

Framing by Wording, Framing by Selection: A Large-Scale Two-Dimensional Audit of French News Headlines, 2022–2025

Amr Sobhy

Le French News Lab

amr@frenchnewslab.org

Abstract

News headlines frame public issues both by *what* they select and by *how* they word it, yet computational framing work typically collapses these operations into a single score. We introduce a two-dimensional framework that separates *salience framing*, measured through four wording devices (loaded vocabulary, blame attribution, threat framing, rhetorical question), from *selection framing*, measured through outlet-level story-form and high-charge distributions. We build a 10,000-headline French supervision set using three LLM annotators with majority-vote resolution and human arbitration, validate the labels against two annotator-independent blind human studies, and apply the strongest classifier to 902,111 deduplicated headlines from 25 French outlets (2022–2025). Three main findings emerge. First, salience and selection divergence are positively correlated yet leave nearly half of outlet-level variance unexplained, populating interpretively distinct off-diagonal cells in a four-cell outlet typology. Second, default classification thresholds systematically inflate corpus-level salience estimates; a precision-floor recalibration protocol corrects this distortion. Third, group-mention analysis reveals sharply unequal salience contexts: headlines mentioning Jews, the Far-right, and Muslims carry the highest detected salience rates, which broad event-context composition does not fully explain (residuals are descriptive, not same-event causal estimates; per-group lexicon precision is reported alongside). To our knowledge, this is the largest framing-focused French *headline* audit to date; we release the supervision set, lexicons, and analysis code.

1 Introduction

Headlines are high-reach editorial artifacts: a majority of shared news URLs are never clicked

(Gabielkov et al., 2016), and experimental work shows that headlines can shape memory and inference even when the article is later read (Ecker et al., 2014). Their compactness makes editorial choice unusually visible. A headline can frame by wording an event as invasion, betrayal, danger, scandal, or responsibility; it can also frame by repeatedly selecting particular event types for particular groups. When a French outlet writes “*Les migrants envahissent*” (“Migrants are invading”) rather than “*Des migrants arrivent*” (“Migrants are arriving”), the word choice is not accidental; it encodes a threat frame, an intergroup opposition, and a causal attribution in three words. When a far-right aggregator fills four in five headlines with crime stories about identifiable social groups, that selection pattern is not random; it systematically concentrates group-threat associations in editorial output. These are two distinct editorial operations.

Computational framing measures often blur them. Sentiment scores and bias classifiers capture how language is used but cannot distinguish neutral wording on a concentrated crime agenda from charged wording on a broader one; topic-distribution measures capture what is selected but not how it is packaged. Auditing harmful or polarizing media ecosystems, as the EU Digital Services Act’s independent-audit regime now requires at platform scale, means knowing whether an outlet’s effect is driven by selection, salience, or their interaction. We study this problem in French news headlines, a setting with rich ideological variation but less computational framing infrastructure than English. We ask: **RQ1** how salience framing varies across French outlets; **RQ2** whether salience and selection are empirically dissociable, taking as the null that they are one construct measured twice, which would show as near-collinear divergences ($r \geq 0.90$) with unpopulated off-diagonal cells; and **RQ3**

which social and political groups appear in the highest-salience headline contexts. Our contributions include the 10,000-headline supervision set, a precision-floor recalibration protocol, and a 902,111-headline audit showing substantial overlap between salience and selection without collapsing them into one axis.

2 Framework and Related Work

Entman (1993)’s definition says framing involves both *selection* and *salience*: selecting aspects of reality and making them more noticeable or meaningful. Entman treats these as coupled within a single framing act; we separate them not as a claim about the act but as a measurement decision at the outlet level, where the two need not coincide and, as §4.1 shows, only partially do. We operationalize this as two measurable streams. **Salience framing** is the wording stream: four primary devices (loaded vocabulary (Stevenson, 1944), blame attribution, threat framing, and rhetorical question) plus a supplementary indicator of ingroup/outgroup construction (us-vs-them). These devices are grounded in problem definition, causal interpretation, moral evaluation, and treatment implication, and in framing theory, critical discourse analysis, and argumentation theory (Entman, 1993; Gamson and Modigliani, 1989; van Dijk, 1991; Reisigl and Wodak, 2001; Walton, 1996; Tankard, 2001). The four primary devices are selected because they are theoretically grounded, observable in headline-length text, and stable enough for supervised measurement at corpus scale. The second criterion is the binding one and excludes devices the same traditions treat as central: personalisation and metaphor, for instance, are recoverable from article context but rarely decidable from a headline alone. **Selection framing** is the distributional stream: which story forms an outlet repeatedly chooses to cover, and how often those story forms are high-charge events such as crime, conflict, scandal, and crisis. This follows agenda-setting and gatekeeping traditions, where influence operates through coverage allocation and accessibility (McCombs and Shaw, 1972; Scheufele and Tewksbury, 2007; Shoemaker and Vos, 2009). Our selection measure is conditional on coverage: it observes story-form distributions *among headlines an outlet chose to publish*, closer to second-level attribute agenda-setting (McCombs, 2005) than to first-level gate-

keeping. Outlets with identical story-form distributions could differ substantially in which events from the world they select to cover at all: a dimension the published-headline corpus cannot observe.

Table 1 summarizes the measurement schema. The four primary devices are grounded in Entman (1993)’s four frame functions: blame attribution (causal interpretation), threat framing (problem definition and treatment implication), loaded vocabulary and rhetorical question (moral evaluation and problem definition), with many-to-many correspondence per headline; us-vs-them marks ingroup/outgroup opposition. Prior work either assigns frame labels from a single typology without decomposing wording from selection (Card et al., 2015; Liu et al., 2019; Mendelsohn et al., 2021), captures a single bias dimension (informational and lexical bias (Fan et al., 2019; Spinde et al., 2021) or persuasion techniques (Piskorski et al., 2023)), or separates issue filtering from slant rather than wording from selection (Budak et al., 2016), and so cannot serve as quantitative baselines for a two-dimensional decomposition (Otmakhova et al., 2024). English-language auditing resources including NELA-GT (Gruppi et al., 2022) and BERT-based political leaning classifiers (Baly et al., 2020) address related questions at scale but do not separate wording from selection mechanisms.

French-language media studies motivate the setting. Benson and Wood (2015) documents that immigration news in France is dominated by governmental and political sources, with limited non-governmental voices (cf. Benson, 2013, for cross-national comparison); Dalibert (2015) shows how minority-movement access to the French public sphere is regulated by “francité,” constraining visibility and legitimacy. Work on migrant representation records oscillation between victim, threat, and voiceless framings, echoing our threat and us-vs-them findings. This literature motivates the 12-group inventory analysed in §4.2, which includes migrants, Muslims, Jews, and the far right. Related francophone work detects opinion/information genre at scale in Québécois and Belgian media (Escoufflaire et al., 2024) and links ownership structures to speaking-time-based political slant in French broadcasting (Cagé et al., 2022). More recent French-media framing and discourse studies remain either small and issue-

Measure	Operational definition
Loaded vocabulary	Evaluative or emotionally charged wording whose neutral paraphrase would preserve the core event claim.
Blame attribution	Wording that assigns causal responsibility for a problem or negative outcome to an actor or group.
Threat framing	Wording that casts an actor, event, or situation as danger, invasion, crisis, or security risk.
Rhetorical question	Headline posed as a question. Operationalised by interrogative form, not by whether the question is judged rhetorical.
Us-vs-them	Explicit ingroup/outgroup contrast through naming, pronouns, identity markers, or oppositional structure.
Story form	10-class event/action category: POLICY, PARLIAMENT, CONFLICT, CRIME, SCANDAL, ELITE, SOCIAL, LAW, ELECTIONS, OTHER.
High-charge	Binary flag for story forms associated with conflict, crime, scandal, crisis, or strong negativity.

Table 1: Measurement schema. Saliency devices are wording-level and non-mutually exclusive; story form and high-charge are selection-level.

specific, such as the 120-article immigration corpus of Song et al. (2024), or topic-bound at moderate scale, such as the 13,795-article AI press corpus of Tsimpoukis (2025). Larger French news corpora, such as Jehle and Le Gallo (2025)’s 400,000-article EU-sentiment study, analyze tone or agenda dynamics rather than explicit framing measurement. Our unit is narrower, focusing on headlines only, but covers more outlets and explicitly separates selection from saliency; to our knowledge, that makes the present corpus unusually large for framing-focused work in the French setting.

The boundary matters most for attribute agenda-setting (McCombs, 2005): repeated CRIME coverage operates through saliency transfer and accessibility, making considerations available in memory (McCombs and Shaw, 1972; Scheufele, 1999), while charged wording operates through applicability, shaping which interpretive schemas audiences invoke (Scheufele, 1999; Scheufele and Tewksbury, 2007). A single “framing intensity” score cannot preserve that distinction.

3 Data and Method

Corpora. We use two datasets with different roles. The supervised development set contains 10,000 France-based news headlines from 25 outlets, designed as a structurally comprehensive cross-section of the French media ecosystem across the political spectrum (Newman et al., 2023; ACPM, 2026). It includes three of France’s leading broadcast and continuous-news brands (*Franceinfo*, *BFMTV*, *TF1 INFO*) and seven major national dailies (*Le Monde*, *Le Fi-*

garo, *Le Parisien*, *Les Echos*, *Libération*, *La Croix*, *L’Humanité*). Selection criteria were reach or circulation rank; format diversity; ideological breadth, explicitly anchoring the range with *La Croix* (Catholic daily), *L’Humanité* (communist historical title), *Fdesouche* (far-right aggregator), *Valeurs actuelles* (conservative-nationalist weekly), and *Causeur* (conservative intellectual magazine); and restriction to outlets covering three shared sections, with headlines jointly balanced across outlets and sections.

Each headline is labeled by three schema-constrained LLM annotators from different model families, following evidence that LLMs can match or outperform crowdworkers on structured annotation tasks (Gilardi et al., 2023; Törnberg, 2025; Ziems et al., 2024)—a task-specific comparison that remains contested, and one our own validation qualifies (Table 4). The annotators, openai/gpt-oss-120b (released August 2025), google/gemma-4-31B (released April 2026), and meta-llama/Llama-3.3-70B-Instruct, produced full-field labels under a shared schema-constrained instruction set; Appendix Table 25 summarizes the released schema; the full annotation prompt is included in the release. Three candidate second annotators were run against the primary annotator at full scale and two were retained; pairwise agreement statistics for all three are released. Final binary labels are resolved by 2/3 majority vote. The 642 three-way story-form conflicts (6.4%) are assigned to a single human arbitrator; post-hoc blind inter-annotator reliability on these cases is reported in Appendix Table 26. Binary conflicts are not arbitrated: the panel is three models with overlapping pretraining rather than three independent coders, so majority vote can propagate shared error (Appendix Table 9), and the blind human studies rather than arbitration are the control for the binary heads; the weakest head, us-vs-them, is reported only as a supplementary indicator. The final split is 6,999 train, 1,500 validation, and 1,501 test headlines.

