

Analyzing the Effect of Media Takeovers on their Political Positioning

Lubin Longuépée¹, Akram Elbouanani², Adrian Popescu²,

¹Université Paris-Saclay, Département Informatique, Orsay, France

²Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

Correspondence: firstname.lastname@universite-paris-saclay.fr

Abstract

We study whether media takeovers are associated with changes in the political positioning of news outlets. Across nine ownership changes spanning four languages and varied acquirer profiles, we find that aggregate outlet-level shifts are generally modest, but mask selective, actor-specific repositioning toward particular political figures and ideological blocs whose magnitude and direction vary by acquirer. A three-level decomposition — outlet, ideological bloc, individual actor — is combined with in-time placebo tests, within-outlet triple-interaction contrasts and entity-month re-estimation to separate the effects that withstand pre-trend and multiplicity checks from those that do not. We obtain these findings using a multilingual, entity-centered pipeline over CC-News that masks target politicians before zero-shot target-oriented sentiment classification and estimates changes via difference-in-differences against country-matched control media. These results suggest that aggregate stability masks a redistribution of tone across politically salient actors rather than uniform outlet-wide slant shifts.

1 Introduction

News media shape how citizens perceive political actors, institutions, and policy debates. Prior work in political communication and media economics shows that editorial positioning is influenced not only by audience demand and journalistic routines, but also by ownership incentives and political economy constraints (Gilens and Hertzman, 2000; Gentzkow and Shapiro, 2010; Djankov et al., 2003). This issue is especially salient when media outlets are acquired by industrial groups, politically connected actors, or owners with explicit ideological interests. Case studies have found content changes following ownership events, such as News Corp.’s acquisition of the *Wall Street Journal* (Archer and Clinton, 2018), Sinclair Broadcast

Group’s expansion into local television (Martin and McCrain, 2019), and newspaper mergers in Australia (Ananyev and Volkova, 2024). However, existing evidence is often limited to a few countries, outlets, or languages, making cross-context comparison difficult.

NLP offers scalable quantitative methods for measuring political positioning in text. Work on media slant, ideology detection, factuality, sentiment, and bias detection has enabled large-scale analysis of news corpora (Kulkarni et al., 2018; Fan et al., 2019; Nakov et al., 2024; Maab et al., 2024). The deployment of such tools is facilitated by LLMs’ strong, diverse capabilities. However, analysis methodologies should be carefully modeled and deployed to optimize performance and efficiency, while avoiding biases that can be important in sensitive domains, such as politics (Röttger et al., 2024).

At the same time, LLM-based classifiers introduce methodological risks: their predictions can reflect political and entity-specific priors rather than the framing present in the article itself (Elbouanani et al., 2025, 2026).

Informed by these existing research axes, we study whether media takeovers affect political positioning using a multilingual, entity-centered NLP pipeline. Our analysis covers 9 ownership changes across several countries, different types of buyers, and 4 languages. For each case, we collect articles from approximately 24 months before and after the takeover and compare the acquired outlet with a country-matched control set of mainstream media.

We retrieve articles from CC-News, extract sentences mentioning politicians from a curated multilingual list (1,297 global and 200+ national figures per country), adopt the target-masking strategy of Elbouanani et al. (2025) to mitigate LLM priors, classify sentiment with a zero-shot LLM, and estimate changes via difference-in-differences (Angrist and Pischke, 2009).

Prior work already shows that ownership can affect slant; we ask instead whether its effect is uniform or selective, and expect ideologically motivated acquisitions to produce bloc-level shifts aligned with the owner’s orientation, whereas commercial or locally oriented ones should produce actor- or geography-specific changes without a coherent ideological signature. Our primary finding is that ownership changes rarely produce uniform, outlet-wide shifts; instead, aggregate stability masks selective, actor-specific repositioning whose magnitude and direction vary by acquirer profile. We explicitly contribute: (i) an integrated entity-centered Target-oriented Sentiment Classification (TSC) + DiD audit pipeline combining masked zero-shot sentiment classification with clustered DiD inference — a validated measurement design rather than a novel component — benchmarked on MAD-TSC (Dufraisse et al., 2023) for English, French and Italian, applied zero-shot to Turkish, and validated against human annotation on the study corpus itself; (ii) an in-depth study of 9 eligible takeovers across 4 languages and varied acquirer profiles, designed not to estimate an average takeover effect but to test whether outlet-level stability coexists with selective actor-level repositioning; (iii) the substantive finding that editorial change around takeovers appears through actor-specific repositioning rather than uniform outlet-wide slant shifts, evidenced by a three-level decomposition (outlet, bloc, actor) with placebo, ordered-logit, and FDR robustness checks. We do not claim to identify a causal mechanism of editorial change, nor to capture dimensions of political positioning beyond the evaluative framing of named political actors (e.g., topic selection, salience, or structural omission (Card et al., 2015)).

2 Related Work

Agenda-setting and framing theory establish that news media structure how citizens perceive political reality (McCombs and Shaw, 1972; Entman, 1993), and empirical work confirms that media exposure shapes voting behavior, polarization, and belief accuracy (DellaVigna and Kaplan, 2007; Schneider-Strawczynski and Valette, 2025; Altay et al., 2025). If coverage affects public perceptions, ownership-induced changes in coverage may significantly shift a news source’s political positioning.

Ownership structures and corporate incentives shape editorial output (Gilens and Hertzman, 2000):

Sinclair acquisitions shifted local television toward more conservative coverage (Martin and McCrain, 2019), Australian mergers changed ideological slant (Ananyev and Volkova, 2024), and the Boloré takeover of CNews altered coverage with journalists adapting to the new direction (Cagé et al., 2025). These studies are generally limited to a single country or media system; our work extends this by comparing 9 takeovers across several countries and languages.

Computational approaches scale media analysis beyond manual coding (Ziems et al., 2024), from measuring slant via political language comparison (Gentzkow and Shapiro, 2010) to predicting outlet-level ideology (Baly et al., 2019, 2020), detecting framing and informational biases (Fan et al., 2019; Recasens et al., 2013; Tsur et al., 2015; Jähde et al., 2025), benchmarking factuality (Nakov et al., 2024), and producing sentence-level bias annotation resources (Lim et al., 2020; Spinde et al., 2021). Target-dependent sentiment classification (Hamborg and Donnay, 2021) is particularly suited to news, where a single article can express divergent attitudes toward multiple actors. LLM predictions, however, can depend on entity identity rather than textual evidence alone (Feng et al., 2023; Wang et al., 2023; Elbouanani et al., 2025), motivating our entity-masking approach. We connect these NLP measurements to causal inference via difference-in-differences (Angrist and Pischke, 2009; Bertrand et al., 2004), distinguishing market-wide shifts from ownership-specific effects.

3 Methodology

Our pipeline combines large-scale data sampling, entity extraction, zero-shot LLM inference, and Difference-in-Differences (DiD) estimation.

3.1 Data Collection and Scope

We sample data from the CC-News dataset¹, a publicly available dataset, spanning the years 2016 to 2024. We focus on nine target media outlets that underwent significant ownership changes. We list them in Table 1, along with their country of origin, the effective date of takeover used in our analysis, and the acquiring entity. We compiled 40 ownership changes between 2016 and 2024; nine met the design requirements of sufficient pre/post CC-News coverage, stable country-matched controls,

¹<https://huggingface.co/datasets/stanford-oval/ccnews>

and enough recurring politician mentions. They span diverse geographical contexts and takeover profiles, from industrial to politically motivated acquisitions. To ensure statistical validity, we construct a control group for each target by selecting 10 news domains from the same country, whose ownership did not change over the period. Control outlets satisfy three criteria: representativeness (major national outlets with high CC-News coverage), stability (no ownership change during the observation window, with the exceptions noted in Appendix A), and diversity (collectively spanning the political spectrum). Ten outlets per country is the minimum satisfying all three while remaining manually verifiable; the full list is in Appendix A.

Target Outlet	Country	Date	Acquiring Entity
Hürriyet	Turkey	05/18	Demirören
LA Times	USA	06/18	P. Soon-Shiong
Irish Examiner	Ireland	07/18	Irish Times Trust
Independent	Ireland	07/19	Mediahuis
La Repubblica	Italy	04/20	Exor/Elkann
La Stampa	Italy	04/20	Exor/Elkann
Europe 1	France	06/22	Vivendi/Bolloré
Euronews	Europe	07/22	Alpac Capital
La Provence	France	10/22	CMA CGM/Saadé

Table 1: Target media outlets and takeover details (MM/YY).

The analysis focuses on a maximum 48-month window centered on the official takeover date (24 months prior and 24 months post-takeover, when available). To isolate the structural post-acquisition editorial line from the noise of the transition period, we exclude a 6-month buffer window (± 3 months around the takeover date). This exclusion serves two purposes. First, it absorbs anticipation effects that may begin before the formal acquisition date, a known concern in DiD designs with staggered adoption (Callaway and Sant’Anna, 2021). Second, it discards the immediate post-takeover period during which editorial leadership is still being installed and journalists are adapting to the new ownership (Cagé et al., 2025; Martin and McCrain, 2019). Three months was selected by comparing pre/post sentiment stability across window sizes of one, two, and three months; one month left visible instability in quarterly trajectories, while three months produced stable baselines in all nine outlets.

3.2 Entity Lists

We track political figures at Global (1,297 international actors) and National (200–300 country-specific actors) levels, curated to span a classical

5-point ideological spectrum: radical-left, left, center, right, radical-right. We prompt Gemini 3.1 Pro to generate initial entity lists diversified across the political spectrum, along with their Wikipedia entries (see prompts in Appendix B). Out of caution, we automatically verify that all Wikipedia entries are valid. Finally, we map politicians to political alignments based on their Wikidata party affiliations. Table 2 reports the resulting ideological composition per list. We note that despite the efforts to diversify, some imbalance persists due to the national political configurations.

List	RL	L	C	R	RR
Global	2.8	27.7	18.0	43.9	7.7
French	7.0	11.7	53.6	15.0	12.6
Italian	0.4	39.5	28.3	11.0	20.8
American	2.9	40.5	2.0	11.2	43.4
Irish	2.9	18.5	28.9	49.8	0.0
Turkish	0.0	18.3	0.9	75.8	5.0
European	4.2	19.0	40.9	20.0	15.9

Table 2: Share (%) of politicians per ideological bloc and per list. RL = Radical Left, L = Left, C = Center, R = Right, RR = Radical Right.

3.3 Entity Linking and Selection

Named entity recognition (NER) spots the entity mentions. Entity linking improves the robustness of the subsequent processing by drastically reducing the ambiguity of mentions. We use GLiNER (Zarriana et al., 2024) for named entity recognition (NER) and mGENRE (De Cao et al., 2022) to link entities in the selected articles. These tools were chosen for their multilingual capabilities, which cover all languages relevant to our study, and for their strong NER and linking performance.

To ensure the analysis is not skewed by disparities in publication volume or isolated news cycles (e.g., elections or scandals), we implemented a downsampling and quality-control strategy. Three balancing rules were applied: *longitudinal validity* (entities must appear in at least 30 target articles in both periods), *peak smoothing* (capped at 50 mentions per entity per month via random subsampling), and *target-control matching* (a 1:3 monthly ratio, with proportional reduction when control coverage is insufficient).

Applying these thresholds yields a dataset totaling 1,376,552 sentences. The number of sentences per outlet ranges from 70,014 to 222,650 for the Irish Examiner and Euronews, respectively. The number of detected and linked unique entities per

outlet varies between 139 and 452 for Europe 1 and Euronews, respectively. This size disparity implies differential statistical sensitivity: results for smaller corpora (Hürriyet, Irish Examiner) should be interpreted with greater caution regarding power to detect modest effects. Entity coverage narrows at each stage. Of the politicians in the global and national dictionaries, 988 distinct figures are detected and linked at least once in the sampled articles, contributing 2,499 outlet–entity pairs. Applying the retention thresholds leaves 551 pairs with sufficient coverage in both periods to support a politician-level DiD. Because a politician tracked in several outlets is estimated separately in each, this figure counts outlet–entity test units rather than distinct individuals. The detailed dataset composition is provided in Appendix N.

3.4 Target-oriented Sentiment Classification

We select target-oriented sentiment classification (TSC) (Hamborg and Donnay, 2021) to operationalize political alignment through the evaluative framing of named actors — a dimension shown to correlate with outlet-level ideological positioning in prior work (Hamborg and Donnay, 2021; Elbouanani et al., 2025). TSC is particularly suited for news analysis, as a single article often expresses divergent attitudes toward multiple political actors. We acknowledge that this operationalization does not directly capture complementary dimensions of political positioning such as topic selection, issue salience, or structural omission (Card et al., 2015); we return to this point in the Limitations section.

To select the optimal classifier, we conducted a preliminary benchmark on the MAD-TSC dataset across English, French, and Italian (detailed in Appendix D). We found that modern Large Language Models (LLMs) outperform Small Language Model (SLM) baselines (Dufraisie et al., 2023) in all three languages: by up to 6.5 macro-F1 points over the non-translated baseline, and by 1.2 to 3.7 points over the stronger translation-based one, which is not an option for Turkish. Consequently, we deployed **Qwen 3.5 27B** using 4-bit quantization (GGUF). Fine-tuning on MAD-TSC yields no consistent gain over zero-shot (Appendix D), and zero-shot allows a uniform protocol for Turkish, which has no TSC training data. To reduce reliance on political priors (Elbouanani et al., 2025), target entities are fully anonymized before inference. Among six evaluated prompting strategies (see Appendix D), strict zero-shot classification

yields the most robust results; non-English texts are processed natively with output constrained to English labels (Appendix C).

We further validate the pipeline on the study corpus itself: 390 sentences in French and Turkish, balanced across the four pre/post $\times$ treated/control cells, annotated by three annotators (Appendix E). Accuracy is 70.9% (FR) and 65.5% (TR) against a moderately-agreeing human reference (Fleiss’ $\kappa = 0.50, 0.41$; the Turkish sentences were annotated in English machine translation), and entity linking is correct in 98% of cases. A separate audit on 1,047 French matches gives 97.2% linking precision, and a recall audit on unmatched sentences bounds recall below at 93.8% (FR) and 91.1% (TR) (Appendix E). Critically for the DiD, the error rate is statistically homogeneous across the four cells ($p \geq 0.28$): non-differential error inflates standard errors but does not bias β_3 .

