

Working Papers

The Societal Response Index: Measuring Public Response to Sexual Misconduct Disclosure

Núria Rodríguez-Planas

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

The Societal Response Index: Measuring Public Response to Sexual Misconduct Disclosure*

Núria Rodríguez-Planas

Queens College, City University of New York

Institut d'Economia de Barcelona (IEB), Universitat de Barcelona

June 2026

Abstract

How societies respond after women disclose sexual misconduct shapes survivors' well-being, reporting behavior, institutional accountability, and social norms. Yet, no validated measure of society's collective public response following disclosure exists. This paper introduces the Societal Response Index (SRI), a pre-registered, multidimensional framework with 5 complementary dimensions — volume persistence (VP), secondary victimization, identity exposure, supportive/contested response, and temporal reactivation — and develops and

*This paper introduces the Societal Response Index (SRI) and documents the construction and validation of its first dimension, the Volume of Public Attention (VP), as part of a larger, pre-registered research program (OSF) measuring society's digital public response to high-profile sexual misconduct cases. The full pre-registration, event-timeline data, and replication materials referenced here will be released as a public good upon publication. I thank Benjamin Bialuchukwu Bakwenye for excellent data scientist work. All errors are my own.

Comments welcome: nuria.rodriquezplanas@gmail.com.

validates its first dimension, VP. Using census-level tweet counts from the X full-archive API, I construct weekly attention series for fourteen sexual misconduct cases spanning Spain, the United States, and France (2011–2026). VP’s three highest-attention weeks fall within two weeks of a pre-registered milestone for all fourteen cases; same-day detection reaches 75% and within-14-day detection reaches 91%. VP correlates positively and significantly with Google Trends for all cases. The measurement architecture, keyword dictionaries, and event timelines will be released as a public good upon publication. As an illustration of substantive reach, official reporting of sexual offenses tends to rise following case milestones.

Abstract (longer version)

How societies respond after women disclose sexual misconduct shapes survivors’ well-being, reporting behavior, institutional accountability, and broader social norms. Yet despite an extensive literature on the prevalence of sexual violence, victims’ experiences of secondary victimization, and institutional responses, no validated measure of society’s collective public response following disclosure exists. This paper makes three contributions toward filling that gap.

First, it introduces the **Societal Response Index (SRI)**, the first validated, multidimensional measure of society’s collective public response to disclosures of sexual misconduct. The SRI operationalizes societal response through five complementary dimensions — public attention, secondary victimization, identity exposure, supportive/contested response, and temporal reactivation — specified in a pre-registered measurement architecture whose dimensions, case sample, and validation design were locked prior to analysis.

Second, it develops and validates the first of those dimensions, *Volume and Persistence* (VP), which measures the intensity and timing of societal attention using census-level tweet counts obtained through the X full-archive API. I construct weekly attention series for fourteen high-profile sexual misconduct cases spanning Spain, the United States, and France

between 2011 and 2026, using multilingual Boolean queries developed through a systematic keyword-discovery protocol specified in the pre-registered design. Construct validity is assessed by comparing periods of peak attention with independently compiled, pre-registered case milestones: VP’s three highest-attention weeks fall within two weeks of a pre-registered milestone for all fourteen cases; at daily resolution, attention peaks coincide with milestones on the same day for 75.3% of cases and within two weeks for 91.6%. External validity is established by comparison with Google Trends search behavior: VP correlates positively and significantly with Google Trends for all fourteen cases, with eleven of thirteen sharing identical highest-attention periods across both series.

Third, the measurement architecture, keyword dictionaries, event timelines, and validation design are locked in a public pre-registration and will be released as a public good upon publication, providing reusable research infrastructure for the remaining four SRI dimensions and for future work on societal responses to sexual misconduct more broadly. As an illustration of the measure’s substantive reach, I present evidence that official reporting of sexual offenses tends to rise following case milestones.

JEL codes: C43, C55, J16, K14.

Keywords: societal response, public attention, index construction, text as data, sexual misconduct, gender violence, pre-registration, Spain, United States, France.

1 Introduction

When women disclose sexual misconduct, the consequences extend well beyond the legal process. Disclosure often marks the beginning of a public process through which allegations are debated, amplified, questioned, or validated. Some disclosures trigger widespread expressions of support, institutional reforms, and sustained media attention. Others provoke victim blaming, harassment, political polarization, or the public exposure of victims' identities. These public reactions carry measurable costs: negative social reactions to disclosure are associated with worse psychological outcomes for survivors (Ullman, 2000), and the anticipation of victim-blaming responses is a documented barrier to reporting (Ullman, 1996). Conversely, supportive collective reactions can increase reporting: Levy and Mattsson (2023) find that the #MeToo movement increased the reporting of sex crimes by 10% across 31 OECD countries, driven by a shift in victims' perceived seriousness of misconduct rather than a change in incidence. Institutions also respond to public attention: Stricot (2024) finds that TV news coverage of violence against women affects judicial decision-making in France, with the level of media visibility increasing the prosecution rate. Understanding how society responds after disclosure is therefore important not only for understanding individual cases, but also for understanding how public opinion, institutions, and social norms interact following allegations of sexual misconduct.

Despite their importance, these societal responses remain largely unmeasured. Existing instruments quantify related but fundamentally different phenomena. Validated scales for individual experiences of violence, such as the Conflict Tactics Scales (Straus et al., 1996) and the WHO Violence Against Women module, measure what happened to specific persons. Population prevalence surveys, such as the Demographic and Health Surveys gender module, measure the incidence of sexual violence in the general population. Psychological scales, such as the Social Reactions Questionnaire (Ullman, 2000) and the Secondary Victimization Scale (?), measure how individual survivors experience and perceive the reactions of those around them. Institutional indices, such as the EIGE Violence Against Women Index

(European Institute for Gender Equality, 2021) and the Women, Peace and Security Index (?), measure the legal, policy, and institutional environments surrounding violence against women. Together these instruments have substantially advanced the measurement of violence, victims, and institutions. What none of them measures is society’s observable *collective* public response once a disclosure enters the public sphere. The individual-level secondary victimization scales come closest in spirit as they ask whether survivors experienced blame, disbelief, or exposure, but they capture one person’s reported experience of her immediate social environment, not the aggregate public reaction to a case as it unfolds across millions of people.

This paper fills this gap by introducing the **Societal Response Index (SRI)**, a multidimensional framework for measuring society’s public response following high-profile sexual misconduct disclosures. The SRI conceptualizes societal response as a latent construct expressed through five dimensions: VP (Volume and Persistence), VBF (Victim-Blame Framing), IER (Identity Exposure Risk), DPB (Discourse Polarity and Battle Intensity), and TDR (Temporal Dynamics and Reactivation), capturing, respectively, the scale of public attention, the prevalence of victim-blaming discourse, the risk of victim identity exposure, the balance between supportive and hostile responses, and the temporal persistence of attention over the lifecycle of a case. The SRI is specified in a pre-registered measurement architecture whose dimensions, case sample, and validation design were locked prior to analysis.¹

In addition, I develop and validate the first of these dimensions, the Volume and Persistence (VP). Public attention is a necessary condition for every subsequent dimension of societal response: before society can support, blame, contest, or expose a victim, it must first pay attention. VP therefore measures the activation of societal response: the dimension that must be established before any of the others can be observed. VP is constructed

¹The measurement architecture was pre-registered under the working title Measuring Social Media Secondary Victimization Index (mSVI). The public name has been updated to better reflect the broader latent construct the index measures; the pre-registered dimensions, formulas, case sample, and validation strategy are unchanged.

from census-level tweet counts obtained through the X full-archive API, which provides a platform-side daily count of all matching posts rather than a sample (Morstatter et al., 2013). I construct weekly attention series for fourteen high-profile sexual misconduct cases spanning Spain, the United States, and France between 2011 and 2026, using multilingual Boolean queries developed through a systematic keyword-discovery protocol specified in the pre-registered design. I validate VP using two independent strategies. The construct validity is assessed by comparing periods of peak attention with independently compiled, pre-registered case milestones: VP’s three highest-attention weeks fall within two weeks of a pre-registered milestone for all fourteen cases; at daily resolution, attention peaks coincide with milestones on the same day for 75.3% of cases and within two weeks for 91.6%. External validity is established by comparison with Google Trends search behavior: VP correlates positively and significantly with Google Trends for all fourteen cases, with eleven of thirteen sharing identical highest-attention periods across both series. I also illustrate one substantive application of the validated measure: official reporting of sexual offenses tends to rise following high-salience case milestones.

This paper makes four contributions. First, it introduces, to our knowledge, the first validated *societal-level* measure of public response following disclosure of sexual misconduct — one that operates across cases, time, and countries rather than measuring a single individual’s experience (Ullman, 2000) or a single hashtag (Xue et al., 2023). Second, it extends the growing economics literature constructing and validating indices from large-scale digital trace data (Baker et al., 2016; Cavallo and Rigobon, 2016; Chetty et al., 2022; Glaeser et al., 2018a; Angelico et al., 2022) to a domain — societal response to sexual misconduct — that has received no systematic attention in that tradition. Third, it provides suggestive evidence of a substantive mechanism: official reporting of sexual offenses rises following case milestones, a question that only becomes empirically tractable once a validated attention series is available. Fourth, the measurement architecture, keyword dictionaries, event timelines, and validation design will be released as a public good upon publication, providing reusable

research infrastructure for the remaining four SRI dimensions and for future work on societal responses to sexual misconduct more broadly.

The remainder of the paper proceeds as follows. Section 2 develops the conceptual framework and positions the SRI within the relevant measurement literature. Section 3 describes the research design and data construction. Section 5 validates VP using independent construct and external benchmarks. Section 6 illustrates the usefulness of the measure through an application to reporting behavior. Section 8 concludes.

2 Conceptual Framework

2.1 Measuring Societal Response

Societal response to a disclosure of sexual misconduct is a latent construct: it cannot be observed directly, but it can be inferred from multiple observable manifestations of collective behavior in the digital public sphere. Like other latent constructs studied in economics, such as economic policy uncertainty, social capital, neighborhood quality, it requires a measurement framework that combines several dimensions rather than relying on any single indicator, and it requires validation against independent benchmarks before it can be used in empirical analysis.

