Politicized Attention Shifts Amplify Polarization in the Information Ecosystem during California Wildfires (Supplementary Information)

Yiheng Chen[1†], Alina Hagen[2†], Fan Yang[3], Ratna B. Dougherty[4], Zihui Ma[5], Lingyao Li[2*], Runlong Yu[1*]

[1]Department of Computer Science, University of Alabama, Tuscaloosa, AL, 35487, USA.
[2]School of Information, University of South Florida, Tampa, FL, 33620, USA.
[3]School of Journalism and Mass Communications, University of South Carolina, Columbia, SC, 29208, USA.
[4]School of Public Affairs, University of South Florida, Tampa, FL, 33620, USA.
[5]Department of Aerospace Engineering, University of Maryland, College Park, MD, 20742, USA.

*Corresponding author(s). E-mail(s): lingyaol@usf.edu; ryu5@ua.edu;
†These authors contributed equally to this work.

Module A: Data, preprocessing, spatial context, and LLM annotation

S1. Brandwatch data fields and query execution

We retrieved public posts from X (formerly Twitter) via the Brandwatch API. Brandwatch provides access to historical X posts via its API integration and returns posts matching the specified query and filters. The user-level attributes extracted include gender, U.S. region, interests, account type, and verification status. These characteristics are inferred by Brandwatch rather than directly self-reported in structured form. According to the company's documentation, the inference approach differs by attribute. For example, geographic location is primarily derived from explicit information provided in users' public X profiles, while attributes such as gender and interests are predicted using machine learning models that analyze publicly available signals, including first names and profile descriptions. We note that the specific algorithms underlying these demographic inferences are proprietary, and detailed methodological information is not publicly disclosed.

S2. Preprocessing details

Bot filtering. To reduce automated amplification, we removed posts authored by accounts on a curated bot list. Candidate accounts were discovered by scanning author identifiers for the substring `bot` (case-insensitive) and manually reviewed to remove false positives; filtering was applied by exact match on normalized author strings. (normalization: case-folding and whitespace trimming).

Text normalization. Post text was standardized by removing retweet-style prefixes (e.g., `RT @user:`) and URLs (`http(s)://...`), normalizing whitespace, and dropping records with empty text after cleaning. These steps reduce boilerplate without altering semantic content.

Deduplication and identity. Throughout the analysis, `Mention Id` is treated as the post-level identifier. To reduce structural redundancy in high-volume periods, we additionally collapsed *exact* duplicates (identical cleaned text strings), retaining the first occurrence.

S3. Wildfire perimeter data, mapping rules, and severe-fire windows

Perimeter retrieval and filtering. Wildfire perimeter polygons were obtained from CAL FIRE's Fire and Resource Assessment Program (FRAP) historical fire perimeters feature service. We queried incidents within 2016-2025 and retained perimeters with mapped area $\geq$ 1,000 acres. To accommodate service record limits, requests were paginated using a fixed record count and increasing offsets. County boundaries were obtained from the U.S. Census Bureau cartographic boundary county dataset; all layers were reprojected to a common coordinate reference system prior to overlay.

Mapping and highlighting for Fig.1a. All perimeters meeting the area threshold are rendered as translucent contextual overlays. A hatched subset highlights "major named fires" curated from CAL FIRE incidents identified as among the most destructive, deadliest, or largest during the study period (Table S1). Major incidents were matched to FRAP perimeters by incident year and name; ambiguous cases were resolved by conservative matching rules documented here to ensure consistent highlighting.

Severe-fire windows for Fig.1d. For the context interaction analysis, we define statewide severe-fire windows by consolidating overlapping major incidents into intervals (Table S8). A day t is labeled *During fire* if it falls within any interval in Table S8 (unioned with no double counting), and *Non-fire* otherwise. Each government entity-day observation inherits the indicator from its calendar date.