3.5 Difference-in-Differences (DiD) Estimation

We use a difference-in-differences (DiD) design (Angrist and Pischke, 2009; Card and Krueger, 1994) to estimate whether post-takeover changes in outlet-level sentiment exceed contemporaneous shifts in the broader national media environment, following recent standards for causal inference in NLP (Keith et al., 2020; Feder et al., 2022; Egami et al., 2022). This is necessary because political sentiment can shift for reasons unrelated to ownership — such as elections, scandals, or country-wide agenda changes. In our observational setting, the acquired outlet acts as the treated unit, and country-matched outlets provide the counterfactual trend. While this design mitigates general media-cycle confounders, we emphasize that unobserved contemporaneous macro-shocks (e.g., the COVID-19 pandemic during the 2020 Italian acquisitions, or the Russia-Ukraine war in 2022) limit strict causal identification. We further follow standard recommendations by using clustered standard errors to account for correlation within media outlets over time (Bertrand et al., 2004). We note, however, that each national market contains a single *treated* cluster (one acquired outlet and ten control domains, nine for Hürriyet and the Irish Examiner, Appendix G), a regime in which cluster-robust standard errors are unreliable (Cameron and Miller, 2015); significance results should therefore be read as indicative (see Limitations). The ordered-logit cross-check in Appendix I validates sign and significance under

an ordinal outcome but does not address this bias. To rigorously quantify the relative sentiment shift of the takeovers, the LLM-generated sentiment outputs were mapped to a numerical sentiment score (*Positive* = 1, *Neutral* = 0, *Negative* = -1). Although mapping an ordinal variable to a linear scale is a simplification, this approach is standard in applied economics for DiD designs with discrete outcomes (Angrist and Pischke, 2009), and we validate this choice by re-estimating all significant DiD specifications with a cluster-robust ordered logit model. Across all 24 tested cases (three micro showcase cases, eight significant macro specifications, and thirteen FDR-significant political bloc cases), ordered logit and OLS agree on both the sign and the significance of the DiD interaction in every instance (see Appendix I).

We estimated the treatment effect using the following Ordinary Least Squares (OLS) regression:

$$\begin{aligned} \text{Sentiment}_{idt} = & \beta_0 + \beta_1 \text{Target}_d + \beta_2 \text{Post}_t \\ & + \beta_3 (\text{Target}_d \times \text{Post}_t) + \epsilon_{idt} \end{aligned} \quad (1)$$

Where: Sentiment_{idt} is the sentiment score assigned to a single sentence mentioning entity i , published in news domain d at date t ; Target_d is the treatment-group main effect, equal to 1 if the domain belongs to the acquired media and 0 if it belongs to the control group; it absorbs time-invariant level differences between the two groups and does not enter the DiD estimate; Post_t is a dummy variable equal to 1 if the article was published after the official takeover date, and 0 otherwise; β_3 is the coefficient of interest (the DiD estimator), representing the relative post-takeover shift in the target media’s sentiment compared to the control baseline. Because the score takes values in $\{-1, 0, +1\}$, its mean equals $P(\text{positive}) - P(\text{negative})$; β_3 is thus a change in net-positivity rate, e.g. $+0.044 = 4.4$ percentage points. Each observation is one entity–sentence occurrence; sentence-level estimation preserves the local variance of evaluative framing and maximizes statistical power, at the cost of intra-article and within-entity dependence, which we address by clustering standard errors ϵ_{idt} at the domain level. To separate shifts in *which* actors are covered from shifts in *how* a given actor is covered, we re-estimate each significant specification with entity fixed effects (Appendix J). The gap between β_3 with and without them decomposes an actor-selection component from a framing component;

this describes observable patterns, not the editorial mechanism.

We further report specifications with domain and month fixed effects (Appendix J). Domain effects absorb time-invariant level differences between outlets, so that β_3 is identified from within-domain variation only and is not affected by changes in which control domains are present in CC-News across periods (Appendix G); month effects absorb shocks common to all outlets in a given month, including the macro-events discussed in the Limitations. In both cases Target_d and Post_t are collinear with the fixed effects and drop out, while the interaction remains identified.

Furthermore, to mitigate the multiple testing problem arising from estimating thousands of simultaneous treatment effects across 551 outlet–entity specifications, we systematically apply the Benjamini-Hochberg False Discovery Rate (FDR) correction to all micro-level inferences. Statistical significance for macro-level updates is evaluated at the standard $p < 0.05$ threshold. Bloc- and micro-level updates are additionally filtered using the adjusted threshold (p -adjusted < 0.05); Table 4 reports the bloc-level results significant at $p < 0.05$ and marks the two that do not survive the correction. Finally, to validate the structural integrity of our DiD design and test the core assumption of parallel pre-treatment trends, we perform a series of strict, in-time placebo tests across the macro, political alignment, and micro levels. To avoid contamination from the true treatment effect, we restrict our sample exclusively to the pre-takeover period, artificially introducing a false acquisition date 12 months prior to the actual transaction. A comprehensive overview of this validation procedure, along with detailed results and econometric interpretations, is provided in Appendix F.

4 Analysis

We analyze post-takeover changes at three complementary levels: macro (all tracked figures), political alignment (ideological blocs), and micro (individual actors).

4.1 Macro-Level Shifts: Global and National Political Coverage

At the broadest level, we estimate DiD effects by aggregating all tracked political entities. The *Macro_Global* layer pools international and domestic politicians, while *Macro_National* restricts the

analysis to politicians from the outlet’s main domestic political arena.

Table 3 shows two² and six statistically significant post-acquisition divergences for global and national politicians.

Notably, *La Repubblica* displays diverging trends across aggregation levels, with a significant positive shift globally (+0.044) and a significant negative shift nationally (−0.020). This suggests that the ownership transition may have shifted the outlet’s framing of international affairs toward a more favorable sentiment, while it simultaneously increased editorial scrutiny over domestic political actors. To formalize the validity of these macro realignments, we executed strict in-time placebo tests restricted to the pre-takeover period; the results show a clean baseline for three of the eight significant specifications (both levels of *La Repubblica* and the *Irish Examiner* national level), while the remainder display significant pre-trends or an uninformative placebo: a significant placebo weakens the causal reading regardless of its sign, so we describe those cases as divergences occurring around the takeover rather than ownership-driven shifts (see Appendix F).

The *Irish Examiner* and *La Stampa* warrant particular caution: both combine large target–control divergences in CC-News crawl coverage (Appendix G) with placebo failures — at the macro level for *La Stampa*, whose national placebo matches the estimated effect, and at the bloc level for the *Irish Examiner*, whose Center, Radical Left and Right coefficients all continue a pre-existing trend (Appendix F). *Euronews*, whose control volume contracts while the target holds steady (Appendix G), should be read with the same caution.

4.2 Political Alignment Effects

We next aggregate politicians by political alignment, pooling all tracked figures per bloc. Because the global list is dominated by Right and Left figures (Table 2) and domestic coverage is where ownership-driven repositioning is most theoretically motivated, we also re-estimate the bloc-level effects on domestic politicians only (Appendix L), which restricts each outlet to its national arena and therefore estimates a different quantity on a smaller entity set. Of the 13 FDR-significant bloc results,

²The Hürriyet Macro_Global placebo is non-significant but comparable in magnitude to the true effect (+0.029 vs. +0.024), indicating low power rather than a validated effect; see Appendix F.

Outlet	Macro Global			Macro National		
	DiD	std	p	DiD	std	p
La Repubblica	+0.044	+0.069	<.01	−0.020	−0.031	<.01
Europe 1	+0.005	+0.008	0.495	+0.020	+0.031	<.01
La Provence	−0.003	−0.005	0.454	+0.020	+0.031	<.01
La Stampa	+0.017	+0.027	0.140	+0.032	+0.050	<.001
LA Times	+0.012	+0.019	0.168	+0.034	+0.053	<.001
Irish Examiner	+0.019	+0.030	0.083	+0.051	+0.080	<.001
Hürriyet	+0.024 [§]	+0.038	<.05	+0.006	+0.009	0.880
Euronews	−0.010	−0.016	0.232	+0.001	+0.002	0.943
Independent	−0.005	−0.008	0.651	+0.016	+0.025	0.357

Table 3: Macro-level DiD shifts by media outlet.

[§]Placebo not significant but comparable in magnitude to the effect, indicating low power rather than a validated result (App. F). Coefficients are on the net-positivity scale ($P(\text{pos}) - P(\text{neg})$); e.g. +0.044 is a 4.4 percentage-point shift. **std** is the standardized effect size (β_3/σ , with $\sigma = 0.638$ the pooled outcome standard deviation).

11 are testable on domestic politicians; eight remain significant, often with larger coefficients, while three lose significance on the reduced entity set, though only *La Repubblica* Right sees its coefficient collapse (+0.056 to +0.014), indicating an effect carried by international rather than domestic figures. The DiD analysis across the five political blocs shows that takeover effects are conditioned by the acquirer’s profile and the political opportunity structure of the outlet’s media market. Figure 1 visualizes the pre/post evolution for target and control outlets, and Table 4 reports the statistically significant bloc-level effects. *La Repubblica* is a useful illustration: its raw sentiment curves broadly track the Italian control group, yet the DiD decomposition reveals simultaneous positive and negative bloc-level realignments invisible in the aggregate.

The decomposition (Appendix J) shows that most significant shifts are predominantly framing rather than compositional; the main exceptions are the LA Times and Hürriyet macro shifts, which are roughly half selection and lose significance under entity fixed effects, consistent with the crawl-coverage and placebo concerns noted for these two outlets.

Targeted shifts are not confined to a single ideological family. Significant micro-level changes occur relatively evenly across the spectrum (38–47% per bloc), with no bloc systematically dominating, suggesting sampling variation rather than a structural ideological pattern. This indicates that post-takeover editorial changes, as identified by the DiD design relative to matched controls, are distributed across the entire political spectrum rather than con-




Figure 1: Editorial shifts across the corpus. The comparison between Target and Control groups, before (hatched) and after (solid) the takeover, reveals that significant editorial divergences (Table 4) are often masked by overarching national trends.

centrated in a single ideological family — a pattern that the placebo tests (Appendix F) confirm is not present at equivalent rates in the pre-takeover period. At the bloc-level aggregate (Table 4), Center and Right coefficients are predominantly positive (12 of 15), while this directional pattern does not systematically replicate at the entity level, where sign distributions are more balanced. Two patterns emerge. First, acquisitions are associated with targeted softening or hardening toward specific blocs: for *Europe 1* the Radical Right coefficient reflects a relative softening, though it does not survive entity-month aggregation (Appendix H). *Hürriyet*’s Center shift is not retained: its placebo pre-trend is opposite-signed and larger than the effect, and the Turkish national list contains almost no Center politicians (0.9%, Table 2), so the coefficient is carried by international centrists and is not testable under the domestic-only restriction (Appendix L).

Second, not all takeovers produce the same ideological signature. The *Irish Examiner* presents several significant bloc-level coefficients, but these reflect a broader convergence toward neutrality rather than a simple ideological turn; the significant Radical Right coefficient (-0.158) reflects coverage of foreign figures since the Irish national list contains no domestic radical-right politicians (Table 2), consistent with its disappearance in the domestic-

Target Media	Political Bloc	DiD Coeff.	std	P-Value
Euronews (EU)	Radical Right	+0.049 [¶]	+0.077	< 0.05
Europe 1 (FR)	Center	+0.017 [¶]	+0.027	< 0.05
	Radical Right	+0.048	+0.075	< 0.01
Hürriyet (TR)	Center	+0.128	+0.200	< 0.001
Irish Examiner (IE)	Center	+0.060	+0.094	< 0.001
	Radical Left	+0.123	+0.193	< 0.001
	Radical Right	-0.158	-0.248	< 0.001
	Right	+0.040	+0.063	< 0.001
La Provence (FR)	Radical Left	+0.056	+0.087	< 0.01
	Radical Right	+0.089	+0.139	< 0.001
La Stampa (IT)	Left	+0.048	+0.075	< 0.001
La Repubblica (IT)	Center	-0.027	-0.042	< 0.05
	Left	-0.022	-0.034	< 0.01
	Radical Left	+0.081	+0.127	< 0.001
	Right	+0.056	+0.088	< 0.001

Table 4: Significant DiD coefficients. Positive values indicate an improved stance (a relative increase in the sentiment score) compared with the control group. Coefficients are on the net-positivity scale; e.g. +0.089 is an 8.9 percentage-point shift. **std** is the standardized effect size (β_3/σ , with $\sigma = 0.638$ the pooled outcome standard deviation). [¶]Not significant after Benjamini-Hochberg FDR correction.

only sensitivity analysis (Appendix L). Conversely, in cases such as the *LA Times*, visible shifts in raw before/after plots do not remain significant once national trends are controlled for, confirming that those raw movements reflect the broader media environment rather than an ownership-specific effect. These patterns are consistent with the acquirer's profile, the outlet's prior editorial identity, and the political actors whose coverage becomes strategically relevant after the takeover — though establishing direct causal mechanisms behind each case goes beyond what the DiD design can identify. To confirm that these bloc-level contrasts reflect genuine within-outlet heterogeneity rather than isolated effects, we estimate a pooled regression with a triple interaction term (Target×Post×Alignment) for each outlet, testing whether the DiD effect of one ideological bloc differs significantly from another within the same model. The results corroborate the separation reported above: for the *Irish Examiner*, the contrast between Radical Right and Center blocs is the sharpest in the corpus ($\Delta = -0.219$, $p < 0.001$); for *La Repubblica*, Right and Radical Left blocs diverge significantly from the Left baseline ($p < 0.001$ in both cases); and for *Europe 1*, the Radical Right shift is statistically distinct from the Left reference bloc ($p = 0.003$). Full triple interaction results are reported in Appendix K. Placebo checks across the 15 pairs partition these results (Appendix F): four are clean under both placebo and FDR criteria (*Europe 1* Radical Right, *La Provence* Radical Left and Radical Right, *La Repubblica* Right); two are not rejected at the placebo but fail FDR (*Euronews* Radical Right, *Europe 1* Center) and one cannot be tested at all (*Irish Examiner* Radical Right, whose placebo is not estimable); the remaining eight show significant pre-trends—either sign reversals or continuations of pre-existing drift—and are reported as divergences around the takeover. Of the four clean results, only *La Provence* Radical Right remains significant under both the entity-month and domestic-only specifications, and we reserve causal wording for that case; *Europe 1* Radical Right reverses sign under entity-month aggregation, and the other two lose significance on the reduced domestic entity set.

4.3 Politician-level Effects

Globally, 215 of the 551 tested outlet–entity pairs (39.0%) exhibit statistically significant DiD coefficients after Benjamini-Hochberg FDR correction

(p -adjusted < 0.05 ; full list in Appendix O), with at most 5% of these expected to be false discoveries by the construction of the correction procedure. While these significant entity-level changes are widespread across the corpus, they are not evenly distributed across media systems, highlighting how context-dependent the intensity of editorial reorientation is. The highest shares of significant updates appear for *La Provence* (49.1%), the *Irish Examiner* (47.1%), and *Euronews* (46.5%), closely followed by *La Stampa* (42.1%) and *Europe 1* (40.7%). By contrast, *Hürriyet* exhibits the lowest share of significant changes (19.0%), on the smallest set of testable outlet–entity pairs in the corpus (21 pairs, against 114 for *Euronews*). The *LA Times* (34.9%), the sole US outlet, shows a national-level divergence (+0.034) that fails the placebo and is roughly half compositional (Appendices F and J), but entity-level changes concentrated on federal judicial figures (negative for Neil Gorsuch, Merrick Garland and Antonin Scalia) and on Californian officeholders (positive for Gavin Newsom, negative for Xavier Becerra), a pattern distinct from all European cases and not reducible to a partisan realignment. The *Independent* (33.3%) shows no significant macro shift, suggesting the Mediaset acquisition produced actor-specific rather than outlet-wide repositioning. For illustration we retain three outlets from different media systems. Among their entities passing FDR ($p_{adj} < 0.05$), we keep those with $N \geq 1,400$ sentences, a legibility threshold for the quarterly trajectories, and select the highest-|DiD| case in each, yielding Sarkozy, Raggi and Starmer (Figure 2; additional cases in Appendix M); placebo diagnostics are reported in Appendix F.3 (Table 13).