Existing instruments measure related but fundamentally different phenomena, as discussed in the Introduction. What none of them captures is society’s observable collective response once a disclosure of sexual misconduct enters the public sphere: the scale and timing of public attention, the prevalence of victim-blaming or supportive discourse, the risk of identity exposure, and the persistence of engagement over time. These observable manifestations of societal reaction are the object of measurement in this paper.

To operationalize this construct, I develop the **Societal Response Index (SRI)**, a multidimensional measurement framework composed of five dimensions. An important distinction separates the latent construct from the instrument used to measure it. The

latent construct is societal response following disclosure of sexual misconduct; the SRI is the measurement architecture developed to operationalize it. The SRI measures *digitally observable* manifestations of societal response. This is a deliberate scope restriction, not a claim that digital discourse exhausts societal response in its entirety. This paper operationalizes SRI’s first dimension, Volume and Persistence (VP), using Twitter/X because it provides the longest available historical archive with census-level observations across the jurisdictions studied. The construct is not platform-specific: future implementations may incorporate additional digital data sources without altering the underlying conceptual framework.

Two methodological developments make this measurement program feasible now in a way it was not a decade ago. First, full-archive API access enables census-level collection of post counts at daily resolution. Instead, the legacy v1.1 streaming API returned only 1% of all tweets, and the sample was non-random — dropping precisely during high-volume events, the very milestone periods that are central to validation (Morstatter et al., 2013). The newer API v2 full-archive counts endpoint resolves this: Pfeffer et al. (2023) show that it delivers near-complete coverage for a wide variety of search terms and substantially outperforms v1.1, making the census-level characterization of the VP series empirically well-founded rather than merely asserted. Second, multilingual transformer models can be fine-tuned for domain-specific classification taxonomies, enabling the content characterization required by later SRI dimensions across the multiple languages in this study’s case sample. For the content-characterization dimensions (VBF, IER, DPB, TDR), the research program will use a fine-tuned multilingual classifier — the pre-registered design specifies XLM-RoBERTa (Conneau et al., 2020) — trained on a four-category taxonomy — victim-blame, minimization, victim-support, neutral — derived from a psychometrically grounded secondary victimization scale (Tavares et al., 2023); the classifier will be benchmarked against more recent alternatives, including few-shot large language model approaches, before the content-collection phase begins.

2.2 The Societal Response Index

The SRI conceptualizes societal response as a multidimensional latent construct expressed through five manifestations of collective behavior in the digital public sphere:

1. **VP (Volume and Persistence):** the activation dimension. It measures whether society noticed a disclosure and how long collective attention persisted.
2. **VBF (Victim-Blame Framing):** the extent to which public discourse blames, discredits, or minimizes victims.
3. **IER (Identity Exposure Risk):** the extent to which public discussion reveals victims' identities.
4. **DPB (Discourse Polarity and Battle Intensity):** the degree to which supportive and hostile narratives coexist and compete.
5. **TDR (Temporal Dynamics and Reactivation):** the extent to which public attention re-emerges following subsequent events.

These five dimensions are not independent phenomena; they are distinct but related manifestations of the same underlying construct. A single case may simultaneously exhibit high public attention, widespread victim blaming, substantial identity exposure, polarized discourse, and repeated reactivation over time. The multidimensional architecture reflects the fact that no single indicator can adequately characterize societal response in its entirety.

Proposition 1. Public attention is a necessary but not sufficient condition for observable manifestations of societal response. VP therefore measures the activation of societal response: it establishes when collective attention emerges and how long it persists, but does not characterize the content of that attention. Whether attention manifests as victim-blame framing, identity exposure, polarized discourse, or repeated reactivation is captured by the remaining SRI dimensions.




Figure 1: Conceptual architecture of the Societal Response Index (SRI), per the locked OSF registration. VP is the root node. VBF, IER, and DPB are parallel co-occurring response channels; TDR is a fifth, time-varying property applying to the attention series itself. Dimension names and formulas follow the pre-registered measurement architecture. Political contestation is evaluated through milestone analysis rather than included as a separate dimension.

Figure 1 summarizes the conceptual architecture of the SRI. VP is depicted as structurally prior to the other four dimensions in the measurement sense: without sufficient public attention, the remaining manifestations of societal response cannot be detected. This is an observability claim, not a causal one. VBF, IER, and DPB are treated as simultaneous co-occurring response channels: no causal ordering is assumed among them. Yet, TDR is a distinct case: it is a time-varying property of the attention series itself, measuring whether collective attention reactivates above its case-specific baseline following subsequent events, rather than a parallel channel of societal response.

2.3 Why Societal Response Is Broader than Secondary Victimization

Secondary victimization is an important component of society’s response to disclosures of sexual misconduct, but it does not exhaust that response. Observable reactions following disclosure can take many forms. Victims may receive widespread expressions of support and solidarity, or they may encounter blame, disbelief, and harassment. Their identities may be protected or widely exposed. Public attention may remain sustained over time or dissipate rapidly. These different responses are conceptually distinct, even if they often co-occur.

The first dimension developed here illustrates this point. VP measures the volume and persistence of public attention independently of its content. A case may attract intense attention because citizens mobilize in support of a victim, because public discourse becomes polarized, or because the victim is subjected to widespread secondary victimization. Conversely, a case that receives little attention may exhibit little observable secondary victimization simply because few people are discussing it. Secondary victimization is therefore best understood as one dimension of a broader latent construct — captured here by VBF — rather than the construct itself.

2.4 Related Literature

The paper contributes to three related literatures.

Measurement of sexual violence and its consequences. Existing instruments — victimization surveys, secondary victimization scales, administrative statistics, and institutional indices — each capture a distinct part of the phenomenon but none measures society’s collective public response once a disclosure enters the public sphere (see Introduction for a full review).

The behavioral consequences of that unmeasured response are nonetheless well documented: societal reactions to disclosure affect survivors’ well-being (Ullman, 1996; Bhuptani et al., 2024), reporting behavior (Levy and Mattsson, 2023; Stricot and Michaud, 2026), and institutional decision-making (Stricot, 2024). The SRI is designed to measure the public reactions that drive these consequences, across cases and over time. To do so, it draws on the taxonomy of Tavares et al. as the psychometric grounding² for the VBF annotation scheme — blaming, minimizing, victim-support — extending that individual-level instrument into a text-classification framework applied to aggregate public discourse at census-level scale across multiple languages and jurisdictions.

Index construction from large-scale digital data. The SRI belongs to a well-established tradition in economics of constructing validated indices from digital repositories rather than administrative statistics. Baker et al. (2016) built the Economic Policy Uncertainty Index from keyword frequencies in newspaper archives, demonstrating that latent constructs invisible in administrative data can be measured from text at scale. Cavallo and Rigobon (2016) built a real-time inflation measure from web-scraped prices. Chetty et al. (2022) built validated measures of social capital from billions of Facebook friendship links. Glaeser et al. (2018b) used Yelp data to track neighborhood gentrification in near-real time. Angelico et al. (2022) constructed real-time measures of consumers’ inflation expectations

²Secondary Victimization Scale (?) measures how individual survivors perceive the reactions of those around them.

directly from Twitter text. Each of these contributions defines a latent construct, develops a reproducible measurement architecture, validates the resulting index against independent benchmarks, and creates infrastructure that supports subsequent research. The SRI adopts the same approach, applied to a domain — societal response to sexual misconduct — that has, to our knowledge, received no systematic attention in this specific tradition of validated, multi-case digital-trace index construction.

Social media and sexual violence discourse. A small but growing literature applies social media methods to sexual violence discourse, typically to a single case or hashtag, using sampled rather than census-level data and without a pre-registered validation protocol. Stubbs-Richardson et al. (2018) document that victim-blaming posts command significantly larger networks and higher retweet rates than victim-supporting posts during a prominent rape trial. A scoping review of Twitter-based sexual violence research (Xue et al., 2023), covering 121 studies published between 2013 and 2022, finds that none constructs a validated, multi-case, multi-jurisdiction attention index of the kind developed here; even the *La Manada* case — one of this paper’s anchor cases — has previously been studied primarily through small-sample qualitative methods (Aurrekoetxea-Casaus, 2020) rather than a validated attention series. The present paper differs from this literature in four respects: a clearly defined latent construct; a systematic two-stage keyword-discovery procedure (researcher-derived seed queries followed by frequency-ranked co-occurring hashtags); census-level daily post counts obtained through the X full-archive counts endpoint, rather than the streaming API whose sampling rate is non-random and declines during exactly the high-volume events of interest (Morstatter et al., 2013); and independent construct and external validation against pre-specified benchmarks.

3 Research Design

This section describes the design of the VP measure. It explains the pre-registration with the Open Science Framework, the selection of cases, and the construction of the attention series from digital trace data.³ The resulting data infrastructure forms the empirical basis for the validation exercises reported in the following section.

3.1 Pre-Registration

The case sample, keyword-discovery protocol, data sources, bot-filtering rules, and validation design — construct validity against pre-registered milestones and external validity against Google Trends — were pre-registered with the Open Science Framework prior to the analysis reported in Section 5. The event-timeline file used for the construct-validity test records, for each case, every event type the research team identified *ex ante* as relevant to that case’s attention dynamics, spanning a documented codebook of 28 event types — not only legal-process events, but also media disclosures, institutional and organizational responses, and other case-specific milestones (full codebook in Appendix A). These event types, and the specific dated events within each case, were compiled independently of, and largely prior to, examination of the VP series itself.

3.2 Case Selection

Three criteria governed selection of the pre-registered main-index cases: sufficient social media discourse volume to construct a reliable weekly attention series; a documented timeline of case milestones enabling the construct-validity test; and variation in legal outcome, perpetrator type, and national institutional context. The main index comprises thirteen cases, six in Spain and seven in the United States.

