commerce. The same lesson applies wherever a salient event's economic footprint depends on how widely it is noticed rather than on the action alone.

References

Advan Research. Foot traffic / weekly patterns plus [dataset]. Dewey Data, 2025. doi: 10.82551/C103-N851.

Umair Ali, Jessica H. Brown, and Chris M. Herbst. Secure communities as immigration enforcement: How secure is the child care market? *Journal of Public Economics*, 233: 105101, 2024. doi: 10.1016/j.jpubeco.2024.105101.

Marcella Alsan and Crystal S. Yang. Fear and the safety net: Evidence from secure communities. *Review of Economics and Statistics*, 106(6):1427–1441, 2024. doi: 10.1162/rest_a_01250.

Catalina Amuedo-Dorantes and Francisca M. Antman. De facto immigration enforcement, ICE raid awareness, and worker engagement. *Economic Inquiry*, 60(1):373–391, 2022. doi: 10.1111/ecin.13041.

Catalina Amuedo-Dorantes, Esther Arenas-Arroyo, and Almudena Sevilla. Immigration enforcement and economic resources of children with likely unauthorized parents. *Journal of Public Economics*, 158:63–78, 2018. doi: 10.1016/j.jpubeco.2017.12.004.

Catalina Amuedo-Dorantes, Esther Arenas-Arroyo, and Chunbei Wang. Is immigration enforcement shaping immigrant marriage patterns? *Journal of Public Economics*, 190: 104242, 2020. doi: 10.1016/j.jpubeco.2020.104242.

Asad L. Asad. *Engage and Evade: How Latino Immigrant Families Manage Surveillance in Everyday Life*. Princeton University Press, 2023. doi: 10.1515/9780691249049.

Scott R. Baker, Robert A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. How does household spending respond to an epidemic? consumption during the 2020 COVID-19 pandemic. *Review of Asset Pricing Studies*, 10(4):834–862, 2020. doi: 10.1093/rapstu/raaa009.

Jiafeng Chen and Jonathan Roth. Logs with zeros? Some problems and solutions. *The Quarterly Journal of Economics*, 139(2):891–936, 2024.

Raj Chetty, John N. Friedman, and Michael Stepner. The economic impacts of COVID-19: Evidence from a new public database built using private sector data. *Quarterly Journal of Economics*, 139(2):829–889, 2024. doi: 10.1093/qje/qjad048.

Alberto Ciancio and Camilo García-Jimeno. Anticipating state action: Risk perceptions and consumption under immigration enforcement. Discussion Paper 18446, IZA Institute of Labor Economics, 2026.

Elizabeth Cox and Chloe N. East. Labor market impacts of ICE activity in trump 2.0. Working Paper 35129, National Bureau of Economic Research, 2026.

Uma De Balanzó, Núria Rodríguez-Planas, and Jennifer Roff. Immigration enforcement visibility and consumer spending. Discussion Paper 18620, IZA Institute of Labor Economics, 2026.

Chloe N. East, Annie Laurie Hines, Philip Luck, Hani Mansour, and Andrea Velásquez. The labor market effects of immigration enforcement. *Journal of Labor Economics*, 41(4): 957–996, 2023. doi: 10.1086/721152.

Thomas Eisensee and David Strömberg. News droughts, news floods, and U.S. disaster relief. *Quarterly Journal of Economics*, 122(2):693–728, 2007. doi: 10.1162/qjec.122.2.693.

Global Disaster Awareness and Coordination System. Gdacs event feed [database]. https://www.gdacs.org, 2026. Accessed via GDACS API.

Austan Goolsbee and Chad Syverson. Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *Journal of Public Economics*, 193:104311, 2021. doi: 10.1016/j.jpubeco.2020.104311.

Exequiel Hernandez. ICE-ing the economy: Immigration enforcement under trump 2.0 and local economic activity. Technical report, SSRN, 2026.

Jordan Herring and Burt S. Barnow. Indirect effects of immigration enforcement on health care utilization among lawfully present older Hispanics. *Social Science & Medicine*, 384: 118540, 2025. doi: 10.1016/j.socscimed.2025.118540.