S4. LLM prompting, JSON schema, parsing, failure handling and model settings

We performed institutional entity extraction and per-entity stance classification using a single closed-world prompt with a strict JSON output contract. The prompt restricts entities to {*government*, *non-profit*, *news*} and requires one JSON object per detected entity with fields `entity_name`, `category`, `representation`, `level` (government only), and `stance`. The full prompt text is provided in S11.

Model settings. Corpus-scale annotation was conducted with `gpt-5-mini`, configured with `reasoning_effort=medium` and `max_completion_tokens=4000`. All model identifiers and token usage statistics were systematically logged to facilitate reproducibility.

Parsing/validation and failure handling. Outputs were deterministically parsed and validated against the JSON contract. Malformed or contract-violating outputs were logged and excluded from aggregation. Transient API failures were retried with exponential backoff; contract violations were not repaired post hoc to preserve a fully contract-compliant pipeline.

S5. Human annotation protocol and validation

Gold standard. We validated entity detection and stance classification on an adjudicated set of $N =$ 200 posts. Two annotators independently applied the closed-world rules in S4 to record all eligible entities and assign stance (*positive*/*negative*/*neutral*) for each entity; disagreements were resolved by adjudication.

Entity matching. Evaluation is conducted at the post–entity instance level. Gold and predicted entity strings are normalized (case-folding and whitespace normalization) with an optional alias map for common variants. Entities are paired using (i) exact match on normalized strings and (ii) greedy one-to-one fuzzy matching above a string-similarity threshold ($\geq$ 0.80). Unmatched gold entities are counted as misses; unmatched predictions are counted as false positives. Stance metrics are computed conditional on matched entities.

Metrics. Let $\mathcal{D}$ denote evaluation posts. For each post $d \in \mathcal{D}$, let G_d and P_d be the sets of gold and predicted entities after normalization, and let $M_d \subseteq G_d \times P_d$ be the resulting one-to-one matches. Entity detection rate (micro-recall) is

$$\mathrm{EDR} = \frac{\sum_{d\in\mathcal{D}} |M_d|}{\sum_{d\in\mathcal{D}} |G_d|}.$$

Let $y(g)$ and $\hat{y}(p)$ denote the gold and predicted stance labels for a matched pair $(g, p) \in M_d$. Stance accuracy is

$$\mathrm{Acc} = \frac{\sum_{d\in\mathcal{D}} \sum_{(g,p)\in M_d} \mathbb{I}(\hat{y}(p) = y(g))}{\sum_{d\in\mathcal{D}} |M_d|}.$$

Stance macro-F1 is computed for $\mathcal{C} = \{\mathrm{pos}, \mathrm{neg}, \mathrm{neu}\}$ over all matched pairs $\bigcup_{d\in\mathcal{D}} M_d$.

Model selection and corpus-scale summary. Table S2 reports validation performance on the adjudicated set; we select `gpt-5-mini` for corpus-scale annotation based on overall performance under the same prompt and strict output constraints. Across the full corpus, 670,488 posts contain at least one eligible entity; in total, 1,120,831 entity instances were extracted, of which 730,337 were linked to posts with county information and thus usable for county-level analyses (Table S6).

Module B: Actor taxonomy, amplification measures, event detection, and ITS

S6. 9-way author taxonomy rules and override governance

We classified author accounts into a mutually exclusive 9-way taxonomy: *Politicians*, *Government agencies*, *News media*, *Journalists*, *Organizations*, *Artists & Sportspersons*, *Business & Executives*, *Professionals*, and *General public*. Classification followed a fixed priority order: (i) exact handle matching to curated lists, (ii) profession/job-title cues from metadata, (iii) keyword/regex rules applied to handles and display names (conditioned on organizational-account signals where applicable), and (iv) default assignment to *General public*. Keyword/regex cues are listed in Table S3.

Manual overrides. After automated classification, a curated override list was applied via exact (case-insensitive) handle matching to correct audited edge cases; overrides supersede automated assignments and are logged with a `manual_override` tag. Override counts by target class are reported in Table S4.

