Analyzing the Effect of Media Takeovers on their Political Positioning

Anonymous ACL submission

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 — combined with in-time placebo tests and within-outlet triple-interaction contrasts isolates these patterns from contemporaneous national media trends. 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 ownership changes reshape the distribution of tone across politically salient actors rather than producing 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 curate diverse lists of politicians, including a 1,297-strong global list and national-level lists comprising at least 200 country-specific politicians and use them to support the analysis.

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).

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 TSC + DiD audit pipeline combining masked zero-shot sentiment classification with clustered DiD inference, benchmarked on MAD-TSC (Dufraisse et al., 2023) for English, French and Italian and applied zero-shot to Turkish; (ii) a multi-country empirical study of 9 takeovers across 4 languages and varied acquirer profiles; and (iii) the substantive finding that ownership-driven editorial change operates 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 Bolloré 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 dataset[1], 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. These specific outlets are selected because they are sufficiently represented in the CC-News corpus and offer a broad diversity in terms of geographical coverage and takeover profiles, ranging 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), and diversity (collectively spanning the political spectrum). Ten outlets per country is the minimum satisfying

[1] https://huggingface.co/datasets/stanford-oval/ccnews

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). The choice of 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 (Zaratiana 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 reservoir sampling), 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 analyzed 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. The detailed dataset composition is provided in Appendix J.

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 cor-

relate 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) substantially outperform Small Language Models (SLMs) (Dufraisse et al., 2023) by up to 6.5 macro-F1 points, justifying their use despite higher computational overhead. Consequently, we deployed **Qwen 3.5 27B** using 4-bit quantization (GGUF).

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).

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 with only 11 domains per national market (one treated outlet and ten controls), conventional cluster-robust standard errors may be downward biased under the few-cluster regime (Cameron and Miller, 2015); significance results should therefore be interpreted with appropriate caution. The ordered-logit cross-check in Appendix F validates sign and significance under an ordinal outcome but does not address the few-cluster bias, which we leave to future work (wild-cluster bootstrap, multiway clustering). 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 F).

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

$$Sentiment_{idt} = \beta_0 + \beta_1 Target_d + \beta_2 Post_t + \beta_3 (Target_d \times Post_t) + \epsilon_{idt} \quad (1)$$

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 a variable equal to 1 if the domain belongs to the acquired media, and 0 if it belongs to the control group; $Post_t$ 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.

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. Furthermore, to mitigate the multiple testing problem

arising from estimating thousands of simultaneous treatment effects across 551 individual political figures, 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, while micro-subject updates are strictly filtered using the adjusted threshold (p-adjusted < 0.05). 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 E.

3.6 Temporal Smoothing and Political Bloc Analysis

Bloc-level analysis examines pre/post stance differences by political alignment, pooling all tracked politicians. Because the global list is dominated by Right and Left figures (Table 2), we report a domestic-only sensitivity analysis (Appendix H) as the primary robustness benchmark.

4 Analysis

We analyze post-takeover changes at three complementary levels: macro (all tracked figures), political alignment (ideological blocs), and micro (individual actors). Ownership change is not primarily visible as a homogeneous outlet-wide shift; rather, the most interpretable effects emerge once coverage is decomposed by ideological family and individual actor — a pattern that is itself one of the paper's substantive findings.

4.1 Macro-Level Shifts: Global and National Political Coverage

Macro-level results serve primarily as a diagnostic baseline: their scarcity of significant effects is itself evidence that ownership operates through selective rather than uniform repositioning, as the two following subsections demonstrate. 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. This distinction separates general changes in the outlet's political tone from shifts in the national political space where editorial repositioning is most likely to be salient.

Table 3 shows two[2] and six statistically significant post-acquisition divergences for global and national politicians. For example, *La Repubblica* exhibits a significant *Macro_Global* shift (+0.044, $p < 0.01$), while *Europe 1* (+0.020, $p < 0.01$) and *La Stampa* (+0.032, $p < 0.001$) show significant *Macro_National* shifts. 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 confirm a clean baseline for half of the significant outlets (including *La Repubblica* and the *Irish Examiner*) and reveal substantial trend reversals for the remaining media, strongly validating the presence of non-random, ownership-driven shifts (see Appendix E for full macro estimates).

Outlet	**Macro Global**		**Macro National**	
	DiD	p	**DiD**	p
La Repubblica	+0.044	<.01	−0.020	<.01
Europe 1	+0.005	0.495	+0.020	<.01
La Provence	−0.003	0.454	+0.020	<.01
La Stampa	+0.017	0.140	+0.032	<.001
LA Times	+0.012	0.168	+0.034	<.001
Irish Examiner	+0.019	0.083	+0.051	<.001
Hürriyet	+0.024 †	<.05	+0.006	0.880
Euronews	−0.010	0.232	+0.001	0.943
Independent	−0.005	0.651	+0.016	0.357

Table 3: Macro-level DiD shifts by media outlet. †Fails placebo robustness (App. E).

4.2 Political Alignment Effects

We next aggregate politicians by political alignment. This level is more informative than a single global average because ownership influence may operate through selective shifts within ideological families rather than through a uniform move-

[2]The Hürriyet Macro_Global result fails the placebo robustness check (placebo +0.029, p=0.159, comparable in magnitude to the true effect of +0.024); see Appendix E.