Kalev Leetaru and Philip A. Schrodt. GDELT: Global data on events, location, and tone, 1979–2012. In *International Studies Association Annual Convention*, San Francisco, CA, 2013.

T. William Lester, Matthew Wilson, and Eli Knaap. The "chilling effect" of ice enforcement: Evidence from high-frequency mobility and spending data in the los angeles region. Working paper, January 26, 2026, 2026. URL https://bpb-us-e2.wpmucdn.com/sites.uci.edu/dist/6/5832/files/2026/03/Chilling-Effect-of-ICE-in-LA-LWK.pdf.

Media Cloud. Media cloud online news archive [database]. https://search.mediacloud.org, 2025. Accessed via Python API client (`mediacloud>=4.3.0`); US Spanish Language Collection, id 196136637.

Ashesh Rambachan and Jonathan Roth. A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591, 2023. doi: 10.1093/restud/rdad018.

SafeGraph. Spend patterns [dataset]. Dewey Data, 2022. doi: 10.82551/NSF5-R186.

Tara Watson. Inside the refrigerator: Immigration enforcement and chilling effects in Medicaid participation. *American Economic Journal: Economic Policy*, 6(3):313–338, 2014. doi: 10.1257/pol.6.3.313.

Appendix A: Robustness of the timing evidence

Identification rests on the cross-event variation, but the pooled event study is still informative about the timing of the response and can be stress-tested against violations of parallel trends. Figure A1 plots the pooled attention interaction δ_τ from a single regression on the stacked CBG-week panel.

Figure A1: Pooled event-time attention interaction

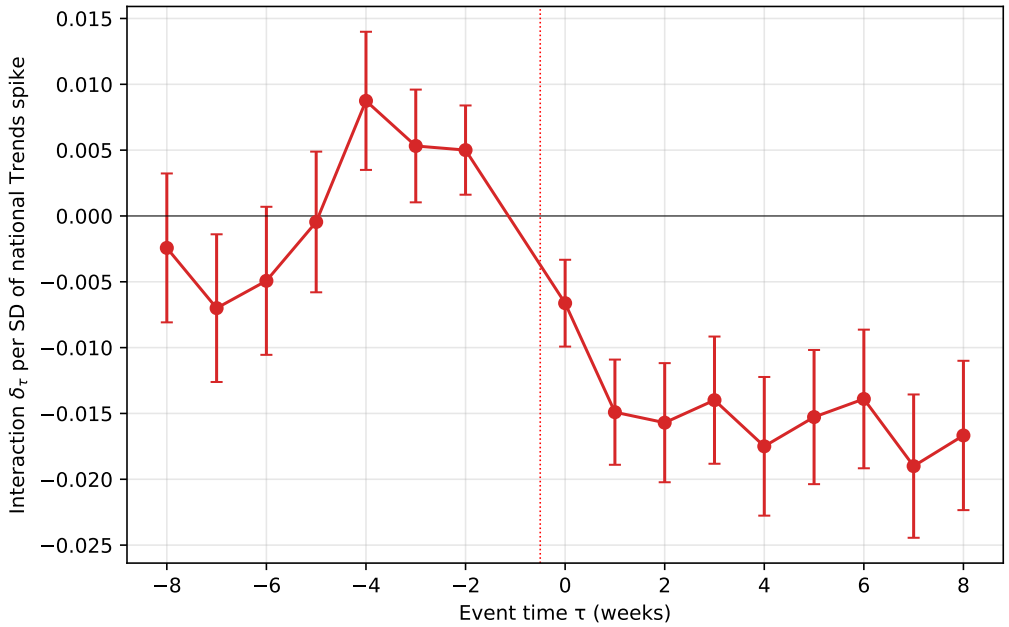

Notes. Estimates from a single pooled regression on the stacked CBG-week panel. $\log(1+y) = \alpha_{c,e} + \lambda_{w,e} + \sum_{\tau \neq -1}[\gamma_\tau + \delta_\tau \text{Spike}_e^z] \cdot \text{HighFB}_c^e \cdot \mathbf{1}[t - T_0^e = \tau] + \varepsilon$, where Spike_e^z is the standardized event-level national Google Trends spike. The figure plots δ_τ. Standard errors are clustered at the CBG level. Whiskers are $\pm 1.96\times$ SE. Reference week is $\tau = -1$. $N = 204{,}510$ CBG-week observations. The pre-event coefficients are not perfectly flat, so the figure is treated as descriptive rather than as an identifying test.