Display mapping for Fig.4. For visualization, sparse categories (*Artists & Sportspersons*, *Business & Executives*, *Professionals*) were merged into *General public*, yielding the 6-way display set used in Fig.4.

S7. Field availability and smoothing choices

Field availability. Engagement and impression fields are not consistently populated prior to 2020 in our export (often recorded as zero). Accordingly, engagement-weighted and impression-weighted analyses are restricted to 2020–2025. Engagement is defined in the main Methods.

Smoothing. We apply rolling-window smoothing matched to each figure's temporal resolution to reduce high-frequency noise while preserving event-scale dynamics. For Fig.2, weekly series are smoothed using a 4-week rolling mean (`min_periods`=1). For Fig.4, monthly series are smoothed using a centred 6-month rolling mean (`center=True`, `min_periods`=2) after linear interpolation over short gaps (up to four consecutive months). For ITS panels derived from daily data, we use 14-day centred rolling averages of daily engagement-weighted negativity/positivity to stabilize day-to-day fluctuations within event windows.

S8. Peak detection and topic attribution

Peak detection. Peaks are detected on weekly smoothed engagement- and impression-weighted series (2020–2025). Candidate peaks are extracted using a prominence threshold of 0.03 and a minimum separation of 3 weeks. The two peak sets are merged, treating impression peaks within ± 2 weeks of any engagement peak as duplicates (retaining the engagement peak). To handle rising endpoints, we apply a tail rule based on prominence within the final 8 weeks and select the peak week as the within-tail week with the largest post count, subject to the same duplicate-resolution logic. Peaks are ranked by prominence and retained for labeling as reported in Table S5.

Topic-assisted event labeling. To support interpretable event naming, we apply BERTopic separately to negative and positive government posts in 2020–2025 after lightweight text cleaning and excluding very short texts (< 20 characters). The pipeline uses SentenceTransformer embeddings (`all-MiniLM-L6-v2`), a count vectorizer with English stop-words and bi-grams, UMAP reduction (fixed random seed), and BERTopic with `min_topic_size`=15 and `nr_topics=auto`. For each peak, topics are scored within a ± 14-day window using within-window enrichment (lift relative to baseline) and excess volume above expected counts; BERTopic outputs assist labeling, and final event names are curated and checked manually for interpretability.

Peak alignment for plotting. For Fig.2, detected peak dates are snapped to the nearest local maximum on the plotted engagement curve within a ± 6-week search window; events mapping to curve peaks within 4 weeks are deduplicated by retaining the higher peak. The full event list (peak date, onset date, label, and event type) is reported in Table S5.

S9. Interrupted time-series (ITS) specification and sensitivity

Dependent series. ITS analyses are conducted on daily aggregates of government posts. For each day t, we compute engagement-weighted negativity and positivity rates using engagement weights as defined in the main Methods. The dependent series used for plotting and estimation is a 14-day centred rolling mean (minimum 7 days) of the daily share.

Event onset and window. The interruption date is the curated onset date reported in Table S5. We extract a symmetric window of ± 36 days around onset and drop days without rolling estimates near the window edges. If the onset date does not coincide with an available daily observation, the interruption index is set to the closest observed day within the window. Windows with fewer than 20 usable observations are treated as insufficient for ITS estimation.

Segmented OLS specification. Within each event window, we fit a segmented linear regression by ordinary least squares:

$$y_t = \beta_0 + \beta_1 T_t + \beta_2 D_t + \beta_3 P_t + \varepsilon_t,$$

where y_t is the smoothed negativity (or positivity) rate, T_t is a linear time index within the window, $D_t = \mathbb{I}(T_t \geq t_0)$ is a post-onset indicator, and $P_t = D_t (T_t - t_0)$ is the post-onset time-since-interruption term. Under this parameterization, β_2 captures the immediate level change at onset and β_3 captures the change in slope relative to the pre-onset trend. Counterfactual trajectories under no interruption are obtained by setting $(D_t, P_t) = (0, 0)$ and projecting the pre-onset trend through the full window.