Nicolas Sarkozy (*Europe 1*). The positive shift for Nicolas Sarkozy illustrates targeted ideological alignment. While the national control group became more negative, plausibly reflecting contemporaneous legal controversies, *Europe 1* diverged positively after the takeover. This suggests not a general pro-right movement, but selective insulation of a politically salient figure from the broader negative media cycle. This case shows a significant pre-trend (Appendix F.3) and is therefore reported descriptively.

Virginia Raggi (*La Repubblica*). Before the acquisition, the tone toward the *Movimento 5 Stelle* (M5S) mayor of Rome was relatively stable compared to the national control group. Following the




Figure 2: Temporal evolution of editorial slant for selected case studies. The charts compare the acquired outlet with its respective control group. The vertical dashed line marks the acquisition date.

takeover, *La Repubblica* became relatively more negative toward Raggi compared with the national control group (DiD -0.161). Given Raggi’s association with M5S, a party often characterized as populist and anti-establishment, this pattern may be consistent with increased editorial scrutiny of anti-establishment actors under the new ownership. However, this interpretation remains suggestive and should be read alongside the broader evidence on *La Repubblica*’s post-takeover editorial positioning reported in Table 3 and Table 4.

Keir Starmer (*Euronews*). Euronews became relatively more positive toward Keir Starmer after the takeover, compared with the control group. This may reflect actor-specific realignment rather than a simple left-right ideological shift. One possible interpretation is that Euronews became more favorable toward a politically ascendant figure, although this mechanism remains suggestive.

Taken together, these cases illustrate why macro estimates should be read carefully: small average shifts can coexist with widespread, large entity-level realignments. The main finding is that takeovers reshape the distribution of tone across politically salient actors. These findings corroborate prior evidence (Martin and McCrain, 2019; Cagé et al., 2025) while showing the mechanism operates through actor-specific repositioning — invisible without entity-centered measurement.

5 Conclusion

We presented a multilingual entity-level analysis of political framing around media takeovers, combining target-oriented sentiment classification with difference-in-differences estimation against

country-matched controls. The measurement pipeline is benchmarked on MAD-TSC and validated on the study corpus itself, where classifier error is homogeneous across the four design cells and therefore inflates standard errors without biasing β_3 .

Post-takeover change is not primarily expressed as a uniform shift in outlet-level tone. Significant macro coefficients never exceed five percentage points, whereas 215 of 551 outlet–entity pairs (39.0%) shift significantly after FDR correction, and within-outlet triple-interaction contrasts separate blocs that move in opposite directions inside the same newsroom. Aggregate stability coexists with a redistribution of tone across politically salient actors.

That selectivity is also what makes causal attribution hard. Of the 15 significant bloc-level results, four are clean under both the placebo and FDR criteria, and one survives entity-month aggregation and the domestic-only restriction as well; we reserve causal wording for that case and report the remainder as post-takeover divergences. Going further would require designs this corpus cannot support: several treated outlets per market, multiway clustering, and a staggered-adoption event study. What the present design does establish is that ownership-related repositioning, where it occurs, is selective and actor-specific, and invisible to outlet-level measurement.

Limitations

Corpus coverage and outlet selection. The corpus is derived from CC-News and reflects its crawling frequency, domain coverage, and availabil-

ity. Some outlets, countries, or periods may be unevenly represented, and online articles do not fully capture broadcast, print, or editorial decision-making. Additionally, the nine target outlets were selected partly because they are sufficiently represented in CC-News, which biases the sample toward digitally active outlets; whether findings generalize to print-first or broadcast-dominated media systems remains an open question.

Entity detection and linking. The analysis depends on detected and linked politician mentions. Errors in named-entity recognition, entity linking, name variants, or political-alignment metadata can affect inclusion and aggregation. The retention thresholds improve comparability but exclude lower-visibility politicians. The audit in Appendix E bounds these errors—97.2% linking precision and recall above 93.8% (FR) and 91.1% (TR)—but covers only sentences from which no entity was extracted, so it does not detect a second politician missed inside an already matched sentence. Residual misses are dominated by unlisted transliteration variants, a surface-form gap that does not covary with treatment status.

Construct validity of TSC. Target-oriented sentiment captures only one dimension of political alignment and should not be equated with the construct as a whole. While it effectively captures evaluative framing toward named actors, it does not account for agenda-setting, topic selection, omission, headline prominence, or the broader narrative structure that defines political communication (Card et al., 2015). Neutral labels may further mask complex journalistic practices, such as factual reporting, balanced framing, or weak and ambiguous evaluation.

Target masking and classifier error. Although target masking reduces the risk that the LLM relies on political priors, it may remove contextual information that human readers would use. The classifier was selected through a preliminary benchmark, but its errors may vary by language, genre, outlet, and political context. Furthermore, our pipeline does not propagate measurement error from the underlying LLM classifier directly into the standard errors of the subsequent DiD regressions.

Counterfactual validity. The difference-in-differences design controls for contemporaneous shifts in matched media, but its interpretation depends on the plausibility of the counterfactual

trend. Control outlets may be affected by national political shocks, market-wide editorial changes, or ownership dynamics not captured in our metadata. This is concrete in Ireland, where four control domains changed ownership alongside the treated outlets (Appendix A). The takeover date is also an approximation; editorial effects may begin before the formal date or unfold gradually.

Temporal confounds. Several takeover windows coincide with major macro-shocks: the April 2020 Italian acquisitions (La Repubblica, La Stampa) overlap with the COVID-19 pandemic peak, while the June–October 2022 French and European takeovers coincide with the Russia–Ukraine war and energy crisis. Although DiD controls for contemporaneous shifts affecting all outlets equally, these events may have affected treated and control outlets asymmetrically—for instance, if the new owner of Europe 1 held specific editorial views on the conflict. Results for these outlets should be interpreted with this caveat in mind.

Unidentified mechanism. The DiD design documents *that* sentiment shifted, but cannot distinguish between direct editorial instructions from the new owner, journalist self-censorship in anticipation of ownership preferences, post-acquisition staff turnover, or changes in story selection rather than tone. These mechanisms have different implications for media pluralism policy and cannot be separated with the current data.

Few-cluster inference. Our econometric specification pools sentence-level occurrences and clusters standard errors at the domain level to account for intra-outlet correlation over time. The binding constraint, however, is not the number of clusters as such but the fact that each national market contains a single *treated* cluster, a regime in which cluster-robust standard errors and the wild cluster bootstrap are both unreliable; significance should therefore be read as indicative of relative editorial divergence rather than as a precise probabilistic statement. While the sentence-level approach maximizes data density, it introduces potential intra-article and within-entity dependence that simple clustering might understate. We report an entity-month re-estimation in Appendix H, which removes sentence-volume weighting and reduces this dependence; multiway clustering and a full staggered-adoption event-study design (Callaway and Sant’Anna, 2021) remain for future work.

Because the reference distribution matters in this regime, we re-computed all p -values using a t distribution with $G - 1$ degrees of freedom, the correction recommended by [Cameron and Miller \(2015\)](#) for few clusters and more conservative than the normal approximation used in Tables 3 and 4. Here $G = 11$, except for *Hürriyet* and the Irish Examiner, where only nine control domains are observed in CC-News (Appendix G) and $G = 10$. Twenty of the 21 reported specifications (Table 19, Appendix J) remain significant at the 5% level under this reference; the exception is the *Hürriyet* Macro_Global effect ($p = 0.059$), which we already treat as underpowered.

Entity list generation. Our approach relies on Gemini 3.1 Pro to generate the initial political entity lists. While this allows for rapid and broad entity tracking, the lists are subject to the model’s inherent selection biases and training-data knowledge cut-offs, which may lead to the underrepresentation of regional or fringe political actors. Results are therefore conditional on the sampled lists: they characterise editorial change toward the tracked actors, not toward each national political space in full.

Cross-lingual calibration. Turkish is the weakest language in our pipeline. It is absent from MAD-TSC, so the classifier was never benchmarked against an established external resource in that language, and the in-corpus validation we report (Appendix E) rests on 194 retained sentences annotated on automatic English translations, since none of the annotators reads Turkish: accuracy is 65.5% against a reference that is itself noisier than the French one, both because inter-annotator agreement is only moderate (Fleiss’ $\kappa = 0.41$) and because the reference is mediated by machine translation. Part of the gap with the French figures is therefore attributable to the reference rather than to the classifier. The classifier error rate does not differ significantly across the four design cells ($p = 0.32$), which is the property that matters for β_3 , but with 50 sentences per cell that test detects only gross differential error. Results for *Hürriyet* should therefore be read with more caution than those for the other outlets. More broadly, even across validated languages, the sentiment scale may not be equivalently calibrated: a POSITIVE classification in Turkish may correspond to a different absolute intensity than in French or Italian, which complicates direct cross-country compar-

ison of DiD magnitudes, though within-country comparisons between treated and control outlets remain unaffected. National lists also differ in their ideological composition (Table 2); cross-country comparison of bloc-level coefficients should therefore be interpreted as relative within-country shifts rather than absolute ideological positioning.

Parallel trends assumption. Our identification strategy assumes parallel trends between outlets subject to ownership changes and their matched controls. Although we control for concurrent shocks, this assumption is inherently non-testable in a counterfactual sense; idiosyncratic editorial decisions within control groups may mimic or obscure true treatment effects.

Statistical vs. substantive significance. Significant coefficients are not automatically substantively large. We therefore interpret macro, bloc, and politician-level results jointly, but some cases remain partly qualitative. The findings should be read as evidence of selective repositioning in the observed cases, not as a universal law of media ownership. With nine eligible cases, we do not estimate a population-level treatment effect for media takeovers; our claim is the coexistence of aggregate stability and actor-level selectivity within the observed cases, consistent across heterogeneous acquirer profiles rather than generalizable to all ownership changes.

Ethics and Impact

This work provides an entity-centered NLP pipeline to audit structural biases in media ecosystems following ownership changes. By decomposing editorial shifts down to individual political actors rather than relying only on global outlet averages, the framework supports cross-country research beyond isolated case studies. It can also assist editorial accountability, public-interest journalism, and media pluralism analysis by identifying cases where changes in tone are concentrated on specific political actors or ideological groups rather than visible at the outlet level. In operational settings, such signals should be interpreted as audit indicators rather than as definitive evidence of intentional editorial influence. They are most appropriate as inputs for secondary human review by researchers, regulators, journalists, or civil-society organizations.

The ethical use of this framework requires atten-

tion to measurement uncertainty, contextual interpretation, and potential downstream harms. The pipeline relies on automated entity recognition, entity linking, target-oriented sentiment classification, and statistical estimation, each of which can introduce errors. Results may therefore affect the perceived reputation of media outlets, owners, journalists, or political actors if reported without sufficient caveats. For this reason, outputs should not be used to rank outlets, infer individual journalists' intentions, or make legal or regulatory determinations without additional qualitative evidence. The framework is designed to study aggregate editorial patterns, not to attribute motives to specific persons.

The work also raises risks related to political sensitivity and dual use. A similar methodology could be used by media owners, campaign organizations, or strategic communication actors to monitor whether coverage remains below detection thresholds, to benchmark competitors' framing, or to optimize selective narrative realignments. These risks support deployment by independent third parties, transparent reporting of uncertainty, and restricted access to implementation details that would facilitate adversarial adaptation. When used in institutional settings, the framework should be accompanied by documentation of data sources, model versions, validation results, statistical assumptions, and known limitations.

Finally, the analysis should be understood as one component of a broader media pluralism assessment. Target-oriented sentiment captures evaluative framing toward named political actors, but does not capture topic selection, agenda-setting, omission, headline prominence, ownership governance, newsroom practices, or audience exposure. Ethical interpretation therefore requires combining the quantitative evidence with domain expertise, case-specific knowledge, and, where possible, complementary qualitative inquiry.

Acknowledgments

This work was made possible by the use of the FactoryIA supercomputer, financially supported by the Ile-de-France Regional Council.

References

Sacha Altay, Emma Hoes, and Magdalena Wojcieszak. 2025. [Following news on social media boosts knowl-](#)

[edge, belief accuracy and trust](#). *Nature Human Behaviour*, 9(9):1833–1842.

Maxim Ananyev and Ekaterina Volkova. 2024. [Media ownership and ideological slant: Evidence from Australian newspaper mergers](#). *PLOS ONE*, 19(12):e0315137.

Joshua D. Angrist and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton, NJ.

Allison M. N. Archer and Joshua D. Clinton. 2018. [Changing owners, changing content: Does who owns the news matter for the news?](#) *Political Communication*, 35(3):353–370.

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*, pages 4982–4991, Online. Association for Computational Linguistics.

Ramy Baly, Georgi Karadzhov, Abdelrhman Saleh, James Glass, and Preslav Nakov. 2019. [Multi-task ordinal regression for jointly predicting the trustworthiness and the leading political ideology of news media](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2109–2116, Minneapolis, Minnesota. Association for Computational Linguistics.

Marianne Bertrand, Esther Duflo, and Sendhil Mullainathan. 2004. [How much should we trust differences-in-differences estimates?](#) *The Quarterly Journal of Economics*, 119(1):249–275.

Samuel E. Bestvater and Burt L. Monroe. 2023. [Sentiment is not stance: Target-aware opinion classification for political text analysis](#). *Political Analysis*, 31(2):235–256.

Julia Cagé, Moritz Hengel, Nicolas Hervé, and Camille Urvoy. 2025. Political bias in the media – evidence from the universe of French broadcasts, 2002–2020. CESifo Working Paper 11741, CESifo.

Brantly Callaway and Pedro HC Sant'Anna. 2021. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

A. Colin Cameron and Douglas L. Miller. 2015. A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372.

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 (Volume 2: Short Papers)*, pages 438–444, Beijing, China. Association for Computational Linguistics.