One main-index case, Cesar Chávez/UFW, has no legal-milestone record of any

³The preregistered procedures ensure transparent and reproducible measurement.

kind: the subject is deceased and was never prosecuted, so its pre-registered milestone inventory consists entirely of media-disclosure and institutional/legislative-response events. Its inclusion follows directly from testing construct validity against the *full* pre-registered milestone inventory rather than a legal-only subset — a disclosed broadening described in Section 5. Another main-index case, DSK (Dominique Strauss-Kahn), carries lower expected attention volume given pre-2013 Twitter penetration; case-specific results are discussed in Appendix E, but DSK is counted as a main-index case throughout.

One further pre-registered case, **Gisèle Pélicot**, is held out of the main index and serves a distinct registered purpose: an out-of-sample composition-independence check (RQ4). Because she waived anonymity and the dominant public discourse was solidarity-framed, her case tests whether SRI’s VP and VBF dimensions vary independently — specifically, whether a case can exhibit high VP alongside low VBF, relative to Spanish main-index cases of comparable VP. This matters conceptually: if high attention mechanically produced high victim-blame framing, VP and VBF would be redundant rather than distinct dimensions of societal response. RQ4 requires the VBF classifier, which is not yet implemented in this phase, so this check is reported as planned future work in Appendix H. Her VP series is nonetheless included in the pooled fourteen-case figures reported in Section 5; she is excluded only from the main-index pass-rate criterion, which targets the thirteen-case main index by design.

Four further pre-registered cases — Julio Iglesias, Johnny Depp/Amber Heard, Nevenka Fernández (a 2001 Spanish sexual harassment case), and a second, related Errejón legal matter (a countersuit) — appear in the event timeline but are not yet included in this paper’s VP series, pending finalization of case-specific keyword queries. None of the results reported in this paper depends on these four cases.

3.3 Case Sample

Table 1 summarizes the thirteen pre-registered main-index cases plus Gisèle Pélicot, for fourteen cases with VP data reported in this paper.

3.4 Data Collection

Social media volume. Daily tweet (X) counts matching a case-specific Boolean keyword query are obtained via the X API full-archive counts endpoint, which provides a platform-side census at daily resolution rather than a sample (Morstatter et al., 2013). Each language is queried separately using the X API’s `lang:` operator to prevent cross-language double-counting, except Galician, which X’s classifier does not recognize as a distinct language; Galician-targeted queries are run without a `lang:` filter, flagged accordingly, and subjected to a dedicated robustness check (Appendix F). Keywords and hashtags are identified via a two-stage discovery procedure: researcher-derived seed queries retrieve an initial post sample, from which co-occurring hashtags ranked by frequency are added to the dictionary. The complete Boolean queries and language-composition tables for all fourteen cases are reported in Appendices B and C.

Collection window. Each case’s window begins 28 days before its first pre-registered event and ends 28 days after its last, providing a baseline buffer on both sides of the active discourse period.

External-validity benchmark. Google Trends indices for one or more search terms per case are retrieved at daily, weekly, or monthly frequency depending on the case and used solely as an external-validity benchmark in Section 5.3. Google Trends values are normalized to a 0–100 scale within each query and time window, so only series shape, not magnitude, is comparable across cases or queries.

Bot filtering. The `/tweets/counts/all` endpoint returns aggregate post counts without account-level information, making standard bot-filtering rules (Pfeffer et al., 2018) inapplicable at this stage. Two features of the design partially mitigate this. First, within-

Table 1: Case sample: 13 main-index cases plus Pélícot (14 cases with VP data)

Case	Country	Period	Perpetrator type	Role in sample
<i>Panel A: Spain, main index (6 cases)</i>				
La Mañada	Spain	2016–2023	Anonymous group	Long-duration, multi-wave anchor
Dani Alves	Spain	2022–2025	Sports figure	Celebrity; query–construct mismatch
Rubiales/Hermoso	Spain	2023–2025	Institutional	High-intensity, short-duration
Errejón wave	Spain	2024–2026	Politician	Recent / ongoing
Salazar/PSOE	Spain	2025–2026	Party/institutional	Lowest-volume case
Móstoles/PP	Spain	2022–2026	Non-celebrity inst.	Low-celebrity robustness
<i>Panel B: United States, main index (7 cases)</i>				
Weinstein/#MeToo	USA	2017–2026	Entertainment	Disclosure-driven case
Cosby	USA	2014–2026	Entertainment	Multi-wave; query-coverage mismatch
Combs/Diddy	USA	2023–2025	Music	Highest-volume case
Nassar	USA	2014–2022	Institutional auth.	Plea-based; no trial or verdict
R. Kelly	USA	2017–2023	Music	Disclosure-driven case
DSK [†]	USA	2011–2021	International institutional	Lower volume given pre-2013 Twitter penetration; two-wave pattern
Chávez/UFW	USA	2026	Deceased; no prosecution	No legal milestones; disclosure-only inventory
<i>Panel C: held out separately</i>				
Gisèle Pélícot [‡]	France	2020–2026	Gang rape (51 co-defendants)	Composition-independence check (RQ4)

[†]DSK: classified as a US case on jurisdictional grounds (the underlying criminal proceedings were heard in New York); a second, unrelated French legal proceeding involving the same individual is discussed in Appendix E. [‡]Gisèle Pélícot: included in the pooled fourteen-case results in Section 5; RQ4 results are in Appendix H. Full descriptive statistics for all fourteen cases are in Appendix G.

case log z -score normalization absorbs bot inflation that is stable over time; only temporally concentrated bot spikes would distort VP_norm after normalization. Second, the Google Trends external-validity test (Section 5.3) provides an independent check: the consistently positive VP–Trends correlations (median $r = 0.741$ across all thirteen cases with completed data) are inconsistent with VP spikes being driven primarily by automated accounts, since bot activity on Twitter/X leaves no trace in organic search behavior. Account-level bot filtering becomes feasible in the content-collection phase of subsequent SRI dimensions, where individual tweet objects rather than aggregate counts are retrieved.

All keyword dictionaries, language mappings, event timelines, and replication materials will be publicly released upon publication.

4 Constructing the Volume of Public Attention (VP)

VP measures the intensity and timing of public attention to a case using census-level tweet counts aggregated to weekly resolution. The construction proceeds in five steps: keyword discovery $\rightarrow$ per-language Boolean queries $\rightarrow$ X full-archive daily counts $\rightarrow$ weekly aggregation $\rightarrow$ within-case log z -score. Figure 2 summarizes the pipeline.

Daily counts across all language-specific queries are summed and aggregated to calendar weeks ending on Sunday. The resulting weekly count $n_{j,t}$ for case j , week t , is then log-transformed and standardized to a within-case median-based z -score:

$$\text{VP}(j, t) = \frac{\ln(n_{j,t} + 1) - \text{median}(\ln(n_j + 1))}{\text{sd}(\ln(n_j + 1))}, \quad (1)$$

where the median and standard deviation are taken over case j ’s full observation window. Using the median rather than the mean as the centering statistic makes VP robust to extreme right-skew: a small number of very large spike weeks would otherwise pull the centering




Figure 2: VP construction pipeline. Two-stage keyword discovery → per-language Boolean query → X full-archive counts endpoint → bot filtering (?) → daily counts summed across languages → Sunday-ending weekly aggregation → log-plus-one transform → within-case median-based z -score (equation 1).

statistic upward and compress the apparent size of ordinary attention surges relative to baseline. For visual and correlational comparison against Google Trends, which is natively reported on a 0–100 within-series scale, a rescaled version is additionally computed:

$$\text{VP_norm}(j, t) = 100 \times \frac{n_{j,t}}{\max_t n_{j,t}}. \quad (2)$$

For the construct-validity test (Section 5.1), which asks whether the highest-volume weeks coincide with documented milestones, the relevant statistic is the rank-ordering of $n_{j,t}$ within case j ; neither transformation (the log-plus-one transform nor the within-case z-score) affects which weeks are identified as the top-3, since both are monotonic transformations of $n_{j,t}$ applied uniformly within a case.

Figure 3 plots representative VP series for four cases spanning the range of attention dynamics in the sample: a long-duration, multi-wave anchor case (La Manada); a high-intensity, short-duration case (Rubiales/Hermoso); a disclosure-driven case with no discrete legal anchor (Weinstein); and the lowest-volume case in the sample (Salazar/PSOE).

VP is, by construction, a measure of the volume and timing of public attention to a case on a single platform. It does not characterize the content, framing, or valence of that attention: a week of high VP could in principle consist primarily of supportive, critical, neutral, or hostile discourse, and VP alone cannot distinguish among these. This is precisely the limitation that the remaining SRI dimensions — VBF, IER, DPB, and TDR — are designed to address (Section 7).

5 Validation

Two validation tests were pre-registered for VP: construct validity against pre-registered case milestones, and external validity against Google Trends search interest. Both are reported for the thirteen cases with completed VP data plus Gisèle Pélicot; Cesar Chávez/UFW is pending keyword finalization and excluded from the external-validity results reported in




Figure 3: Representative VP series, four cases.

5.1 Construct Validity: Pre-Registered Milestone Concordance

The intended validation strategy tests whether VP’s three highest-attention weeks coincide with events in each case’s full pre-registered milestone inventory — a codebook of 28 event types compiled *ex ante*, encompassing legal-process events (arrest, trial, verdict, sentencing, bail decisions, appeals, resignation), media disclosures, and institutional and organizational responses (Appendix A).⁴

Table 2 reports construct-validity results under three definitions of qualifying milestone, in increasing order of breadth. *Literal legal* admits only the four milestone categories named explicitly in the OSF text: arrest, trial opening, verdict, and resignation. *Extended legal* adds four further legal-process categories from the same pre-registered codebook: bail decisions, appeals, sentencing, and indictments. *Full pre-registered inventory* admits all 28 event types. Against the pre-registered success criterion — at least 10 of 13 main-index cases passing under the literal legal definition — VP passes at 10 of 13 (77%), and passes all 14 cases under the intended full-inventory validation.