Module C: Network construction and layout alignment

S10. Network construction, community detection and diagnostics

Network construction. Yearly user–user interaction networks are built from the deduplicated post corpus. Nodes represent unique authors. A directed edge $i \to j$ is added when author i replies to or reshared author j; edge weights equal annual interaction counts.

Undirected projection and filtering. For community detection and visualization, the directed network is projected to an undirected weighted graph by summing reciprocal interaction weights for each unordered pair. We retain the giant connected component and restrict to the top 2,000 nodes by degree to improve readability and maintain comparable density across years.

Community detection and polarization metrics. Communities are detected using the Leiden algorithm with modularity optimization on the undirected weighted graph. For each yearly network we report modularity (Q) and the E-I index contrasting inter-community versus intra-community connectivity, alongside bridging-node share and the bimodality coefficient (definitions in main Methods).

Layout alignment across years. To enable cross-year visual comparison, layouts are aligned by assigning the largest communities to fixed angular sectors and anchoring sector identity using stable high-centrality accounts (identified by PageRank) as reference points. Within-community coordinates are generated using a seeded force-directed layout and mapped into the assigned sector; remaining nodes are placed near the centre using the same seeded procedure.

Visualization and reproducibility. Network panels were rendered with NetworkX and Matplotlib using a fixed global random seed (`seed=42`) for community detection and layout to ensure reproducible output. To eliminate hash-order dependence, nodes and edges are lexicographically sorted before conversion to igraph; degree-rank ties are broken alphabetically by username.

S11. Prompt used for entity and stance extraction

```
Identify and categorize all mentions of governmental agencies, non-profit organizations, or news outlets in the
↪  following X post related to California wildfires. Do all analysis and reasoning silently (do not reveal it).
↪  Return only the valid JSON with no other comments.

# TASK
1) Read the post carefully in its entirety.
2) Identify all distinct entities explicitly mentioned.
3) Restrict entities to these categories only: government, non-profit, news.
4) For each entity, emit one JSON object with:
   - entity_name (string)
   - category ("government" | "non-profit" | "news")
   - representation ("individual" | "organizational")
   - level ("federal" | "state" | "local" | null)  // government only; otherwise null
   - stance ("positive" | "negative" | "neutral")

IMPORTANCE:
- Do not return any entities if they are not present in the post.
- Note the distinction between similar words: CAFire and CAL FIRE are not the same.
- No inference beyond canonical renaming is permitted.

# SCOPE
- Types: government, non-profit, news (include political parties as government)
- Geographic focus: U.S. entities only (primarily California)
- Non-profit news organizations -> category="non-profit" (not "news")
- EXCLUDE: places mentioned solely as locations, events/incidents, for-profit corporations/brands, metaphorical
↪  wildfire uses, non-U.S. entities
```

```
# NAMING CONVENTIONS (entity_name)
- If an @handle appears in the post for an entity, record the @handle exactly as written (do not expand it).
- Only expand to official names when the surface name appears without a handle.
- Merge duplicate mentions of @handles and surface names into one entity (default to the @handle when present).
- Canonicalization examples (surface names only):
  - FEMA -> Federal Emergency Management Agency
  - USFS -> United States Forest Service
  - Trump -> Donald Trump
  - Newsom -> Gavin Newsom

# CATEGORY DEFINITIONS
- government: public agencies (federal/state/local), political offices, public officials, political parties, and generic public roles explicitly referenced (e.g., firefighters, police, authorities).*
- non-profit: NGOs, charities/foundations, advocacy groups, research orgs, coalitions, cultural/educational institutions, and non-profit media. **, ***
- news: for-profit mass media outlets and individual journalists working for for-profit media. ****

# REPRESENTATION
- individual: a specific named person whose actions/statements are referenced.
- organizational: an agency, office, group, outlet, or generic public body/title without a specific named person.

# LEVEL (government only)
- Infer level when possible using context (agency acronyms, titles, explicit jurisdiction).
- organizational:
  - federal: U.S. agencies (e.g., FEMA, USFS)
  - state: California state agencies (e.g., CAL FIRE)
  - local: county/city agencies (e.g., Sheriff's Office, Fire Department)
- individual:
  - federal: President/VP, U.S. Senators/Reps, Cabinet, heads of federal agencies
  - state: Governor, state legislators, statewide officials, state agency directors
  - local: mayors, councilmembers, county supervisors, sheriffs, local fire chiefs
- Political parties: assign state/local only if explicitly scoped; otherwise level=null.
- If level cannot be determined from text, use null.

# STANCE (per entity; independent)
- positive: explicit praise/thanks/support; favorable evaluation; calls for donation/support/volunteering; describes helpful actions/aid/positive impact.
- negative: criticism/blame; questioning competence; highlighting failures/inaction/harm; calls for investigations/removals/changes; describes harmful actions/negative impact.
- neutral: factual mention, info-sharing, announcements, logistical updates, citations without evaluative language.
- Irony/quotes do not flip polarity without explicit evaluative language.

# RECALL-FIRST HEURISTICS
- Favor inclusion when uncertain; default stance="neutral" and level=null if needed.
- Include generic public bodies only when explicitly referenced as actors/targets.
- Include @handles only if they plausibly belong to {government|non-profit|news} and are actors/targets (not mere "via"/"h/t" credit).
- Do NOT include:
  - places used only as locations/context
  - event/incident names (e.g., wildfire names) as events
  - for-profit corporations/brands
  - non-U.S. entities
  - formulaic credit mentions (e.g., "via @X", "h/t @X")

# OUTPUT FORMAT (STRICT)
Return the results in valid JSON as a single minified array (single-line). No markdown, no commentary.
- If no eligible entities are present, output exactly: []
- All enum values must be lowercase.
- level must be "federal"|"state"|"local" or null (JSON null).

Example output schema (format only):
[{"entity_name":"<name or @handle>","category":"government|non-profit|news","representation":"individual|organizational","level":"federal|state|local|null","stance":"positive|negative|neutral"}]

Here is the X post: {text}
```