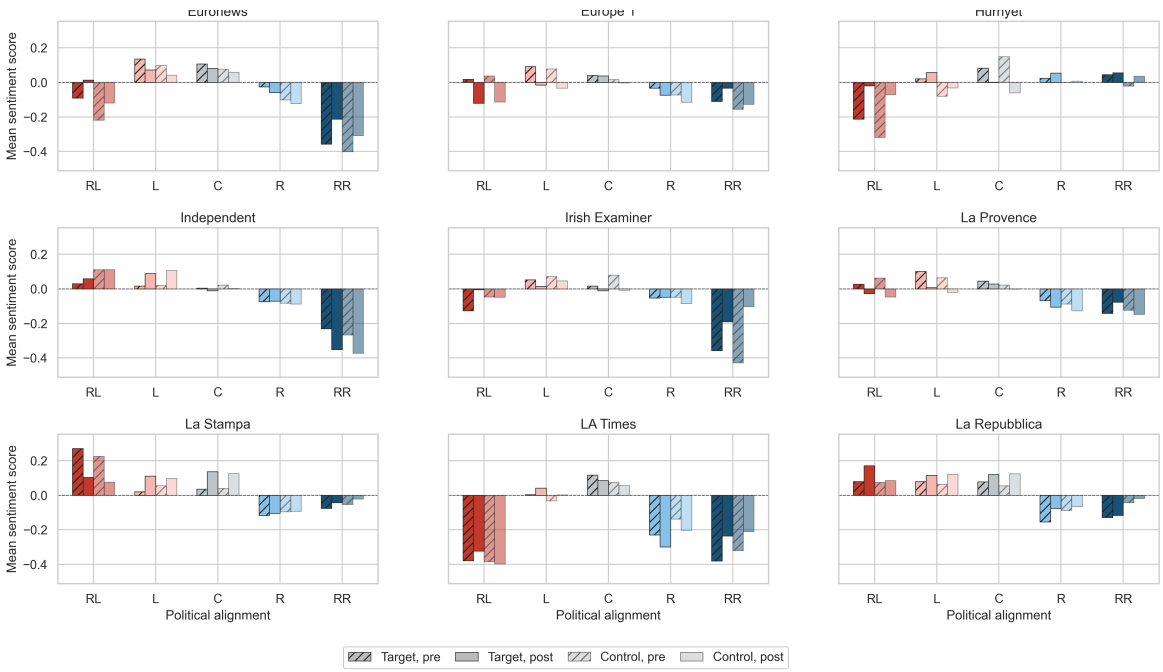

Figure 1: Editorial shifts across the corpus. The comparison between Target (solid) and Control (hatched) groups reveals that significant editorial divergences (Table 4) are often masked by overarching national trends.

ment across all actors. Since domestic coverage is where ownership-driven repositioning is most theoretically motivated, we treat the domestic-only restriction (Appendix H) as the primary robustness benchmark; results below hold or strengthen under that restriction. 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.

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 concentrated in a single ideological family — a pattern that the placebo tests (Appendix E) 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 (11 of 13), 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, while Hürriyet's positive Center shift captures a buffering effect against a sharper control group decline, though the small Center share in the Turkish list (0.9%, Table 2) warrants caution.

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-only sensitivity analysis (Appendix H). 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. This distinction is central to the paper's contribution: ownership matters, but its effect is not generic. These patterns are consistent with

Target Media	Political Bloc	DiD Coeff.	P-Value
Euronews (EU)	Radical Right	+0.049	< 0.05
Europe 1 (FR)	Center	+0.017	< 0.05
	Radical Right	+0.048	< 0.01
Hürriyet (TR)	Center	+0.127	< 0.001
Irish Examiner (IE)	Center	+0.060	< 0.001
	Radical Left	+0.122	< 0.001
	Radical Right	−0.158	< 0.001
	Right	+0.039	< 0.001
La Provence (FR)	Radical Left	+0.055	< 0.01
	Radical Right	+0.089	< 0.001
La Stampa (IT)	Left	+0.048	< 0.001
La Repubblica (IT)	Center	−0.026	< 0.05
	Left	−0.021	< 0.01
	Radical Left	+0.081	< 0.001
	Right	+0.056	< 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.

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. Across the 15 significant bloc-level pairs, restricting attention to the 13 Center and Right bloc coefficients, 11 are positive, while at the micro-level, entity-by-entity sign decomposition shows that significant shifts for Radical Left politicians are predominantly negative — a pattern not directly visible in the bloc-level aggregates of Table 4, where small positive averages can coexist with a majority of individually negative entity-level changes. 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 ($\Delta = -0.219$, $p < 10^{-29}$); for *La Repubblica*, Right and Radical Left blocs diverge significantly from the Left baseline ($p < 10^{-7}$ in both cases); and for *Europe 1*, the Radical Right shift is statistically distinct from other blocs ($p = 0.003$). Full triple interaction results are reported in Appendix G. This structural heterogeneity is further supported by strict pre-treatment placebo checks across the 15 pairs: key ideological shifts—such as the softening of *Europe 1* and *La Provence* toward the Radical Right—show no significant pre-existing trend in the placebo period, consistent with these changes emerging after rather than before the takeover, while other outlets display more complex ongoing transitional dynamics detailed in Appendix E.

4.3 Politician-level Effects

Globally, 215 out of 551 tested political figures (39.0%) exhibit statistically significant DiD coefficients after Benjamini-Hochberg FDR correction (p-adjusted < 0.05; full list in Appendix K), 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%), which suggests either a more concentrated set of targeted figures, stronger background volatility in the Turkish control baseline, or lower statistical power due to a smaller number of tracked entities. The *LA Times* (34.9%), the sole US outlet, shows a significant national-level shift (+0.034) but entity-level changes concentrated on judicial and gubernatorial figures (e.g., positive shifts for Gavin Newsom and John Kerry, negative for Neil Gorsuch and Merrick Garland), consistent with a California-focused editorial reorientation under its new owner rather than a partisan realignment — a pattern distinct from all European cases. The *Independent* (33.3%) shows no significant macro shift, suggesting the Mediahuis acquisition produced actor-specific rather than outlet-wide repositioning. Among entities passing FDR ($p_{\text{adj}} < 0.05$) with $N \geq 1{,}400$ sentences, we select the highest-|DiD| case per outlet, yielding Sarkozy, Raggi, and Starmer (Figure 2; additional cases in Appendix I); placebo diagnostics are reported in Appendix E.3 (Table 11).