Because the pre-event coefficients in Figure A1 are not flat, this pooled time path is treated as descriptive, and its sensitivity to violations of parallel trends is quantified. Table A1 applies the relative-magnitudes sensitivity analysis of Rambachan and Roth [2023] to the late-window average of the pooled event study. The robust confidence set for the interaction δ_τ already includes zero once the allowed post-event deviation from parallel trends reaches roughly one-tenth of the largest pre-event deviation (a breakdown $\bar{M} \approx 0.1$). The pooled time path therefore does not support identification on its own. The main evidence instead comes from the precision-weighted cross-event regression, which the design-validity checks in Appendix C show is not explained by differential pre-trends, and from the

instrumental-variables estimates.

Table A1: HonestDiD pre-trend sensitivity of the pooled event study

	Robust 95% CI (late-window avg)
Panel A: Interaction δ_τ (effect of news coverage, Figure A1), est. $= -0.0165$	
Parallel trends ($\bar{M} = 0$)	[-0.0212, -0.0118]
$\bar{M} = 0.1$	[-0.0277, -0.0057]
$\bar{M} = 0.2$	[-0.0356, 0.0021]
$\bar{M} = 0.3$	[-0.0441, 0.0103]
$\bar{M} = 0.4$	[-0.0479, 0.0187]
$\bar{M} = 0.5$	[-0.0479, 0.0272]
$\bar{M} = 1.0$	[-0.0479, 0.0479]
$\bar{M} = 1.5$	[-0.0479, 0.0479]
$\bar{M} = 2.0$	[-0.0479, 0.0479]
Breakdown $\bar{M}$	0.10
Panel B: Main γ_τ (average response), est. $= -0.0023$	
Parallel trends ($\bar{M} = 0$)	[-0.0076, 0.0029]
$\bar{M} = 0.1$	[-0.0110, 0.0063]
$\bar{M} = 0.2$	[-0.0156, 0.0109]
$\bar{M} = 0.3$	[-0.0207, 0.0161]
$\bar{M} = 0.4$	[-0.0261, 0.0215]
$\bar{M} = 0.5$	[-0.0315, 0.0268]
$\bar{M} = 1.0$	[-0.0535, 0.0535]
$\bar{M} = 1.5$	[-0.0535, 0.0535]
$\bar{M} = 2.0$	[-0.0535, 0.0535]
Breakdown $\bar{M}$	0.00

Notes: Rambachan–Roth (2023) relative-magnitudes sensitivity for the pooled event study (national Google Trends measure, 7 pre-periods, 9 post-periods). The target is the late-window average ($\tau \in [4, 8]$). $\bar{M}$ bounds the largest post-treatment deviation from parallel trends by $\bar{M}$ times the largest observed pre-treatment deviation; $\bar{M} = 0$ is the standard CI under exact parallel trends. The breakdown $\bar{M}$ is the largest value at which the robust 95% confidence set still excludes zero. Confidence sets use the conditional least-favorable method.

Figure A2 examines the dispersion of the response over event time. At each event time, a random-effects meta-analysis of the 51 per-event coefficients estimates the between-event standard deviation of true effects net of sampling noise. Because each γ_τ^e is measured relative to the $\tau = -1$ reference week, dispersion accumulates mechanically with distance from the reference in both directions, so the informative comparison is at matched distance. Post-event

dispersion exceeds pre-event dispersion at all seven matched distances, by a factor of 1.4 on average, and rises to 0.035 log points by $\tau = +8$. The lower panel estimates the effect of news coverage separately at each event time. The effect of news coverage is indistinguishable from zero in every pre-event week and emerges after the event, stabilizing near -0.02 log points per standard deviation over the late window. The cross-event heterogeneity is therefore consistent with being triggered by the event rather than being pre-existing, and its alignment with news coverage has no counterpart before the event.