Supplementary Tables

- **Table S1.** Major named fires used for hatched overlays in Fig.1a.
- **Table S2.** Model performance on entity extraction and stance classification.
- **Table S3.** 9-way author taxonomy keyword/regex cues.
- **Table S4.** Manual override counts for the author taxonomy.
- **Table S5.** Peak events used in Fig.2.
- **Table S6.** Corpus-scale annotation summary (entity and stance distributions; two-panel table).
- **Table S7.** Mixed-effects model EMMs and planned contrasts for Fig.1c–d.
- **Table S8.** Consolidated severe-fire windows used to define the intervention context in Fig.1d.

Table S1: Major California wildfire events (2016-01-01 – 2025-02-12) identified by CAL FIRE as among the most destructive, deadliest, or largest incidents. An asterisk (*) indicates that the incident appears on CAL FIRE's Top 20 list for that category.

Incident Name	**Created**	**Extinguished**	**Destructive**	**Deadliest**	**Largest**
Tubbs Fire	2017-10-08	2017-10-31	*	*	
Redwood Valley Fire	2017-10-08	2017-10-28		*	
NUNS Fire	2017-10-08	2017-10-31	*		
Atlas Fire	2017-10-09	2017-11-17		*	
Thomas Fire	2017-12-04	2018-01-12	*		*
Carr Fire	2018-07-23	2018-08-30	*	*	*
Mendocino Complex Fire	2018-07-27	2019-01-04			*
Camp Fire	2018-11-08	2018-11-25	*	*	
Woolsey Fire	2018-11-08	2018-11-21	*		
CZU Lightning Complex Fire	2020-08-16	2020-09-22	*		
SCU Lightning Complex Fire	2020-08-16	2020-10-01			*
August Complex	2020-08-16	2020-11-11			*
North Complex Fire	2020-08-17	2020-12-03	*	*	*
LNU Lightning Complex Fire	2020-08-17	2020-10-02	*	*	*
Glass Fire	2020-09-17	2020-10-20	*		
Dixie Fire	2021-07-13	2021-10-25	*		*
Monument Fire	2021-07-30	2021-10-26			*
Caldor Fire	2021-08-14	2021-10-21	*		*
Eaton Fire	2025-01-07	2025-01-31	*	*	
Palisades Fire	2025-01-07	2025-01-31	*	*	