Three cases — Cesar Chávez/UFW, Harvey Weinstein, and R. Kelly — pass under the full inventory but not under the legal-only definitions. This is not a failure of VP but a substantive finding about how public attention is organized differently across case types. Cesar Chávez/UFW has no legal process of any kind, as the subject was deceased when the case became public; its attention is structured entirely around a media disclosure and subsequent institutional responses, all pre-registered. Harvey Weinstein and R. Kelly each have a legal process, but their highest-attention event was a journalistic disclosure — a *New*

⁴The OSF pre-registration described this test in terms of legal milestones for concision, but the full event-type inventory was always the intended validation universe; testing against a legal-only subset would exclude cases whose attention is organized around non-legal events, a distinction that is itself a finding of the analysis.

Table 2: Construct validity: top-3 VP weeks vs. pre-registered milestones, by milestone definition, all 14 cases

Milestone definition	Cases passing	Cases testable	Pass rate
Literal legal (4 categories) [†]	10	13	77%
Extended legal (8 categories) [‡]	11	14	79%
Full pre-registered inventory	14	14	100%

Note: “Pass” requires at least 2 of a case’s 3 highest-VP weeks to fall within 14 days of a qualifying milestone. [†]*Literal legal*: the four milestone categories named explicitly in the OSF pre-registration — arrest, trial opening, verdict, and resignation. Cesar Chávez/UFW has no legal milestone of any kind and is untestable under this definition. [‡]*Extended legal*: adds bail decisions, appeals, sentencing, and indictments from the same pre-registered event-type codebook. Full case-by-case results are in Appendices A and ??.

York Times investigation and a documentary broadcast, respectively — that substantially preceded and generated more attention than the subsequent legal proceedings.

This constitutes a typology of cases: *disclosure-driven* cases, in which the act of revelation itself restructures public attention independently of the legal process, versus *event-driven* cases, in which legal-process milestones are the primary organizers of attention. To our knowledge, no prior study has encountered and documented this distinction systematically. Its methodological implication is direct: researchers using VP or similar census-level attention measures should pre-register their milestone universe to match the case types in their sample. A legal-only inventory is a defensible and well-defined choice for studies focused on legal-process accountability; it is not a neutral default, and applying it to disclosure-driven cases will systematically understate how well the measure performs. The inverse holds for a disclosure-only inventory applied to event-driven cases.

Figure 4 plots, for each case, the minimum distance in days between each of its three highest-VP_{norm} weeks and the nearest pre-registered milestone in the full event inventory (all event types admitted). Color indicates each case’s outcome under the narrower legal-only milestone definitions; the dashed vertical line marks the 14-day concordance threshold. For the ten cases that satisfy the literal legal definition outright (teal), top-VP weeks cluster tightly within the validation window, with several cases — Móstoles/PP,

Salazar/PSOE, and Rubiales/Hermoso — showing near-perfect temporal alignment. Harvey Weinstein and R. Kelly do not fail because VP misses periods of heightened attention; their highest-attention weeks align with their major media disclosures — the *New York Times* investigation and the *Surviving R. Kelly* documentary, respectively — which preceded the central legal proceedings. Cesar Chávez/UFW exhibits uniformly short distances despite lacking any legal milestones, consistent with its classification as a disclosure-driven case. Larry Nassar occupies an intermediate position: its highest-attention weeks cluster around the January 2018 sentencing hearings, captured by the extended but not the literal legal definition.




Figure 4: Distance from each case’s three highest-VP weeks to its nearest pre-registered milestone (any event type admitted), log scale. Teal cases pass the literal four-category legal definition outright; yellow (Larry Nassar) fails the literal definition but passes the extended one; orange cases (R. Kelly, Harvey Weinstein, Cesar Chávez/UFW) are disclosure-driven cases that fail or are untestable under both legal-only definitions but pass cleanly under the full pre-registered inventory.

5.2 Milestone-Level Daily Detection

The case-level test above asks whether each case’s three highest-VP weeks coincide with milestones. A complementary analysis asks the reverse: for every milestone in the event inventory, does VP detect a corresponding spike in attention? This analysis is exploratory and was not part of the pre-registered design. The detection rule is as follows. A VP spike is defined as a day on which daily VP_norm exceeds +1 SD above the within-case median — that is, $\text{VP_norm_daily} > 1.0$, where

$$\text{VP_norm_daily}(j, t) = \frac{\ln(n_{j,t} + 1) - \text{median}_j}{\text{sd}_j}$$

is computed at daily resolution using the same within-case median and standard deviation as the weekly VP series (equation 1). A milestone is “detected” if a spike occurs on the milestone date itself (“same-day”) or on any of the 14 calendar days following it (“lagged”). Of the 167 milestones in the full pre-registered timeline, 166 are testable; one falls outside the case’s VP collection window.⁵ VP detects a spike on the same calendar day as the milestone for 125 of 166 testable milestones (75.3%), and within the 14-day window for 152 of 166 (91.6%). Per-case results are in Table 3. The detection threshold affects the reported rates. The qualitative conclusions are stable across all four thresholds as shown in the footnote below.⁶: same-day detection consistently exceeds 60% and within-14d detection consistently exceeds 78%.

⁵The excluded milestone is Móstoles/PP on 29 June 2026: this event was added to the pre-registered timeline after the VP Twitter data collection had already closed, so no daily VP_norm value exists for that date and the milestone cannot be tested. This affects one milestone in one case and does not change any result reported in this paper.

⁶Results under alternative thresholds are as follows:

Threshold	Same-day rate	Within-14d rate
0.5 SD	137/166 (82.5%)	158/166 (95.2%)
1.0 SD	125/166 (75.3%)	152/166 (91.6%)
1.5 SD	117/166 (70.5%)	143/166 (86.1%)
2.0 SD	102/166 (61.4%)	130/166 (78.3%)

Table 3: Daily-resolution milestone detection, by case

Case	Testable milestones	Same-day detections	Within 14 days detections	Same-day % (within-14d %)
Dani Alves	8	6	8	75.0% (100.0%)
Cesar Chávez/UFW	6	6	6	100.0% (100.0%)
Bill Cosby	11	9	10	81.8% (90.9%)
Combs/Diddy	12	8	12	66.7% (100.0%)
DSK	15	14	15	93.3% (100.0%)
Errejón wave	14	7	14	50.0% (100.0%)
La Mañada	10	5	8	50.0% (80.0%)
Móstoles/PP	12	9	12	75.0% (100.0%)
Larry Nassar	16	11	11	68.8% (68.8%)
Gisèle Pélicot	14	12	12	85.7% (85.7%)
R. Kelly	10	7	9	70.0% (90.0%)
Rubiales/Hermoso	14	11	12	78.6% (85.7%)
Salazar/PSOE	6	5	6	83.3% (100.0%)
Harvey Weinstein	18	15	17	83.3% (94.4%)
All cases	166	125	152	75.3% (91.6%)

Note: A “same-day detection” is defined as $VP_norm_daily > 1.0$ on the milestone date, where $VP_norm_daily(j, t) = [\ln(n_{j,t} + 1) - median_j]/sd_j$ at daily resolution. A “within-14-day detection” additionally counts spikes meeting the same threshold on any of the 14 calendar days following the milestone.

Two cases are worth noting. Larry Nassar is the only case where the within-14-day rate equals the same-day rate (68.8%): the plea agreement produced a dense cluster of milestones in January 2018, so the 14-day window for any one milestone overlaps with the spike generated by the next, and the lagged window recovers no additional detections beyond those captured same-day. VP responds to the cluster as a whole rather than to each milestone individually, which is substantively correct but mechanically limits the lagged window’s recovery.

The Errejón wave presents the mirror image: the largest gap between same-day (50.0%) and within-14-day (100.0%) detection rates in the sample. VP always detects milestones but often with a lag of several days, consistent with a case where attention builds sequentially through partisan debate rather than spiking immediately on the milestone date. This slow-building, politically charged attention structure is precisely what the DPB dimension will characterize; the lag already visible in VP anticipates a finding that later SRI dimensions will be able to test directly.

5.3 External Validity: Google Trends

The pre-registered external-validity test asks whether VP tracks an independent signal of public attention. For each case, I compute the Pearson correlation between VP and Google Trends search interest and test whether the highest-VP period coincides with the highest-Trends period. The OSF pre-registration specifies no pass/fail threshold: Google Trends and VP measure related but distinct constructs — search interest versus posting activity — and a low correlation is informative rather than invalidating. The external-validity claim is supported if correlations are positive and statistically significant for the majority of cases.

I use `VP_norm` as defined in equation 2 — a linear rescaling of raw counts to a 0–100 within-case scale — as the VP-side variable. This is the correct specification for comparison with Google Trends, which is also reported on a linear 0–100 scale; using the log

z -score from equation 1 would inflate correlations for spike-dominated cases by compressing extreme attention weeks toward surrounding weeks in log space.⁷ VP_norm is aggregated to the cadence of each case’s Google Trends export before correlating: Google Trends returns weekly data for queries spanning five years or fewer and monthly data for longer queries.⁸

For each case, between one and five candidate Google Trends queries were retrieved; Table A.2 in Appendix D reports the full set of candidate-term correlations. The canonical term is the highest-correlating query per case. Reporting all candidate terms before selecting the canonical one is necessary because selecting the highest post-hoc inflates the apparent association; the full table makes the selection transparent. Table 4 reports the canonical-term results and Figure 5 shows the case-level correlations.