Figure A2: Between-event dispersion and the effect of news coverage over event time

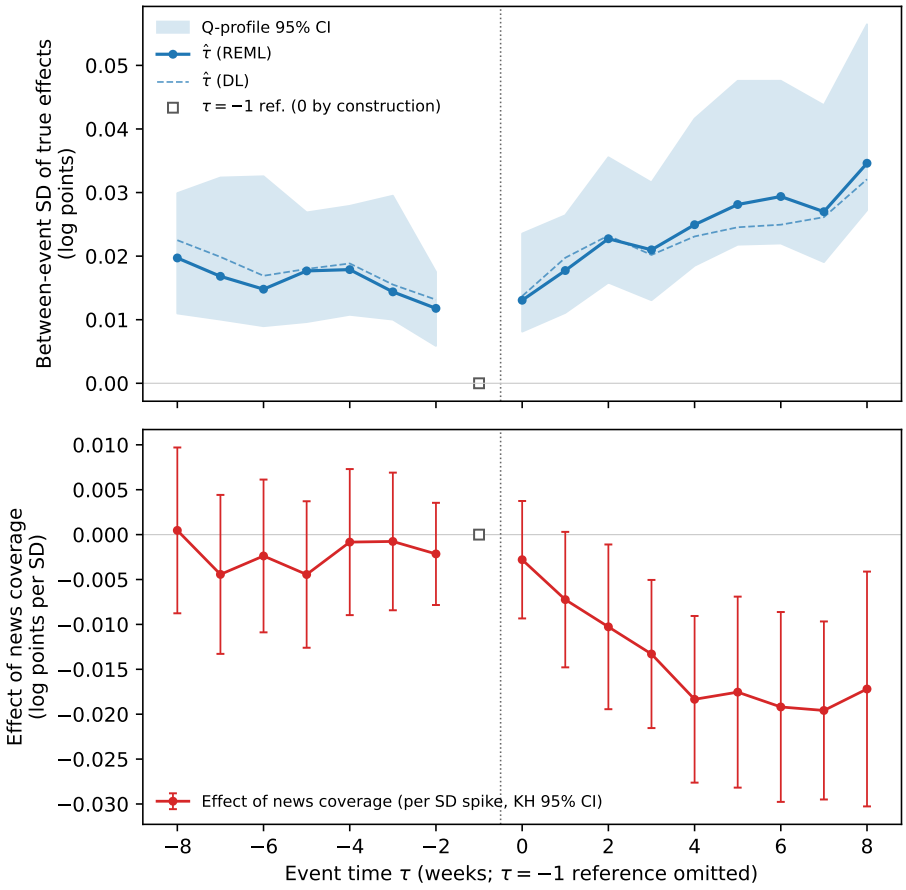

Notes. The top panel reports, at each event time τ, the estimated between-event standard deviation of true effects from a random-effects meta-analysis of the 51 per-event coefficients γ_τ^e, each weighted by its CBG-clustered sampling variance, with REML point estimates and Q-profile 95 percent confidence bands. The bottom panel reports the coefficient from a precision-weighted regression of γ_τ^e on the standardized Spanish-language news-coverage spike at each τ, with Knapp–Hartung 95 percent confidence intervals. The reference week $\tau = -1$ is zero for every event by construction and is marked separately. Estimates are pointwise across τ. The underlying numbers appear in the Online Appendix.

Appendix B: Instrument validity

Two checks support the exclusion restriction behind the news-shock instrument: a balance check against predetermined event characteristics and a direct test of the home-region channel.

Table B1 addresses the most direct timing concern. The question is whether foreign-disaster exposure is proxying for observable differences in the timing of events. The timing measure is related to the U.S.-disaster measure and to reported arrests, which is why the specification with both instruments and controls is reported. It is not related to pre-event foot-traffic trends or the placebo response.