The corpus-level analysis set approximates real editorial output. We drew a proportionally stratified sample targeting one million headlines from the 25-outlet panel across outlet, section, and year for 2022–2025, then removed exact duplicate URLs and within-outlet duplicate headlines, yielding 902,111 unique headlines. Eligibility was restricted to the Politics, Economy, and Society sections used throughout corpus construction. Per-

outlet counts range from 646 (Blast) to 108,307 (Le Figaro); full counts appear in Appendix Table 13. Identical headlines across outlets are retained (syndication is part of the observable editorial field); classification is restricted to headline text, the standalone framing unit. The unit also strips context that disambiguates attribution: quotation, irony, and blame voiced by a cited source rather than by the outlet are often unrecoverable, which bounds the blame-attribution and threat-framing devices in particular.

Labels. Salience is multi-label: loaded vocabulary, blame attribution, threat framing, rhetorical question, and us-vs-them construction. Selection is modeled through a binary high-charge label (Galtung and Ruge, 1965; Harcup and O’Neill, 2017) and a 10-class story-form taxonomy (Table 1), developed iteratively during pilot annotation on a 500-headline exploratory sample and guided by French press section conventions and prior news-values classifications (Harcup and O’Neill, 2017); unlike topic-oriented systems such as IPTC Media Topics, story form classifies the journalistic action a headline reports rather than just the subject domain, so one immigration headline may be POLICY and another CRIME (Appendix Table 12). Target-group analysis is handled separately through a hand-curated lexicon of 219 French surface-form terms spanning 12 groups (Appendix Table 7 reports the full v6 label distribution).

Models. We compare three supervised measurement stacks on the final v6 split: a TF-IDF + Logistic Regression baseline, camembert-base, and xlm-roberta-large (Table 3). For transformer models, a multi-task encoder predicts the five salience devices and charge band through six independent binary heads on a shared headline representation, while a separate encoder predicts story form through a single 10-class head. We use xlm-roberta-large for corpus inference because it delivers the strongest overall held-out performance: mean salience+charge F1 is .742 (XLM-R), .726 (CamemBERT), .476 (TF-IDF); XLM-R also leads on story-form macro-F1 (.731). CamemBERT scores higher on rhetorical questions alone (.900 vs. .841), likely reflecting French-specific interrogative morphosyntax; for corpus inference we therefore use XLM-R for all heads except rhetorical question, where CamemBERT is substituted. The ensemble was

Head	MinP	Thr.	P	R	F1
loaded	.70	.67	.733	.698	.715
blame	.70	.83	.719	.650	.683
threat	.60	.84	.723	.712	.718
rhet. q.	.85	.89	.957	.809	.877
us-vs-them [‡]	.60	.89	.600	.625	.612
charge	.70	.64	.779	.785	.782
any sal.	—	—	.838	.780	.808

Table 2: Final v6 production thresholds for the mixed-encoder ensemble, selected on the validation set only. [‡]Us-vs-them is a supplementary indicator, not a primary prevalence estimate.

fixed on validation-set metrics before corpus inference; no post-hoc adjustments were made. Exhaustive encoder benchmarking is not the contribution; the substantive cross-outlet findings are bounded against encoder choice by the TF-IDF replication in §4.3 ($r = 0.667$ vs. 0.736).

Both models use a 48-token maximum sequence length (<2% of headlines exceed this; details in Appendix Table 24). Three-seed evaluation confirms robust model selection (Appendix Table 24); results in Table 3 are from seed 42. Thresholds are selected on validation only via a precision-floor-constrained recalibration policy: for each binary head, we search a 0.10–0.99 grid and retain the threshold maximizing validation F1 subject to pre-specified per-head precision floors, assigned from annotator agreement tiers before validation metrics are observed: rhetorical question ($\kappa=.770$, highest) receives a .85 floor; the mid-agreement heads (loaded vocabulary, blame attribution, charge band) receive .70; and the lowest-agreement heads (threat framing, us-vs-them) receive .60 (per-head floors and thresholds in Table 2).

Agreement patterns (Appendix Table 9) are consistent with the v6 merge design: binary fields show high unanimous agreement while story form is harder, matching the need for arbitration; us-vs-them is weakest ($\kappa = .408$).

Independent blind validation. Two annotators with no affiliation to this work independently re-annotated a stratified 499-headline held-out sample (primary validation) using a bilingual annotation guide derived from theoretical construct definitions, with no access to LLM labels or expected label distributions (Table 4). The sample is temporally representative of the held-out test split (39.5% post-October 7, 2023, vs. 41.3% of date-resolved test headlines), directly covering the

Model	Loaded	Blame	Threat	Rhet.Q	UvT [‡]	Charge	Story (macro-F1)
TF-IDF + Logistic Regression	.518	.468	.429	.491	.381	.572	.543
CamemBERT-base	.770	.657	.642	.900	.617	.769	.723
XLM-RoBERTa-large [†]	.758	.672	.749	.841	.633	.801	.731

Table 3: Held-out test-set F1 by model (final v6 split, seed 42, production thresholds). [†]Production model; CamemBERT substituted for rhetorical question only. [‡]Us-vs-them is a supplementary indicator.

Device	H-H κ	Maj-LLM κ	F1	Prec.	Rec.
Loaded vocabulary	.542	.608	.768	.644	.951
Blame attribution	.601	.718	.675	.579	.809
Threat framing	.613	.668	.564	.420	.857
Rhetorical question	.894	.869	.862	.820	.909
Us-vs-them [†]	.510	.596	.471	.369	.649
Any salience [‡]	.680	.463	.780	.646	.984

Table 4: Independent blind re-annotation of the 499-headline validation sample by two unaffiliated annotators. Maj-LLM = majority-vote human labels vs. LLM consensus; conflict rows excluded, effective N per head 383–488.

period most exposed to possible LLM pretraining overlap; within this subset, any-salience human-LLM agreement declines only modestly ($\kappa = .424$ vs. .463 full-sample). Majority-vote human-LLM κ ranges .608–.869 across the four primary devices; recall is consistently high ($\geq .65$) while precision is lower (.37–.82), confirming that the merged labels are systematically liberal (see Limitations). Human-human $\kappa = .680$ on the any-salience composite confirms that annotators agree on overall salience judgements even where device attributions diverge. A corroborating independent study on a separate $N=350$ sample yields consistent results (Appendix Table 20). Appendix Table 19 provides a qualitative error analysis illustrating principal failure modes.

Group analysis uses a lexicon of explicit anchor terms mapping French surface forms to 12 canonical groups. The lexicon contains 219 terms and detects 134,046 headlines (14.9% of the corpus), covering explicit mentions only; indirect or paraphrastic references are missed by design. Groups were selected by theoretical relevance (targeting identities whose framing asymmetries are documented in French and European media research (Benson and Wood, 2015; Dalibert, 2015; van Dijk, 1991)) and by a minimum corpus-frequency criterion of at least 1,000 detectable headline mentions in a pilot pass. A blinded two-human audit of a 200-headline sample confirms high agreement: $\kappa = .94$ –.98 on the representative tier ($n=140$) and $\kappa = .83$ –.85 on a harder stress-test tier ($n=60$). The lexicon shows precision $\geq .83$

for 9 of 12 groups; LFI (.69), Jews (.70), and Unions (.82, marginal) are lower-precision and should be treated as indicative (Appendix Table 8).

Analysis plan. We summarize salience and selection distinctiveness with Jensen–Shannon divergence from the corpus baseline, following the use of Jensen–Shannon measures for media-agenda comparison (Pinto et al., 2019), and test dissociability with Pearson and Spearman correlations against the conventional $r < 0.90$ criterion (RQ2). We build the four-cell typology from median splits on salience and selection divergence as a cartographic convenience rather than a claim of categorical differences; continuous outlet positions are visible in Figure 1. Robustness checks compare focal outlets against all other outlets within the same story form using χ^2 and ϕ as a coarse descriptive stress test; story form is partly editorially mediated and therefore is not an exogenous same-event control. Salience distinctiveness measures divergence in the distribution of rhetorical devices relative to the corpus baseline rather than raw device frequency, so an outlet may exhibit frequent charged language while remaining low in Sal.JS if its device mixture mirrors the broader media distribution. Because Sal.JS is computed over a 5-bin device distribution and Sel.JS over a 10-bin story-form distribution, the two divergences are not directly comparable in magnitude; the typology uses within-axis rankings rather than cross-axis absolute values.

4 Results

Across the full corpus, 34.6% of headlines contain at least one salience device (loaded vocabulary most frequent, us-vs-them rarest; Appendix Table 7). High-charge story forms account for 31.3%. These rates are lower than default-threshold estimates would produce, reflecting the precision-floor recalibration; the 34.6% figure is best interpreted relative to the French corpus baseline rather than as a universal prevalence benchmark.




Figure 1: **Selection vs. salience divergence across French news outlets.** Note that *Ouest-France* sits high in (b) because it is far from the baseline overall while its headlines deploy framing devices unusually *rarely*, not because they deploy them aggressively. (a) Each outlet’s distance from the corpus baseline decomposed into selection divergence (D_{sel}^{JS} , what stories are emphasised) and salience divergence (D_{sal}^{JS} , which framing devices are deployed in the headline). Both quantities are Jensen–Shannon divergences against the cross-outlet mixture distribution. Marker colour encodes the four-class typology defined in §4.1. (b) Same outlets re-expressed as a salience share ($D_{sal}^{JS} / (D_{sal}^{JS} + D_{sel}^{JS})$) against summed divergence. The dashed line at 0.5 separates outlets whose share falls on the selection side from those on the salience side. Vertical position in (b) is summed distance from the baseline, not salience intensity; outlets close to the bottom edge are close to the corpus baseline on both dimensions. Both divergences share the same bounded scale, so the share is well defined, but their expected magnitudes differ with support size (5 vs. 10 bins); both coordinates in (b) are therefore read ordinally — as which axis dominates within an outlet, and as rank distance from the baseline — not as an equivalence of magnitudes.