Nicolas Sarkozy (*Europe 1*). The positive shift for Nicolas Sarkozy illustrates targeted ideological alignment. While the national control group

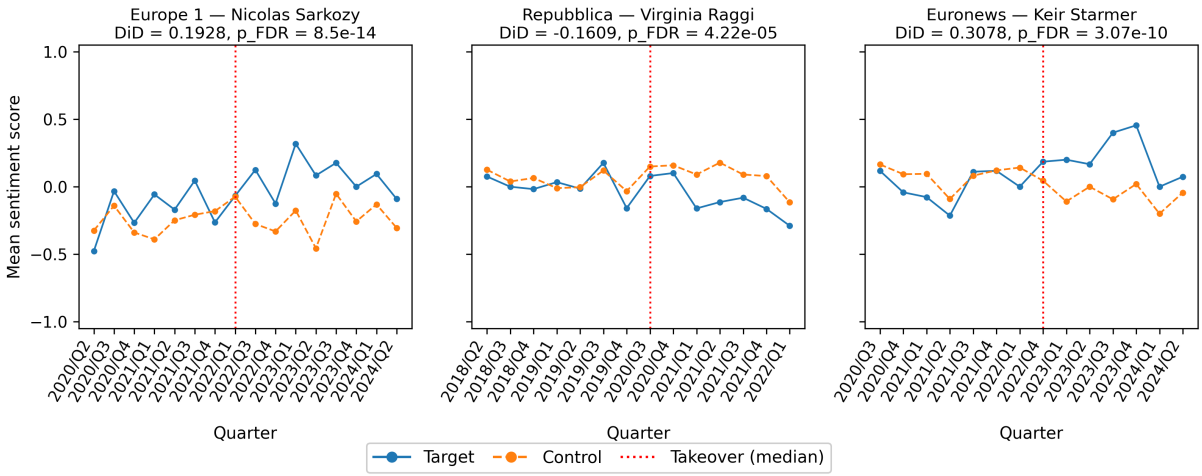

Figure 2: Temporal evolution of editorial slant for selected case studies. The charts compare the acquired outlet with its respective control group. The vertical dashed line marks the acquisition date.

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.

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 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.

Micro-level placebo tests (Table 11, Appendix E) confirm that shifts for Starmer and Raggi occur exclusively after the formal acquisitions, while Sarkozy's trajectory shows a clean post-takeover trend reversal. 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. Using target-oriented sentiment classification and difference-in-differences estimation, we compared acquired outlets with country-matched controls before and after ownership changes. The results show that takeover effects are not primarily expressed as uniform shifts in average outlet-level tone. Macro-level coefficients are often modest, while political-alignment and politician-level analyses reveal more substantial and interpretable changes.

This pattern indicates that ownership-related repositioning is selective and context-dependent. It may affect particular ideological blocs, domestic political arenas, or individual figures without producing a simple global change in slant. The findings support the value of entity-centered NLP for media analysis and suggest that future work on ownership and editorial change should combine aggregate estimates with finer-grained political and contextual interpretation.

Limitations

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

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. 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. However, because each national market contains a limited number of domains (one treated and ten controls), conventional cluster-robust standard errors may be downward biased. While this sentence-level approach maximizes data density, it introduces potential intra-article and within-entity dependence that simple clustering might understate. Future work should address this by aggregating to entity-month averages, deploying multiway clustering, and adopting a full staggered-adoption event-study design (Callaway and Sant'Anna, 2021).

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 under-representation of regional or fringe political actors.

Cross-lingual calibration. While our sentiment analysis pipeline has been validated for the major

languages in our sample, we acknowledge the lack of a human-annotated validation set specifically for Turkish (*Hürriyet*). Consequently, performance metrics for this language are inferred from cross-lingual zero-shot capabilities rather than direct domestic ground truth. 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 comparison 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.

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 attention 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.

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, Jisun An, Haewoon Kwak, Reem Yoosuf, Mohamed Ali, Marie-Francine Moens, 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 2103–2116.

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. Technical report, Sciences Po, Department of Economics.

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. Boydstun, 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 LLMs. *Preprint*, arXiv:2601.12374.

Robert M. Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4):51–58.

Lisa Fan, Marshall White, Eva Sharma, Ruisi 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, Yuhan 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, Junting 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 any significant ownership changes or takeovers during the respective 48-month observation windows.

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

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*.

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. To maintain anonymity during the double-blind review process, the complete dictionary of the 551 tracked political figures, along with their associated parties and assigned 5-point ideological labels, has been de-identified. The full, disaggregated dataset will be made publicly available in a public repository upon publication.

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 substantially outperform SLMs. Specifically, the best LLM configuration (**Qwen 3.5 27B GUID Q4KM**) yields improvements of 2.9, 6.5, and 6.4 macro-F1 points for English, Italian, and French, respectively. This performance gap is particularly pronounced for non-English languages, where pre-trained SLMs are often less effective (Dufraisse et al., 2023). While SLMs offer computational efficiency, we prioritized these significant performance gains and the ease of cross-lingual deployment across the four languages in our corpus. Consequently, we deployed the Qwen 3.5 27B model for the primary media outlet analysis.

Model	EN	IT	FR
RoBERTa SPC	72.3	65.8	70.8
RoBERTa SPC TRAN	72.3	65.8	70.8
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

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 showed only negligible variations: 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. This confirms that anonymization incurs at worst a very slight performance penalty, and occasionally 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 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 provides strong evidence that the parallel trends assumption holds and that the treated and control outlets were following similar trajectories prior to the acquisition. When a placebo coefficient is significant, comparing its direction (sign) to the true effect allows us to determine whether the acquisition accelerated an existing trend or diverged from it.