Table S2: Model performance on entity extraction and stance classification ($N = 200$ adjudicated posts). **Bold** = best; <u>underline</u> = second-best per column.

Model	**Entity detection rate**	**Stance Acc**	**Stance F1**
Gemini-2.5-Flash	93.1298	83.9695	86.1953
Gemini-2.5-Pro	89.3130	80.9160	87.1097
GPT-4o-mini	53.4351	46.5649	72.9951
GPT-4o	17.5573	16.0305	63.0640
GPT-4.1-mini	67.1756	54.9618	73.7420
GPT-4.1	68.7023	64.8855	84.0847
GPT-o3-mini	80.9160	74.0458	86.2024
GPT-o3	<u>94.6565</u>	<u>86.2595</u>	87.4664
GPT-o4-mini	93.1298	**87.0229**	**90.6004**
GPT-5-nano	74.0458	64.8855	79.5151
GPT-5-mini	**95.4198**	<u>86.2595</u>	<u>88.3214</u>
GPT-5	92.3664	85.4962	88.2161

Table S3: Keyword/regex cues used in the 9-way author taxonomy. Patterns are case-insensitive and applied to author handle/display name and profession tags as described in S6.

<table>
<tr><th>Assigned class</th><th>Keyword/regex list and conditions (case-insensitive)</th></tr>
<tr><td>Politicians</td><td>Profession tag: `profession=politician`.
Job-title regex:
`politician|governor|senator|mayor|congress|representative|council\s?member|legislat|assembl|commissioner|supervisor|secretary of state|attorney general|white house|president(?!.*university)`.
Executive safeguard: when `profession=executive`, `president` alone does not trigger politician assignment; a non-`president` politician keyword is required.</td></tr>
<tr><td>Journalists</td><td>Profession tag: `profession=journalist`.
Job-title regex:
`journalist|reporter|editor(?!.*video)|anchor|correspondent|columnist|news director|bureau chief|managing editor`.</td></tr>
<tr><td>News media</td><td>Legacy media regex:
`news|tv\b|radio|times\b|post\b|tribune|herald|gazette|chronicle|press\b|wire|nbc|abc\b|cbs|fox\b|cnn|bbc|ap\b|reuters|bloomberg|broadcast|journal\b|dispatch`.
Applied to organisational accounts, and used as a backstop for outlets mislabeled as non-organisational.</td></tr>
<tr><td>Government agencies</td><td>Gov/emergency-service regex:
`department|dept\.?|agency|bureau|commission|authority|police|sheriff|fire dept|calfire|usfs|fema|nws\b|city of\b|county of\b|state of\b|district\b|emergency|public works|parks?\s?(?:and|&)\s?rec`.</td></tr>
<tr><td>Organizations</td><td>Assigned if the handle matches a curated organisational entity list (e.g., non-profit/organisation accounts), or as the organisational fallback when not assigned to Government agencies or News media.</td></tr>
<tr><td>Artists & Sportspersons</td><td>`profession=artist` OR `profession=sportpersons`.</td></tr>
<tr><td>Business & Executives</td><td>`profession=executive`.</td></tr>
<tr><td>Professionals</td><td>`profession=scientist` OR `profession=teacher` OR `profession=health` OR `profession=legal`.</td></tr>
<tr><td>General public</td><td>Assigned when no higher-priority rule fires.</td></tr>
</table>

Table S4: Manual override counts for the 9-way author taxonomy. Overrides are applied *after* automated classification via exact (case-insensitive) handle matching and recorded with a `manual_override` tag ($N = 925$ accounts).