4.1 Outlet Profiles and Dissociability

Figure 1 maps all 25 outlets in the two-dimensional divergence space, Table 5 summarizes the four-cell membership, and Appendix Table 13 reports the full outlet-level values. Salience divergence from the corpus baseline varies sharply: *Slate.fr* is the most salience-distinctive outlet, *Fdesouche* carries the clearest threat/blame profile, *Causeur* is distinctive through loaded and interrogative wording, and *Ouest-France* is double-distinctive through unusually low device and high-charge rates (7.2% any-salience, 4.6% high-charge, both lowest in the panel). The Double-distinctive cell spans 7.2%–72.6% any-salience: cell membership reflects position relative to thresholds, not editorial similarity within the cell. Selection divergence produces a different ranking, with *Fdesouche*, *Causeur*, and *Blast* among the clearest agenda outliers (*Causeur*’s ELITE-heavy selection profile should be read with caution because ELITE-category conflict cases show lower arbitration agreement; Appendix Table 26) and high-

charge coverage concentrated especially in *Fdesouche*, *Blast*, and *Valeurs actuelles*.

Device composition reveals three outlet signatures: *threat-blame*, where threat and blame co-elevate (*Fdesouche* most clearly, with *Valeurs actuelles* at lower intensity); *interrogative-evaluative*, where rhetorical questions and loaded vocabulary dominate without the broad threat/blame profile (*Slate.fr*, *Causeur*, *L’Express*; rhetorical-question rates index headline form, see Limitations); and *institutional blame*, where accountability language rises without a strong threat register (*Mediapart*, *L’Humanité*, *Le Parisien*; blame rates are similarly directional). (Loaded vocabulary: est. precision 0.643 post-prior-shift, –5.7 pp below the 0.70 floor; Appendix Table 10.) Similar any-salience rates can therefore reflect distinct interpretive mechanisms.

Rate estimates widen materially only for the smallest outlets; *Blast* (646 headlines) should be read as an illustrative case profile (Appendix Table 15).

Dissociability rests directly on the populated

Cell	n	Sal.JS	Sel.JS	Sal.%
Double-distinctive	9	0.027	0.081	.484
Saliency-dominant	3	0.017	0.018	.358
Selection-dominant	3	0.008	0.046	.510
Low-distinctiveness	10	0.004	0.014	.376

Table 5: Four-cell typology. Sal.JS/Sel.JS = saliency/selection Jensen–Shannon divergence from the corpus baseline (§3); cells assigned by median splits. Sal.% = mean any-saliency prevalence.

off-diagonal cells of the bivariate distribution (Figure 1; four-cell summary in Table 5): the cells are perturbation-stable (only seven near-boundary outlets shift under $\pm 20\%$ median changes, Appendix Table 27), and concrete cases anchor the corners (JDD high on selection yet low on saliency, TF1 INFO the reverse), holding independently of any correlation threshold. The correlation only quantifies the overlap: saliency distinctiveness explains 54.2% of outlet-level selection variance ($r=0.736^1$, LOO-corrected $r=0.721$; 95% bootstrap CI [.503, .936], $N=25$; reliably non-zero, permutation $p=.0001$), leaving 45.8% unexplained. Both estimates fall within the conventional $r < 0.90$ criterion; with $N=25$ the interval is wide and its upper bound (.936) exceeds it, so the claim does not rest on the threshold. This residual maps onto interpretable off-diagonal profiles: selection-dominant outlets exhibit the *highest* raw saliency prevalence (any-saliency = 0.510) despite the *lowest* rhetorical divergence (Sal.JS = 0.008), while double-distinctive outlets combine high saliency (any-saliency = 0.484) with substantially greater divergence (Sal.JS = 0.027). Frequent charged language therefore does not by itself imply rhetorical distinctiveness.

To test whether this association is ecological aggregation, we estimate a longitudinal outlet $\times$ month panel (1,184 observations; Appendix Table 29). Under two-way outlet+month fixed effects with cluster-robust inference (outlet clusters, $G=25$), the within-outlet coupling remains positive ($\beta = 0.758$, 95% CI [−.11, 1.63]; $p = .084$); we treat this as confirmatory rather than a stand-alone test. Excluding Fdesouche lowers the cross-sectional correlation to $r = 0.688$ ($\rho = 0.711$), and the unexplained variance ($1-r^2 = 0.53$) remains sufficient to populate distinct off-diagonal cells. Because three

¹Split-half reliability: $r_{xx}=0.979$, $r_{yy}=0.988$; Spearman-corrected latent $r=0.748$ (see Limitations for the leave-one-out baseline correction; $r=0.721$).

high-volume low-divergence outlets (Le Figaro, Le Parisien, Franceinfo) contribute $\approx 32\%$ of the corpus baseline, their low divergence is partly self-referential; a leave-one-out correction yields $r=0.721$, $\rho=0.727$, preserving the dissociability criterion (full LOO analysis in Limitations). A tone-only audit would therefore miss selection-dominant outlets, while a topic-only audit would miss saliency-dominant ones.

4.2 Group-Mention Saliency Contexts

van Dijk (1991) establishes that group framing operates through both event-context selection and adversarial wording devices; our two-dimensional framework keeps those pathways analytically separate. Explicit group mentions occur in highly asymmetric model-estimated saliency contexts (Table 6). Group saliency rates reflect headline rhetorical intensity, not editorial sentiment or targeting; interpretive caveats in §Ethics apply throughout. Headlines mentioning Jews have the highest observed any-saliency rate, driven predominantly by antisemitism reporting and Israel/Gaza security coverage (lexicon precision for this group is .70, among the lowest of the twelve, though recall is .97, so the detector is liberal rather than blind): 78.7% carry at least one detected saliency device and 79.8% are high-charge. A year $\times$ story-form reweighting estimates that 49.6% saliency would be expected from their coarse event-context mix alone, leaving a descriptive +29.1 pp residual; this residual is evidence that broad story-form composition does not fully explain the pattern, not evidence of same-event causal editorial targeting. Far-right mentions show a similar raw/residual profile (77.7% raw; +27.7 pp residual), reflecting electoral, parliamentary, and conflict coverage of a political movement rather than targeting of its members. Muslims (66.4%; +19.8 pp residual) and Migrants (58.4%; +17.7 pp) concentrate in policy and security event contexts; Police (57.6%; +11.1 pp) coverage is dominated by blame attribution in accountability-focused reporting; and Unions (53.8%; +3.6 pp) reflect the conflict-heavy 2023–2024 pension-reform mobilization cycle. Seniors remain below baseline, and Workers fall below the saliency rate expected from their year $\times$ story-form mix. The 2024–2025 saliency rise for Jews (77.2% to 84.1%; Appendix Table 23) is consistent with post-October-2023 crisis coverage; high saliency reflects event-

Group	N	Sal%	Exp%	Res.	Top form
Jews	4,646	78.7	49.6	+29.1	CRIME
Far-right	3,981	77.7	50.0	+27.7	CONFLICT
Muslims	4,720	66.4	46.6	+19.8	CRIME
Migrants	9,606	58.4	40.7	+17.7	CRIME
Police	14,313	57.6	46.5	+11.1	CRIME
Unions	6,877	53.8	50.3	+3.6	CONFLICT
Workers	15,456	39.8	52.8	-13.0	CONFLICT
Seniors	21,987	32.1	34.6	-2.5	POLICY

Table 6: Group-mention salience contexts (top 8 of 12 groups). Sal% = observed any-salience; Exp% = expected under year $\times$ story-form reweighting; Res. = observed - expected (pp).

selection and wording contexts, not a uniform targeting pattern (caveats in §Ethics).

All differences from baseline are significant under BH-FDR-corrected two-proportion z -tests (12 hypotheses; $p < 0.001$; intervals do not propagate classifier uncertainty; see Limitations for causal-identification disclosures).

Device profiles differ by group (Appendix Table 17): Jews show the strongest multi-device signal (blame, loaded vocabulary, threat); Police are blame-dominated; and the Far-right shows broad elevation across all three adversarial devices. Selection adds a second layer: Police and Migrants concentrate in CRIME; Workers and Unions concentrate in CONFLICT (56.9% and 56.8%, reflecting the 2023–2024 pension-reform cycle); and RN concentrates in ELECTIONS and PARLIAMENT.

Outlet-by-group patterns (Appendix Table 17) add an important caution: Fdesouche’s elevation spans multiple groups, indicating a broad high-intensity editorial register rather than targeted group framing, while Mediapart and L’Humanité peak on Police and the Far-right through an accountability lens. Removing Fdesouche confirms most group effects are robust (median exclusion effect -0.5 pp; Appendix Table 22). Results for the supplementary us-vs-them indicator are reported in Appendix Table 14.

4.3 Robustness

A dual-ablation study confirms that predictions are not driven by entity memorization: named-entity masking (spaCy) yields negligible confidence drops ($|\Delta p| \leq 0.085$) and low label-flip rates (5.3%–11.4%), while masking the single most salient semantic token (Integrated Gradients) produces substantially larger effects (e.g., $|\Delta p| = 0.266$, 31.6% flip rate for rhetorical questions). Within-story-form comparisons show

that salience differences persist inside broad story-form bins (Appendix Table 18): Fdesouche and Valeurs actuelles remain elevated within CRIME, and Slate.fr retains its interrogative-evaluative profile within SOCIAL; these are descriptive stress tests rather than causal controls because story-form assignment is itself partly shaped by editorial judgment. A within-outlet headline bootstrap confirms near-perfect rank stability ($\rho = .996/.998$ for salience/selection JS), and selection-divergence rankings are robust to taxonomy granularity (Appendix Table 28). The top group pattern is temporally persistent: Jews, the Far-right, and Muslims remain the three highest-salience groups in every year from 2022 to 2025, with annual lifts in the $\times 1.78\text{--}\times 2.35$ range (Appendix Table 23).