E.1 Macro-Level Placebo Estimates

Table 9 reports the placebo results exclusively for the 8 macro-level models that were statistically significant in our main analysis.

The results confirm the robustness of several key findings. Half of the models (4 out of 8),

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.)	$+0.020^{**}$	$-0.019^{\dagger}$	0.007
La Provence (Nat.)	$+0.020^{**}$	$-0.024^{\dagger}$	<0.001
LA Times (Nat.)	$+0.034^{***}$	$-0.017^{\dagger}$	<0.001
La Stampa (Nat.)	$+0.032^{***}$	$+0.032^{\dagger}$	0.001

Table 9: 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 ($p < 0.05$). The top four models validate a clean treatment effect, while the bottom four reveal pre-takeover trend reversals.

including both levels for *La Repubblica* and the national level for the *Irish Examiner*, show completely non-significant placebo coefficients, validating a clean treatment effect. Furthermore, for *Europe 1*, *La Provence*, and the *LA Times*, the placebo coefficients are statistically significant but *negative*. Since their true post-takeover DiD coefficients are *positive*, this demonstrates that the media acquisitions did not merely follow a pre-existing editorial drift, but actively reversed a prior negative trend toward domestic political figures. 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.

E.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 10).

The bloc-level analysis reveals that 6 of the 14 testable pairs display no significant pre-trends. Most notably, the ideological realignments observed in the French media landscape—specifically the softening toward the Radical Right by both *Europe 1* and *La Provence* — are validated as remarkably clean, immediate shocks with placebo coefficients near zero ($p > 0.10$). Similarly, *La Repubblica*'s shift toward the Right and *Euronews*'s shift toward the Radical Right withstand the placebo check. Where placebo effects are significant (e.g., *Irish Examiner*), they often suggest gradual anticipation dynamics or broader editorial transitions rather than abrupt post-takeover shocks.

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)	$+0.055^{**}$	$+0.020$	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)	$+0.127^{***}$	$-0.339^{\dagger}$	<0.001
Irish Examiner (Center)	$+0.060^{***}$	$+0.050^{\dagger}$	0.042
Irish Examiner (Rad. Left)	$+0.122^{***}$	$+0.188^{\dagger}$	0.030
Irish Examiner (Right)	$+0.039^{***}$	$+0.031^{\dagger}$	<0.001
La Stampa (Left)	$+0.048^{***}$	$+0.051^{\dagger}$	0.003
La Repubblica (Center)	-0.026^{*}	$-0.026^{\dagger}$	0.025
La Repubblica (Left)	-0.021^{**}	$+0.037^{\dagger}$	<0.001
La Repubblica (Rad. Left)	$+0.081^{***}$	$+0.084^{\dagger}$	<0.001

Table 10: 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.

E.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 ($\beta_3 = -0.047$,

Outlet (Subject)	Real DiD	Placebo DiD	Placebo p
Euronews (Keir Starmer)	$+0.310^{***}$	$+0.021$	0.528
La Repubblica (Virginia Raggi)	-0.150^{***}	-0.005	0.886
Europe 1 (Nicolas Sarkozy)	$+0.190^{***}$	$-0.047^{\dagger}$	0.047

Table 11: 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.

$p = 0.047$). However, the true treatment effect is overwhelmingly positive ($\beta_3 = +0.190$, $p < 10^{-15}$). Rather than invalidating the result, this indicates that *Europe 1* was on a trajectory of increasing criticism toward Sarkozy alongside the control group, and the Bolloré takeover abruptly reversed this dynamic to produce a massive positive realignment.

F 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 12 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.

G 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:

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

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 13 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.

H 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 14 reports the comparison. The core findings are robust: all Irish Examiner bloc effects, La Stampa Left, La Provence Radical Right, and La Repubblica (Center, Left, Radical Left) remain significant and directionally consistent, often 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, some borderline results from the primary analysis do not survive the restriction: Euronews Radical Right, Europe 1 Center and Radical Right, and La Repubblica Right lose significance when

Table 12: Robustness check: OLS vs. cluster-robust ordered logit on the DiD interaction term. Coefficients are not comparable in magnitude (different scales). Sign ✓ means both estimators agree on direction; Sig. ✓ 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}
(a) Micro showcase cases					
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
(b) Macro Global / National (OLS-significant)					
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
(c) Political bloc (FDR-significant)					
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

Table 13: 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

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.

I 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., Virginia Raggi), 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).

J 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.