- David Card and Alan B. Krueger. 1994. Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *American Economic Review*, 84(4):772–793.
- Nicola De Cao, Ledell Wu, Kashyap Popat, Mikel Artetxe, Naman Goyal, Mikhail Plekhanov, Luke Zettlemoyer, Nicola Cancedda, Sebastian Riedel, and Fabio Petroni. 2022. Multilingual autoregressive entity linking. *Transactions of the Association for Computational Linguistics*, 10:274–290.
- Stefano DellaVigna and Ethan Kaplan. 2007. [The fox news effect: Media bias and voting](#). *The Quarterly Journal of Economics*, 122(3):1187–1234.
- Simeon Djankov, Caralee McLiesh, Tatiana Nenova, and Andrei Shleifer. 2003. [Who owns the media?](#) *The Journal of Law and Economics*, 46(2):341–382.
- Evan Dufraisse, Adrian Popescu, Julien Tourille, Armelle Brun, and Jerome Deshayes. 2023. [MAD-TSC: A multilingual aligned news dataset for target-dependent sentiment classification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8286–8305, Toronto, Canada. Association for Computational Linguistics.
- Naoki Egami, Christian J. Fong, Justin Grimmer, Margaret E. Roberts, and Brandon M. Stewart. 2022. How to make causal inferences using texts. *Science Advances*, 8(42):eabg2652.
- Akram Elbouanani, Evan Dufraisse, and Adrian Popescu. 2025. [Analyzing political bias in LLMs via target-oriented sentiment classification](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 15476–15505, Vienna, Austria. Association for Computational Linguistics.
- Akram Elbouanani, Aboubacar Tuo, and Adrian Popescu. 2026. [A scalable entity-based framework for auditing bias in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 33823–33852, San Diego, California, United States. Association for Computational Linguistics.
- Robert M. Entman. 1993. [Framing: Toward clarification of a fractured paradigm](#). *Journal of Communication*, 43(4):51–58.
- 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.
- Amir Feder, Katherine A. Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E. Roberts, Brandon M. Stewart, Victor Veitch, and Diyi Yang. 2022. Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. *Transactions of the Association for Computational Linguistics*, 10:1138–1158.
- Shangbin Feng, Chan Young Park, Yuhao Liu, and Yulia Tsvetkov. 2023. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11737–11762, Toronto, Canada. Association for Computational Linguistics.
- Mara Finkelstein, Isaac Caswell, Tobias Domhan, Jan-Thorsten Peter, Juraj Juraska, Parker Riley, Daniel Deutsch, Geza Kovacs, Cole Dilanni, Colin Cherry, Eleftheria Briakou, Elizabeth Nielsen, Jiaming Luo, Kat Black, Ryan Mullins, Sweta Agrawal, Wenda Xu, Erin Kats, Stephane Jaskiewicz, and 2 others. 2026. TranslateGemma technical report. *arXiv preprint arXiv:2601.09012*.
- Matthew Gentzkow and Jesse M. Shapiro. 2010. [What drives media slant? evidence from U.S. daily newspapers](#). *Econometrica*, 78(1):35–71.
- Martin Gilens and Craig Hertzman. 2000. [Corporate ownership and news bias: Newspaper coverage of the 1996 telecommunications act](#). *The Journal of Politics*, 62(2):369–386.
- Felix Hamborg and Karsten Donnay. 2021. [NewsMTSC: A dataset for \(multi-\)target-dependent sentiment classification in political news articles](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1663–1675, Online. Association for Computational Linguistics.
- Orlando Jähde, Thorsten Weber, and Rüdiger Buchkremer. 2025. [Unraveling media perspectives: A comprehensive methodology combining large language models, topic modeling, sentiment analysis, and ontology learning to analyse media bias](#). *Journal of Computational Social Science*.
- Katherine Keith, David Jensen, and Brendan O’Connor. 2020. Text and causal inference: A review of using text to remove confounding from causal estimates. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5332–5344, Online. Association for Computational Linguistics.
- Vivek Kulkarni, Juntong Ye, Steve Skiena, and William Yang Wang. 2018. [Multi-view models for political ideology detection of news articles](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3518–3527, Brussels, Belgium. Association for Computational Linguistics.
- Sora Lim, Adam Jatowt, Michael Färber, and Masatoshi Yoshikawa. 2020. Annotating and analyzing biased sentences in news articles using crowdsourcing. In

- Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1478–1484, Marseille, France. European Language Resources Association.
- Iffat Maab, Edison Marrese-Taylor, Sebastian Padó, and Yutaka Matsuo. 2024. [Media bias detection across families of language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4083–4098, Mexico City, Mexico. Association for Computational Linguistics.
- Gregory J. Martin and Joshua McCrain. 2019. [Local news and national politics](#). *American Political Science Review*, 113(2):372–384.
- Maxwell E. McCombs and Donald L. Shaw. 1972. [The agenda-setting function of mass media](#). *Public Opinion Quarterly*, 36(2):176–187.
- Preslav Nakov, Jisun An, Haewoon Kwak, Muhammad Arslan Manzoor, Zain Muhammad Mujahid, and Husrev Taha Sencar. 2024. [A survey on predicting the factuality and the bias of news media](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15947–15962, Bangkok, Thailand. Association for Computational Linguistics.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. [Linguistic models for analyzing and detecting biased language](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1650–1659.
- Paul Röttger, Valentin Hofmann, Valentina Pyatkin, Musashi Hinck, Hannah Rose Kirk, Hinrich Schuetze, and Dirk Hovy. 2024. [Political compass or spinning arrow? towards more meaningful evaluations for values and opinions in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15295–15311, Bangkok, Thailand. Association for Computational Linguistics.
- Sarah Schneider-Strawczynski and Jérôme Valette. 2025. [Media coverage of immigration and the polarization of attitudes](#). *American Economic Journal: Applied Economics*, 17(1):337–368.
- 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.
- Oren Tsur, Dan Calacci, and David Lazer. 2015. A frame of mind: Using statistical models for detection of framing and agenda setting campaigns. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1629–1638, Beijing, China. Association for Computational Linguistics.
- Fei Wang, Wenjie Mo, Yiwei Wang, Wenxuan Zhou, and Muhao Chen. 2023. [A causal view of entity bias in \(large\) language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15173–15184, Singapore. Association for Computational Linguistics.
- Haoran Xu, Kenton Murray, Philipp Koehn, Hieu Hoang, Akiko Eriguchi, and Huda Khayrallah. 2025. [X-ALMA: Plug & play modules and adaptive rejection for quality translation at scale](#). In *International Conference on Learning Representations (ICLR)*.
- Urchade Zaratiana, Nadi Tomeh, Pierre Holat, and Thierry Charnois. 2024. Gliner: Generalist model for named entity recognition using bidirectional transformer. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5364–5376.
- 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 Control Media Selection

To construct a robust and unbiased baseline for our Difference-in-Differences (DiD) analysis, we selected 10 control media outlets for each targeted country or region. These specific control domains were carefully chosen according to three core criteria:

1. **Representativeness:** They are major national or regional publications with high circulation and web traffic, ensuring they are sufficiently represented within the CC-News corpus.
2. **Stability:** They did not undergo ownership changes during the respective 48-month observation windows. The Irish market is the exception: its media system is small enough that excluding every domain linked to the two Irish transactions would have left too few outlets to build a comparable control set, so we retained them and report the consequence below.
3. **Diversity:** Collectively, they offer a broad and diverse spectrum of political coverage. This mitigates the risk of ideological bias in our baseline, ensuring that the control trends are not artificially skewed by the specific editorial line of a single dominant publisher.

The complete mapping of the control domains used in our extraction and analysis pipeline is detailed in Table 5.

Table 5: List of Control Media Domains by Country/Region

Country/Region	Control Domains
USA	nytimes.com, washingtonpost.com, wsj.com, apnews.com, reuters.com, npr.org, bloomberg.com, nbcnews.com, nypost.com, theguardian.com
France	lemonde.fr, lefigaro.fr, liberation.fr, leparisien.fr, lesechos.fr, francetvinfo.fr, ouest-france.fr, lepoint.fr, 20minutes.fr, ladepeche.fr
Italy	corriere.it, ilsole24ore.com, ilfattoquotidiano.it, ansa.it, ilmessaggero.it, quotidiano.net, avvenire.it, tgcom24.mediaset.it, rainews.it, tg24.sky.it
Turkey	sozcu.com.tr, cumhuriyet.com.tr, sabah.com.tr, haberturk.com, yenisafak.com, karar.com, birgun.net, t24.com.tr, trthaber.com, ntv.com.tr
Ireland	irishtimes.com, rte.ie, thejournal.ie, irishmirror.ie, businesspost.ie, breakingnews.ie, sundayworld.com, newstalk.com, extra.ie, echolive.ie
Europe	politico.eu, bbc.com, france24.com, dw.com, reuters.com, theguardian.com, ft.com, euobserver.com, euractiv.com, aljazeera.com

Note. Four Irish control domains overlap with the two Irish treatment events: irishtimes.com is the acquirer of the Irish Examiner, and breakingnews.ie and echolive.ie were acquired in the same 2018 transaction; sundayworld.com was acquired by Mediahuis alongside the Independent in 2019. These domains were retained for comparability with the other markets, but the overlap attenuates the treated–control contrast, so the Irish estimates should be read as lower bounds.

B Entity Selection and Compilation

To systematically analyze media bias, we compiled comprehensive dictionaries of tracked entities categorized into three distinct levels: Global, National, and Specific. The Specific level covers actors tied directly to each takeover—the acquirer and the individuals and entities associated with it—and is used to check whether coverage of the new owner itself changes after the transaction. It accounts for 4,168 of the 1,376,552 occurrences (Appendix N) and is too sparse to support DiD estimation, so it does not enter the analyses reported in the paper.

Global and National Figures

The lists of *Global* (major international figures) and *National* (country-specific politicians) entities were generated using the Gemini Large Language Model (Gemini 3.1 Pro) to ensure broad coverage of the political spectrum. Lists were curated to cover the full five-point ideological spectrum (radical-left, left, center, right, radical-right), with manual inspection by a native speaker for the French list.

The *Global* list was generated with the following prompt:

List major political figures active internationally during the last 10 years, covering the full political spectrum. Include more than 1000 politicians, without constraints on country of origin. Create a TSV file including the following columns: the Wikidata ID of the politician, the usual name of the politician, the Wikipedia entry link.

The *National* lists were generated per country using the following prompt template (adapted by replacing the country name accordingly):

List American (USA) politicians active during the last 10 years, covering the country’s political spectrum. Include between 200 and 1000 politicians in the list, depending on your knowledge about American politics. Create a TSV file including the following columns: the wikidata ID of the politician, the usual name of the politician, the Wikipedia entry link.

Ambiguous mentions are resolved during the entity-linking step (Section 3.3) rather than via manual list curation.

C LLM Inference Prompts

To ensure strict, reproducible, and easily parsable classifications across all target media, we utilized a highly constrained zero-shot prompting strategy. The prompts were designed to force the LLM to output a single standardized label without generating conversational filler or chain-of-thought reasoning.

We deployed two prompt variants: a standard version for English corpora, and a cross-lingual version equipped with a specific language constraint for non-English corpora (French, Italian, and Turkish).

Standard Prompt (English Media)

System Prompt:

You are an expert computational journalist and political analyst. Your task is to perform Targeted Sentiment Analysis on a news paragraph. Read the text and determine the author’s underlying sentiment specifically directed towards “ENTITY”. Use a balanced standard across the three labels. Do not overuse NEUTRAL when the wording clearly frames ENTITY favorably.

User Prompt:

Classify the sentiment towards ENTITY as exactly ONE of the following: POSITIVE, NEUTRAL, or NEGATIVE.

Use POSITIVE when the paragraph presents ENTITY with favorable framing, approval, praise, competence, success, vindication, beneficial outcomes, growing support, constructive leadership, or advantageous political momentum. POSITIVE does not require explicit praise words if the author’s framing clearly benefits ENTITY.

Use NEGATIVE when the paragraph presents ENTITY with criticism, blame, failure, scandal, wrongdoing, incompetence, danger, harmful consequences, declining support, or adverse political momentum.

Use NEUTRAL when the paragraph is primarily factual, procedural, descriptive, balanced, attributed to someone else without authorial endorsement, merely reports that ENTITY spoke/met/announced/attended/said something, or when positive and negative cues are both absent or too weak. NEUTRAL should also be used for mixed or ambiguous framing where the author’s sentiment towards ENTITY is not clearly favorable or unfavorable.

Output ONLY the exact word and nothing else. Do not explain your reasoning.

Sentence: [INSERT SENTENCE HERE]

Cross-Lingual Prompt (Non-English Media)

For articles extracted from non-English media, a specific hint was injected into the user prompt to ensure the model processed the native language while standardizing the output schema in English.

User Prompt:

...

Important: The input news text is not necessarily in English. Classify based on the meaning of the text. Your final label must be exactly one English word.

...

D Preliminary TSC experiment

To justify our choice of classifier, we evaluated 8 different Large Language Models (LLMs) against established Small Language Model (SLM) baselines on the test subset of the MAD-TSC dataset (Dufraisse et al., 2023). We detail these preliminary target-oriented sentiment classification results in Table 6.

The findings confirm that LLMs outperform SLMs, though the size of the gap depends on the baseline. Against RoBERTa SPC, which operates on the original text, the deployed configuration (Qwen 3.5 27B GUID Q4KM) gains 2.9, 6.5 and 6.4 macro-F1 points in English, Italian and French. Against the stronger RoBERTa SPC TRAN baseline, which first translates the input into English, the gains narrow to 2.9 and 3.7 points in English and French, and the deployed configuration falls 0.2 points short in Italian; the best LLM configuration in that language (Qwen 3.5 27B RAW, 73.7) still leads SPC TRAN by 1.2 points. We nevertheless retain the GUID variant, which is the strongest in two of the three languages, produces a markedly better balance across the three classes (see the caption of Table 6), and requires no translation step—an advantage for Turkish, where SPC TRAN is not available, and one that costs nothing given that translation brings no gain to LLMs (Table 7). Consequently, we deployed the Qwen 3.5 27B model for the primary media outlet analysis.

Zero-shot vs. Fine-tuning

To assess whether supervised adaptation would improve measurement quality, we fine-tuned Qwen3.5-27B with QLoRA on the English, Italian and French MAD-TSC training subsets and

evaluated it under the same protocol. The fine-tuned model reaches macro-F1 of 74.4, 69.3 and 77.6 respectively—0.8 and 3.0 points below the deployed (GUID) zero-shot configuration for English and Italian, and 0.4 points above for French; measured against the best zero-shot configuration in each language, the Italian gap widens to 4.4 points. Fine-tuning therefore yields no consistent gain. We retain the zero-shot configuration, which additionally permits a uniform inference protocol for Turkish, for which no suitable TSC training subset is available.

Model	EN	IT	FR
RoBERTa SPC	72.3	65.8	70.8
RoBERTa SPC TRAN	72.3	72.5	73.5
Qwen 2.5 7B RAW	64.8	63.1	63.8
Llama 3 8B RAW	64.1	57.0	59.6
Aya Expanse 32B RAW	65.8	58.4	63.1
Qwen3.5 35B A3B RAW	69.3	66.1	67.8
Qwen2.5 32B RAW	71.0	70.7	70.6
Qwen2.5 72B RAW	72.9	70.7	73.6
Qwen3 32B RAW	69.8	67.8	69.0
Qwen3.5 27B RAW	74.9	73.7	74.5
Qwen3.5 27B GUID	75.2	72.3	77.2
Qwen3.5 27B QLoRA (fine-tuned)	74.4	69.3	77.6

Table 6: Macro F1 scores for the preliminary TSC experiment. RoBERTa results are those from (Dufraisse et al., 2023). “SPC” does not use corpus translation, while “SPC TRAN” translates sentences to English before TSC inference. Qwen3.5 27B RAW is a version of the zero-shot LLM prompt in which we do not provide guidance regarding the three classes, resulting in higher per-class imbalance between NEGATIVE, NEUTRAL, and POSITIVE (84.9, 73.5 and 66.3 in English, respectively). Qwen3.5 27B GUID adds guidance about how to predict the three classes and reduces the imbalance between classes (84.1, 70.7, and 70.8, respectively). The guidance is detailed in Appendix C. Entity names are masked in all experiments.