An encoder-independence check recomputes Sal.JS/Sel.JS with a TF-IDF story-form model (macro-F1=0.543): r drops from 0.736 to 0.667 (still < 0.90); ρ is stable (0.741 vs. 0.747), bounding the shared-encoder confound at $\Delta r \approx 0.07$. Excluding the rhetorical-question head entirely lowers the corpus any-salience rate from 34.6% to 30.6% and r from 0.736 to 0.561 (0.650 excluding Fdesouche), flipping four near-median outlets (Appendix Table 28). Cook’s D analysis shows that excluding three high-leverage outlets raises r to 0.865, still leaving $\sim 25\%$ unexplained variance; RQ-head exclusion flips four near-boundary outlets (Appendix Table 28). Broadcast-cluster format artifacts are null ($R^2=0.0003$), and XLM-R confidence gaps across ideological clusters are negligible (≤ 0.015).

Temporal trend. Any-salience rises monotonically, 31.7% (2022) to 38.1% (2025). A rate-composition decomposition assigns $+4.86$ pp of the $+6.34$ pp change to within-story-form wording and $+1.20$ pp to composition. It holds in nine of ten story forms, under joint outlet $\times$ section $\times$ story-form standardisation (31.8% $\rightarrow$ 37.4%), and with the October–December 2023 and June–July 2024 windows excluded; 23 of 25 outlets rise. One classifier scores all years, so drift is excluded by construction; standardisation controls event category, not intensity.

5 Discussion

The two-dimensional framework resolves a structural ambiguity in single-axis audits: outlets similarly scored may differ sharply in mechanism, as the outlet profiles illustrate. TF1 INFO com-

bines moderate raw salience (38.3%) with elevated rhetorical divergence ($\text{Sal.JS}=0.029$) but near-baseline selection divergence ($\text{Sel.JS}=0.011$), consistent with salience-led broadcast packaging. JDD shows the inverse pattern: substantial raw salience (41.9%) but almost no rhetorical divergence ($\text{Sal.JS}=0.002$) alongside clearer selection divergence ($\text{Sel.JS}=0.037$), indicating agenda concentration without comparable wording distinctiveness. Fdesouche’s dual-distinctiveness (72.6% any-salience; 82.1% high-charge; $\text{Sal.JS}=0.038$, $\text{Sel.JS}=0.147$) therefore reflects not mere bias but a structural editorial register where accessibility and applicability amplify together (Scheufele and Tewksbury, 2007; McCombs and Shaw, 1972; Scheufele, 1999; van Dijk, 1991). The group-mention findings corroborate prior unequal-salience research (van Dijk, 1991); separating wording-device load from story-form concentration adds resolution unavailable to single-axis designs. Whether the pattern reflects stable editorial stance, event-context concentration, or their interaction requires richer annotation; the present corpus enables such follow-on work (Benson and Wood, 2015; Dalibert, 2015). The precision-floor recalibration protocol, with its fully released threshold log, extends to any longitudinal multi-label framing audit; whether default-threshold inflation would accumulate to distort findings remains an empirical question.

Limitations

Scope of the audit. The audit measures editorial output patterns; whether those patterns produce accessibility or applicability effects in audiences requires experimental work (Price and Tewksbury, 1997; Scheufele, 2004). The study measures headlines only and does not capture article-body framing or cross-platform dynamics. The selection framing measure is further constrained to published-headline story-form concentration and cannot detect first-level event-selection gatekeeping; outlets indistinguishable on Sel.JS may differ in which real-world events they choose to cover. The outlet typology is two-dimensional by construction; finer-grained axes could further separate mobilisation-register from adversarial-register outlets within the double-distinctive cell.

Device performance and ensemble characterisation. The rhetorical-question head indexes headline form, not rhetorical intent: the coding

guide marks any interrogative headline, the annotator prompt excludes purely informational ones, and both fire on nearly all question-marked headlines (of 56 in the 499-item sample, coders mark 96.4% and 94.6%, annotators 85.7%). Its $\kappa = .869$ therefore reflects a near-syntactic judgement, bounding the interrogative-evaluative signature in §4.1; its loaded-vocabulary component is unaffected. Across the three remaining primary heads, blame attribution achieves the strongest independent validation (Maj-LLM $\kappa = .718$; Table 4), while loaded vocabulary shows the lowest human-human agreement ($\kappa_{\text{HH}} = .542$), indicating genuine task ambiguity rather than model failure. Threat framing and us-vs-them reach moderate-to-substantial Maj-LLM agreement ($\kappa = .668$ and $.596$; Table 4), with a corroborating independent study on a separate 350-headline sample yielding consistent results (Appendix Table 20). Isolated French headlines remain difficult for implicit-causality and rhetorical-register inference: even where annotators agree on device presence, the v6 ensemble operates as a recall-oriented liberal annotator ($\text{recall} \geq .65$ across all heads), so corpus-level rates should be interpreted as upper-bound prevalence estimates. Cross-cluster confidence gaps are negligible (≤ 0.015 ; §4.3), suggesting that the recall-biased training signal does not produce outlet-differential annotation error sufficient to distort distributional rankings. A French-specific alternative (mistral-small3.2), one of the candidate annotators not retained, underperformed the ensemble on all five devices ($\Delta\kappa = -0.03$ to -0.25 ; Appendix Table 11), suggesting the pragmatic ceiling is task-inherent rather than an English-pretraining artefact, consistent with broader evidence that pragmatic understanding in LLMs remains highly sensitive to training strategy (Ruis et al., 2023).

Story-form classification and outlet support.

Story-form classification is imperfect around the SOCIAL/OTHER/ELITE boundary. The supervised development set is uniformly balanced across outlets while the production corpus is proportionally stratified, introducing an estimated precision drift of -4.8 pp for the XLM-R any-salience baseline under prior shift (full per-device estimates in Appendix Table 10; the final ensemble corpus any-salience rate is slightly higher because the rhetorical-question head is patched from CamemBERT). Outlet-level estimates are descrip-

tive; smaller outlets, especially Blast (646 headlines), should be read as case profiles rather than equally precise population estimates.

Group detection and event-context confounds.

Group detection is lexicon-based and covers explicit mentions only. We chose explicit surface-form matching over NER or LLM-based entity extraction for reproducibility and to eliminate hallucination risk; the lexicon is a transparent lower bound, not an exhaustive account of group references. Group salience rates are descriptive: all reported differences from the corpus baseline are significant after Benjamini-Hochberg correction, but extreme baseline volume differences (e.g., Jews 0.52% vs. Farmers 2.45%) mean that high salience for rare groups is heavily driven by narrow, high-intensity news events, with the yearly breakdown (Appendix Table 23) reducing sub-annual clustering concern. The year $\times$ story-form reweighting in Appendix Table 16 is a coarse event-context normalization, not a same-event causal design: story forms are useful bins, but assigning a headline to CRIME, POLICY, CONFLICT, or SOCIAL is itself partly editorially mediated. A true causal separation of event intensity from outlet framing would require matched same-event corpora, article-level event clustering, or parallel cross-national coverage. A blinded two-human audit confirms high agreement with the released group labels ($\kappa = .83-.98$), though residual ambiguity concentrates in boundary-heavy cases (workers, broader far-right references).

Annotator release-date and temporal contamination.

Two LLM annotators (openai/gpt-oss-120b, released August 2025; google/gemma-4-31B, released April 2026) have training data that may overlap with the 2024–2025 portion of the annotation period; their pre-training could introduce systematic annotation priors for high-salience events (Israel/Gaza coverage, French elections) beyond pure device detection. We directly tested this concern with a period-stratified agreement analysis comparing 2022–2023 vs. 2024–2025 annotations. The analysis finds no evidence of systematic contamination-inflated annotations: any-salience κ deltas across the three annotator pairs range from +0.003 to +0.032, and several binary heads show equal or lower cross-model agreement in the later period, inconsistent with the contamination hypothesis. The κ delta test is an aggregate check and cannot rule out systematic annotation

priors on specific high-salience events where LLM pretraining coverage is densest. The two highest-risk events, Israel/Gaza coverage from October 2023 onward and the June 2024 French legislative elections, are exactly the events most likely to carry model-specific framing priors beyond the annotation schema. The independent human validation sample ($N=499$) is 39.5% post-October 7, 2023, providing annotation-independent ground truth directly for the contamination-exposed period; within this subset, any-salience human–LLM agreement declines only modestly ($\kappa=.424$ vs. .463 full-sample), bounding the contamination effect at the aggregate level. A residual risk remains for specific event clusters within this window; downstream users who require contamination-free labels for Israel/Gaza or election coverage should use the human-validated $N=499$ subset only. Contamination is also distinct from homogeneity: priors shared across the three annotators would not register in a period-stratified comparison. Appendix Table 9 tests this separately and finds annotator–annotator agreement materially exceeding annotator–human agreement for loaded vocabulary and threat framing, (LLM–LLM minus LLM–human κ gaps of +.177 and +.139, against +.067 or less for the other three heads), so corpus rates for those two heads are the least independent of the five.

Typology sensitivity. Causeur’s Double-distinctive typology assignment rests partly on its ELITE-heavy selection profile, where ELITE-category arbitration agreement is lower ($\kappa = .276-.417$); post-hoc arbitration agreement is otherwise well calibrated, higher for HIGH-confidence ($\kappa = .729$) than MEDIUM-confidence ($\kappa = .586$) cases; a sensitivity analysis merging ELITE into OTHER confirms Causeur remains Double-distinctive under this taxonomy perturbation (BFMTV and Marianne shift cells; no focal outlet affected).

Baseline self-reference and precision-floor shortfalls.

The Sal.JS and Sel.JS corpus baseline is partially self-referential for three high-volume low-divergence outlets (Le Figaro, Le Parisien, Franceinfo; collectively $\approx 32\%$ of corpus): their low observed divergence partially reflects their own weight in the baseline rather than absolute editorial similarity to the panel median. A leave-one-out (LOO) baseline confirms rankings are robust: LOO-

corrected $r = 0.721$, $\rho = 0.727$ (vs. standard $r = 0.736$, $\rho = 0.741$), with only two near-boundary outlets changing typology cells (Valeurs actuelles: double→selection-dominant; L’Humanité: selection→double-distinctive); all focal outlets discussed in the text are stable, and the dissociability criterion ($r < 0.90$) holds comfortably under LOO correction. Two device heads fall below their stated production-precision floors after prior-shift correction (Elkan 2001 method, Appendix Table 10): loaded vocabulary (est. 0.643, floor 0.70, -5.7 pp) and us-vs-them (est. 0.550, floor 0.60, -5.0 pp); blame attribution is marginally above floor (est. 0.701); threat framing (est. 0.715) and rhetorical question (est. 0.929) remain comfortably above their floors. Corpus-level loaded-vocabulary and us-vs-them rates are therefore lower-precision estimates relative to the other heads. See also the inline note in §4.1 for loaded vocabulary.