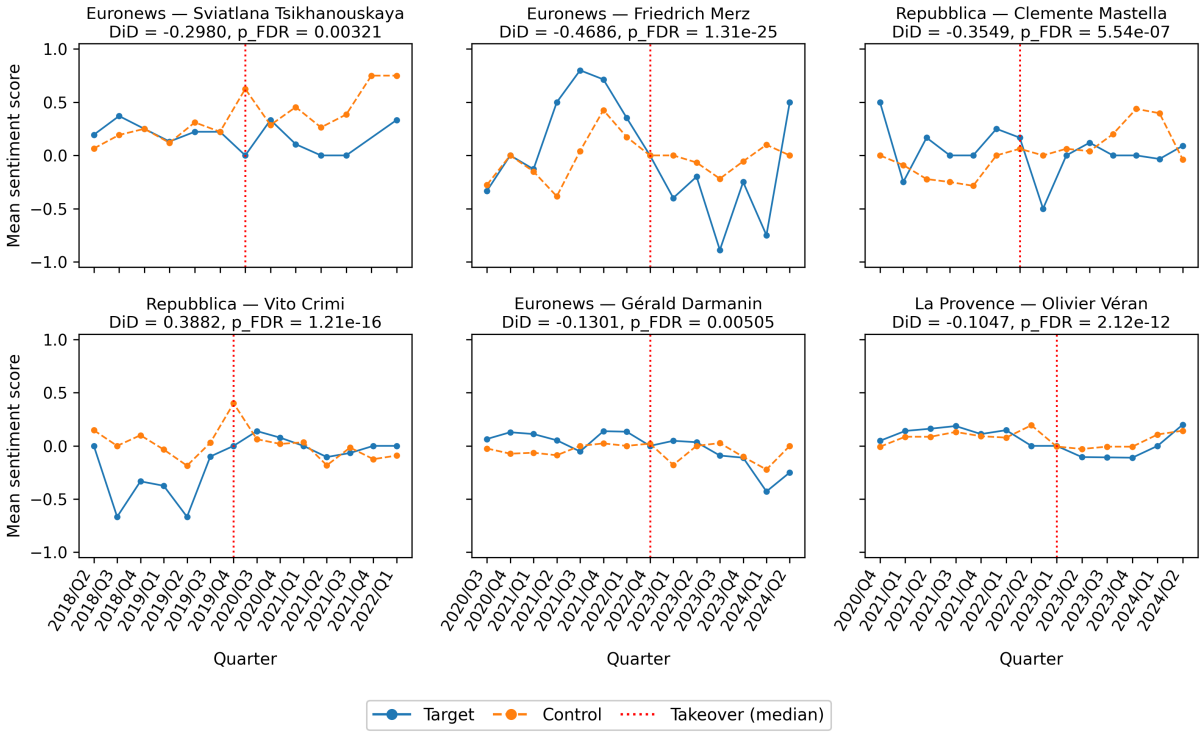

Figure 3: Supplementary case studies of significant micro-level editorial shifts following media acquisitions (Sviatlana Tsikhanouskaya, Friedrich Merz, Clemente Mastella, Virginia Raggi, Gérald Darmanin, and Olivier Véran). The charts compare the acquired outlet with its respective control group.

Unit of observation. Each row in the dataset is an *entity–sentence occurrence*: a single detected mention of a political figure in a target sentence, together with metadata (`occurrence_id`, `role`, `period`, `date`, `match_variant`, etc.). When several entities appear in the same target sentence, the pipeline emits one row per entity; hence the total row count exceeds the number of distinct `occurrence_id` keys (approximately 277,000 rows list more than one entity slug in the `entities` field). The headline total reported in the main text—**1,376,552**—is therefore the number of *entity–sentence occurrences*, not the number of unique sentences or articles.

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 entity–sentence occurrence 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` occurrences.

Table 14: Bloc-level DiD: all politicians vs. domestic only. Sig. = significant at p-adj < 0.05 after Benjamini–Hochberg FDR correction.

Outlet	**Bloc**	β_{all}	**Sig.**	$\beta_{\text{nat.}}$	**Sig.**	**Dir.**
Euronews	Rad. Right	+0.050	✓	+0.018	×	✓
Europe 1	Center	+0.018	✓	+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	×	✓

Table 15: Composition of `ccnews_linked` by target outlet (nine media clusters). Occurrences = entity–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
`europe1_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 (`europe1_fr`) and 452 (`euronews_eu`).

K 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 16: Complete Significant Results (Macro and Micro Levels)

Outlet	Level	Entity	DiD Coeff.	P-value
Euronews	Micro_subject	Alberto Fernández	+0.381	< 0.001
Euronews	Micro_subject	Alexander De Croo	+0.103	0.015
Euronews	Micro_subject	Alexander Lukashenko	-0.152	0.007
Euronews	Micro_subject	Angela Merkel	-0.082	< 0.001
Euronews	Micro_subject	Bashar al-Assad	+0.328	< 0.001
Euronews	Micro_subject	Bill Gates	+0.133	0.009
Euronews	Micro_subject	Boris Johnson	-0.195	< 0.001
Euronews	Micro_subject	Bruno Le Maire	-0.070	0.024
Euronews	Micro_subject	Charles Michel	-0.047	0.033
Euronews	Micro_subject	Cyril Ramaphosa	-0.092	< 0.001
Euronews	Micro_subject	David Cameron	+0.222	< 0.001
Euronews	Micro_subject	Didier Reynders	-0.157	< 0.001
Euronews	Micro_subject	Dmitry Peskov	+0.036	0.002
Euronews	Micro_subject	Dmytro Kuleba	-0.073	< 0.001
Euronews	Micro_subject	Donald Tusk	+0.209	< 0.001
Euronews	Micro_subject	Ebrahim Raisi	-0.301	< 0.001
Euronews	Micro_subject	Frank-Walter Steinmeier	-0.094	< 0.001
Euronews	Micro_subject	Frans Timmermans	-0.139	0.003
Euronews	Micro_subject	Friedrich Merz	-0.469	< 0.001
Euronews	Micro_subject	Gabriel Attal	-0.139	0.003
Euronews	Micro_subject	Grant Shapps	-0.087	0.046
Euronews	Micro_subject	Greta Thunberg	+0.040	0.004
Euronews	Micro_subject	Guillermo Lasso	+0.223	< 0.001
Euronews	Micro_subject	Gérald Darmanin	-0.130	0.003
Euronews	Micro_subject	Jacinda Ardern	+0.211	< 0.001
Euronews	Micro_subject	Jair Bolsonaro	+0.113	0.046
Euronews	Micro_subject	Janet Yellen	-0.183	< 0.001
Euronews	Micro_subject	Josep Borrell	-0.050	0.005
Euronews	Micro_subject	Justin Trudeau	+0.091	0.025
Euronews	Micro_subject	Kamala Harris	-0.121	0.034
Euronews	Micro_subject	Keir Starmer	+0.308	< 0.001
Euronews	Micro_subject	Klaus Iohannis	-0.356	< 0.001
Euronews	Micro_subject	Kristalina Georgieva	-0.258	< 0.001
Euronews	Micro_subject	Kyriakos Mitsotakis	-0.069	0.006
Euronews	Micro_subject	Linda Thomas-Greenfield	-0.103	0.005
Euronews	Micro_subject	Maia Sandu	-0.335	< 0.001
Euronews	Micro_subject	Mario Draghi	+0.107	0.004
Euronews	Micro_subject	Markus Söder	+0.249	< 0.001
Euronews	Micro_subject	Maroš Šefčovič	-0.073	0.002
Euronews	Micro_subject	Milorad Dodik	-0.177	< 0.001