Table B1: Balance check: foreign disaster news shocks and event characteristics

Covariate	N	Coef. on foreign disasters (cond.)	p
Log arrests	49	+0.296	0.034
Pre-event beta	51	−0.002	0.624
Backward placebo beta	51	+0.004	0.674
Cross-state disaster days (std.)	51	+0.486	0.003
Joint test of listed covariates	49	$F = 4.35$	0.006

Notes: Each row reports the coefficient on standardized foreign-disaster days in a regression of the listed predetermined event characteristic on foreign-disaster days, second-Trump-era, event-publicity, and broad-sector controls. The foreign-disaster measure counts days in the ICE event's post-event window that fall within the first seven days of a GDACS red-alert sudden disaster outside the United States, Latin America, and the Caribbean. The joint test reverses the regression and asks whether the listed covariates jointly predict foreign-disaster days after the same era, publicity, and sector controls. Heteroskedasticity-robust inference is used. The table is a balance diagnostic, not a proof of exclusion.

Table B2 pressure-tests the exclusion restriction directly. The restriction requires that foreign-disaster news shocks shift the neighborhood response only by crowding out coverage of the enforcement action. Because the outcome is a within-catchment differential, U.S.-wide disaster shocks difference out, and the main residual concern is a home-region channel in which a disaster in a country tied to the treated population affects those neighborhoods directly. The qualifying disasters are predominantly Asia-Pacific (73 percent), with the remainder elsewhere outside Latin America and the Caribbean, none with a material tie to the Mexican or Central-American-origin population. Restricting the instrument to Asia-Pacific disasters alone leaves the estimate close to the headline with a first-stage F of 47, and dropping disasters in the countries with the largest non-Latino immigrant populations in the U.S. preserves a significant estimate. The Africa-and-Europe-only slice has too few exposed

events to support a first stage. The overidentified specification that adds U.S. cross-state disasters, which would violate the exclusion through a different channel, does not reject the Sargan test ($p = 0.83$). The instrument is unrelated to the pre-event coefficient and the placebo response, and it correlates with reported arrests and, mechanically, with the U.S.-disaster measure, which is why arrests are controlled for and the two disaster instruments are reported jointly.

Table B2: Exclusion-restriction pressure tests for the foreign-disaster IV

	Coef.	(SE)	First-stage F	Disasters
Panel A: Geographic restriction of the instrument				
Baseline (ex-LatAm/Caribbean)	-1.059^{***}	(0.338)	52.7	113
Asia-Pacific only	-1.240^{***}	(0.402)	46.6	83
Africa / Europe / MENA only	-0.520	(0.857)	5.6	27
Drop largest US (non-Latino) diaspora origins	-0.889^{**}	(0.439)	25.8	69
Panel B: Overidentification and balance				
Foreign + US disasters (over-id)	-1.052^{***}	(0.340)	25.8	
Sargan over-id p	0.830			
Balance joint F (covariates $\rightarrow$ instrument)	$F = 4.35$, $p = 0.006$			
OLS reference	-0.919^{***}	(0.256)		

Notes: The endogenous regressor is the Mediacloud Spanish-ICE news coverage spike and the outcome is per-event β^{e}_{late}. Panel A re-estimates the foreign-disaster IV restricting the instrument to disasters by region of the affected country. Of the 113 qualifying red-alert non-Latin-American sudden disasters, 83 (73%) are in the Asia-Pacific and the remainder elsewhere outside Latin America and the Caribbean, none with a material tie to the Mexican or Central-American-origin population. The diaspora row drops the largest US non-Latino foreign-born origin countries (the Philippines, India, China, Vietnam, Korea). Panel B reports the overidentified specification using the foreign-disaster and US-disaster instruments jointly with its Sargan test, and the joint test that predetermined event characteristics predict the instrument. First-stage statistics are robust Wald F tests. Heteroskedasticity-robust standard errors. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Appendix C: Robustness of the effect of news coverage

Table C1 reports the full random-effects decomposition of the cross-event variation summarized in Section 4.2, together with the precision-weighted regression on the standardized news coverage spikes.