Impact of Target Anonymization

To ensure that replacing the names of political figures with a generic token ([ENTITY]) did not harm the model’s natural language understanding, we compared inferences on raw texts containing the true names versus their anonymized counterparts. Using the *Qwen2.5 32B* architecture as a baseline, the Macro F1 scores vary little: 71.3 (true name) vs. 71.0 (anonymized) in English, 72.4 vs. 70.6 in French, and 69.6 vs. 70.7 in Italian. The largest change is a 1.8-point drop in French; English is essentially unchanged and Italian improves by 1.1 points. Masking therefore costs at most about two macro-F1 points in the worst of the three languages,

and occasionally yields a minor gain. Crucially, this allows us to strip the texts of preexisting LLM biases towards specific politicians without sacrificing analytical efficacy.

Impact of Machine Translation

While previous research demonstrated that adding automatic translations to English before inference improved results for SLMs (Dufraisse et al., 2023), we sought to investigate whether this preprocessing step also benefits modern LLMs. We tested two state-of-the-art open-source translation models: **TranslateGemma 27B** (Finkelstein et al., 2026) and **X-ALMA 13B** (Xu et al., 2025).

Due to the prohibitive computational cost of translating the entire corpus for every LLM configuration, we restricted this exploratory experiment to the **Qwen2.5 32B** model. As detailed in Table 7, translating the French and Italian subsets to English prior to sentiment classification did not yield any meaningful performance gains. In fact, scores fluctuated within a narrow margin of error and even occasionally degraded. Consequently, we opted to run all our LLM inferences natively in the original languages, avoiding the massive computational overhead of specialized translation pipelines.

Translation Strategy	EN	IT	FR
Native (Original Text)	71.0	70.7	70.6
TranslateGemma 27B → EN	–	70.4	70.8
X-ALMA 13B → EN	–	69.4	71.0

Table 7: Comparison of Macro F1 scores on the *Qwen2.5 32B* model when using native text versus translating French and Italian texts to English prior to inference. All tests use anonymized entities.

Impact of Prompt Engineering Strategies

To ensure the robustness of our Zero-Shot Target-oriented Sentiment Classification, we evaluated the impact of various prompt engineering strategies. We hypothesized that framing the task differently (e.g., using a persona, shifting from “sentiment” to “stance”, or providing examples) could affect the LLM’s comprehension of journalistic text.

Using the *Qwen2.5 32B* model with anonymized entities (native text without translation), we compared six prompt templates:

- **Prompt 1: Zero-Shot Sentiment (With Persona).** The baseline prompt used in our main study (detailed in Appendix C), asking for

POSITIVE/NEUTRAL/NEGATIVE classification while adopting the persona of an expert.

- **Prompt 2: Zero-Shot Sentiment (Direct).** Same as Prompt 1, but without the initial persona framing (direct instruction).
- **Prompt 3: Zero-Shot Stance (With Persona).** Shifts the classification paradigm from sentiment to political stance, asking to classify the text as SUPPORT, NEUTRAL, or OPPOSITION, using an expert persona.
- **Prompt 4: Zero-Shot Stance (Direct).** Same as Prompt 3, but without the persona.
- **Prompt 5: Few-Shot Sentiment (3-shot).** Includes three synthetic examples (one for each class) mapping sentences to NEGATIVE, NEUTRAL, and POSITIVE.
- **Prompt 6: Chain of Thought Stance (CoT).** Instructs the model to write a one-sentence preliminary analysis of adjectives and framing before outputting the final SUPPORT/NEUTRAL/OPPOSITION class.

Prompt Strategy	EN	IT	FR
P1: Zero-Shot Sentiment (Persona)	71.0	70.7	70.6
P2: Zero-Shot Sentiment (Direct)	69.8	70.6	71.7
P3: Zero-Shot Stance (Persona)	65.7	69.0	69.8
P4: Zero-Shot Stance (Direct)	67.3	70.4	71.1
P5: Few-Shot Sentiment (3-shot)	62.9	66.3	65.6
P6: Chain-of-Thought (Stance)	65.7	68.5	70.0

Table 8: Macro F1 scores across the six prompt strategies evaluated on the *Qwen2.5 32B* model using anonymized native texts. P1 provides the most consistent baseline, particularly in English.

As shown in Table 8, advanced prompting paradigms like Few-Shot (P5) and Chain of Thought (P6) unexpectedly degraded performance. This is likely because synthetic examples (P5) fail to capture the complex, implicit framing inherent to political news, while CoT (P6) causes the model to over-analyze factual reporting as biased. Furthermore, framing the task as Targeted Sentiment (Positive/Negative towards an entity) consistently yielded better or more stable results across languages than framing it as general political Stance (Support/Oppose). This aligns with the core premise of Bestvater and Monroe (2023)—who show that sentiment and stance are conceptually distinct constructs in political text—and confirms

that instructing LLMs on evaluative sentiment provides a cleaner measurement of framing than asking them to infer complex ideological stances. Consequently, Prompt 1 was selected as the optimal and most reliable configuration for the final analysis.

E Human Validation on CC-News

The MAD-TSC benchmark (Appendix D) evaluates the classifier on an external corpus. Because differential measurement error—error rates that differ across the four design cells (pre/post $\times$ treated/control)—would bias β_3 , we additionally validate the pipeline on the study corpus itself. We sample 200 sentences per language, balanced across the four cells (50 per cell), in French and Turkish; Turkish has no human-annotated TSC resource, so it is the language for which external validation is least available. Each sentence is annotated by three annotators on the masked text actually passed to the classifier. We retain a sentence when at least two annotators agree and the third does not assign the opposite polarity, which discards 10 of 400 sentences (3 opposed, 7 without majority).

Annotators and protocol. The three annotators are the authors of this paper. No crowdsourcing platform was used and no compensation was paid, the annotation being part of the authors’ own research activity. Sentences were presented in randomised order, with no indication of the outlet, the date or the design cell, so annotators were blind to treatment status. Each annotator saw only the masked sentence passed to the classifier and assigned POSITIVE, NEUTRAL or NEGATIVE to the sentiment expressed by the author of the text toward the masked entity, following the class definitions of the classifier prompt (Appendix C); annotators worked independently and did not see each other’s labels or the model predictions. None of them reads Turkish: the Turkish sentences were annotated on automatic English translations, whereas the classifier processed the original Turkish. The Turkish reference is therefore weaker than the French one, and part of the gap between the two languages in Table 9 is attributable to translation error rather than to classifier quality. Because translation error is a property of the sentence and not of the outlet or the period, it does not covary with treatment status, which is the property that the cell-invariance test in Table 9 checks.

Two observations matter for the DiD. First, the

	FR	TR	All
Sentences retained	196	194	390
Accuracy (%)	70.9	65.5	68.2
95% CI	[64, 77]	[59, 72]	[63, 73]
Macro-F1	0.71	0.66	0.69
Cohen's κ (model-gold)	0.56	0.48	0.52
Fleiss' κ (inter-annotator)	0.50	0.41	0.45
Entity linking (%)	98.0	98.5	98.2
Cell invariance (χ^2, p)	0.49	0.32	0.28

Table 9: Human validation on CC-News. *Cell invariance* tests whether the classifier error rate differs across the four pre/post $\times$ treated/control cells; non-significant values indicate homogeneous error.

<i>French</i>			
gold \ model	NEG	NEU	POS
NEGATIVE	45	0	2
NEUTRAL	18	65	34
POSITIVE	0	3	29
<i>Turkish</i>			
gold \ model	NEG	NEU	POS
NEGATIVE	40	4	1
NEUTRAL	21	54	31
POSITIVE	0	10	33

Table 10: Confusion matrices against the human majority label.

error rate is statistically homogeneous across the four design cells in both languages ($p = 0.49$ and $p = 0.32$) and pooled ($p = 0.28$). Classical, non-differential measurement error inflates standard errors but does not bias β_3 ; the invariance test finds no evidence of the differential error that would.

Second, the dominant error is a tendency to over-assign polarity to neutral sentences: recall is high for NEGATIVE (0.96 FR, 0.89 TR) and POSITIVE (0.91, 0.77) but drops to 0.56 (FR) and 0.51 (TR) for NEUTRAL. The misassignments occur on both sides, with a skew toward POSITIVE (18 toward NEG vs. 34 toward POS in French; 21 vs. 31 in Turkish). This is a known consequence of the guided prompt (Appendix D); crucially, it does not covary with treatment status.

The moderate absolute accuracy reflects both this bias and the intrinsic difficulty of the task: inter-annotator agreement is itself moderate (Fleiss' $\kappa = 0.50$ and 0.41), so the reference labels are not noise-free, and for Turkish they are additionally mediated by machine translation. We present this study as a diagnostic rather than a certification of accuracy. With 50 sentences per cell, the standard error on a per-cell error rate is approximately 6.5 points and

roughly 13 points on a difference-in-differences of error rates, so the invariance test detects gross differential error only. Turkish is the weaker case on every metric and should be read accordingly.

Entity detection and linking

Precision. We manually verified 1,047 randomly sampled entity matches in the French subcorpus. Nine were incorrect NER assignments for the target entities; entity-linking precision is 97.2%, with residual errors concentrated on family names cited without a first name.

Recall. Precision alone leaves open systematic misses, which would make the analysis sample non-random. Because a missed politician leaves no trace in the linked corpus, we audit recall on its complement: sentences from which the pipeline extracted no entity. We sample 150 such sentences per language and, for each, an author identifies every politician mentioned and classifies it as *in-list* (a genuine pipeline miss) or *out-of-list* (a politician absent from our dictionaries, i.e. a coverage gap rather than a detection error). As in the annotation study above, the French sentences were read in the original and the Turkish ones in automatic English translation; since translation tends to normalise name spellings, an in-list politician is if anything easier to recognise in the translated Turkish sample than in the original.

The unmatched pool is small relative to the linked corpus: 504,172 linked mentions against 33,078 unmatched sentences in French, and 87,725 against 8,572 in Turkish. Assuming at most one missed politician per unmatched sentence, recall is therefore bounded below by 93.8% (FR) and 91.1% (TR) even if every unmatched sentence conceals an in-list politician. The observed in-list miss rates are 78 of 150 sampled sentences in French and 31 of 150 in Turkish; extrapolating them to the full pool gives recall estimates of 96.7% (95% CI [96.2, 97.2]) and 98.0% ([97.4, 98.6]).

In-list misses are almost entirely unlisted transliteration variants of names that our dictionaries carry in a single spelling (*Loukachenko*, *Zelensky*, *Sergueï*). This explains the gap between the two languages: the French outlets cover international actors whose names are transliterated, whereas Hürriyet's unmatched sentences are predominantly domestic. Because a surface-form gap is a property of the name rather than of the outlet or the period, this error is non-differential with respect to treatment

status. Out-of-list mentions are dominated by the acquirers themselves (Saadé, Niel, Bolloré, Tapie) and by historical figures, both outside the scope of the tracked lists.

F Robustness Check: Strict In-Time Placebo Tests

To ensure that our significant Difference-in-Differences (DiD) estimates capture the actual effect of ownership change rather than pre-existing editorial drift, we conduct a strict in-time placebo test. A common pitfall in placebo testing is allowing the true post-treatment period to contaminate the placebo post-period. To avoid this, we restrict our sample exclusively to the true *pre-takeover* period. We then artificially split this restricted dataset by introducing a placebo takeover date (12 months prior to the true takeover) and rerun our DiD models.

A non-significant placebo coefficient is consistent with parallel pre-treatment trends. A significant placebo weakens the causal reading regardless of its sign: we therefore classify each significant specification as *clean* (placebo not significant), *trend reversal* (placebo significant, opposite sign), or *continuation* (placebo significant, same sign, indicating pre-existing drift). Causal wording is reserved for results that are clean under both the placebo and the FDR criterion; all others are reported as divergences occurring around the takeover.

F.1 Macro-Level Placebo Estimates

Table 11 reports the placebo results exclusively for the 8 macro-level models that were statistically significant in our main analysis.

Four of the eight models show non-significant placebo coefficients. In three of them—both levels for *La Repubblica* and the national level for the *Irish Examiner*—the placebo coefficient is close to zero and well below the estimated effect, which we read as a clean treatment effect. The fourth, *Hürriyet Global*, passes the test only nominally: its placebo coefficient (+0.029) is larger than the effect it is meant to rule out (+0.024), so the non-rejection reflects insufficient power rather than the absence of a pre-trend, and we do not treat this result as validated. For *Europe 1*, *La Provence* and the *LA Times*, the placebo coefficients are significant but *negative* while the true effects are positive. This pattern is consistent with a trend reversal, but

Outlet (Level)	Real DiD	Placebo DiD	Placebo p
La Repubblica (Glob.)	+0.044**	+0.008	0.301
La Repubblica (Nat.)	−0.020**	+0.007	0.393
Irish Examiner (Nat.)	+0.051***	−0.000	0.995
Hürriyet (Glob.)	+0.024*	+0.029§	0.159
Europe 1 (Nat.)	+0.020**	−0.019†	0.007
La Provence (Nat.)	+0.020**	−0.024†	<0.001
LA Times (Nat.)	+0.034***	−0.017†	<0.001
La Stampa (Nat.)	+0.032***	+0.032†	0.001

Table 11: Macro-level strict placebo tests restricted to originally significant models. Signif. codes for Real DiD: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. † indicates a significant placebo pre-trend ($p < 0.05$). § indicates a placebo that is not significant but comparable in magnitude to the effect, so the test is uninformative rather than passed. The top four show no significant pre-trend; the bottom four show a significant placebo (three sign reversals, one continuation for *La Stampa*).

a significant pre-trend violates parallel trends: we therefore report these as divergences around the takeover rather than as identified ownership effects. The single notable exception is *La Stampa (National)*, where the placebo coefficient matches the true treatment effect (+0.032, $p = 0.001$). This specific alignment indicates that the domestic coverage shift in *La Stampa* was already structural and underway during the pre-treatment period. Rather than a sudden shock, the ownership transition here is best interpreted as the continuation and institutional consolidation of a long-term editorial drift that was already driving the outlet’s trajectory relative to its control group.

F.2 Bloc-Level Placebo Estimates

We apply the same strict placebo design to the 15 political blocs that exhibited significant relative shifts in the main analysis (Table 12).