Override target class	Count
Politicians	64
Government agencies	34
News media	38
Journalists	242
Organizations	281
Artists & Sportspersons	215
Business & Executives	0
Professionals	0
General public	51
Total	**925**

Table S5: Full list of detected peaks and attributed events used in Fig.2. Peak dates are discourse peaks; onset dates record the anchor date associated with the named event. Event type is assigned by comparing wildfire vs. political keyword overlaps in the BERTopic label (ties default to political/metaphorical).

Polarity	Rank	Peak date	Onset date	Event label	Event type
Negative	1	2025-02-11	2025-01-07	LA Palisades & Eaton Fires	Wildfire
Negative	2	2024-04-23	2024-03-01	CA Insurance Market Collapse	Political
Negative	3	2022-05-03	2022-04-22	Shellenberger vs Newsom / Budget Cuts	Wildfire
Negative	4	2024-11-19	2024-11-05	Post-2024 Election Backlash	Political
Negative	5	2021-02-09	2021-01-28	MTG "Jewish Space Lasers"	Political
Negative	6	2021-07-06	2021-07-04	Recall Newsom / Noem Border Stunt	Political
Negative	7	2023-05-30	2023-05-26	State Farm Exits CA / Debt Ceiling	Wildfire
Negative	8	2020-11-17	2020-11-03	2020 Post-Election / COVID Surge	Wildfire
Negative	9	2023-09-26	2023-09-15	PG&E Undergrounding Failures	Political
Negative	10	2020-08-04	2020-07-15	Emergency Budget Cuts / Fire Season	Wildfire
Positive	1	2022-12-27	2022-12-26	CalFire Tesla Rescue Miracle	Political
Positive	2	2023-12-05	2023-12-05	CAL FIRE Grants & Women in Fire	Wildfire
Positive	3	2024-08-13	2024-07-24	Smokey Bear 80th / Park Fire Containment	Political
Positive	4	2024-12-31	2025-01-07	LA Fires Early Response / FEMA Aid	Wildfire
Positive	5	2023-08-01	2023-08-08	Maui Wildfire Aid / Hawaii Response	Wildfire
Positive	6	2022-11-22	2022-11-08	Favorable Fire Season / Forest Investment	Wildfire
Positive	7	2021-11-30	2021-11-15	Build Back Better / Wildfire Infra	Political
Positive	8	2021-01-12	2021-01-09	CA Wildfire Budget / Erbes Fire Response	Wildfire
Positive	9	2020-09-15	2020-08-17	Western Fire Siege / Incarcerated FF	Wildfire
Positive	10	2021-03-30	2021-03-22	Fire Season Readiness / Rx Burns	Wildfire

Table S6: Corpus-scale annotation summary across the full corpus. **A,** Entity distribution by category, representation and (for government) administrative level. **B,** Stance distribution by entity category. Percentages in Panel A are within category. In Panel A, *Unknown* indicates that the government level could not be determined; An em dash (—) indicates not applicable. Percentages in Panel B are within category.

A. Entity distribution

Category	Representation	Level	Count	%
Government				
	Organizational	Federal	109,042	13.4
		State	237,944	29.3
		Local	135,161	16.6
		Unknown	137,487	16.9
	Individual	Federal	99,033	12.2
		State	70,589	8.7
		Local	15,455	1.9
		Unknown	7,409	0.9
News				
	Organizational	—	130,024	68.5
	Individual	—	59,875	31.5
Non-profit				
	Organizational	—	107,591	90.6
	Individual	—	11,221	9.4
		Total	1,120,831	

B. Stance distribution

Category	Negative	Neutral	Positive	Total
Government	154,617 (19.0%)	484,116 (59.6%)	173,387 (21.3%)	812,120
News	13,525 (7.1%)	156,954 (82.7%)	19,420 (10.2%)	189,899
Non-profit	4,223 (3.6%)	61,327 (51.6%)	53,262 (44.8%)	118,812
Total	172,365	702,397	246,069	1,120,831

Table S7: Estimated marginal means (EMMs) and planned contrasts from mixed-effects models.