Table C1: Cross-event variation in neighborhood responses

Panel A: Cross-event response variation	
Events	51
Random-effects pooled mean response	-0.0049
95% CI	[-0.0145, 0.0047]
Share of variance due to true cross-event differences	69.7%
SD of true event effects	0.0287
Prediction interval for a new event	[-0.0568, 0.0470]

Panel B: Predictors of event-level response

	Coef. per SD	SE	True variance explained
Spanish-ICE spike	-0.0184	0.0051	30.1%
Spanish-ICE | arrests	-0.0192	0.0053	25.1%
English-GDELT spike	-0.0154	0.0050	18.5%

Notes: Random-effects meta-analysis of the 51 per-event late-window (weeks 4–8) estimates, using the standard error of the late-window average from the CBG-clustered covariance matrix. A formal homogeneity test rejects a common effect across events (Cochran's $Q = 164.9$, $p < 0.001$). The share-of-variance row reports the portion of observed cross-event dispersion attributed to true differences in the underlying response rather than sampling error. The SD of true event effects is the restricted-maximum-likelihood estimate of τ; the DerSimonian–Laird estimate is 0.0254. The prediction interval is the Higgins–Thompson–Spiegelhalter interval for the true effect of a new event. Panel B reports precision-weighted regressions with Knapp–Hartung adjusted standard errors. "True variance explained" is the share of between-event variance explained by the covariate. Spikes are standardized to the 51-event SD; the arrests specification has $N = 49$.

Figure C1 plots the 51 per-event late-window estimates against the random-effects pooled mean and prediction interval reported in Table C1 above.

Figure C1: Per-event late-window estimates with the random-effects pooled mean and prediction interval