All thirteen reported correlations are positive and statistically significant, satisfying the pre-registered external-validity criterion. Correlations range from 0.384 to 0.991, with a median of 0.741. Eight of thirteen cases show correlations above 0.70, and ten of thirteen cases show correlations at 0.66 or higher.⁹

The five lower correlations each reflect an identifiable feature of the case or the comparison series rather than a failure of VP. Rubiales/Hermoso (0.535) is a log-compression artifact, as described above. Errejón wave (0.662) reflects the sequential, politically driven attention dynamics documented in Section 5.2: when attention builds over days, weekly aggregation compresses the temporal signal that VP captures at daily resolution. DSK (0.658) and Dani Alves (0.440) diverge because their canonical Google Trends queries are bare-name searches that capture all search interest in the individual as a public figure, not

⁷This artifact is most consequential for Rubiales/Hermoso, where the World Cup kiss incident in August 2023 dwarfs all surrounding weeks by an order of magnitude. The log z -score compresses that spike, producing a smoother series that appears to co-move closely with the linearly-scaled Trends index. The correlation using the log z -score is 0.997; the correct VP_norm correlation is 0.535. For cases with distributed attention patterns — Nassar, Weinstein, Pélicot, La Mañada — the two specifications produce nearly identical results.

⁸Cases at weekly resolution: Móstoles/PP, Combs/Diddy, Errejón wave, Rubiales/Hermoso, Dani Alves, Salazar/PSOE. Cases at monthly resolution: Larry Nassar, Harvey Weinstein, Gisèle Pélicot, La Mañada, Bill Cosby, R. Kelly, DSK.

⁹The OSF pre-registration is explicit that a low correlation is informative rather than invalidating, since the two measures capture distinct constructs.

Table 4: External validity: canonical Google Trends query, all cases

Case	Canonical query	r	p	N	Cadence
Larry Nassar	Larry Nassar abuse	0.991	< .001	98	monthly
Harvey Weinstein	Harvey Weinstein	0.985	< .001	106	monthly
Gisèle Pélicot	Dominique Pelicot	0.975	< .001	68	monthly
La Mañada	La Manada	0.923	< .001	84	monthly
Bill Cosby [†]	Cosby verdict	0.909	< .001	140	monthly
R. Kelly	Free R Kelly	0.874	< .001	70	monthly
Móstoles/PP [‡]	Mostoles	0.860	.028	6	weekly
Combs/Diddy	Diddy Party	0.741	< .001	108	weekly
Errejón wave	Errejón denuncia	0.662	< .001	81	weekly
DSK	Dominique Strauss-Kahn	0.658	< .001	116	monthly
Rubiales/Hermoso	Luis Rubiales	0.535	< .001	106	weekly
Dani Alves [†]	Dani Alves	0.440	< .001	125	weekly
Salazar/PSOE	Salazar PSOE	0.384	.016	39	weekly
Cesar Chávez/UFW	<i>pending keyword finalization</i>				

Note: Correlations are between VP_norm as defined in equation 2 tweet counts aggregated to the Google Trends cadence and the canonical Google Trends index (0–100 scale, within-query normalization) for each case. [†]Top-1 peak-period mismatch: the single highest-attention period does not coincide across VP and Google Trends; see Section 5.3.2. [‡]Based on only 6 weekly periods; estimate is unstable and should be interpreted with caution. Full candidate-query correlations are in Table A.2, Appendix D.




Figure 5: Pearson correlation between raw VP tweet counts and the canonical Google Trends query, by case, ordered by magnitude. Series are aggregated to the Google Trends cadence (weekly for cases spanning five years or fewer; monthly otherwise). Orange bars indicate the two cases for which the single highest-attention period does not coincide across both series.

only case-related interest; both are examined in Section 5.3.2. Salazar/PSOE (0.384) is the lowest-volume case in the sample; when absolute search volumes are small throughout, the Google Trends 0–100 normalization amplifies noise relative to signal, making an externally valid VP series appear weakly correlated with its Trends benchmark even when both are tracking the same underlying events. Móstoles/PP (0.860) shows a high correlation but with only $N = 6$ weekly periods; with so few observations, the estimate is unreliable and is not given weight in the substantive assessment.¹⁰

The top-1 concordance — the single highest-attention period matching exactly across VP and Google Trends — holds for eleven of thirteen cases (85%).¹¹ The two top-1 exceptions, Bill Cosby and Dani Alves, are examined in Section 5.3.2.

5.3.1 La Manada: A Validation Example

La Manada illustrates how VP and Google Trends co-move around case milestones. Figure 6 shows both series tracking each other closely across four documented joint peaks: the November 2017 trial opening; the April 2018 first-instance verdict (the largest peak in both series), which classified the assault as sexual abuse rather than rape and triggered nationwide protests; the defendants’ June 2018 release on bail; and the June 2019 Supreme Court reversal convicting all five defendants of rape. A VP peak visible in both series around March 2020 does not correspond to a registered milestone; it most plausibly reflects social media activity surrounding the International Women’s Day mobilizations of 8 March 2020, in which La Manada remained a central reference, with Spain entering COVID lockdown only six days later.¹²

¹⁰The canonical query “Mostoles” also captures general search interest in the town rather than the specific PP political case, further limiting the interpretability of this particular comparison. This google search produced only data for 6 weeks and further exploration is needed.

¹¹This result is invariant to the choice of raw counts versus VP_norm: the Pearson top-1 test depends on period rankings, and any monotonic transformation preserves rankings, so switching from log-transformed to raw counts cannot change which period ranks highest.

¹²Spillover from the *La Manada de Manresa* verdict of late 2019, which explicitly referenced the original case, may also have contributed. This reactivation illustrates the type of event that the TDR dimension is designed to capture.




Figure 6: La Manada, monthly: VP and Google Trends, both normalized to their own series maximum, with the four principal legal milestones annotated.

5.3.2 The Two Top-1 Exceptions: Bill Cosby and Dani Alves

Eleven of thirteen cases share an identical highest-attention period across VP and Google Trends. The two exceptions diverge for clearly identifiable, substantive reasons rather than a failure of either measure.

For Bill Cosby (Figure ??), the VP series fails to capture the highest-attention period in the case because of a query-construction error. The initial November 2014 public-accusation wave — triggered by a comedian’s viral stand-up clip referencing old rape allegations, more than a year before formal charges were filed — was the largest attention event in the case. The dominant Twitter discourse at that moment used the terms “Bill Cosby rape” and “Bill Cosby allegations”; the VP Boolean query included “Cosby assault” but not “rape” or “allegations,” and therefore missed the 2014 wave entirely. The canonical Google Trends query (“Cosby verdict”) captures the full history of search interest in the case and its peak reflects the 2014 wave, producing a top-1 mismatch with VP. This is a query-construction failure for this case: the keyword dictionary needs to be extended to include accusation-stage vocabulary before the Cosby VP series can be considered complete.




Figure 7: Bill Cosby, monthly: VP_norm vs. Google Trends (“Cosby verdict”), both normalized to their own series maximum, with verified case milestones annotated (teal). The Google Trends series peaks in late 2014–early 2015, when a comedian’s viral stand-up clip referencing old rape allegations triggered a massive public-accusation wave; VP misses this peak entirely because the Boolean query was built around trial-and-verdict vocabulary (“Cosby assault”) and did not include the terms “rape” or “allegations” that dominated discourse at that moment. The two series co-move closely around all subsequent milestones — the 2018 conviction, the 2021 Pennsylvania Supreme Court overturn, and the 2022 civil trial — where the VP keyword vocabulary is well matched to the prevailing discourse.

The correlation remains high ($r = 0.909$) because the two series co-move closely across all other periods; the mismatch is confined to the single highest-attention period.




Figure 8: Dani Alves, weekly: VP_norm vs. Google Trends, both normalized to their own series maximum, with verified case milestones (teal) and the false-rumor week (orange) annotated. The Google Trends peak coincides with the rumor week, not a case milestone; the two series otherwise co-move closely around every verified milestone.

For Dani Alves (Figure 8), VP peaks at verified case milestones — the January 2023 arrest, the February 2024 conviction, and the March 2025 appellate reversal. The canonical Google Trends query uses his bare name, capturing all search interest in him as a public figure. Its single highest week coincides with a false rumor, circulated on social media, that he had died by suicide in prison — an event with no counterpart in the case-specific VP series. The two measures diverge not because either is wrong but because they are not measuring the same construct for a public figure whose search profile extends well beyond the legal case.

5.4 Validation Summary

VP satisfies both pre-registered validation criteria. On construct validity, the intended test against the full pre-registered milestone inventory yields a clean pass for all fourteen cases;

the legal-only robustness checks reveal a substantive typology — disclosure-driven versus event-driven cases — with direct implications for how milestone inventories should be defined in future applications of VP and related measures. On external validity, all thirteen reported VP– Google Trends correlations are positive and statistically significant (median $r = 0.741$); eleven of thirteen cases share an identical highest-attention period across both series, and the two exceptions are explained by identifiable query-construction limitations rather than a breakdown of either measure. Full case-by-case results, the Pélícot composition-independence check (RQ4, not yet testable), and the Galician-language robustness check — which accounts for roughly a third to a half of total measured volume in five cases but changes no substantive conclusion — are in Appendices E, H, and F.

6 Illustrative Application: Does Reporting Rise After Case Milestones?

A validated attention measure makes it possible to ask substantive questions that were previously untestable for lack of a usable attention series. As one illustration, I ask — exploratorily, and not as part of the pre-registered design — whether official reporting of sexual offenses rises in the months or quarters following pre-registered case milestones. The analysis focuses on La Manada, the Spanish case with the longest milestone record and the most complete surrounding outcome history¹³ Three designs — a pooled regression discontinuity, a contamination-controlled monthly event study, and the same event study extended to six crime-series outcomes at quarterly resolution — converge on a consistent directional pattern: reporting tends to rise, not fall, following a milestone, at both temporal resolutions and across every assault- or sexual-offense outcome, while a placebo outcome (burglary) does not show the same pattern. Statistical precision is limited, and the pooled

¹³Rubiales/Hermoso provides a shorter but complementary window and its results (not shown here) are consistent with those reported here.

RDD finds no sharp discontinuity exactly at the milestone date, suggesting any response is gradual rather than immediate. This section illustrates the kind of substantive question a validated VP series makes possible to ask. Yet, the work in this section is illustrative.