Ethics, Reproducibility, and Data Availability

This study audits public editorial output and does not infer protected attributes of private individuals. The study period covers the first enforcement year of the EU Digital Services Act (DSA, February 2024 onward), which introduced requirements for independent audits of very large online platforms’ algorithmic systems; our methodology is a descriptive research audit and is not a regulatory compliance instrument, but the transparent lexicon, public dataset release, and reproducible threshold protocol are designed to be consistent with independent audit principles. Group labels refer to explicit lexical mentions in headlines and should not be interpreted as sentiment, endorsement, or hostility labels. Outlet-level and group-level results describe aggregate measurement patterns and should not be used for individual profiling or content moderation decisions.

The group salience rates carry a specific dual-use risk: high salience for Jews-mention headlines reflects antisemitism reporting and post-October-2023 Israel/Gaza security coverage, not editorial targeting; Far-right salience reflects electoral, parliamentary, and conflict coverage of a political movement; elevated Muslim rates concentrate in policy and security contexts. These numbers describe the rhetorical and event contexts in which groups appear in headlines (the story types and

wording conventions surrounding their mention), not editorial sentiment, intent, or uniform targeting. The Fdesouche exclusion check (Appendix Table 22) confirms elevated salience for Jews, Far-right, and Muslims is distributed across the broader outlet panel rather than concentrated in a single far-right source (Migrants salience drops 10.3 pp without Fdesouche, remaining 13.5 pp above baseline).

Specific political misuse scenarios warrant explicit counter-framing: (1) antisemitic actors may cite “Jews: highest salience lift” as evidence of Jewish over-representation in media discourse, inverting the finding, which measures the rhetorical intensity of event-driven coverage contexts surrounding group mentions, not claims about the communities themselves; (2) immigration-restriction actors may cite Migrants and Muslims figures as evidence of legitimizing media concern, ignoring that these rates are event-driven and concentrated in security and policy story forms; (3) Far-right actors may frame their high salience as evidence of media “persecution” rather than coverage density of a contested political movement in an election-intensive period. In all three scenarios, the interpretive key is the same: high salience in this audit measures the charge level of news contexts, driven by event type and outlet editorial register, not editorial hostility or targeting of the named groups. LLM-assisted annotations are not treated as ground truth: the final supervision set uses majority-vote resolution and selective human arbitration.

The final v6 supervision set, comprising 10,000 French headlines with binary salience labels, story-form labels, split assignments, and outlet/section/date metadata, is released as a manifest-only artifact: headline text is distributed as (outlet, date, section, URL) pointers to respect publisher copyright, so each record is sufficient for authorized reconstruction but does not redistribute verbatim content. The human-validation samples are the one disclosed exception: 516 headlines are released verbatim, because a blind double-annotation study cannot be re-run without the text that was annotated. The group-mention lexicons (219 terms, 12 groups) and the full annotation schema are included in full. The 902,111-headline corpus inference layer is released as a predictions-only artifact: headline identifier, outlet, section, date, predicted labels, and raw classifier confidence scores at six-decimal precision,

with no headline text. The annotator-panel materials are included in full: pairwise agreement tables for all three candidate second annotators against the primary annotator, and per-annotator device labels for the human-validated 499-headline subset, which together make the panel homogeneity check in Appendix Table 9 independently reproducible. Analysis and visualization scripts are released, except those requiring headline text. Model training and threshold-calibration code is not released: it is bound to the licensed 902,111-headline corpus and to the trained checkpoints, neither of which can be redistributed. All materials are available at <https://github.com/lefrenchnewslab/framing-wording-selection>.

References

- ACPM. 2026. [Classement audience OneNext — presse quotidienne nationale 2026 S1](#). Alliance pour les Chiffres de la Presse et des Médias (ACPM).
- Ramy Baly, Giovanni Da San Martino, James Glass, and Preslav Nakov. 2020. [We can detect your bias: Predicting the political ideology of news articles](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4982–4991, Online. Association for Computational Linguistics.
- Rodney Benson. 2013. *Shaping Immigration News: A French-American Comparison*. Cambridge University Press, Cambridge.
- Rodney Benson and Timothy Wood. 2015. [Who says what or nothing at all? speakers, frames, and frameless quotes in unauthorized immigration news in the united states, norway, and france](#). *American Behavioral Scientist*, 59(7):802–821.
- Ceren Budak, Sharad Goel, and Justin M. Rao. 2016. [Fair and balanced? quantifying media bias through crowdsourced content analysis](#). *Public Opinion Quarterly*, 80(S1):250–271.
- Julia Cagé, Moritz Hengel, Nicolas Hervé, and Camille Urvoy. 2022. [Hosting media bias: Evidence from the universe of french broadcasts, 2002–2020](#). Technical Report hal-03878119, Sciences Po / HAL.
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. [The media frames corpus: Annotations of frames across issues](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444, Beijing, China. Association for Computational Linguistics.
- Marion Dalibert. 2015. [Médias et mouvements sociaux minoritaires : un accès à la sphère publique régulé par la “francité” ?](#) *Sciences de la société*, 94:15–29.
- Ulrich K. H. Ecker, Stephan Lewandowsky, Ee Pin Chang, and Rekha Pillai. 2014. [The effects of subtle misinformation in news headlines](#). *Journal of Experimental Psychology: Applied*, 20(4):323–335.
- Charles Elkan. 2001. The foundations of cost-sensitive learning. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 973–978.
- Robert M. Entman. 1993. [Framing: Toward clarification of a fractured paradigm](#). *Journal of Communication*, 43(4):51–58.
- Louis Escoufflaire, Antonin Descampe, Antoine Venant, and Cédric Fairon. 2024. La subjectivité dans le journalisme québécois et belge : transfert de connaissance inter-médias et inter-cultures. In *Actes de JEP-TALN-RECITAL 2024, 31ème Conférence sur le Traitement Automatique des Langues Naturelles, volume 2 : traductions d’articles publiés*, pages 12–13. Translation of article originally published at JADT 2024 (17th International Conference on Statistical Analysis of Textual Data). Data is francophone (Québec + Belgium), not France.
- Lisa Fan, Marshall White, Eva Sharma, Ruishi Su, Prafulla Kumar Choubey, Ruihong Huang, and Lu Wang. 2019. [In plain sight: Media bias through the lens of factual reporting](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6343–6349, Hong Kong, China. Association for Computational Linguistics.
- Maksym Gabielkov, Arthi Ramachandran, Augustin Chaintreau, and Arnaud Legout. 2016. Social clicks: What and who gets read on twitter? In *ACM SIGMETRICS / IFIP Performance 2016*, Antibes Juan-les-Pins, France.
- Johan Galtung and Mari Holmboe Ruge. 1965. The structure of foreign news. *Journal of Peace Research*, 2(1):64–91.
- William A. Gamson and Andre Modigliani. 1989. Media discourse and public opinion on nuclear power. *American Journal of Sociology*, 95(1):1–37.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [ChatGPT outperforms crowd workers for text-annotation tasks](#). *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.
- Mauricio Gruppi, Benjamin D. Horne, and Sibel Adali. 2022. [NELA-GT-2022: A large multi-labelled news dataset for the study of misinformation in news articles](#). *Preprint*, arXiv:2203.05659.
- Tony Harcup and Deirdre O’Neill. 2017. [What is news? news values revisited \(again\)](#). *Journalism Studies*, 18(12):1470–1488.

- Camille Jehle and Florian Le Gallo. 2025. [Europe in the headlines: What two decades of french news reveal about EU sentiment](#). Technical Report Working Paper 1008, Banque de France.
- J. Richard Landis and Gary G. Koch. 1977. [The measurement of observer agreement for categorical data](#). *Biometrics*, 33(1):159–174.
- Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019. [Detecting frames in news headlines and its application to analyzing news framing trends surrounding u.s. gun violence](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 504–514, Hong Kong. Association for Computational Linguistics.
- Maxwell E. McCombs. 2005. [A look at agenda-setting: Past, present and future](#). *Journalism Studies*, 6(4):543–557.
- Maxwell E. McCombs and Donald L. Shaw. 1972. [The agenda-setting function of mass media](#). *Public Opinion Quarterly*, 36(2):176–187.
- Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. [Modeling framing in immigration discourse on social media](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263.
- Nic Newman, Richard Fletcher, Craig T. Robertson, Kirsten Eddy, and Rasmus Kleis Nielsen. 2023. [Reuters Institute Digital News Report 2023](#). Technical report, Reuters Institute for the Study of Journalism, University of Oxford.
- Yulia Otmakhova, Shima Khanehzar, and Lea Frermann. 2024. [Media framing: A typology and survey of computational approaches across disciplines](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15407–15428, Bangkok, Thailand. Association for Computational Linguistics. Outstanding Paper Award, ACL 2024.
- Sebastián Pinto, Federico Albanese, Claudio O. Dorso, and Pablo Balenzuela. 2019. [Quantifying time-dependent media agenda and public opinion by topic modeling](#). *Physica A: Statistical Mechanics and its Applications*, 524:614–624.
- Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. [SemEval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multi-lingual setup](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2343–2361, Toronto, Canada. Association for Computational Linguistics.
- Vincent Price and David Tewksbury. 1997. News values and public opinion: A theoretical account of media priming and framing. In G. A. Barnett and F. J. Boster, editors, *Progress in the Communication Sciences*, volume 13, pages 173–212. Ablex, New York.
- Martin Reisigl and Ruth Wodak. 2001. *Discourse and Discrimination: Rhetorics of Racism and Antisemitism*. Routledge, London.
- Laura Ruis, Akbir Khan, Stella Biderman, Sara Hooker, Tim Rocktäschel, and Edward Grefenstette. 2023. [The Goldilocks of pragmatic understanding: Fine-tuning strategy matters for implicature resolution by LLMs](#). In *Advances in Neural Information Processing Systems*, volume 36.
- Bertram Scheufele. 2004. [Framing-effects approach: A theoretical and methodological critique](#). *Communications*, 29(4):401–428.
- Dietram A. Scheufele. 1999. [Framing as a theory of media effects](#). *Journal of Communication*, 49(1):103–122.
- Dietram A. Scheufele and David Tewksbury. 2007. [Framing, agenda setting, and priming: The evolution of three media effects models](#). *Journal of Communication*, 57(1):9–20.
- Pamela J. Shoemaker and Tim P. Vos. 2009. *Gatekeeping Theory*. Routledge, New York.
- Wenlong Song, Bo Pang, Zihan Wang, Yilang Xu, Zetong Liu, and Murong Tan. 2024. [Differentiated secularism: discourse shaping of immigrants from different religious backgrounds in the french media](#). *Humanities and Social Sciences Communications*, 11:893.
- Timo Spinde, Manuel Plank, Jan-David Krieger, Terry Ruas, Bela Gipp, and Akiko Aizawa. 2021. [Neural media bias detection using distant supervision with babe - bias annotations by experts](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1166–1177, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Charles L. Stevenson. 1944. *Ethics and Language*. Yale University Press, New Haven.
- James W. Tankard. 2001. The empirical approach to the study of media framing. In Stephen D. Reese, Oscar H. Gandy, and August E. Grant, editors, *Framing Public Life*, pages 95–106. Lawrence Erlbaum, Mahwah, NJ.
- Petter Törnberg. 2025. [Large language models outperform expert coders and supervised classifiers at annotating political social media messages](#). *Social Science Computer Review*, 43(6):1181–1195.
- Panos Tsimpoukis. 2025. [Contesting dominant AI narratives on an industry-shaped ground: public discourse and actors around AI in the french press and social media \(2012–2022\)](#). *Journal of Science Communication*, 24(2):A10.