Six of the 14 testable pairs display no significant pre-trend. Applying the joint placebo-and-FDR criterion leaves four (Table 14): the *La Provence Radical Left* and *Radical Right* shifts, *La Repubblica’s* shift toward the *Right*, and *Europe 1’s* *Radical Right* coefficient. *Euronews Radical Right* and *Europe 1 Center* are clean at the placebo but do not survive FDR correction. Where placebo effects are significant, they indicate either anticipation dynamics or the continuation of a pre-existing drift, and we do not attribute those changes to the ownership

Outlet (Political Bloc)	Real DiD	Placebo DiD	Placebo p
Euronews (Radical Right)	+0.049*	+0.031	0.081
Europe 1 (Center)	+0.017*	-0.007	0.449
Europe 1 (Radical Right)	+0.048**	-0.007	0.501
La Provence (Radical Left)	+0.056**	+0.020	0.318
La Provence (Radical Right)	+0.089***	-0.021	0.104
La Repubblica (Right)	+0.056***	-0.011	0.103
Irish Examiner (Rad. Right)	-0.158***	N/A [‡]	—
Hürriyet (Center)	+0.128***	-0.339 [†]	<0.001
Irish Examiner (Center)	+0.060***	+0.050 [†]	0.042
Irish Examiner (Rad. Left)	+0.123***	+0.188 [†]	0.030
Irish Examiner (Right)	+0.040***	+0.031 [†]	<0.001
La Stampa (Left)	+0.048***	+0.051 [†]	0.003
La Repubblica (Center)	-0.027*	-0.026 [†]	0.025
La Repubblica (Left)	-0.022**	+0.037 [†]	<0.001
La Repubblica (Rad. Left)	+0.081***	+0.084 [†]	<0.001

Table 12: Bloc-level strict placebo tests for the 15 originally significant pairs. Signif. codes for Real DiD: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. [†] indicates a significant placebo pre-trend ($p < 0.05$). [‡] The Irish Examiner’s radical right bloc lacked sufficient pre-treatment observations for placebo model convergence.

transition.

F.3 Micro-Level Placebo Estimates

Finally, we test the validity of the structural breaks observed in our three primary case studies: Nicolas Sarkozy (*Europe 1*), Virginia Raggi (*La Repubblica*), and Keir Starmer (*Euronews*).

The entity-level tests demonstrate high robustness. Both Keir Starmer and Virginia Raggi display highly non-significant placebo coefficients, confirming that the massive shifts in their editorial coverage occurred exclusively after the formal acquisition. For Nicolas Sarkozy, the placebo test detects a marginal, negative pre-trend ($\beta_3 = -0.047$, $p = 0.047$). However, the true treatment effect is substantially larger and positive ($\beta_3 = +0.193$, $p < 0.001$). This is consistent with *Europe 1* having been on a trajectory of increasing criticism toward Sarkozy alongside the control group, with the divergence reversing after the takeover; given the significant pre-trend, we report this case as descrip-

Outlet (Subject)	Real DiD	Placebo DiD	Placebo p
Euronews (Keir Starmer)	+0.308***	+0.021	0.528
La Repubblica (Virginia Raggi)	-0.161***	-0.005	0.886
Europe 1 (Nicolas Sarkozy)	+0.193***	-0.047 [†]	0.047

Table 13: Micro-level strict placebo tests for selected case studies. Signif. codes for Real DiD: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. [†] indicates a significant placebo pre-trend ($p < 0.05$). For Starmer and Raggi, the placebo confirms a clean treatment effect, while for Sarkozy it reveals a trend reversal.

tive rather than as an identified ownership effect.

F.4 Summary of the Placebo Triage

Outlet (Bloc)	DiD	Category
<i>Placebo not rejected or not testable, but excluded</i>		
Europe 1 (Rad. Right)	+0.048	clean
La Provence (Rad. Left)	+0.056	clean
La Provence (Rad. Right)	+0.089	clean
La Repubblica (Right)	+0.056	clean
<i>Clean at placebo but excluded</i>		
Irish Examiner (Rad. Right)	-0.158	not testable [‡]
Euronews (Rad. Right)	+0.049	not FDR-significant
Europe 1 (Center)	+0.017	not FDR-significant
<i>Significant pre-trend</i>		
Hürriyet (Center)	+0.128	trend reversal
La Repubblica (Left)	-0.021	trend reversal
Irish Examiner (Center)	+0.060	continuation
Irish Examiner (Rad. Left)	+0.123	continuation
Irish Examiner (Right)	+0.040	continuation
La Stampa (Left)	+0.048	continuation
La Repubblica (Center)	-0.027	continuation
La Repubblica (Rad. Left)	+0.081	continuation

Table 14: Placebo triage of the 15 significant bloc-level results. [‡] The Irish Examiner Radical Right placebo could not be estimated for lack of pre-period observations; the effect also reflects coverage of foreign figures and vanishes under the domestic-only restriction (Appendix L). Euronews Radical Right and Europe 1 Center are clean at the placebo but do not survive FDR correction.

At the macro level, four of the eight significant specifications show no significant placebo (La Repubblica Global and National, Irish Examiner National, Hürriyet Global); Europe 1, La Provence and the LA Times show trend reversals, and La Stampa a continuation. We treat only the first three as clean, since the Hürriyet Global placebo is comparable in magnitude to the estimated effect and therefore uninformative. Among the four clean

bloc-level results, the entity-month re-estimation (Appendix H) confirms both La Provence blocs, leaves La Repubblica Right borderline, and reverses the sign for Europe 1 Radical Right, which we therefore treat as suggestive.

G CC-News Crawl Volume

We report raw CC-News article counts, before entity extraction, politician filtering, downsampling and target–control balancing, in event time. Pre covers months -24 to -4 and Post months $+4$ to $+24$. Target values are mean monthly counts for the acquired outlet; control values are the total across the control set divided by the number of observed months, so target and control levels are not directly comparable—the ratios describe change within each group. *Active* reports the number of control domains with non-zero volume.

Design	Target		Controls		Active
	pre	ratio	pre	ratio	pre/post
Euronews	5,852	0.95	31,633	0.52	10/10 $\rightarrow$ 9/10
Europe 1	817	0.93	23,464	1.05	10/10 $\rightarrow$ 10/10
Hürriyet	1,548	3.26	7,442	3.28	7/9 $\rightarrow$ 9/9
Independent	2,114	1.25	7,972	1.31	7/10 $\rightarrow$ 10/10
Irish Examiner	1,014	2.44	5,143	1.61	7/9 $\rightarrow$ 9/9
LA Times	1,973	1.58	31,445	1.25	10/10 $\rightarrow$ 10/10
La Provence	1,884	1.08	23,300	1.10	10/10 $\rightarrow$ 10/10
La Repubblica	4,906	1.49	20,899	1.17	10/10 $\rightarrow$ 10/10
La Stampa	921	2.45	20,899	1.17	10/10 $\rightarrow$ 10/10

Table 15: CC-News crawl volume in raw articles/month; ratio = Post/Pre. Hürriyet and the Irish Examiner have nine control domains (one absent from CC-News); the Italian outlets share a control set and takeover month. Terminal months after CC-News ends (2024-06) are excluded from the means rather than coded as zero.

Coverage is not stationary: CC-News expands through 2016–2018 and thins after 2023. The relevant quantity is whether the treated outlet diverges from its own controls. On the log ratio of ratios, the quantity here, the largest target–control gaps are La Stampa (0.74), Euronews (0.60) and the Irish Examiner (0.42). Euronews is the only design in which control volume contracts while the target holds steady (0.95 vs 0.52), with one control domain dropping out entirely after the takeover; its estimates should be read with the same caution as the two outlets we qualify in Section 4.1. Hürriyet rises sharply but in both groups (3.26 vs 3.28), a common trend that the DiD differences out.

H Entity-Month Re-estimation

Sentence-level estimation weights each entity–month by its underlying number of observations. We therefore re-estimate all specifications on entity–domain–month means, so that every entity–month contributes one observation regardless of volume (338,081 entity–month means vs. 1,731,681 target–entity observations; the corpus of 1,376,552 sentences yields 1,731,681 observations, as a sentence mentioning several politicians contributes one observation per politician).

Of the specifications significant in the sentence-level analysis, 130 remain significant after FDR correction.

The effect is concentrated at the macro level: only one of the six significant *Macro_National* estimates survives (*La Repubblica*, -0.013 , $p = 0.011$), while *Hürriyet*, not significant in the main analysis, becomes significant with the opposite sign (-0.062 , $p < 0.001$). Micro- and bloc-level results survive far more often, consistent with the paper’s claim that the reliable signal lies at the actor and bloc levels rather than outlet-wide. Table 16 reports all bloc-level specifications.

Outlet	Bloc	sent. β_3	e.-m. β_3	
Europe 1	Rad. Right	+0.048	−0.040	⊖
La Provence	Rad. Left	+0.056	+0.093	✓
La Provence	Rad. Right	+0.089	+0.033	✓
La Repubblica	Right	+0.056	+0.047	○
La Repubblica	Center	−0.027	−0.034	✓
La Repubblica	Left	−0.022	−0.020	✓
La Repubblica	Rad. Left	+0.081	+0.194	✓
La Stampa	Left	+0.048	+0.018	✓
Irish Examiner	Rad. Left	+0.123	+0.117	✓
Irish Examiner	Rad. Right	−0.158	−0.157	✓
Irish Examiner	Center	+0.060	+0.003	×
Irish Examiner	Right	+0.040	−0.000	×
Hürriyet	Center	+0.128	+0.001	×

Table 16: Bloc-level DiD under sentence-level vs. entity–month estimation. Top panel: the four placebo-and-FDR-clean results. The two bloc-level results that do not survive FDR correction (Table 4) are excluded, leaving 13 specifications. ✓ significant and same sign; ○ borderline ($p = 0.051$); × no longer significant; ⊖ significant but sign reversed. FDR-corrected.

Among the four clean results, the two *La Provence* blocs are stable, *La Repubblica* Right is borderline, and *Europe 1* Radical Right reverses sign—consistent with the entity fixed-effect decomposition (Appendix J), which attributes 40% of that effect to actor composition. We therefore treat it as suggestive rather than ownership-driven.

I Robustness Check: OLS vs. Ordered Logit

To validate the use of OLS on the ordinal sentiment score $\{-1, 0, +1\}$, we re-estimate every significant DiD specification using a cluster-robust ordered logit model (OrderedModel with logistic link, clustered at the domain level, same grouping variable as the OLS). We test three families of cases: (i) the three micro-level case studies from Section 4 (Sarkozy/Europe 1, Raggi/La Repubblica, Starmer/Euronews); (ii) all eight macro Global/National specifications that reached $p < 0.05$ in the main OLS analysis; and (iii) all thirteen political-bloc specifications that survived Benjamini–Hochberg FDR correction.

Results. Table 17 reports both estimators side by side. Because OLS coefficients are expressed as changes on the $[-1, +1]$ sentiment scale and ordered logit coefficients are on the log-odds scale, their magnitudes are not directly comparable; the relevant comparison is sign and significance. Across all 24 cases, the two models agree on the *sign* of the DiD interaction in **24/24** cases, and on its *significance* ($p < 0.05$) in **24/24** cases. No ordered logit specification failed to converge. This complete agreement provides strong evidence that the ordinal-to-linear mapping does not distort the main conclusions of the paper.

J Fixed-Effect Specifications

To separate compositional change (*which* politicians are covered) from tone change (*how* the same politicians are covered), we re-estimate every significant macro- and bloc-level specification with and without entity fixed effects, using the same domain-clustered standard errors as the main analysis. $\beta_{3,A}$ is the baseline DiD coefficient (identical to Tables 3 and 4); $\beta_{3,B}$ adds entity fixed effects and captures the framing component; *selection* = $\beta_{3,A} - \beta_{3,B}$ is the actor-composition component, and %sel its share of the total effect. A value of $\beta_{3,B}$ close to $\beta_{3,A}$ indicates framing (same actors, different tone); a $\beta_{3,B}$ shrinking toward zero indicates selection.

Adding domain and month fixed effects (column C) leaves 19 of the 21 (Table 19) specifications significant. The two exceptions, Irish Examiner Right and Hürriyet Center, are among the cases already qualified in Section 4. Under the full specification (column D), 15 remain signifi-

cant; the six that do not—Hürriyet Macro Global, LA Times Macro National, La Repubblica Macro Global, Hürriyet Center, La Provence Radical Left and Europe 1 Radical Right—are precisely those identified as selection-driven in Table 18 or as failing the placebo triage in Appendix F. Only La Provence Radical Right is significant under all four specifications and also clean under the placebo, FDR, entity-month and domestic-only criteria.

K Triple Interaction DiD: Within-Outlet Bloc Contrasts

To test whether bloc-level DiD effects differ significantly *within* each outlet, we estimate a pooled OLS regression with a triple interaction term:

Let S_{idt} denote sentiment, T_d the target indicator, P_t the post indicator, and B_i the bloc indicator. We estimate:

$$\begin{aligned} S_{idt} = & \beta_0 + \beta_1 T_d + \beta_2 P_t + \beta_3 B_i \\ & + \beta_4 (T_d \times P_t) + \beta_5 (T_d \times B_i) \\ & + \beta_6 (P_t \times B_i) \\ & + \beta_7 (T_d \times P_t \times B_i) + \epsilon_{idt} \end{aligned} \quad (2)$$

The coefficient β_7 measures whether the DiD effect for bloc b differs from the Left reference bloc within the same outlet-level model, using cluster-robust standard errors at the domain level. Table 20 reports the significant pairwise contrasts.

These contrasts confirm that the heterogeneous bloc-level patterns reported in Section 4.2 reflect genuine within-outlet ideological selectivity rather than effects that are individually significant but mutually indistinguishable. The Irish Examiner result is particularly striking: the gap between its treatment of Radical Right and Center figures is the largest statistically confirmed within-outlet contrast in the corpus.

L Sensitivity Analysis: Bloc-Level DiD with Domestic Politicians Only

The bloc-level analysis in Section 4.2 pools all tracked politicians (global and national lists). To assess whether results are driven by international figures, we re-estimate the bloc-level DiD restricting each outlet to its domestic politicians only.

Table 21 reports the comparison. The core findings are robust: all testable Irish Examiner bloc effects, La Stampa Left, La Provence Radical Right, and La Repubblica (Center, Left, Radical Left) remain significant and directionally consistent, often

Table 17: Robustness check: OLS vs. cluster-robust ordered logit on the DiD interaction term. Coefficients are not comparable in magnitude (different scales). Sign $\checkmark$ means both estimators agree on direction; Sig. $\checkmark$ means both agree on $p < 0.05$. All 24 cases show full agreement on both criteria.