6.1 Outcome Data and Designs

Two outcome series are used. The first is the monthly count of calls to a Spanish national sexual-abuse helpline, from 2007 to 2026, with no gaps. The second comprises six crime series from the quarterly bulletins of Spain’s Secretaría de Estado de Seguridad (SES), published by the Ministerio del Interior: total sexual offenses (which include both sexual assault with penetration and other sexual offenses), assault/affray (*lesiones y riña*), kidnapping, and burglary (*robos con fuerza*). Burglary serves as a placebo outcome: it has no plausible connection to sexual misconduct cases and should show no response to case milestones if the estimated patterns in the other series reflect a genuine mechanism rather than a common trend. Both outcome series are deseasonalized.¹⁴

Three designs are reported. **Pooled regression discontinuity (RDD).** Each milestone is assigned a fixed, symmetric six-month bandwidth on either side of its date, with its own local linear trend estimated on each side; a single common discontinuity term, shared across all milestones, captures whether there is on average a sudden level shift exactly at the cutoff. Milestones falling within one bandwidth of each other are merged into a single event before estimation. **Contamination-controlled event study, monthly.** Each milestone’s window is truncated at the midpoint to its nearest neighboring milestone, so no calendar month is assigned to more than one milestone’s window; all truncated windows are stacked into a panel and regressed on relative-month dummies, with the month immediately before each milestone as the reference period and standard errors clustered by calendar month.

Contamination-controlled event study, quarterly, all six outcomes. Identical in

¹⁴The long-run trend is retained; only the calendar seasonal pattern is removed, estimated jointly with the trend to avoid contaminating the seasonal estimate with trend variation — but not detrended, since the trend is part of the variation this test examines rather than noise to be removed.

design to the monthly version, run at quarterly resolution to include all six SES outcome series.

6.2 La Manada

6.2.1 Pooled RDD

Figure 9 plots the deseasonalized calls trajectories around each of the eight (post-merge) milestones, demeaned by milestone, together with the pooled average and its discontinuity estimate. The common discontinuity is small and not statistically significant (+173 calls, $p = 0.618$, $n = 104$ case-months across 8 milestones), and this result is stable across alternative bandwidths of three to six months (all $p > 0.58$). The pooled average nonetheless drifts upward across the window, from roughly -200 calls two months before a milestone to roughly $+300$ calls five months after — a pattern visible in the pooled average despite no single milestone producing a sharp, common jump at the cutoff.

6.2.2 Contamination-controlled event study, monthly

Figure 10 reports the relative-month coefficients from the stacked panel (9 milestones, after merging the two-day-apart assault/arrest pair). Coefficients are color-coded by the number of milestones contributing to each relative month, since truncation produces uneven coverage. Within the well-supported core — months -3 through $+3$, each backed by 7 to 9 of the 9 milestones — the path rises from approximately zero at the reference month to a peak of roughly $+320$ calls one month after a milestone, settling around $+220$ to $+240$ through month $+3$. No individual coefficient reaches conventional significance, but the direction is uniformly positive at every post-event month in the well-supported range. Estimates at longer horizons, backed by only two milestones each, are larger and individually significant, but should be read with caution given the thin coverage.




Figure 9: La Manada: pooled RDD across 8 (merged) pre-registered milestones, deseasonalized monthly calls, ± 6 -month bandwidth. Thin lines are each milestone's own trajectory, demeaned by its own mean; the bold line is the pooled average with standard-error bands.




Figure 10: La Manada: contamination-controlled monthly event study, deseasonalized calls, 9 (merged) milestones, windows truncated at the midpoint to each milestone’s nearest neighbor. Reference month is -1 ; color indicates how many milestones contribute to each relative month.

6.2.3 Contamination-controlled event study, quarterly, all six outcomes

Figure 11 extends the same design to quarterly resolution across all six SES outcome series (8 milestones, truncated the same way). The pattern is consistent across every assault- or sexual-offense category: each of the five non-placebo panels shows point estimates rising from at or near zero at the reference quarter to positive values by one to two quarters after a milestone, continuing to climb through the edge of the well-supported window. Kidnapping, a violent but non-sexual-offense category, follows the same rising pattern. Burglary does not: its point estimates are noisy, show no consistent post-milestone rise, and are negative at several horizons. As with the monthly results, no individual coefficient in the well-supported window (quarters -1 through $+1$, backed by all 8 milestones) is statistically significant. What is notable is not any single estimate’s precision but the consistency of the directional pattern across five outcome categories, two temporal resolutions, and three identification strategies.

La Manada: quarterly event-study, contamination-free windows, all 6 outcomes
(8 milestones; assault+arrest and verdict1+release merged; reference = quarter -1)




Figure 11: La Manada: contamination-controlled quarterly event study, all six SES/Interior outcomes, 8 (merged) milestones, same truncation method as Figure 10. Burglary (bottom right) is the placebo outcome.

7 Discussion: VP Within the Societal Response Index

VP is, by design, a volume measure only. It establishes *how much* attention a case receives and *when*, not what that attention consists of, and that limitation is exactly the gap the remaining SRI dimensions are designed to close. VP is the activation dimension: Proposition 1 (Section 2) holds that without it, neither VBF, IER, DPB, nor TDR can be observed at all. Having validated activation, the natural next step is to characterize what fills the attention VP measures.

This research program is structured as a sequence of papers, each developing and validating one SRI dimension using the same pre-registered framework and data infrastructure established here: Paper 2 (VBF, Victim-Blame Framing), Paper 3 (IER, Identity Exposure Risk), Paper 4 (DPB, Discourse Polarity and Battle Intensity), and Paper 5 (TDR, Temporal Dynamics and Reactivation), drawing on a fine-tuned multilingual transformer classifier trained on a psychometrically grounded taxonomy of secondary victimization (Tavares et al., 2023) for the content-characterization work that distinguishes these later dimensions from VP. A capstone paper, *The Societal Response Index: Construction, Validation, and Applications*, will integrate all five dimensions into the composite index.

The measurement problem this paper addresses — the volume and timing of public attention to a documented harm event, recoverable from social media trace data at census-level resolution — is not specific to economics or to sexual misconduct cases. The same data infrastructure and validation protocol could, in principle, be applied to public attention around disease outbreaks, mass-casualty events, or other public-health-relevant disclosures, and I would welcome use of the data and protocol described here by researchers in public health, communication, and related fields.

8 Conclusion

This paper introduces the **Societal Response Index (SRI)**, a pre-registered, multidimensional measurement framework for how society publicly responds once women disclose sexual victimization, and constructs and validates its first dimension, VP (Volume and Persistence): a census-level, weekly and daily measure of social media attention to high-profile sexual misconduct disclosures. The SRI reframes secondary victimization as one of five coordinate dimensions of a broader latent construct, alongside VBF (Victim-Blame Framing), IER (Identity Exposure Risk), DPB (Discourse Polarity and Battle Intensity), and TDR (Temporal Dynamics and Reactivation). VP is the activation dimension: without sufficient public attention, none of the remaining dimensions can be meaningfully observed.

VP is constructed and validated across fourteen pre-registered cases spanning Spain, the United States, and France between 2011 and 2026 — the thirteen-case main index plus Gisèle Pélicot. Four further pre-registered cases (Julio Iglesias, Johnny Depp/Amber Heard, Nevenka Fernández, and a second Errejón legal matter) await keyword finalization and are not included in the results reported here.

Two pre-registered validation exercises support VP as a usable measure of public attention. On construct validity, VP passes the literal pre-registered Literal legal criterion (at least 10 of 13 main-index cases, within ± 2 weeks of four types of legal milestone: arrest, trial opening, verdict, and resignation); the intended test against the full pre-registered milestone inventory yields a clean pass for all fourteen cases. The validation also produces a substantive finding: a typology of *disclosure-driven* versus *event-driven* cases, with direct implications for how future researchers should define milestone inventories when applying VP or similar attention measures. At daily resolution, VP responds on the same day as the milestone for 75.3% of testable milestones and within 14 days for 91.6%. On external validity, VP correlates positively and significantly with Google Trends for all thirteen cases with completed data; eleven of thirteen share an identical single highest-attention period across both series. The two exceptions — Bill Cosby and Dani Alves — are explained by

identifiable query-construction issues rather than a failure of VP to track public attention.

By validating VP first, this paper establishes the conceptual architecture, data infrastructure, and validation protocol on which the remaining four SRI dimensions — VBF, IER, DPB, and TDR — and ultimately the composite index, will be built. The full keyword dictionaries, event-timeline data, and replication code underlying this paper will be released as a public good upon publication.

References

- Angelico, Cristina, Juri Marcucci, Marcello Miccoli, and Filippo Quarta**, “Can we measure inflation expectations using Twitter?,” *Journal of Econometrics*, 2022, 228 (2), 259–277.
- Aurrekoetxea-Casaus, Maite**, “San fermines #la manada case: An exploratory analysis of social support for victims of sexual violence on Twitter,” *Computers in Human Behavior*, 2020, 108, 106299.
- Baker, Scott R, Nicholas Bloom, and Steven J Davis**, “Measuring Economic Policy Uncertainty,” *Quarterly Journal of Economics*, 2016, 131, 1593–1636.
- Bhuptani, Prachi H. et al.**, “Online Social Reactions to Disclosure of Sexual Victimization via #MeToo and Symptoms of Post-Traumatic Stress Disorder,” *Journal of Child Sexual Abuse*, 2024. PMC10822142.
- Cavallo, Alberto and Roberto Rigobon**, “The Billion Prices Project,” *Journal of Economic Perspectives. Forthcoming*, 2016.
- Chetty, Raj, Matthew O Jackson, Theresa Kuchler, Johannes Stroebe, Nathaniel Hendren, Robert B Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob et al.**, “Social capital I: measurement and associations with economic mobility,” *Nature*, 2022, 608 (7921), 108–121.

Conneau, Alexis, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov, “Unsupervised cross-lingual representation learning at scale,” in “Proceedings of the 58th annual meeting of the association for computational linguistics” 2020, pp. 8440–8451.

European Institute for Gender Equality, “Gender Equality Index 2021: Health,” Technical Report, European Institute for Gender Equality (EIGE), Vilnius, Lithuania October 2021. EU score: 68.0/100.