Teun A. van Dijk. 1991. *Racism and the Press*. Routledge, London.

Douglas N. Walton. 1996. *Argumentation Schemes for Presumptive Reasoning*. Lawrence Erlbaum, Mahwah, NJ.

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. [Can large language models transform computational social science?](#) *Computational Linguistics*, 50(1):237–291.

A Supplementary Material

Component	Count	Share
Politics	3,339	33.4%
Economy	3,335	33.4%
Society	3,326	33.3%
Any salience	4,093	40.9%
Loaded vocabulary	2,391	23.9%
Blame attribution	1,435	14.4%
Threat framing	871	8.7%
Us-vs-them	698	7.0%
Rhetorical question	749	7.5%
Charge = HIGH	2,745	27.5%
OTHER	1,980	19.8%
POLICY	1,883	18.8%
SOCIAL	1,571	15.7%
ELITE	1,159	11.6%
CONFLICT	756	7.6%
CRIME	753	7.5%
LAW	646	6.5%
ELECTIONS	607	6.1%
SCANDAL	425	4.2%
PARLIAMENT	220	2.2%

Table 7: Final v6 label distribution in the 10,000-headline development set after majority-vote merging and story-form arbitration.

Group	Terms	Prec.	Rec.	F1
Migrants	26	1.000	.674	.805
Police	29	.873	.674	.761
RN	14	.831	.736	.780
LFI	9	.691	1.000	.818
Workers	14	.860	.491	.625
Farmers	40	.877	.934	.904
Women	17	.857	.640	.733
Muslims	20	.940	.808	.869
Jews	14	.696	.970	.810
Unions	12	.821	.979	.893
Far-right	13	.901	.941	.921
Seniors	11	.967	.725	.829

Table 8: Held-out validation of the 219-term group-mention lexicons: precision, recall, and F1 for each of the 12 canonical groups.

(a) Reliability summary ($N = 10,000$)

Field	3-way agree.	Mean κ
Loaded vocabulary	81.1%	.661
Blame attribution	85.4%	.639
Threat framing	90.0%	.623
Rhetorical question	94.8%	.770
Us-vs-them	88.3%	.408
Charge band	82.2%	.699
Story form	55.7%	.635

(b) Homogeneity check ($N = 499$)

Device	LLM-LLM	LLM-hum.	H-H
Loaded vocabulary	.650	.474	.542
Blame attribution	.521	.510	.601
Threat framing	.623	.485	.613
Rhetorical question	.809	.807	.894
Us-vs-them	.505	.438	.510

Table 9: (a) Reliability summary for the final v6 stack. (b) Homogeneity check: LLM-LLM = mean pairwise κ among annotators; LLM-hum. = annotators vs. coders; H-H = the two coders.

Head	Threshold	Val. Prev.	Val. Prec.	Prod. Prev.	Est. Prec.	Δ pp
Loaded vocabulary	0.67	23.9%	73.3%	17.0%	64.3%	−9.0
Blame attribution	0.83	14.7%	71.9%	13.6%	70.1%	−1.8
Threat framing	0.84	8.8%	72.3%	8.5%	71.5%	−0.8
Rhetorical question	0.89	7.3%	95.7%	4.4%	92.9%	−2.8
Us-vs-them	0.89	6.4%	60.0%	5.3%	55.0%	−5.0
Any-salience	OR	41.1%	83.8%	33.7%	79.0%	−4.8

Table 10: Estimated precision drift under class-prior shift. Expected production precision is computed from validation-set TPR and FPR at each head’s chosen threshold, applying production corpus prevalence via Elkan (2001).

Field	Mistral-human κ	Ensemble-human κ	Δ
Loaded vocabulary	.553	.684	−.131
Blame attribution	.351	.379	−.028
Threat framing	.517	.614	−.097
Rhetorical question	.324	.400	−.076
Us-vs-them	.298	.546	−.248
Any salience	.435	.559	−.124
Story form	.449	.546	−.097

Table 11: French-specific Mistral comparison on the 499-headline human validation sample. Δ is Mistral-human κ minus ensemble-human κ ; negative values indicate lower agreement than the released ensemble.

Device (context)	Example and rationale
Loaded vocabulary (Jews)	“ <i>Au Dîner du Crif, François Bayrou dénonce «la bête délirante et meurtrière» de l’antisémitisme</i> ” (“At the Crif dinner, François Bayrou denounces the ‘deranged and murderous beast’ of antisemitism”) — <i>Le Figaro</i> , id 3789940, test split: the noun phrase <i>bête délirante et meurtrière</i> is an evaluative monstrosity metaphor with no neutral paraphrase of the same event. The loaded lexicon is voiced by a quoted speaker rather than by the outlet, illustrating the attribution limit of the headline unit noted in §3.
Blame attribution (Police)	“ <i>Mort de Blessing Matthew : les gendarmes mis en cause</i> ” (“Death of Blessing Matthew: the gendarmes implicated”) — <i>Mediapart</i> , id 3561734, validation split: causal responsibility for a death is assigned to a named institutional actor in the outlet’s own voice, with no attributing source.
Threat framing (Migrants)	“ <i>«Impression de submersion» migratoire : Yaël Braun-Pivet «gênée» par les propos de François Bayrou</i> ” (“Migratory ‘sense of submersion’: Yaël Braun-Pivet ‘uncomfortable’ with François Bayrou’s remarks”) — <i>Le Parisien</i> , id 1454100, train split: the submersion/overload lexicon reframes a demographic process as a security emergency, consistent with the threat-device definition in Table 1; as above, the phrase is quoted rather than asserted.
Rhetorical question	“ <i>Dépenses de santé : faut-il rembourser les soins en fonction des revenus ?</i> ” (“Health spending: should care be reimbursed according to income?”) — <i>Le Point</i> , id 436456, validation split: an interrogative headline that advances an evaluative proposition rather than seeking information.
Story form vs. topic (immigration)	“ <i>Immigration : vers plus de contrôle des mariages des personnes étrangères en situation irrégulière</i> ” (“Immigration: towards greater scrutiny of marriages of irregular foreign nationals”; <i>La Croix</i> , id 1376080) is POLICY; “ <i>Naufrage de migrants dans la Manche : quatre personnes mises en examen et incarcérées</i> ” (“Migrant shipwreck in the Channel: four charged and jailed”; <i>JDD</i> , id 2310601) is CRIME. Both are train-split rows. Same subject domain, different journalistic action.

Table 12: One example per primary salience device from the v6 supervision set, each carrying its device label by unanimous agreement of all three LLM annotators.

Outlet	N	Sal.JS	Sel.JS	High-Chg%	Typology
Fdesouche	16,124	0.038	0.147	82.1	Double-distinctive
Slate.fr	3,907	0.051	0.071	23.1	Double-distinctive
Causeur	3,056	0.033	0.099	30.4	Double-distinctive
Blast [†]	646	0.025	0.095	57.9	Double-distinctive
Ouest-France	90,423	0.021	0.082	4.6	Double-distinctive
Mediapart	7,092	0.016	0.075	48.7	Double-distinctive
Valeurs actuelles	14,286	0.011	0.056	54.6	Double-distinctive
Les Echos	55,694	0.031	0.065	10.4	Double-distinctive
L’Express	11,795	0.015	0.034	27.5	Double-distinctive
TF1 INFO	44,014	0.029	0.011	29.6	Salience-dominant
20 Minutes	62,832	0.011	0.019	42.6	Salience-dominant
BFMTV	57,617	0.012	0.023	40.8	Salience-dominant
Marianne	9,412	0.010	0.056	41.9	Selection-dominant
JDD	17,598	0.002	0.037	35.9	Selection-dominant
L’Humanité	17,911	0.011	0.043	33.2	Selection-dominant
Le Figaro	108,307	0.001	0.004	31.2	Low-distinctiveness
Le Parisien	93,877	0.004	0.016	43.8	Low-distinctiveness
Franceinfo	83,574	0.002	0.004	28.4	Low-distinctiveness
Le Monde	39,017	0.005	0.018	28.3	Low-distinctiveness
La Croix	35,980	0.002	0.010	23.0	Low-distinctiveness
Libération	33,921	0.004	0.021	34.5	Low-distinctiveness
Le Point	32,522	0.008	0.011	29.3	Low-distinctiveness
CNews	27,252	0.005	0.016	43.1	Low-distinctiveness
Le HuffPost	17,737	0.004	0.025	31.2	Low-distinctiveness
Le Nouvel Obs	17,517	0.002	0.015	34.5	Low-distinctiveness

Table 13: Full four-cell outlet typology with per-outlet headline counts. Median split thresholds: Sal.JS=0.0110, Sel.JS=0.0249. Displayed values are rounded; cell assignments use full precision.

Group	UvT lift	UvT rate	Interpretation
Far-right	×10.60	56.0%	elevated directional signal only; not a primary prevalence estimate
Muslims	×6.83	36.1%	elevated directional signal only; explicit mentions only
LFI	×2.97	15.7%	weaker directional signal in party/mobilisation coverage

Table 14: Supplementary us-vs-them analysis. Given only moderate human–human ($\kappa = .510$) and human–model ($\kappa = .596$) agreement, these values are qualitative signals, not point estimates.