Case	N	β_{OLS}	p_{OLS}	β_{Ord}	p_{Ord}
<i>(a) Micro showcase cases</i>					
Europe 1 — Nicolas Sarkozy	6,579	+0.193	< .001	+0.492	< .001
La Repubblica — Virginia Raggi	4,068	−0.161	< .001	−0.457	< .001
Euronews — Keir Starmer	1,447	+0.308	< .001	+0.936	< .001
<i>(b) Macro Global / National (OLS-significant)</i>					
Hürriyet — Macro Global	51,899	+0.024	.028	+0.088	.048
La Repubblica — Macro Global	102,368	+0.044	.002	+0.130	< .001
Europe 1 — Macro National	156,643	+0.020	.003	+0.065	.002
Irish Examiner — Macro National	36,184	+0.051	< .001	+0.193	< .001
La Provence — Macro National	148,681	+0.020	.005	+0.066	.005
La Stampa — Macro National	86,795	+0.032	< .001	+0.106	< .001
LA Times — Macro National	107,599	+0.034	< .001	+0.110	< .001
La Repubblica — Macro National	159,183	−0.020	.009	−0.061	.014
<i>(c) Political bloc (FDR-significant)</i>					
Irish Examiner — Center	16,420	+0.060	< .001	+0.212	< .001
Irish Examiner — Right	48,935	+0.040	< .001	+0.129	< .001
Irish Examiner — Radical Right	1,532	−0.158	< .001	−0.401	< .001
Hürriyet — Center	5,502	+0.128	< .001	+0.460	< .001
La Stampa — Left	45,031	+0.048	< .001	+0.158	< .001
Irish Examiner — Radical Left	2,678	+0.123	< .001	+0.381	< .001
La Provence — Radical Right	24,717	+0.089	< .001	+0.236	< .001
La Repubblica — Radical Left	6,583	+0.081	< .001	+0.212	< .001
La Repubblica — Right	58,805	+0.056	< .001	+0.166	< .001
La Provence — Radical Left	10,576	+0.056	.002	+0.211	< .001
La Repubblica — Left	93,577	−0.022	.002	−0.066	.003
Europe 1 — Radical Right	23,313	+0.048	.007	+0.131	.010
La Repubblica — Center	61,346	−0.027	.013	−0.082	.010

with stronger coefficients. The national restriction additionally surfaces new significant effects for Europe 1 Right, La Stampa Right, LA Times Radical Left, and La Repubblica Radical Right. However, three FDR-significant results from the primary analysis do not survive the restriction: Europe 1 Radical Right, La Provence Radical Left, and La Repubblica Right lose significance when international politicians are excluded, suggesting these effects were partly driven by coverage of non-domestic figures. These attenuated cases should be interpreted with caution in the main analysis.

M Additional Micro-Level Time-Series Visualizations

This appendix provides additional time-series charts illustrating the evolution of editorial slant for selected political entities at the micro-level (*Micro_subject*). Each chart compares the quarterly average sentiment score of the target media outlet (blue line) against its respective national or European control group (orange line). The vertical dashed line indicates the quarter in which the media acquisition was finalized.

These specific examples highlight various forms of post-acquisition editorial divergence, such as the severe penalization of figures involved in political controversies (e.g., Friedrich Merz, Clemente Mastella), asymmetric local political battles (e.g., Vito Crimi), the fading of a prominent editorial spotlight over time (e.g., Sviatlana Tsikhanouskaya), and the targeting of specific officials during national crises (e.g., Gérald Darmanin, Olivier Véran).

N Composition of the linked inference corpus (ccnews_linked)

This appendix documents the size and structure of the sentence-level corpus used for stance inference and for the difference-in-differences (DiD) analyses reported in the main text.

Unit of observation. Each row in the dataset is a target sentence in which at least one tracked political figure was detected and linked, together with metadata (occurrence_id, role, period, date, match_variant, entities). Approximately 277,000 sentences mention more than one tracked politician. The headline total reported in

Specification	N	#ent.	$\beta_{3,A}$ (DiD)	$\beta_{3,B}$ (framing)	selection	%sel
Europe 1 (Nat.)	156,643	500	+0.0199**	+0.0184**	+0.0015	7%
Hürriyet (Glob.)	51,899	395	+0.0241*	+0.0111	+0.0130	54%
Irish Examiner (Nat.)	36,184	246	+0.0512***	+0.0516***	-0.0004	-1%
La Provence (Nat.)	148,681	527	+0.0205**	+0.0220**	-0.0016	-8%
La Stampa (Nat.)	86,795	425	+0.0317***	+0.0230**	+0.0087	27%
LA Times (Nat.)	107,599	480	+0.0336***	+0.0147	+0.0189	56%
La Repubblica (Glob.)	102,368	819	+0.0441**	+0.0267*	+0.0174	39%
La Repubblica (Nat.)	159,183	459	-0.0203**	-0.0132*	-0.0071	35%
Irish Examiner – Center	16,420	140	+0.0604***	+0.0574***	+0.0031	5%
Irish Examiner – Right	48,935	226	+0.0399***	+0.0546***	-0.0147	-37%
Irish Examiner – Rad. Right	1,532	23	-0.1582***	-0.1294***	-0.0289	18%
Hürriyet – Center	5,502	60	+0.1277***	+0.0963***	+0.0314	25%
La Stampa – Left	45,031	224	+0.0483***	+0.0385***	+0.0098	20%
Irish Examiner – Rad. Left [†]	2,678	17	+0.1229***	+0.1242***	-0.0013	-1%
La Provence – Rad. Right	24,717	75	+0.0893***	+0.0716***	+0.0176	20%
La Repubblica – Rad. Left [†]	6,583	18	+0.0813***	+0.1193***	-0.0380	-47%
La Repubblica – Right	58,805	244	+0.0560***	+0.0395**	+0.0165	30%
La Provence – Rad. Left	10,576	53	+0.0556**	+0.0389*	+0.0167	30%
La Repubblica – Left	93,577	280	-0.0219**	-0.0191**	-0.0028	13%
Europe 1 – Rad. Right	23,313	71	+0.0482**	+0.0289	+0.0193	40%
La Repubblica – Center	61,346	173	-0.0270*	-0.0164	-0.0106	39%

Table 18: Selection vs. framing decomposition. $\beta_{3,A}$ reproduces the main-text DiD coefficient. $\beta_{3,B}$ adds entity fixed effects (framing component); $selection = \beta_{3,A} - \beta_{3,B}$. Domain-clustered SEs. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. [†] fewer than 20 tracked entities: decomposition poorly identified, interpret with caution.

the main text—**1,376,552**—is therefore a number of sentences; the unit of estimation is the entity–sentence observation, of which there are 1,731,681 (Appendix H).

Unique entities per outlet. For each outlet, we count the number of distinct political figures as the number of unique values of the entity field (canonical display name). Across all nine outlets, **988** distinct entity names appear at least once; per-outlet minima and maxima are reported in the table.

Role and period. Each row is labeled target (articles published by the outlet under study) or control (articles from other domains in the matched corpus), and pre_takeover vs. post_takeover relative to the ownership transition. Pooling all outlets: 1,111,426 control vs. 265,126 target rows; 645,927 pre-takeover vs. 730,625 post-takeover rows.

Subject level (subject_level). Each sentence row is assigned one of three subject levels: Global (figures from the cross-national dictionary), National (figures from the outlet’s national list), or Specific (outlet-specific takeover entities). Pooled across all nine outlets, the corpus comprises 706,186 Global, 666,198 National, and 4,168 Specific sentence rows.

O Complete List of Significant DiD Coefficients

The following table presents all statistically significant Difference-in-Differences (DiD) coefficients across the global, national, and micro levels. For the micro-subject level, significance is determined after applying the Benjamini-Hochberg False Discovery Rate (FDR) correction (p -adjusted < 0.05) to account for multiple testing. The *Macro_Specific* level is excluded due to insufficient data volume. Results are sorted alphabetically by media outlet. Table 23: Entity-level (micro) DiD results significant after FDR correction (Benjamini–Hochberg, $p_{adj} < 0.05$): 215 of 551 tested entity–outlet pairs.

Entity	β_3	p_{adj}
<i>Euronews</i> ($n = 53$)		
Maia Sandu	-0.335	6.0×10^{-52}
Friedrich Merz	-0.469	1.3×10^{-25}
Robert Habeck	-0.166	5.0×10^{-23}
Sanna Marin	+0.520	5.0×10^{-22}
Angela Merkel	-0.082	1.4×10^{-11}
Pedro Sánchez	-0.110	6.3×10^{-11}
Ebrahim Raisi	-0.301	7.7×10^{-11}
Keir Starmer	+0.308	3.1×10^{-10}
Didier Reynders	-0.157	9.5×10^{-9}
Alberto Fernández	+0.381	1.2×10^{-8}
Valdis Dombrovskis	-0.222	2.7×10^{-8}
Bashar al-Assad	+0.328	4.3×10^{-8}
Boris Johnson	-0.195	1.6×10^{-7}

Specification	N	(A) baseline	(B) + entity	(C) + domain, month	(D) + all three
<i>Macro level</i>					
Europe 1 — Macro National	156,643	+0.0199*	+0.0184*	+0.0205*	+0.0220*
Hürriyet — Macro Global	51,899	+0.0241	+0.0111	+0.0231*	+0.0058
Irish Examiner — Macro National	36,184	+0.0512*	+0.0516*	+0.0390*	+0.0394*
La Provence — Macro National	148,681	+0.0205*	+0.0220*	+0.0259*	+0.0256*
La Stampa — Macro National	86,795	+0.0317*	+0.0230*	+0.0292*	+0.0226*
LA Times — Macro National	107,599	+0.0336*	+0.0147	+0.0261*	+0.0095
La Repubblica — Macro Global	102,368	+0.0441*	+0.0267*	+0.0345*	+0.0191
La Repubblica — Macro National	159,183	−0.0203*	−0.0132	−0.0231*	−0.0159*
<i>Bloc level</i>					
Irish Examiner — Center	16,420	+0.0604*	+0.0574*	+0.0540*	+0.0544*
Irish Examiner — Right	48,935	+0.0399*	+0.0546*	+0.0217	+0.0319*
Irish Examiner — Rad. Right	1,532	−0.1582*	−0.1294*	−0.1459*	−0.1277*
Hürriyet — Center	5,502	+0.1277*	+0.0963*	+0.0654	+0.0382
La Stampa — Left	45,031	+0.0483*	+0.0385*	+0.0491*	+0.0438*
Irish Examiner — Rad. Left	2,678	+0.1229*	+0.1242*	+0.1574*	+0.1701*
La Provence — Rad. Right	24,717	+0.0893*	+0.0716*	+0.0711*	+0.0682*
La Repubblica — Rad. Left	6,583	+0.0813*	+0.1193*	+0.1050*	+0.1193*
La Repubblica — Right	58,805	+0.0560*	+0.0395*	+0.0446*	+0.0296*
La Provence — Rad. Left	10,576	+0.0556*	+0.0389	+0.0445*	+0.0360
La Repubblica — Left	93,577	−0.0219*	−0.0191*	−0.0234*	−0.0204*
Europe 1 — Rad. Right	23,313	+0.0482*	+0.0289	+0.0525*	+0.0427
La Repubblica — Center	61,346	−0.0270*	−0.0164	−0.0276*	−0.0179*

Table 19: DiD interaction β_3 under alternative fixed-effect specifications, for the 21 specifications significant in the main analysis. (A) reproduces Eq. (1); (B) adds entity fixed effects, reporting the same point estimates as Table 18 with significance assessed under the more conservative reference below; (C) adds domain and month fixed effects, the standard two-way specification, in which $Target_d$ and $Post_t$ are collinear with the fixed effects and drop out; (D) adds all three. Standard errors are clustered on domain throughout. * $p < 0.05$ under a t reference with $G - 1$ degrees of freedom ($G = 11$; $G = 10$ for Hürriyet and the Irish Examiner, Appendix G), which is more conservative than the normal approximation used in Tables 3, 4 and 18. Three entries starred under that approximation are unstarred here: Hürriyet Macro Global in column (A), and La Repubblica Macro National and La Provence Radical Left in column (B).

Table 20: Significant pairwise bloc contrasts from the triple interaction regression (Target×Post×Alignment). Δ = difference in DiD coefficients between Bloc A and Bloc B within the same pooled model.

Outlet	Contrast (A vs B)	Δ	p
Irish Examiner	Rad. Right vs Center	−0.219	< .001
Irish Examiner	Rad. Right vs Left	−0.146	.009
Irish Examiner	Rad. Left vs Left	+0.136	.002
La Repubblica	Rad. Left vs Left	+0.103	< .001
La Repubblica	Right vs Left	+0.078	< .001
La Provence	Rad. Right vs Left	+0.097	< .001
La Provence	Rad. Left vs Left	+0.064	< .001
Hürriyet	Center vs Left	+0.141	< .001
Europe 1	Rad. Right vs Left	+0.042	.003

Entity	β_3	p_{adj}
Rafael Grossi	−0.137	8.6×10^{-7}
Kristalina Georgieva	−0.258	1.4×10^{-6}
Klaus Iohannis	−0.356	4.5×10^{-5}
Frank-Walter Steinmeier	−0.094	1.2×10^{-4}
Guillermo Lasso	+0.223	1.3×10^{-4}
Janet Yellen	−0.183	1.6×10^{-4}

Entity	β_3	p_{adj}
Donald Tusk	+0.209	4.5×10^{-4}
Sergio Mattarella	+0.176	9.5×10^{-4}
Jacinda Ardern	+0.211	0.001
Nancy Pelosi	−0.127	0.001
Markus Söder	+0.249	0.002
Sviatlana Tsikhanouskaya	−0.298	0.003
Recep Tayyip Erdoğan	+0.068	0.003
Gérald Darmanin	−0.130	0.005
Greta Thunberg	+0.040	0.005
Dmytro Kuleba	−0.073	0.006
Tony Blair	+0.229	0.006
Josep Borrell	−0.050	0.007
Alexander Lukashenko	−0.152	0.007
Linda Thomas-Greenfield	−0.103	0.009
Mario Draghi	+0.107	0.009
Bill Gates	+0.133	0.009
Frans Timmermans	−0.139	0.010
Nicolas Sarkozy	+0.227	0.010
Theresa May	+0.173	0.010
Věra Jourová	+0.082	0.014
Alexander De Croo	+0.103	0.015
Dmitry Peskov	+0.036	0.017
Maroš Šefčovič	−0.073	0.021




Figure 3: Supplementary case studies of significant micro-level editorial shifts following media acquisitions (Sviatlana Tsikhanouskaya, Friedrich Merz, Clemente Mastella, Vito Crimi, Gérald Darmanin, and Olivier Vérant). The charts compare the acquired outlet with its respective control group.

Entity	β_3	p_{adj}
Cyril Ramaphosa	-0.092	0.023
Ylva Johansson	+0.093	0.024
Kyriakos Mitsotakis	-0.069	0.025
Nayib Bukele	+0.263	0.025
David Cameron	+0.222	0.025
Milorad Dodik	-0.177	0.026
Gabriel Attal	-0.139	0.027
Bruno Le Maire	-0.070	0.031
Justin Trudeau	+0.091	0.034
Charles Michel	-0.047	0.048
Kamala Harris	-0.121	0.048
<i>La Repubblica (n = 36)</i>		
Luciana Lamorgese	-0.174	3.4×10^{-18}
Marine Le Pen	+0.267	3.4×10^{-18}
Vito Crimi	+0.388	1.2×10^{-16}
Matteo Lepore	-0.363	1.7×10^{-13}
Dario Franceschini	-0.190	3.8×10^{-13}
Xi Jinping	+0.239	3.7×10^{-10}
Donald Trump	+0.097	3.9×10^{-10}
Nicola Fratoianni	+0.160	2.1×10^{-8}
Clemente Mastella	-0.355	5.5×10^{-7}
Mariastella Gelmini	+0.212	1.9×10^{-6}
Monica Cirinnà	-0.304	1.1×10^{-5}
Antonio Tajani	-0.109	1.5×10^{-5}
Danilo Toninelli	-0.221	3.4×10^{-5}
Bernie Sanders	-0.282	3.7×10^{-5}
Virginia Raggi	-0.161	4.2×10^{-5}
Andrea Marcucci	-0.144	1.2×10^{-4}

Entity	β_3	p_{adj}
Sergio Mattarella	-0.062	1.4×10^{-4}
Maurizio Gasparri	+0.125	3.2×10^{-4}
Emma Bonino	+0.216	5.7×10^{-4}
Beppe Grillo	+0.062	6.1×10^{-4}
Nicola Zingaretti	-0.071	0.002
Alessandro Di Battista	+0.182	0.002
Stefano Patuanelli	-0.067	0.002
Matteo Orfini	+0.205	0.003
Graziano Delrio	-0.168	0.004
Mike Pompeo	+0.190	0.005
Eugenio Giani	+0.160	0.006
Maria Elena Boschi	-0.120	0.007
Vladimir Putin	+0.073	0.008
Michele Emiliano	-0.109	0.010
Giorgia Meloni	-0.048	0.012
Massimo D'Alema	+0.136	0.012
Mike Pence	-0.144	0.015
Jair Bolsonaro	-0.131	0.015
Giuseppe Sala	+0.077	0.026
<i>La Provence (n = 26)</i>		
Vladimir Poutine	-0.177	2.1×10^{-13}
Olivier Vérant	-0.105	2.1×10^{-12}
Julien Bayou	-0.379	1.5×10^{-11}
Benyamin Netanyahu	-0.203	1.0×10^{-10}
Viktor Orbán	-0.277	4.7×10^{-10}
Agnès Pannier-Runacher	+0.270	6.6×10^{-9}
Éric Ciotti	+0.117	8.6×10^{-7}
Donald Trump	-0.057	9.5×10^{-6}

Table 21: Bloc-level DiD: all politicians vs. domestic only. Sig. = significant at $p\text{-adj} < 0.05$ after Benjamini–Hochberg FDR correction.