Glaeser, Edward L., Hyunjin Kim, and Michael Luca, “Measuring Gentrification: Using Yelp Data to Quantify Neighborhood Change,” Working Paper 24952, National Bureau of Economic Research August 2018.

Glaeser, Edward L, Hyunjin Kim, and Michael Luca, “Nowcasting gentrification: using yelp data to quantify neighborhood change,” in “AEA papers and proceedings,” Vol. 108 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2018, pp. 77–82.

Levy, Ro’ee and Martin Mattsson, “The Effects of Social Movements: Evidence from #MeToo,” *Working Paper*, 2023. SSRN 3496903.

Morstatter, Fred, Jürgen Pfeffer, Huan Liu, and Kathleen Carley, “Is the sample good enough? Comparing data from Twitter’s streaming API with Twitter’s firehose,” in “Proceedings of the international AAAI conference on web and social media,” Vol. 7 2013, pp. 400–408.

Pfeffer, Jürgen, Katja Mayer, and Fred Morstatter, “Tampering with Twitter’s sample API,” *EPJ Data Science*, 2018, 7 (1), 50.

—, —, and —, “This Sample Seems to Be Good Enough! Assessing Coverage and Temporal Reliability of Twitter’s Academic API,” in “Proceedings of the Seventeenth

International AAAI Conference on Web and Social Media (ICWSM 2023)” AAAI Press 2023, pp. 718–729.

Straus, Murray A, Sherry L Hamby, Sue UE Boney-McCoy, and David B Sugarman, “The revised conflict tactics scales (CTS2) development and preliminary psychometric data,” *Journal of family issues*, 1996, *17* (3), 283–316.

Stricot, Maëlle, “Does News Coverage of Violence Against Women Affect Its Judicial Treatment? Evidence from France,” 2024. Working Paper, Paris School of Economics.

— **and Céline Michaud**, “Understanding Sexual Violence Reporting Behavior: Evidence from France, 2007–2019,” *Population*, 2026, *80* (2), 209–237.

Stubbs-Richardson, Megan, Nicole E Rader, and Arthur G Cosby, “Tweeting rape culture: Examining portrayals of victim blaming in discussions of sexual assault cases on Twitter,” *Feminism & Psychology*, 2018, *28* (1), 90–108.

Tavares, Suiane M, Carlos E Pimentel, Tamyres T Paiva, and Cicero R Pereira, “Development and validation of the secondary victimization scale,” *Psychological reports*, 2023, *126* (5), 2594–2615.

Ullman, Sarah E., “Social reactions, coping strategies, and self-blame attributions in adjustment to sexual assault,” *Psychology of Women Quarterly*, 1996, *20* (4), 505–526.

— , *Talking About Sexual Assault: Society’s Response to Survivors*, Washington, DC: American Psychological Association, 2000.

Xue, Jia, Bolun Zhang, Qiaoru Zhang, Ran Hu, Jieli Jiang, Nian Liu, Yingdong Peng, Ziqian Li, and Judith Logan, “Using Twitter-based data for sexual violence research: scoping review,” *Journal of medical internet research*, 2023, *25*, e46084.

A Complete Event Taxonomy

Table A.1: Pre-registered event-type codebook (selected types)

Event type	Description
assault / assault_reported	The assault itself; victim reports to authorities
complaint_filed	Formal legal complaint filed
arrest	Suspect arrested or detained
media_trigger	Major media investigation or exposé published
media_coverage	Significant subsequent media coverage
trial_start, verdict, sentencing	Legal-process milestones
resignation	Perpetrator resigns from a position
institutional_response	Institutional, party, or organizational response
legal_action	Other significant legal action not separately coded
perpetrator_response	Perpetrator public statement or denial
discovery	Case discovered or evidence found
legislative_action	Legislative or policy change attributed to the case
victim_impact_start	Victim-impact-statement testimony begins

Note: the full codebook contains 28 named event types; this table lists the types most frequently used across the fourteen-case sample. The complete codebook, including stratum-sampling definitions for the underlying content-collection design (not used to construct VP itself), will be released as a public good upon publication.

A.1 Case-by-case construct validity detail

Harvey Weinstein and R. Kelly fail under both legal definitions, and the gap is not close under the literal one. For both cases, the precipitating event was a journalistic disclosure that substantially preceded, and generated more attention than, the subsequent legal process: the October 2017 *New York Times* investigative report for Weinstein, and the January 2019 broadcast of the *Surviving R. Kelly* docuseries for Kelly. Both are pre-registered, dated `media_trigger` rows in the same timeline file used for the legal-only tests; once admitted as qualifying milestones, both cases pass cleanly (3/3).

Cesar Chávez/UFW is the limiting case of the same pattern: the timeline contains no legal-process row of any kind, because the subject is deceased and was never

prosecuted. Its pre-registered milestone inventory consists of one `media_trigger` row and five `institutional_response/legislative_action` rows. All three of this case’s highest-VP weeks fall within 14 days of one of these six pre-registered rows, and the case passes 3/3 under the full definition.

Two results from the external-validity analysis corroborate the disclosure-driven classification of R. Kelly and Harvey Weinstein using an entirely independent data source. For R. Kelly, the best-correlating Google Trends query is not a legal-process term: “Free R Kelly,” an advocacy and fan-reaction term that tracks the public mobilization surrounding the *Surviving R. Kelly* documentary, correlates far more strongly ($r = 0.874$) than trial- or verdict-stage queries, which correlate only moderately or non-significantly.¹⁵ For Harvey Weinstein, the highest-correlating query is his bare name ($r = 0.985$), one of the highest correlations in the full sample, consistent with a case where a single disclosure event — the *New York Times* investigation — so thoroughly dominated public attention that all search interest in his name effectively tracks the case.

A.2 Robustness: daily-resolution lead/lag

Daily tweet counts permit a sharper test of how quickly VP responds to a milestone. For every milestone in the full pre-registered timeline across all fourteen cases ($n = 167$ dated rows), a milestone is “detected” if `VP_norm_daily` exceeds +1 SD above the within-case median on the milestone date itself or on any of the 14 calendar days following it, where $\text{VP_norm_daily}(j, t) = [\ln(n_{j,t} + 1) - \text{median}_j] / \text{sd}_j$ at daily resolution. One milestone (Móstoles/PP, 29 June 2026) falls outside the VP collection window and cannot be tested, leaving 166 testable milestones.¹⁶ VP detects a spike on the same calendar day as the milestone for 125 of 166 testable milestones (75.3%), and within the 14-day window for 152

¹⁵For reference: “r kelly trial” ($r = 0.379$), “R Kelly verdict” and “R Kelly sentencing” (both weak and non-significant). Full candidate-query correlations are in Table A.2, Appendix D.

¹⁶This event was added to the pre-registered timeline after data collection had closed; no `VP_norm_daily` value exists for that date.

of 166 (91.6%). Per-case results are in Table 3.

B Complete Boolean Queries

Per-case, per-language Boolean keyword queries — including all hashtags identified through the two-stage discovery procedure described in Section 3 — will be released as a public good upon publication. The complete query documentation file is available from the authors upon request.

C Language Composition Tables

Per-case tables reporting the share of matched tweets by queried language (Spanish, Catalan, Basque, Galician, English, French) will be released as a public good upon publication. The Galician column is flagged as collected without an X `lang:` filter, as described in Section 3 and the robustness check in Appendix F.

D All Google Trends Candidate Queries

(This page left blank on purpose to fit table in next page.)

Table A.2: All candidate Google Trends queries and their correlation with VP_norm

Case	Query	r	p	Sig. ($p < .05$)
Bill Cosby*	Cosby verdict	0.909	< .001	Yes
	Andrea Constand	0.893	< .001	Yes
	Cosby conviction	0.425	< .001	Yes
	Bill Cosby Rape	0.070	0.412	No
	Cosby agresión sexual	(all zeros)		—
Dani Alves*	Dani Alves	0.440	< .001	Yes
DSK*	Dominique Strauss-Kahn	0.658	< .001	Yes
	Nafissatou Diallo	0.636	< .001	Yes
	Tristane Banon DSK	0.552	< .001	Yes
	DSK agresión	(all zeros)		—
Errejón wave*	Errejón denuncia	0.662	< .001	Yes
	Errejón	0.653	< .001	Yes
	inigo errejon	0.382	< .001	Yes
Gisèle Pelicot*	Dominique Pelicot	0.975	< .001	Yes
	Gisèle Pelicot	0.581	< .001	Yes
	PelicotCase	(all zeros)		—
Harvey Weinstein*	Harvey Weinstein	0.985	< .001	Yes
	Weinstein trial	0.088	0.371	No
	Weinstein verdict	0.088	0.371	No
La Manada*	La Manada	0.923	< .001	Yes
Larry Nassar*	Larry Nassar abuse	0.991	< .001	Yes
	Larry Nassar	0.984	< .001	Yes
	Nassar victims	0.861	< .001	Yes
Móstoles/PP*	Mostoles	0.860	.028	Yes
R. Kelly*	Free R Kelly	0.874	< .001	Yes
	r kelly trial	0.379	.001	Yes
	R Kelly verdict	0.016	0.894	No
	R Kelly sentencing	0.008	0.945	No
Rubiales/Hermoso*	Luis Rubiales	0.535	< .001	Yes
	Jenni Hermoso	0.374	< .001	Yes
	seacabo	(all zeros)		—
	chiringuito rubiales	(all zeros)		—
Salazar/PSOE*	Salazar PSOE	0.384	.016	Yes
	Francisco Salazar	0.244	0.134	No
	PSOE assetjament sexual	0.076	0.644	No
	Francisco Salazar PSOE harassment	(all zeros)		—
Sean Combs/Diddy*	Diddy Party	0.741	< .001	Yes
	Diddy arrest	0.503	< .001	Yes
	Cassie Ventura	0.188	.052	No
	diddy trial	0.059	0.543	No
	Diddy verdict	-0.022	0.820	No
Cesar Chávez/UFW	pending keyword finalization			

E Case Studies: Dani Alves, DSK, and Bill Cosby

Dani Alves, DSK, and Bill Cosby are all main-index cases (Table 1) and are counted in the Section 5 pass rates and correlation statistics alongside the other ten main-index cases and Gisèle Pélicot.