Domain	Case	N	Any-Sal% 95% CI	High-Chg% 95% CI
Outlet	Fdesouche	16,124	72.6 [71.9, 73.3]	82.1 [81.5, 82.6]
Outlet	Blast [†]	646	67.3 [63.6, 70.8]	57.9 [54.1, 61.6]
Outlet	Causeur	3,056	59.1 [57.4, 60.9]	30.4 [28.8, 32.1]
Outlet	Le Figaro	108,307	33.2 [32.9, 33.5]	31.2 [30.9, 31.5]
Group	Jews	4,646	78.7 [77.5, 79.9]	79.8 [78.6, 80.9]
Group	Far-right	3,981	77.7 [76.3, 78.9]	65.8 [64.3, 67.2]
Group	Muslims	4,720	66.4 [65.0, 67.7]	67.3 [65.9, 68.6]
Group	Migrants	9,606	58.4 [57.4, 59.3]	61.6 [60.6, 62.6]
Group	Police	14,313	57.6 [56.8, 58.4]	79.7 [79.0, 80.3]
Group	Workers	15,456	39.8 [39.0, 40.6]	25.4 [24.8, 26.1]

Table 15: Representative 95% Wilson intervals for outlet- and group-level headline rates, recomputed from the canonical paper artifacts.

Group	N	Corpus%	Any-Sal%	Y×Story exp.%	Resid. pp	Dominant context
Jews	4,646	0.52	78.7	49.6	+29.1	CRIME 32.5%
Far-right	3,981	0.44	77.7	50.0	+27.7	CONFLICT 44.3%
Muslims	4,720	0.52	66.4	46.6	+19.8	CRIME 23.4%
Migrants	9,606	1.06	58.4	40.7	+17.7	CRIME 42.4%
Police	14,313	1.59	57.6	46.5	+11.1	CRIME 56.7%
Unions	6,877	0.76	53.8	50.3	+3.6	CONFLICT 56.8%
LFI	13,746	1.52	49.5	42.6	+6.9	CONFLICT 36.0%
Women	10,214	1.13	48.3	40.8	+7.5	SOCIAL 37.9%
RN	18,751	2.08	44.9	39.2	+5.7	ELECTIONS 32.5%
Farmers	22,112	2.45	41.9	37.3	+4.6	SOCIAL 29.4%
Workers	15,456	1.71	39.8	52.8	-13.0	CONFLICT 56.9%
Seniors	21,987	2.44	32.1	34.6	-2.5	POLICY 39.2%

Table 16: Group-level salience contexts across 902,111 headlines. Y×Story exp.% is the expected any-salience rate after reweighting to each group’s year×story-form composition.

Group	Dominant salience signature	Dominant selection signature	Interpretive caution
Jews	blame ×3.87; threat ×2.00; loaded ×2.30	high-charge 79.8%	often antisemitism/security coverage, not sentiment
Police	blame ×3.28; threat ×2.23	CRIME 56.7%; LAW 15.0%	accountability register
Far-right	loaded ×2.76; blame ×2.78; threat ×3.00	CONFLICT 44.3%; ELECTIONS 11.6%	multi-device profile
Muslims	threat ×3.71	CRIME 23.4%; POLICY 8.9%	explicit mentions only
Migrants	blame ×2.60; threat ×2.17	CRIME 42.4%	explicit mentions only
LFI	moderate salience lift	ELECTIONS 26.2%; CONFLICT 36.0%	party + mobilisation coverage
RN	moderate salience lift	ELECTIONS 32.5%; PARLIAMENT 8.1%	party/electoral framing
Workers	near-baseline salience	CONFLICT 56.9%	pension-reform period effect likely

Table 17: Condensed group-level mechanism table. It separates rhetorical device lift from selected story-form context.

Outlet	Form	n_f	Outlet%	Others	χ^2	ϕ
Fdesouche	CRIME any-sal.	9,407	75.8	43.7	3627.4	.167
Fdesouche	CRIME blame	9,407	64.6	30.8	4490.3	.186
Fdesouche	CRIME threat	9,407	41.0	13.8	4846.8	.193
Valeurs actuelles	CRIME any-sal.	4,213	57.6	45.6	234.3	.042
Slate.fr	SOCIAL rhet. q.	1,870	23.4	7.9	596.7	.056
Causeur	ELITE any-sal.	1,079	59.3	33.5	312.8	.075

Table 18: Within-story-form checks. Values are percentages; “Others” pools non-focal outlets within the same story form. All contrasts are significant at $p < .001$.

Head	Type	Headline	Error source
Loaded	FP	<i>Nationalisons la dette ! (Nationalize the debt!)</i>	Surface-form political imperative over-triggers on the charged verb; annotators classified this as advocacy rather than framing rhetoric.
Loaded	FN	<i>Pour Macron (et pour que rien ne change), bienvenue dans la commune la plus riche de France (For Macron (and so that nothing changes), welcome to the richest municipality in France)</i>	Sarcastic loaded expression (<i>pour que rien ne change</i>) requires discourse-level reading beyond surface form; ironic register under-represented in training.
Blame	FP	<i>Nîmes : l'homme armé qui s'est suicidé au tribunal recherchait «des magistrats mafieux» (Nîmes: the armed man who killed himself at the courthouse was looking for "mafia judges")</i>	Reported speech containing explicit blame language is attributed to a third party; model does not distinguish endorsed from reported blame frames.
Blame	FN	<i>Chez TotalEnergies, détruire la planète continue de rapporter gros (At TotalEnergies, destroying the planet is still highly profitable)</i>	Implicit blame through an agentive nominalized construction; model under-weights non-explicit accusation patterns relative to overt subject-verb-object blame frames.
Threat	FP	<i>Que risquent les Français qui partent faire le djihad ? (What do French people who leave to wage jihad risk?)</i>	Co-occurrence of <i>risquent</i> and <i>djihad</i> triggers the threat surface pattern; the actual frame is a juridical inquiry about legal consequences, not a threat directed at an audience.
Threat	FN	<i>Explosion de l'immigration en France : les demandes d'asile ont bondi de plus de 30 % en 2022, les régularisations de clandestins en hausse de 8 % (Explosion in immigration in France: asylum applications jumped by more than 30% in 2022, regularizations of undocumented migrants up 8%)</i>	Threat metaphor (<i>explosion</i>) combined with quantitative surge framing; the statistical register suppresses the model's threat signal relative to the metaphor.
Rhet. q.	FP	<i>De quoi meurt-on à travers le monde ? (What do people die from across the world?)</i>	Interrogative surface form triggers detection; the question is a factual epidemiological inquiry rather than a rhetorical framing device implying a political claim.
Rhet. q.	FN	<i>Avions : pourquoi les prix des billets ont-ils autant augmenté ? (Flights: why have ticket prices increased so much?)</i>	Precision-constrained threshold near-miss: model score 0.884 falls just below the recalibrated ceiling of 0.89; a genuine rhetorical question sacrificed at the precision ceiling.
Us-vs-them	FP	<i>Ces "ultradiscounters" qui lancent la guerre des prix en France (These "ultra-discounters" launching the price war in France)</i>	Commercial competition metaphor (<i>guerre des prix</i>) and categorical group label trigger the opposition pattern; frame lacks the political or social identity dimension required for this head.
Us-vs-them	FN	<i>Charente-Maritime : un maire agressé par des gens du voyage, un homme placé en garde à vue (Charente-Maritime: a mayor assaulted by Travellers, one man taken into custody)</i>	<i>Gens du voyage</i> is a salient social out-group in French discourse but this identity category is under-represented in training data, producing systematic under-detection of such headlines.

Table 19: Qualitative error analysis: ten representative false positives and false negatives from the 1,501-headline test set, two per salience head, drawn from highest-confidence misclassifications.

Device	H-H κ	Interp.	A1-LLM κ	A2-LLM κ	Maj-LLM κ	Interp.	F1	Prec.	Rec.
Loaded vocabulary	.699	substantial	.638	.523	.682	substantial	.766	.629	.980
Blame attribution	.772	substantial	.568	.559	.648	substantial	.653	.510	.907
Threat framing	.588	moderate	.664	.409	.685	substantial	.746	.623	.930
Rhetorical question	.974	almost perfect	.895	.923	.920	almost perfect	.917	.930	.904
Us-vs-them	.564	moderate	.717	.588	.845	almost perfect	.769	.833	.714
Any salience	.636	substantial	.735	.511	.715	substantial	.909	.864	.959

Table 20: Corroborating independent blind study on a separate sample: two unaffiliated annotators ($N=350$). *Note.* Stratified 350-headline sample drawn exclusively from the v6 test split (all HIGH-confidence; seed 42; no overlap with the primary 499-headline sample in Table 4). Annotation protocol matches the primary study (two unaffiliated annotators, blind to LLM labels and expected distributions). H-H = human-human reliability; Maj-LLM = majority-vote human labels vs. LLM consensus (conflict rows, where $A1 \neq A2$, excluded; conflict counts: loaded 53, blame 31, threat 72, rhetorical question 3, us-vs-them 60, any salience 48). F1/Prec./Rec.: human majority vote as gold standard. Interpretation thresholds follow Landis and Koch (1977).

Table 21: Threshold recalibration sensitivity by outlet: any-salience and high-charge rates under default and recalibrated production thresholds, with the change in percentage points.

Outlet	Any-Sal% (default)	Any-Sal% (recal.)	Δ	Chg% (default)	Chg% (recal.)	Δ
<i>Corpus aggregate</i>	41.5	34.6	-6.9	34.6	31.3	-3.3
Fdesouche	81.1	72.6	-8.5	83.8	82.1	-1.7
Blast [†]	72.6	67.3	-5.3	61.8	57.9	-3.9
Causeur	65.6	59.1	-6.5	35.1	30.4	-4.7
Slate.fr	55.3	49.6	-5.7	26.6	23.1	-3.5
Mediapart	67.4	58.8	-8.6	53.5	48.7	-4.8
Valeurs actuelles	60.9	52.0	-8.9	58.1	54.6	-3.5
Marianne	68.3	61.2	-7.1	46.4	41.9	-4.5
TF1 INFO	45.7	38.3	-7.4	33.2	29.6	-3.6
BFMTV	37.9	29.9	-8.0	44.9	40.8	-4.1
Le Figaro	40.1	33.2	-6.9	34.4	31.2	-3.2
Franceinfo	39.1	31.7	-7.4	32.1	28.4	-3.7
La Croix	36.3	29.3	-7.0	26.3	23.0	-3.3
Ouest-France	9.4	7.2	-2.2	5.4	4.6	-0.8

Note. “Default” uses a uniform 0.50 threshold; “Recal.” uses the precision-floor production thresholds from Table 2. Default thresholds inflate corpus-wide any-salience by 6.9 pp and high-charge by 3.3 pp, but the outlet ranking is preserved. [†]Blast ($N=646$): case profile, wide CIs (Table 15).