Continued on next page

Table 16 – continued from previous page

Outlet	Level	Entity	DiD Coeff.	P-value
Euronews	Micro_subject	Nancy Pelosi	-0.127	< 0.001
Euronews	Micro_subject	Nayib Bukele	+0.263	0.002
Euronews	Micro_subject	Nicolas Sarkozy	+0.227	0.003
Euronews	Micro_subject	Pedro Sánchez	-0.110	< 0.001
Euronews	Micro_subject	Rafael Grossi	-0.137	< 0.001
Euronews	Micro_subject	Recep Tayyip Erdoğan	+0.068	< 0.001
Euronews	Micro_subject	Robert Habeck	-0.166	< 0.001
Euronews	Micro_subject	Sanna Marin	+0.520	< 0.001
Euronews	Micro_subject	Sergio Mattarella	+0.176	< 0.001
Euronews	Micro_subject	Sviatlana Tsikhanouskaya	-0.298	< 0.001
Euronews	Micro_subject	Theresa May	+0.173	0.005
Euronews	Micro_subject	Tony Blair	+0.229	0.002
Euronews	Micro_subject	Valdis Dombrovskis	-0.222	< 0.001
Euronews	Micro_subject	Věra Jourová	+0.082	0.003
Euronews	Micro_subject	Ylva Johansson	+0.093	0.003
Europe 1	Macro_National	All (Global/National)	+0.020	< 0.01
Europe 1	Micro_subject	Agnès Pannier-Runacher	+0.155	< 0.001
Europe 1	Micro_subject	Alexeï Navalny	-0.145	< 0.001
Europe 1	Micro_subject	António Guterres	+0.158	< 0.001
Europe 1	Micro_subject	Aurore Bergé	+0.112	0.034
Europe 1	Micro_subject	Aurélien Pradié	+0.180	< 0.001
Europe 1	Micro_subject	David Lisnard	+0.192	0.002
Europe 1	Micro_subject	Donald Trump	-0.068	< 0.001
Europe 1	Micro_subject	Elon Musk	+0.215	< 0.001
Europe 1	Micro_subject	Gabriel Attal	+0.059	< 0.001
Europe 1	Micro_subject	Gérard Larcher	-0.124	0.003
Europe 1	Micro_subject	Jair Bolsonaro	+0.205	< 0.001
Europe 1	Micro_subject	Jean-Luc Mélenchon	+0.034	0.049
Europe 1	Micro_subject	Nicolas Sarkozy	+0.193	< 0.001
Europe 1	Micro_subject	Olivier Véran	+0.036	< 0.001
Europe 1	Micro_subject	Sandrine Rousseau	+0.155	< 0.001
Europe 1	Micro_subject	Stanislas Guerini	+0.203	< 0.001
Europe 1	Micro_subject	Valérie Pécresse	+0.080	0.031
Europe 1	Micro_subject	Viktor Orbán	-0.165	0.002
Europe 1	Micro_subject	Vladimir Poutine	-0.094	< 0.001
Europe 1	Micro_subject	Volodymyr Zelensky	+0.137	< 0.001
Europe 1	Micro_subject	Xavier Bertrand	+0.124	< 0.001
Europe 1	Micro_subject	Xi Jinping	+0.110	< 0.001
Europe 1	Micro_subject	Éric Zemmour	+0.070	0.006
Hürriyet	Macro_Global	All (Global/National)	+0.024	0.028
Hürriyet	Micro_subject	Ahmet Davutoğlu	-0.257	< 0.001
Hürriyet	Micro_subject	António Guterres	+0.129	< 0.001
Hürriyet	Micro_subject	Emmanuel Macron	+0.174	< 0.001
Hürriyet	Micro_subject	Selahattin Demirtaş	-0.139	0.049
Hürriyet	Micro_subject	Sergey Lavrov	+0.045	0.031
Independent	Micro_subject	Arlene Foster	-0.098	< 0.001
Independent	Micro_subject	Benjamin Netanyahu	-0.087	< 0.001
Independent	Micro_subject	Boris Johnson	+0.126	< 0.001