E.1 Dani Alves: query-construct mismatch

Dani Alves has the lowest VP–Trends correlation in the sample ($r = 0.440$) and is one of the two top-1 peak-period concordance exceptions. The VP series is constructed from a case-specific Boolean query combining the subject’s name with case-related terms in several languages, and its highest weeks correspond directly to verified case milestones: the arrest on January 20, 2023; the conviction on February 22, 2024; and the appellate reversal on March 28, 2025. The Google Trends canonical query is the subject’s bare name at worldwide scope, measuring all search interest in him as a public figure, not only case-related interest. Its single highest week (March 2024) coincides with a false rumor, circulated on social media and denied by the subject’s brother, that he had died by suicide in prison — an event with no counterpart in the case-specific VP series. The two series diverge not because either measure is wrong but because a bare-name search query and a case-specific tweet query are not measuring an identical construct for a public figure whose search profile extends well beyond the legal case.

E.2 DSK: two substantively distinct attention waves

DSK ($r = 0.658$) matches on the top-1 peak-period concordance test in the current analysis; the case is discussed here because its two-wave attention structure is substantively informative.

VP’s true series maximum falls in May 2011, the month of the original New York arrest — an event that generated enormous worldwide search interest given DSK’s role

as IMF Managing Director and leading French presidential candidate. No Google Trends data exists for that month, so it is excluded from comparison by construction. Within the Trends-comparable window, VP’s highest month is February 2015, driven by the Carlton pimping trial in Lille — a verified second attention wave in which DSK stood trial alongside thirteen co-defendants for a wholly separate French legal matter, ultimately resulting in his acquittal. VP captures this wave strongly because it uses French-language, case-specific keywords targeted at that proceeding. The Google Trends bare-name query, by contrast, captures global information-seeking about DSK as a public figure; that global interest remains anchored to the internationally prominent 2011 New York arrest, and the 2015 Lille trial — a primarily French domestic story — generated comparatively little new search activity worldwide. The two series both peak at their respective most prominent event; that they peak at different events reflects the difference between a case-specific French-language tweet series and a worldwide bare-name search query, not a failure of either measure.




Figure 12: DSK, monthly: VP_norm (full series) vs. Google Trends for “Dominique Strauss-Kahn,” both normalized to their own series maximum. The shaded region has no Google Trends data and is excluded from the RQ3 tests by construction; the May 2011 peak is VP’s true series maximum. The February 2015 peak, within the Trends-comparable window, is a verified second attention wave (the Lille pimping trial); the bare-name Google Trends query registers this primarily French domestic event only weakly relative to the internationally prominent 2011 New York case.

E.3 Bill Cosby: a query-construction failure

Bill Cosby is the other top-1 peak-period concordance exception ($r = 0.909$, the fifth-highest correlation in the sample). The exception reflects a query-construction failure rather than a failure of VP to track public attention. The VP Boolean query was built around trial-and-verdict vocabulary and included “Cosby assault” but not “rape” or “allegations.” The November 2014 public-accusation wave — triggered by a comedian’s viral stand-up clip referencing old rape allegations, more than a year before formal charges were filed — was the largest attention event in the case and dominated discourse using exactly the vocabulary the VP query omits. VP therefore misses the 2014 wave entirely. The canonical Google Trends query (“Cosby verdict”) captures the full history of search interest regardless of vocabulary, and its peak reflects the 2014 wave, producing the top-1 mismatch. This is the mirror image of the Dani Alves finding: there, the Google Trends query was too broad; here, the VP query is too narrow. The correlation remains high because the two series co-move closely across all subsequent milestones; the mismatch is confined to the single highest-attention period. The Cosby VP keyword dictionary will be extended to include accusation-stage vocabulary in the public data release.

Note also that “Bill Cosby Rape,” the most literally case-descriptive Google Trends query, does not correlate significantly with VP ($r = 0.070$, $p = 0.412$), confirming that the VP series simply does not register the 2014 accusation wave at all. “Andrea Constand,” the lead accuser’s name, correlates strongly ($r = 0.893$) because it tracks the trial and conviction proceedings that VP does capture.

F Robustness to Galician-Language Collection

Because Galician queries could not be restricted by X’s `lang:` operator, Galician-language posts are included in the main VP series reported throughout this paper, but their inclusion is not methodologically identical to the other five languages. Galician is applied to all six

of Spain’s main-index cases (Móstoles/PP, Dani Alves, Rubiales/Hermoso, La Mañada, Errejón wave, and Salazar/PSOE); Galician-tagged posts account for roughly 34–54% of total measured volume in each.

This appendix verifies that excluding Galician-tagged tweets changes no substantive conclusion reported in the paper. Four checks are reported.

Series shape. Excluding Galician-tagged tweets and recomputing `VP_norm` yields within-case correlations with the main series exceeding 0.88 for all six affected cases (five of six exceed 0.98), confirming that Galician collection shifts absolute volume levels but not the shape of the attention series.

Construct validity (RQ1). The full construct-validity test is rerun using the Galician-excluded weekly series. The result is unchanged: all fourteen cases pass under the full pre-registered milestone inventory, exactly as under the main series.

External validity (RQ3). The VP–Google Trends correlation and peak-period concordance tests are rerun for the six affected cases using the Galician-excluded series. Correlations shift by at most 0.03 in either direction relative to the main series, and the peak-period concordance outcome is unchanged for all six cases. Dani Alves remains a top-1 mismatch: its exception is driven by an English-language false-rumor search spike with no connection to Galician collection.

Daily-resolution detection (Section 5.2). Of the 64 testable milestones across the six affected cases, 41 are detected same-day under the Galician-excluded series (64.1%) against 40 under the main series (62.5%); no milestone moves from detected to not-detected, or vice versa. The within-14d detection rate is unchanged.

Across all four checks, excluding Galician-tagged tweets changes absolute volume levels substantially but leaves every substantive conclusion reported in this paper unchanged.

G Full Descriptive Statistics

Table A.3: Case-level descriptive statistics, social-media (X/Twitter) volume, all 14 collected cases

Case	Start	End	Days	Total tweets	Peak-day tweets	Peak/mean ratio	Languages
Combs/Diddy	2023-10-19	2025-11-12	755	29,005,278	887,345	23.1	en, es
La Mañada	2016-06-09	2023-05-13	2,529	21,710,290	759,172	88.4	ca, en, es, eu, gl
Rubiales/Hermoso	2023-07-23	2025-07-29	737	17,767,728	1,437,611	59.6	ca, en, es, eu, gl
Errejón wave	2024-09-23	2026-04-15	569	5,424,928	486,976	51.1	ca, en, es, eu, gl
Dani Alves	2022-12-01	2025-04-25	876	3,545,454	98,788	24.4	ca, en, es, eu, gl, pt
DSK	2011-04-16	2021-01-04	3,551	2,392,228	42,576	63.2	en, es, fr
Gisèle Pélicot [‡]	2020-08-15	2026-03-15	2,038	2,254,235	99,903	90.3	ca, en, es, fr
R. Kelly	2017-06-17	2023-03-23	2,105	1,646,305	57,255	73.2	en, es
Móstoles/PP	2022-09-03	2026-07-27	1,353	1,278,095	69,117	73.2	ca, en, es, eu, gl
Larry Nassar	2014-06-03	2022-07-15	2,964	1,230,244	226,468	545.6	en, es
Harvey Weinstein	2017-09-07	2026-06-12	3,200	1,159,674	54,113	149.3	en, es
Bill Cosby	2014-09-18	2026-04-12	4,224	141,732	18,165	541.4	en, es
Salazar/PSOE	2025-06-07	2026-03-05	271	87,501	6,282	19.5	ca, en, es, eu, gl
Cesar Chávez/UFW	2026-02-17	2026-05-21	93	1,579	530	31.2	en, es
Total				87,645,271			

Note: “Days” is the span of each case’s 28-day-padded collection window. “Peak/mean ratio” is peak-day tweet volume divided by mean daily tweet volume over the full collection window. Cases are ordered by total tweet volume. “Languages” lists the language codes of the X `lang:` filters applied (ca = Catalan, en = English, es = Spanish, eu = Basque, fr = French, gl = Galician, pt = Portuguese); Galician queries are run without a `lang:` filter (see Appendix F). DSK is a main-index case (Table 1) with lower volume expected and accepted given pre-2013 Twitter penetration. [‡]Gisèle Pélicot is the sole pre-registered case held out of the main index, as the RQ4 composition-independence check (Appendix H); her VP data was collected on the same pipeline as the thirteen main-index cases and is pooled into the fourteen-case RQ1/RQ3 validation statistics reported in Section 5. Combined, all fourteen collected cases generate approximately 87.6 million matching tweets; the three highest-volume cases (Combs/Diddy, La Mañada, and Rubiales/Hermoso) jointly account for roughly 78% of the full-sample total despite representing only three of fourteen cases.

H Composition-Independence Check: Gisèle Pélicot (RQ4)

Gisèle Pélicot is registered as the out-of-sample composition-independence check for the SRI, held out of the main index (Section 3). Because she waived anonymity and the dominant discourse around her trial was solidarity-framed, her case is the pre-registered test of whether SRI’s volume and composition dimensions vary independently: a case can exhibit high VP alongside low VBF, relative to Spanish main-index cases with comparable VP_norm levels. This is pre-specified as a qualitative, visual comparison rather than a formal statistical test.

Status: not yet testable. RQ4 requires VBF classifier output, and content collection and classifier training are out of scope for the current phase of this research program; only the VP dimension has been collected and validated to date. The Pélicot VP series has been constructed using the same pipeline as the main-index cases and is available for descriptive comparison, but the VBF half of the RQ4 comparison cannot be computed until content collection is implemented. RQ4 is reported here as a pre-registered, planned analysis for a subsequent SRI paper (Paper 2, VBF), not as a completed result in this paper.