Table 22: Fdesouche exclusion sensitivity: group-level any-salience rates and baseline lifts recomputed with the Fdesouche aggregator removed from the outlet panel.

Group	Fdesouche share	Sal% (all outlets)	Sal% (excl. Fdesouche)	Δ
Jews	3.9%	78.7	78.2	-0.5 pp
Far-right	3.4%	77.7	77.2	-0.4 pp
Muslims	16.7%	66.4	64.4	-2.0 pp
Migrants	27.0%	58.4	48.1	-10.3 pp
Police	8.7%	57.6	55.1	-2.5 pp
Unions	1.3%	53.8	53.5	-0.4 pp
LFI	2.9%	49.5	48.8	-0.7 pp
Women	3.7%	48.3	47.0	-1.3 pp
RN	1.9%	44.9	44.6	-0.3 pp
Farmers	0.9%	41.9	41.6	-0.3 pp
Workers	0.8%	39.8	39.5	-0.3 pp
Seniors	0.9%	32.1	31.7	-0.4 pp

Note. Median exclusion effect across all 12 groups is -0.5 pp. Migrants is the most sensitive case: removing Fdesouche lowers salience by 10.3 pp (58.4% to 48.1%), but the group still sits 13.5 pp above the corpus baseline.

Group	2022	2023	2024	2025
Jews	74.6% ($\times 2.35$)	75.3% ($\times 2.26$)	77.2% ($\times 2.21$)	84.1% ($\times 2.21$)
Far-right	74.3% ($\times 2.34$)	74.7% ($\times 2.24$)	79.6% ($\times 2.27$)	81.4% ($\times 2.14$)
Muslims	63.8% ($\times 2.01$)	66.7% ($\times 2.00$)	65.5% ($\times 1.87$)	67.9% ($\times 1.78$)
Police	58.4% ($\times 1.84$)	52.9% ($\times 1.59$)	58.0% ($\times 1.66$)	63.1% ($\times 1.66$)
Migrants	53.2% ($\times 1.68$)	59.4% ($\times 1.78$)	58.7% ($\times 1.68$)	62.3% ($\times 1.64$)

Table 23: Year-by-year any-salience rates for the five highest-salience groups. Parenthetical values are lifts relative to the corpus-wide any-salience baseline for the same calendar year.

Hyperparameter	CamemBERT-base	XLM-RoBERTa-large
Pretrained model	camembert-base	xlm-roberta-large
Max sequence length (tokens)	48	48
Per-device batch size	16	8
Gradient accumulation steps	2 (eff. batch = 32)	4 (eff. batch = 32)
Learning rate	2×10^{-5}	1×10^{-5}
Weight decay	0.01	0.01
Dropout	0.2	0.1
Warmup ratio	0.06	0.06
Max epochs	20	20
Early stopping patience	3 (validation loss)	3 (validation loss)
Optimizer	AdamW	AdamW
Loss (salience/charge heads)	BCE with class-balanced pos. weights	BCE with class-balanced pos. weights
Loss (story-form head)	Cross-entropy with class weights	Cross-entropy with class weights
Classification pooling	[CLS] token	[CLS] token
Salience architecture	Shared encoder, 6 independent binary heads	Shared encoder, 6 independent binary heads
Story-form architecture	Separate encoder, 10-class head	Separate encoder, 10-class head
Training seed	42	42
Hardware	Google Colab (A100 / T4)	Google Colab (A100 / T4)

Table 24: Training hyperparameters for CamemBERT-base and XLM-RoBERTa-large on the final v6 supervision set. Both train on the same 6,999-headline split; thresholds use the 1,500-headline validation split only.

Device	Definition (YES)	Canonical YES examples	Canonical NO examples
Loaded vocabulary	Strong emotional, exaggerated, metaphorical, or value-laden wording beyond neutral reporting	<i>catastrophe, chaos, désastre, submersion, effondrement</i> ; metaphorical or dramatic phrasing	Neutral descriptive wording: <i>réforme, mesure, discussion</i>
Blame attribution	An actor is clearly presented as causing harm, failure, or a negative outcome	Explicit causation: <i>provoque, entraîne, aggrave</i> ; responsibility: <i>responsable de, à cause de</i> ; accusation: <i>accusé de, mis en cause</i>	Simple criticism or opposition without causation: <i>X critique Y, X s'oppose à Y</i>
Threat framing	Wording casts an actor, event, or situation as dangerous, destabilizing, or harmful	<i>menace, danger, risque, fait craindre</i> ; escalation: <i>au bord de, sous pression</i> ; metaphorical threat: <i>déferlante, explosion de, submerge</i>	Neutral reporting of events without danger framing
Rhetorical question	A question that implies judgment, criticism, or suspicion	Implies evaluation: <i>Faut-il craindre ? Macron a-t-il échoué ? Jusqu'où ira...</i>	Neutral informational questions: <i>Que dit la loi ? Que sait-on ?</i>
Us-vs-them [†]	Explicit opposition between two distinct groups, both holding social, political, ethnic, national, class, or institutional identity	<i>La France face aux migrants</i> ; class opposition: <i>les 1% vs le reste du monde</i> ; institutional: <i>gouvernement vs syndicats</i>	Politician vs. politician (CONFLICT); commercial competition; one actor + vague target; crime report victim/perpetrator
Charge band (HIGH)	HIGH = strong emotional intensity, crisis, violence, confrontation, or scandal	Crisis, attack, scandal, strong confrontation	Routine, neutral reporting (LOW)

Table 25: Condensed annotation schema for the released supervision criteria. [†]Us-vs-them requires both groups to hold explicit social, political, ethnic, national, class, or institutional identity.

Metric	Value	N
Cohen's κ (nominal)	.654	642
Cohen's κ (linear)	.666	642
Krippendorff's α	.653	642
% raw agreement	69.9%	642
<i>By primary annotator confidence</i>		
HIGH	.729	318
MEDIUM	.586	314
LOW	.157	10

Table 26: Post-hoc human-human agreement on the 642 three-way story-form conflict cases resolved by single-annotator arbitration. A second annotator applied the same codebook blind; original gold labels were not modified.

Outlet	Default cell	Perturbed cell(s)	Triggering perturbation(s)
Valeurs actuelles Le HuffPost	double-distinctive low-distinctiveness	selection-dominant selection-dominant	raise salience median by 10% or 20% lower salience median by 10% or 20%; any-salience-only perturbation except raising selection median
Marianne	selection-dominant	double-distinctive	lower salience median by 10% or 20%; any selection-only perturbation
Libération BFMTV	low-distinctiveness salience-dominant	selection-dominant double-distinctive; low-distinctiveness	lower selection median by 20% lower selection median by 10% or 20%; raise salience median by 20%
L'Humanité 20 Minutes	selection-dominant salience-dominant	double-distinctive low-distinctiveness	lower salience median by 10% or 20% raise salience median by 10% or 20%

Table 27: Outlets whose four-cell assignments change under $\pm 20\%$ perturbations of the salience and selection median split points.

Outlet	N	Full Any-Sal%	RQ-excl Any-Sal%	Δ pp	Δ Sal.JS
Fdesouche	16,124	72.6	72.2	0.5	−.026
Slate.fr	3,907	49.6	31.9	17.7	−.032
Causeur	3,056	59.1	50.0	9.1	+.002
Blast [†]	646	67.3	65.9	1.4	−.018
Ouest-France	90,423	7.2	5.8	1.4	−.002
Mediapart	7,092	58.8	56.8	2.0	−.009
Valeurs actuelles	14,286	52.0	49.9	2.1	−.009
Les Echos	55,694	24.2	21.0	3.2	+.004
L'Express	11,795	45.1	38.1	7.0	−.001
TF1 INFO	44,014	38.3	27.4	10.9	−.019
20 Minutes	62,832	39.1	33.5	5.7	−.000
BFMTV	57,617	29.9	27.8	2.1	−.001
Marianne	9,412	61.2	53.1	8.0	+.001
JDD	17,598	41.9	37.9	4.0	−.001
L'Humanité	17,911	49.9	44.1	5.8	+.001
Le Figaro	108,307	33.2	30.3	2.9	−.001
Le Parisien	93,877	40.3	37.3	3.1	−.001
Franceinfo	83,574	31.7	27.0	4.6	−.001
Le Monde	39,017	37.2	34.8	2.5	−.004
La Croix	35,980	29.3	25.6	3.7	+.000
Libération	33,921	45.7	40.9	4.8	+.000
Le Point	32,522	40.2	33.9	6.3	−.001
CNews	27,252	38.5	35.0	3.5	−.001
Le HuffPost	17,737	37.4	32.3	5.0	−.000
Le Nouvel Obs	17,517	42.2	36.8	5.4	−.000
Corpus	902,111	34.6	30.6	4.0	—

Table 28: Rhetorical-question-excluded any-salience robustness: any-salience recomputed as loaded vocabulary OR blame attribution OR threat framing. [†]Blast ($N=646$): case profile, wide CIs.

(a) Sparse-cell exclusion sensitivity					(b) Cluster-robust inference ($G=25$)				
Min. headlines	N cells	β	SE	p	Specification	β	Cluster SE	p	95% CI
5	1,197	0.8793	0.0637	<.001	Outlet FE	0.862	0.443	.063	[−.05, 1.78]
10	1,184	0.7582	0.0660	<.001	Two-way FE	0.758	0.420	.084	[−.11, 1.63]
20	1,161	0.8552	0.0689	<.001					
30	1,151	0.9218	0.0684	<.001					

Table 29: Temporal panel robustness. (a) Sensitivity to sparse-cell exclusion thresholds (two-way outlet+month fixed effects, classical SEs). (b) Cluster-robust standard errors clustered by outlet ($G=25$).