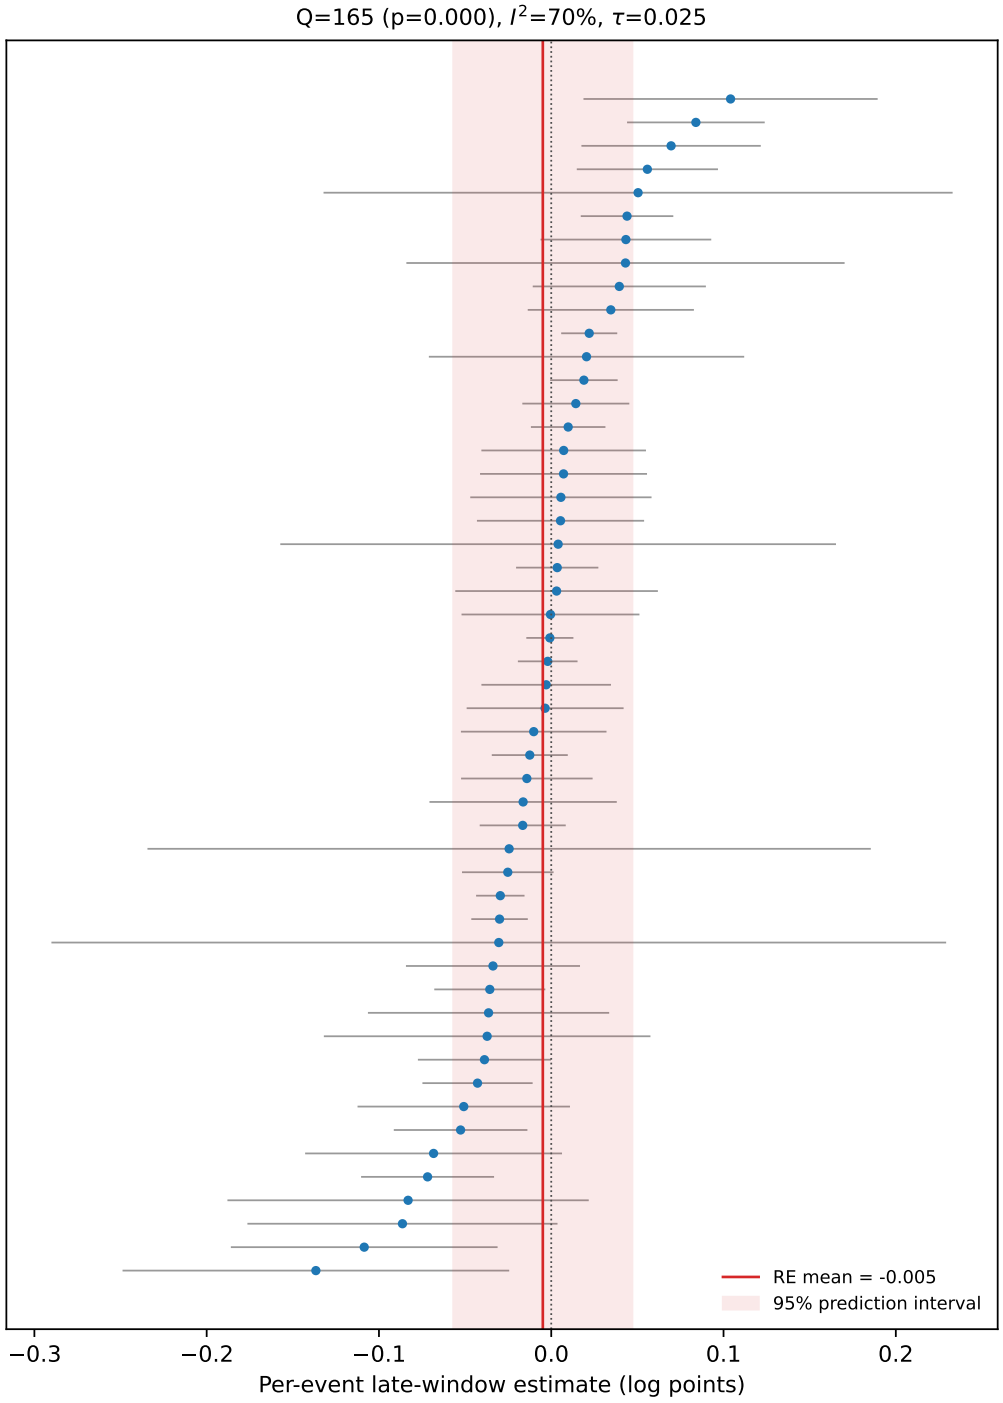

Notes. Each marker is one of the 51 fitted events, ordered by point estimate, with whiskers at ±1.96 times the standard error of the late-window average (the standard error of the linear combination of weekly coefficients, clustered at the CBG level). The vertical line is the random-effects pooled mean and the shaded band is the 95 percent prediction interval for the true effect of a new event.

Table C2 reports the cross-event OLS horse race in per-unit-of-spike units. The per-unit coefficients are larger in magnitude than the per-standard-deviation estimates in Table 1 because the spike is unstandardized. The columns anchor on the Mediacloud Spanish-ICE spike alone and with the arrests control, add a standalone arrests specification, and replace the headline measure with the English-GDELT and national Google Trends variants.

Table C2: News-coverage horse race. Late-window estimate on attention and severity

	(1) Spa. ICE	(2) + arr.	(3) Arr. only	(4) GDELT only	(5) Trends only
MC Spa. ICE spike	-0.9193^{***}	-0.9987^{***}			
	(0.2559)	(0.2649)			
English GDELT spike				-1.3382^{***}	
				(0.4382)	
National Trends spike $\times 10^{-3}$					-0.5254^{***}
					(0.1576)
$\log(1 + \text{Arrests}_e)$		-0.0053	-0.0028		
		(0.0042)	(0.0049)		
Constant	$+0.0193^{**}$	$+0.0386^{**}$	$+0.0015$	$+0.0114$	$+0.0110$
	(0.0096)	(0.0184)	(0.0172)	(0.0088)	(0.0086)
N	51	49	49	51	44
R^2	0.227	0.260	0.006	0.173	0.215

Notes. OLS regressions of the per-event late-window estimate on the post-event news coverage spike (peak$[\tau \in [0, 8]]$ minus baseline$[\tau \in [-8, -5]]$, event-week aggregation) and $\log(1 + \text{Arrests}_e)$. The headline measure in columns (1)–(2) is the Mediacloud US Spanish-language news ICE-spike (407-source collection, share-of-coverage). Column (3) reports the standalone severity specification. Columns (4)–(5) report the same single-regressor attention specification with English-GDELT (full 51-event panel) and US national Google Trends 'ICE raid' (44-event 2024–2026 subsample) as corroborating attention measures. The headline measure closely co-moves with the corroborating measures, so the single-regressor attention columns should be read as cross-source robustness rather than independent channels. Heteroskedasticity-robust (HC1) standard errors in parentheses. Stars denote $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