Continued on next page

Table 16 – continued from previous page

Outlet	Level	Entity	DiD Coeff.	P-value
Independent	Micro_subject	Brendan Howlin	+0.139	< 0.001
Independent	Micro_subject	Danny Healy-Rae	+0.247	0.003
Independent	Micro_subject	Denis Naughten	+0.086	0.035
Independent	Micro_subject	Dominic Raab	+0.196	0.007
Independent	Micro_subject	Donald Trump	-0.074	0.002
Independent	Micro_subject	Emmanuel Macron	+0.081	< 0.001
Independent	Micro_subject	Jair Bolsonaro	-0.161	0.003
Independent	Micro_subject	Joe Biden	+0.311	< 0.001
Independent	Micro_subject	Leo Varadkar	+0.055	0.004
Independent	Micro_subject	Michael Healy-Rae	-0.270	< 0.001
Independent	Micro_subject	Mike Pompeo	-0.105	< 0.001
Independent	Micro_subject	Mitch McConnell	-0.111	0.004
Independent	Micro_subject	Nancy Pelosi	-0.232	< 0.001
Independent	Micro_subject	Pedro Sánchez	+0.188	< 0.001
Independent	Micro_subject	Recep Tayyip Erdoğan	+0.115	0.037
Independent	Micro_subject	Richard Bruton	+0.136	< 0.001
Independent	Micro_subject	Stephen Donnelly	+0.136	0.035
Independent	Micro_subject	Vladimir Putin	+0.089	< 0.001
Irish Examiner	Macro_National	All (Global/National)	+0.051	< 0.001
Irish Examiner	Micro_subject	Donald Trump	+0.111	< 0.001
Irish Examiner	Micro_subject	Donald Tusk	+0.234	< 0.001
Irish Examiner	Micro_subject	Emmanuel Macron	+0.147	< 0.001
Irish Examiner	Micro_subject	Enda Kenny	+0.057	< 0.001
Irish Examiner	Micro_subject	Gerry Adams	-0.118	0.029
Irish Examiner	Micro_subject	Leo Varadkar	+0.082	< 0.001
Irish Examiner	Micro_subject	Mattie McGrath	-0.150	< 0.001
Irish Examiner	Micro_subject	Michael Creed	+0.077	0.003
Irish Examiner	Micro_subject	Micheál Martin	+0.069	< 0.001
Irish Examiner	Micro_subject	Mike Pence	-0.125	< 0.001
Irish Examiner	Micro_subject	Paschal Donohoe	+0.073	< 0.001
Irish Examiner	Micro_subject	Paul Kehoe	+0.279	< 0.001
Irish Examiner	Micro_subject	Pearse Doherty	-0.109	0.028
Irish Examiner	Micro_subject	Pope Francis	+0.102	0.048
Irish Examiner	Micro_subject	Richard Bruton	+0.096	< 0.001
Irish Examiner	Micro_subject	Shane Ross	+0.063	< 0.001
Irish Examiner	Micro_subject	Vladimir Putin	-0.152	< 0.001
LA Times	Macro_National	All (Global/National)	+0.034	< 0.001
LA Times	Micro_subject	Andrew Cuomo	-0.139	< 0.001
LA Times	Micro_subject	Angela Merkel	-0.082	< 0.001
LA Times	Micro_subject	Antonin Scalia	-0.147	0.002
LA Times	Micro_subject	Barack Obama	-0.031	0.023
LA Times	Micro_subject	Bernie Sanders	+0.097	< 0.001
LA Times	Micro_subject	Bill Gates	+0.274	< 0.001
LA Times	Micro_subject	Emmanuel Macron	+0.133	< 0.001
LA Times	Micro_subject	Gavin Newsom	+0.308	0.002
LA Times	Micro_subject	John Kerry	+0.290	< 0.001
LA Times	Micro_subject	Lindsey Graham	-0.100	< 0.001
LA Times	Micro_subject	Lisa Murkowski	-0.070	0.005

Continued on next page

Table 16 – continued from previous page

Outlet	Level	Entity	DiD Coeff.	P-value
LA Times	Micro_subject	Merrick Garland	-0.270	< 0.001
LA Times	Micro_subject	Mick Mulvaney	-0.103	< 0.001
LA Times	Micro_subject	Mitt Romney	+0.123	0.027
LA Times	Micro_subject	Neil Gorsuch	-0.325	< 0.001
LA Times	Micro_subject	Pope Francis	-0.145	< 0.001
LA Times	Micro_subject	Rand Paul	-0.189	< 0.001
LA Times	Micro_subject	Recep Tayyip Erdoğan	+0.259	< 0.001
LA Times	Micro_subject	Rick Perry	+0.199	0.005
LA Times	Micro_subject	Ryan Zinke	+0.119	0.046
LA Times	Micro_subject	Ted Cruz	+0.056	0.005
LA Times	Micro_subject	Vladimir Putin	-0.056	0.002
LA Times	Micro_subject	Xavier Becerra	-0.115	< 0.001
La Provence	Macro_National	All (Global/National)	+0.020	< 0.01
La Provence	Micro_subject	Agnès Pannier-Runacher	+0.270	< 0.001
La Provence	Micro_subject	Alexeï Navalny	+0.083	< 0.001
La Provence	Micro_subject	Benyamin Netanyahou	-0.203	< 0.001
La Provence	Micro_subject	Bruno Le Maire	+0.102	< 0.001
La Provence	Micro_subject	Bruno Retailleau	+0.068	0.042
La Provence	Micro_subject	Charles Michel	-0.165	0.006
La Provence	Micro_subject	Clément Beaune	+0.077	< 0.001
La Provence	Micro_subject	Dmitri Peskov	-0.078	< 0.001
La Provence	Micro_subject	Donald Trump	-0.057	< 0.001
La Provence	Micro_subject	Gérald Darmanin	-0.048	< 0.001
La Provence	Micro_subject	Jair Bolsonaro	+0.091	< 0.001
La Provence	Micro_subject	Josep Borrell	-0.104	< 0.001
La Provence	Micro_subject	Julien Bayou	-0.379	< 0.001
La Provence	Micro_subject	Kamala Harris	+0.122	0.046
La Provence	Micro_subject	Kim Jong-un	-0.232	0.002
La Provence	Micro_subject	Manuel Bompard	+0.202	0.002
La Provence	Micro_subject	Marine Le Pen	+0.116	< 0.001
La Provence	Micro_subject	Marion Maréchal	+0.076	0.002
La Provence	Micro_subject	Olaf Scholz	+0.074	0.034
La Provence	Micro_subject	Olivier Faure (homme politique)	+0.052	0.006
La Provence	Micro_subject	Olivier Véran	-0.105	< 0.001
La Provence	Micro_subject	Valérie Pécresse	-0.117	< 0.001
La Provence	Micro_subject	Viktor Orbán	-0.277	< 0.001
La Provence	Micro_subject	Vladimir Poutine	-0.177	< 0.001
La Provence	Micro_subject	Volodymyr Zelensky	+0.068	0.006
La Provence	Micro_subject	Édouard Philippe	-0.118	0.002
La Provence	Micro_subject	Élisabeth Borne	+0.050	0.024
La Provence	Micro_subject	Éric Ciotti	+0.117	< 0.001
La Repubblica	Macro_Global	All (Global/National)	+0.044	< 0.01
La Repubblica	Macro_National	All (Global/National)	-0.020	< 0.01
La Repubblica	Micro_subject	Alessandro Di Battista	+0.182	< 0.001
La Repubblica	Micro_subject	Andrea Marcucci	-0.144	< 0.001
La Repubblica	Micro_subject	Andrea Orlando	-0.060	0.038
La Repubblica	Micro_subject	Antonio Tajani	-0.109	< 0.001
La Repubblica	Micro_subject	Beppe Grillo	+0.062	< 0.001