Panel a. EMMs: Government level × representation (Fig. 1c)

Level	Representation	EMM	SE	95% CI
Federal	Individual	−0.0457	0.0091	[−0.0636, −0.0279]
Federal	Organizational	0.0946	0.0068	[0.0813, 0.1080]
Local	Individual	0.1830	0.0097	[0.1640, 0.2020]
Local	Organizational	0.2097	0.0045	[0.2008, 0.2186]
State	Individual	0.1291	0.0086	[0.1123, 0.1459]
State	Organizational	0.1954	0.0054	[0.1848, 0.2060]

Type-II LRT (ML) fixed effects

Level	$\chi^2 = 463.27$	$p = 2.53 \times 10^{-101}$
Representation	$\chi^2 = 150.20$	$p = 1.57 \times 10^{-34}$
Level × Representation	$\chi^2 = 58.26$	$p = 2.23 \times 10^{-13}$

Planned contrasts: (Org − Ind) within level (Bonferroni $k = 3$)

Local	$\Delta = +0.0267$	$z = 2.52$	$p_{\text{adj}} = 0.0349$ (*)
State	$\Delta = +0.0663$	$z = 6.58$	$p_{\text{adj}} = 1.34 \times 10^{-11}$ (***)
Federal	$\Delta = +0.1404$	$z = 12.46$	$p_{\text{adj}} = 4.71 \times 10^{-37}$ (***)

Pooled raw means across all government entity-days: Ind $= -0.0308$, *Org* $= 0.1553$, $\Delta = 0.1861$; *representation main-effect* $p_{\text{LRT}} = 1.57 \times 10^{-34}$.

Panel b. EMMs: Representation × intervention (Fig. 1d)

Representation	Intervention	EMM	SE	95% CI
Individual	Non-fire	0.1177	0.0058	[0.1064, 0.1290]
Individual	During fire	0.0183	0.0073	[0.0041, 0.0325]
Organizational	Non-fire	0.1891	0.0035	[0.1822, 0.1960]
Organizational	During fire	0.1648	0.0042	[0.1565, 0.1731]

Type-II LRT (ML) fixed effects

Representation	$\chi^2 = 228.23$	$p = 1.45 \times 10^{-51}$
Intervention	$\chi^2 = 161.79$	$p = 4.60 \times 10^{-37}$
Representation × Intervention	$\chi^2 = 95.02$	$p = 1.88 \times 10^{-22}$

Planned contrasts: (Org − Ind) within context (Bonferroni $k = 2$)

Non-fire	$\Delta = +0.0714$	$z = 10.95$	$p_{\text{adj}} = 1.29 \times 10^{-27}$ (***)
During fire	$\Delta = +0.1465$	$z = 17.80$	$p_{\text{adj}} = 1.29 \times 10^{-70}$ (***)

$n = 151{,}073$ *government entity-day observations (Non-fire: 112,475; During fire: 38,598).*

Table S8: Consolidated severe-fire windows used to define the intervention context in Fig.1d. Windows are derived from Table S1 by grouping co-occurring major incidents into a single acute interval.

Window label	Start	End
Wine Country	2017-10-08	2017-11-17
Thomas	2017-12-04	2018-01-12
Carr+Mendocino	2018-07-23	2018-08-30
Camp+Woolsey	2018-11-08	2018-11-25
Lightning+Glass	2020-08-16	2020-10-20
Dixie+Caldor	2021-07-13	2021-10-25
LA Fires	2025-01-07	2025-01-31