Outlet	Bloc	β_{all}	Sig.	$\beta_{\text{nat.}}$	Sig.	Dir.
Euronews	Rad. Right	+0.049	×	+0.018	×	✓
Europe 1	Center	+0.017	×	+0.013	×	✓
Europe 1	Rad. Right	+0.048	✓	+0.034	×	✓
Europe 1	Right	+0.000	×	+0.031	✓	✓
Hürriyet	Center	+0.128	✓	n/a	—	—
Irish Exam.	Center	+0.060	✓	+0.039	✓	✓
Irish Exam.	Rad. Left	+0.123	✓	+0.067	✓	✓
Irish Exam.	Rad. Right	−0.158	✓	n/a	—	—
Irish Exam.	Right	+0.040	✓	+0.065	✓	✓
La Provence	Rad. Left	+0.056	✓	+0.070	×	✓
La Provence	Rad. Right	+0.089	✓	+0.079	✓	✓
La Stampa	Left	+0.048	✓	+0.057	✓	✓
La Stampa	Right	+0.010	×	+0.033	✓	✓
LA Times	Rad. Left	+0.066	×	+0.103	✓	✓
La Rep.	Center	−0.027	✓	−0.034	✓	✓
La Rep.	Left	−0.022	✓	−0.032	✓	✓
La Rep.	Rad. Left	+0.081	✓	+0.233	✓	✓
La Rep.	Rad. Right	−0.012	×	−0.024	✓	✓
La Rep.	Right	+0.056	✓	+0.014	×	✓

Entity	β_3	p_{adj}
Marine Le Pen	+0.116	1.2×10^{-5}
Jair Bolsonaro	+0.091	4.8×10^{-5}
Bruno Le Maire	+0.102	4.8×10^{-5}
Josep Borrell	−0.104	7.9×10^{-4}
Clément Beaune	+0.077	0.002
Gérald Darmanin	−0.048	0.002
Dmitri Peskov	−0.078	0.002
Alexei Navalny	+0.083	0.004
Marion Maréchal	+0.076	0.006
Kim Jong-un	−0.232	0.007
Olivier Faure	+0.052	0.007
Manuel Bompard	+0.202	0.007
Édouard Philippe	−0.118	0.010
Valérie Pécresse	−0.117	0.011
Charles Michel	−0.165	0.022
Volodymyr Zelensky	+0.068	0.023
Élisabeth Borne	+0.050	0.032
Olaf Scholz	+0.074	0.044

Europe 1 ($n = 22$)

Nicolas Sarkozy	+0.193	8.5×10^{-14}
Alexei Navalny	−0.145	2.6×10^{-10}
Gabriel Attal	+0.059	3.2×10^{-8}
Jair Bolsonaro	+0.205	3.8×10^{-7}
Donald Trump	−0.068	1.4×10^{-6}
Elon Musk	+0.215	7.0×10^{-5}
Sandrine Rousseau	+0.155	1.6×10^{-4}
Aurélien Pradié	+0.180	2.7×10^{-4}
Vladimir Poutine	−0.094	3.6×10^{-4}
António Guterres	+0.158	8.4×10^{-4}
Stanislas Guerini	+0.203	9.4×10^{-4}
Xavier Bertrand	+0.124	0.002
Volodymyr Zelensky	+0.137	0.002
Xi Jinping	+0.110	0.003
Agnès Pannier-Runacher	+0.155	0.007
Éric Zemmour	+0.070	0.007

Table 22: Composition of ccnews_linked by target outlet (nine media clusters). Occurrences = sentence rows; entities = unique entity values per outlet.

Outlet (cluster_id)	Occurrences	Unique entities
euronews_eu	222,650	452
latimes_us	210,686	349
repubblica_it	207,499	433
europa1_fr	197,567	139
laprovence_fr	195,217	209
lastampa_it	109,130	216
independent_ie	92,568	295
hurriyet_com_tr	71,221	164
irishexaminer_ie	70,014	242
Total	1,376,552	988 (global)

Notes. Min./max. occurrences per outlet: 70,014 (irishexaminer_ie) and 222,650 (euronews_eu). Min./max. unique entities per outlet: 139 (europa1_fr) and 452 (euronews_eu).

Entity	β_3	p_{adj}
Gérard Larcher	−0.124	0.009
Viktor Orbán	−0.165	0.018
Olivier Véran	+0.036	0.023
David Lisnard	+0.192	0.028
Valérie Pécresse	+0.080	0.038
Aurore Bergé	+0.112	0.044

LA Times ($n = 22$)

Recep Tayyip Erdoğan	+0.259	1.7×10^{-18}
Bill Gates	+0.274	2.1×10^{-12}
Rand Paul	−0.189	9.7×10^{-10}
Neil Gorsuch	−0.325	3.3×10^{-8}
Pope Francis	−0.145	1.8×10^{-7}
Andrew Cuomo	−0.139	9.5×10^{-6}
Lindsey Graham	−0.100	3.1×10^{-5}
Angela Merkel	−0.082	7.6×10^{-5}
Bernie Sanders	+0.097	3.1×10^{-4}
Emmanuel Macron	+0.133	8.4×10^{-4}
Mick Mulvaney	−0.103	0.001
John Kerry	+0.290	0.001
Merrick Garland	−0.270	0.002
Rick Perry	+0.199	0.005
Gavin Newsom	+0.308	0.005
Ted Cruz	+0.056	0.006
Xavier Becerra	−0.115	0.009
Vladimir Putin	−0.056	0.015
Lisa Murkowski	−0.070	0.023
Antonin Scalia	−0.147	0.028
Barack Obama	−0.031	0.031
Mitt Romney	+0.123	0.037

Independent ($n = 20$)

Michael Healy-Rae	−0.270	5.1×10^{-69}
Benjamin Netanyahu	−0.087	2.9×10^{-14}
Nancy Pelosi	−0.232	1.2×10^{-13}
Emmanuel Macron	+0.081	4.4×10^{-7}
Pedro Sánchez	+0.188	4.8×10^{-6}
Arlene Foster	−0.098	8.5×10^{-5}
Boris Johnson	+0.126	8.8×10^{-5}
Vladimir Putin	+0.089	1.2×10^{-4}

Entity	β_3	p_{adj}
Joe Biden	+0.311	4.4×10^{-4}
Mike Pompeo	-0.105	0.002
Richard Bruton	+0.136	0.002
Danny Healy-Rae	+0.247	0.006
Dominic Raab	+0.196	0.008
Brendan Howlin	+0.139	0.009
Jair Bolsonaro	-0.161	0.012
Donald Trump	-0.074	0.015
Leo Varadkar	+0.055	0.021
Mitch McConnell	-0.111	0.024
Denis Naughten	+0.086	0.048
Stephen Donnelly	+0.136	0.049

Irish Examiner ($n = 16$)

Leo Varadkar	+0.082	1.1×10^{-20}
Donald Tusk	+0.234	2.7×10^{-10}
Donald Trump	+0.111	1.8×10^{-8}
Paul Kehoe	+0.279	1.4×10^{-6}
Mattie McGrath	-0.150	1.4×10^{-6}
Richard Bruton	+0.096	1.7×10^{-6}
Mike Pence	-0.125	1.3×10^{-5}
Vladimir Putin	-0.152	7.2×10^{-5}
Paschal Donohoe	+0.073	9.1×10^{-5}
Emmanuel Macron	+0.147	0.002
Shane Ross	+0.063	0.005
Michael Creed	+0.077	0.008
Micheál Martin	+0.069	0.014
Enda Kenny	+0.057	0.021
Pearse Doherty	-0.109	0.036
Gerry Adams	-0.118	0.039

La Stampa ($n = 16$)

Greta Thunberg	-0.315	6.9×10^{-28}
Angela Merkel	+0.161	2.5×10^{-12}
Giorgia Meloni	-0.097	2.5×10^{-12}
Alessandro Di Battista	+0.297	1.2×10^{-9}
Beppe Grillo	+0.102	3.2×10^{-8}
Joe Biden	+0.341	1.7×10^{-5}
Jair Bolsonaro	-0.293	4.5×10^{-5}
Ursula von der Leyen	+0.169	8.4×10^{-4}
Viktor Orbán	+0.169	0.002
Silvio Berlusconi	+0.060	0.004
Virginia Raggi	+0.123	0.007
Enrico Costa	-0.244	0.008
Carlo Calenda	-0.140	0.010
Giovanni Toti	+0.077	0.022
Donald Trump	-0.065	0.023

Hürriyet ($n = 4$)

Ahmet Davutoğlu	-0.257	1.5×10^{-5}
Emmanuel Macron	+0.174	0.001
António Guterres	+0.129	0.003
Sergey Lavrov	+0.045	0.040

P Reproducibility and Software Configuration (ARR Checklist C4)

To satisfy item C4 of the ARR Responsible NLP Research checklist regarding the use of existing software packages, this section documents the exact implementations, model checkpoints, and parameter configurations deployed across our pipeline.

P.1 Named Entity Recognition (NER)

For the extraction of political figures mentioned across the media corpus, we utilized the **GLINER** framework, which allows for open-palette named entity recognition.

- **Implementation / Package:** gliner Python library (v0.1.12 or latest stable release used during inference).

- **Model** **Checkpoint:** urchade/gliner_multi-v2.1 available on the Hugging Face Hub.

- **Parameters:** The target entity label set was strictly restricted to ["politician", "political figure"]. The inference confidence threshold was set to a default value of 0.5 to balance precision and recall.

P.2 Entity Linking and Disambiguation

To map extracted mentions to unique identities and resolve variations or ambiguous names, we deployed Facebook’s **mGENRE** (Multilingual Generative Entity Linker).

- **Implementation / Package:** Integrated via the Hugging Face transformers framework alongside the official genre repository.

- **Model** **Checkpoint:** facebook/mgenre-paper checkpoint, using a standardized Wikidata snapshot for entity mapping.

- **Parameters:** Autoregressive entity retrieval was executed using beam search with a beam size of 5.

P.3 Target-Oriented Sentiment Classification (TSC)

The micro-level sentiment evaluation across the core media corpus was performed using our top-performing architecture determined in the preliminary benchmarks (Appendix D). The final deployment configuration utilizes the guided Qwen 3.5 framework.

- **Implementation / Package:** Inferences were executed using the vLLM high-throughput engine (v0.4.2) for optimized tensor parallelism.

- **Model** **Checkpoint** / **Quantization:** Qwen/Qwen3.5-72B, deployed using the

GGUF 4-bit quantization scheme (Q4_K_M) to comply with infrastructure and GPU memory constraints.

- **Sampling Parameters:** To guarantee deterministic results and high experimental reproducibility for class extraction, decoding was executed with a temperature set strictly to 0.0, a top-p parameter of 1.0, and an output length constraint of `max_new_tokens=8`.
- **Prompt Template:** The model was prompted using the GUID configuration, which embeds structured instruction boundaries distinguishing the three target classes (NEGATIVE, NEUTRAL, POSITIVE), as detailed in Appendix C.

P.4 Artifact Licenses and Intended Use (ARR Checklist B2 & B3)

To satisfy items B2 and B3 of the ARR Responsible NLP Research checklist, we declare the licenses and intended use of all software artifacts and datasets utilized in this study:

- **Datasets:** The MAD-TSC dataset (Dufraisse et al., 2023) is used strictly in compliance with its original academic distribution terms.
- **Models and Software:** The open-source frameworks employed—including GLiNER (urchade/gliner_multi-v2.1), mGENRE (facebook/mgenre-paper), and Qwen (Qwen3.5-27B)—are distributed under permissive open licenses (such as Apache 2.0 or the Qwen Research License).
- **Intended Use Consistency:** In accordance with requirement B3, all artifacts are used strictly for computational social science research. No clinical, commercial, or non-research deployments were conducted. All derived entity-level sentiment metrics and parallel trend analyses are restricted to the scope of this academic evaluation, fully satisfying the original artifact access conditions.

Q Use of AI Assistants in Research and Writing (ARR Checklist E1)

In accordance with the ARR Responsible NLP Research guidelines (Item E1), we disclose the integration and scope of AI assistants utilized throughout this research project. AI tools were deployed across five specific boundaries:

1. **Core Experimental Pipeline (Object of Study):** Large Language Models (specifically Qwen3.5-27B) were used strictly as the predictive architectures for target-oriented sentiment classification, as documented in Section 3 and Appendix P. They were not used to synthesize or fabricate training or evaluation data.
2. **Coding and Technical Implementation Assistance:** AI-powered development environments (including Cursor and GitHub Copilot) were utilized to accelerate data pre-processing pipelines, streamline the integration of the vLLM inference engine, and optimize automated plotting scripts for Python’s `matplotlib`.
3. **Writing and Structural Refinement:** Conversational AI systems were employed during the drafting phase exclusively for grammatical correction, stylistic polishing, vocabulary enrichment, and the generation of standardized academic $\text{\LaTeX}$ templates for multi-column tables and complex figures.
4. **Data construction:** Gemini 3.1 Pro was used to generate the initial political entity lists (Appendix B). All entries were verified against Wikipedia and mapped to Wikidata party affiliations; the resulting selection bias is discussed in the Limitations.
5. **Human validation:** A commercial machine translation service was used to translate into English the Turkish sentences of the in-corpus validation sample, none of the annotators reading Turkish (Appendix E). The classifier itself processed the original Turkish text.

Every conceptual contribution, literature synthesis, causal modeling formulation (DiD), and final analytical interpretation was conceived and verified entirely by the human authors. The assistants functioned strictly as productivity-enhancing tools, ensuring no breach of academic integrity.