Continued on next page

Table 16 – continued from previous page

Outlet	Level	Entity	DiD Coeff.	P-value
La Repubblica	Micro_subject	Bernie Sanders	-0.282	< 0.001
La Repubblica	Micro_subject	Clemente Mastella	-0.355	< 0.001
La Repubblica	Micro_subject	Danilo Toninelli	-0.221	< 0.001
La Repubblica	Micro_subject	Dario Franceschini	-0.190	< 0.001
La Repubblica	Micro_subject	Donald Trump	+0.097	< 0.001
La Repubblica	Micro_subject	Emma Bonino	+0.216	< 0.001
La Repubblica	Micro_subject	Eugenio Giani	+0.160	0.004
La Repubblica	Micro_subject	Giorgia Meloni	-0.048	0.001
La Repubblica	Micro_subject	Giuseppe Conte (scrittore)	+0.137	< 0.001
La Repubblica	Micro_subject	Giuseppe Sala (politico)	+0.077	0.003
La Repubblica	Micro_subject	Graziano Delrio	-0.168	< 0.001
La Repubblica	Micro_subject	Jair Bolsonaro	-0.131	0.002
La Repubblica	Micro_subject	Luciana Lamorgese	-0.174	< 0.001
La Repubblica	Micro_subject	Maria Elena Boschi	-0.120	0.002
La Repubblica	Micro_subject	Mariastella Gelmini	+0.212	< 0.001
La Repubblica	Micro_subject	Marine Le Pen	+0.267	< 0.001
La Repubblica	Micro_subject	Massimo D'Alema	+0.136	< 0.001
La Repubblica	Micro_subject	Matteo Lepore	-0.363	< 0.001
La Repubblica	Micro_subject	Matteo Orfini	+0.205	< 0.001
La Repubblica	Micro_subject	Maurizio Gasparri	+0.125	< 0.001
La Repubblica	Micro_subject	Michele Emiliano	-0.109	0.005
La Repubblica	Micro_subject	Mike Pence	-0.144	< 0.001
La Repubblica	Micro_subject	Mike Pompeo	+0.190	0.001
La Repubblica	Micro_subject	Monica Cirinnà	-0.304	< 0.001
La Repubblica	Micro_subject	Nicola Fratoianni	+0.160	< 0.001
La Repubblica	Micro_subject	Nicola Zingaretti	-0.071	< 0.001
La Repubblica	Micro_subject	Sergio Mattarella	-0.062	< 0.001
La Repubblica	Micro_subject	Stefano Patuanelli	-0.067	< 0.001
La Repubblica	Micro_subject	Susanna Ceccardi	-0.170	0.046
La Repubblica	Micro_subject	Teresa Bellanova	-0.086	0.045
La Repubblica	Micro_subject	Virginia Raggi	-0.161	< 0.001
La Repubblica	Micro_subject	Vito Crimi	+0.388	< 0.001
La Repubblica	Micro_subject	Vladimir Putin	+0.073	0.003
La Repubblica	Micro_subject	Xi Jinping	+0.239	< 0.001
La Stampa	Macro_National	All (Global/National)	+0.032	< 0.001
La Stampa	Micro_subject	Alessandro Di Battista	+0.297	< 0.001
La Stampa	Micro_subject	Angela Merkel	+0.161	< 0.001
La Stampa	Micro_subject	Beppe Grillo	+0.102	< 0.001
La Stampa	Micro_subject	Carlo Calenda	-0.140	< 0.001
La Stampa	Micro_subject	Donald Trump	-0.065	0.003
La Stampa	Micro_subject	Enrico Costa (politico)	-0.244	0.002
La Stampa	Micro_subject	Giorgia Meloni	-0.097	< 0.001
La Stampa	Micro_subject	Giovanni Toti	+0.077	0.007
La Stampa	Micro_subject	Giuseppe Conte (scrittore)	+0.199	< 0.001
La Stampa	Micro_subject	Greta Thunberg	-0.315	< 0.001
La Stampa	Micro_subject	Jair Bolsonaro	-0.293	< 0.001
La Stampa	Micro_subject	Joe Biden	+0.341	< 0.001
La Stampa	Micro_subject	Silvio Berlusconi	+0.060	< 0.001

Continued on next page

Table 16 – continued from previous page

Outlet	Level	Entity	DiD Coeff.	P-value
La Stampa	Micro_subject	Ursula von der Leyen	+0.169	< 0.001
La Stampa	Micro_subject	Viktor Orbán	+0.169	< 0.001
La Stampa	Micro_subject	Virginia Raggi	+0.123	0.002

L 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.

L.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.

L.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.

L.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/Qwen3.5-27B`, 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.

L.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.

M 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 three 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 L. 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 preprocessing 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 LaTeX templates for multi-column tables and complex figures.

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.