The log transformation of the outcome also does not drive the result. Table C3 re-estimates every first-step event study by Poisson pseudo-maximum likelihood on raw visit counts and re-runs the second step on the resulting estimates. No CBG-week observation in the estimation panels has zero visits, so the concerns raised by Chen and Roth [2024] about log-like transformations with zeros do not arise mechanically, and the PPML check addresses the functional-form question directly.

Table C3: PPML robustness of the effect of news coverage

First-step outcome model	Precision-weighted	OLS	2SLS
$\log(1 + \text{visits})$ (baseline)	-0.0184^{***}	-0.0222^{***}	-0.0256^{***}
	(0.0051)	(0.0062)	(0.0082)
Poisson (PPML), visit counts	-0.0156^{***}	-0.0157^{**}	-0.0158
	(0.0055)	(0.0064)	(0.0104)
Events		51	

Notes: Each cell reports the coefficient on the standardized Spanish-ICE news-coverage spike from the paper's second step, using per-event late-window estimates from the indicated first-step outcome model. The PPML rows re-estimate every per-event distributed-lag DiD by Poisson pseudo-maximum likelihood on raw visit counts with the same CBG and calendar-week fixed effects and CBG-clustered standard errors; PPML coefficients are semi-elasticities, comparable to the log-outcome estimates. The per-event PPML and log estimates correlate at $r = 0.71$ across events. No CBG-week observation in the estimation panels has zero visits, so the two outcome models share an identical estimation sample. Precision-weighted columns use Knapp–Hartung standard errors, OLS columns HC1, and 2SLS columns the foreign-disaster news-shock instrument with robust standard errors. Stars: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Table C4 collects the design-validity checks summarized in the main text. It reports the correlations between per-event design quality and the news coverage spike, and the effect of news coverage with design-quality controls and on the uncontaminated and flat-pre-trend subsamples.

Table C4: Design validity: per-event design quality is orthogonal to attention

Panel A: Design metric regressed on the standardized Spanish-ICE spike (HC1)					
Design metric	N	slope/SD	t	p	R^2
Pre-trend joint F (non-flatness)	51	-217.6642	-1.04	0.301	0.047
|pre-event coef|	51	-0.0056	-1.31	0.191	0.045
|placebo response|	51	-0.0060	-1.25	0.210	0.035
# CBGs in catchment	51	-9.3960	-0.24	0.810	0.001
Per-event SE (imprecision)	51	-0.0005	-0.11	0.915	0.000
# contaminating events	51	0.1667	1.04	0.296	0.025
Frac. catchment shared	51	0.0256	0.51	0.611	0.004
Panel B: Attention coefficient with and without design controls					
	N	coef/SD	SE	t	p
Spanish-ICE spike (per SD), no controls	51	-0.0184	0.0051	-3.65	0.001
+ design controls	51	-0.0185	0.0054	-3.42	0.001
Panel C: Heterogeneity on clean subsamples					
Subsample	N	I^2	τ	spike p	True var. expl.
All events	51	69.7%	0.0287	0.001	30.1%
No contamination	34	61.1%	0.0309	0.027	14.8%
No contam. & flat pre-trends	23	62.8%	0.0301	0.051	16.6%

Notes: Panel A regresses each per-event design-quality metric on the standardized Spanish-ICE news spike with HC1 standard errors; a null slope means design quality does not vary systematically with attention. Panel B shows the precision-weighted regression attention coefficient is unchanged when standardized design metrics (pre-trend F, # CBGs, # contaminating events) are added as controls. Panel C recomputes the heterogeneity statistics and the attention precision-weighted regression on progressively cleaner subsamples (no overlapping-catchment contamination; additionally non-rejected flat pre-trends at $p > 0.10$). τ is in log points.