
Who Gets What Ads? Tail-Aware Metrics and the Composition of Mobile Political Exposure

Sage Open
XX(X):1–17
©The Author(s) 2026
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Chris J. Vargo¹ and Bridget Barrett¹

Abstract

We examine who receives explicitly partisan political ads on mobile social media. We passively metered Facebook, Instagram, X, and TikTok feeds for 379 U.S. adults; each respondent contributed the 30 days before their party-identification survey, and pooled impressions spanned October 4 through November 5, 2024 (UTC). The analysis includes 435,006 ad impressions. We distinguish partisan ad balance, the relative mix of Democratic versus Republican impressions, from partisan volume, the total number of explicitly partisan impressions. Explicitly partisan political-ad load was sparse for most respondents and concentrated for a small subset, especially on Facebook and Instagram when all observed respondent-platform cells were included; platform rankings changed under denominator restrictions. In respondent-level models, party identification was associated with partisan ad balance, while total partisan volume was not reliable after adjustment for activity, platform coverage, and demographics. The results support a composition-first view of delivered ad exposure and tail-aware transparency, while remaining descriptive rather than causal.

Keywords

political advertising; mobile social media; passive metering; ad delivery; microtargeting; political communication

Author Accepted Manuscript — Accepted September 23, 2026 at SAGE Open. No highlights, tracked changes, or markup.

¹University of Colorado Boulder, College of Media, Communication and Information, Boulder, CO, USA

Corresponding author:

a

Introduction

Who Gets What Ads? Tail-Aware Metrics and the Composition of Mobile Political Exposure

Targeted political advertising remains central to debates about microtargeting and democratic accountability, even though field evidence often finds small, heterogeneous, or null average effects on persuasion and turnout (Coppock et al., 2020; Kalla & Broockman, 2018; Aggarwal et al., 2023; Allcott et al., 2026; Barrett & McGregor, 2025). That gap between strategic belief and modest average effects makes delivered exposure a core empirical question: who actually receives political ads in routine use? Public data usually describe what campaigns buy, not whose feeds those buys reach (Edelson et al., 2018; Mehta & Erickson, 2022; Merrill & Tobin, 2019; DeLong, 2021; Silva et al., 2020).

We address this gap by metering mobile feeds across Facebook, Instagram, X, and TikTok during a U.S. general-election period and analyzing the partisan composition of impressions at the respondent level. Two hypotheses follow from the supply-side expectation that campaigns prioritize reachable copartisans and high-yield supporters:

Hypothesis 1a (copartisan Democratic targeting): Net of demographics and political engagement, greater exposure to explicitly Democratic-aligned ads will be associated with lower values on the Republican-coded seven-point PID scale, indicating more Democratic identification.

Hypothesis 1b (copartisan Republican targeting): Net of the same covariates, greater exposure to explicitly Republican-aligned ads will be associated with higher values on the Republican-coded seven-point PID scale.

Campaigns combine voter-file augmentation, list matching, and platform tools such as custom and look-alike audiences to prioritize reachable supporters and high-yield donors (Hersh, 2015; Fowler, 2020; Fowler et al., 2018, 2021, 2023; Ridout et al., 2018, 2021; Barrett & McGregor, 2025; Allcott et al., 2026). These practices extend pre-digital segmentation strategies that target persuadable or mobilizable subconstituencies at reasonable cost (Hillygus & Shields, 2008; Nickerson & Rogers, 2014). Fundraising and list-building objectives can therefore bias delivery toward copartisans and the politically engaged (Allcott et al., 2026; Ridout et al., 2021). If this bias appears in delivered impressions, Democratic advertisers should be overrepresented in Democratic users' feeds and Republican advertisers in Republican users' feeds.

Platform optimization can amplify or redirect campaign intent. Ad delivery unfolds through auctions and pacing systems that balance bids, quality signals, and expected engagement under budget constraints (Edelman et al., 2007; Varian, 2007). These systems can skew realized audiences relative to stated targets and produce demographic or political disparities in delivery and price (Ali et al., 2021; Bär et al., 2024; Lambrecht & Tucker, 2019). Work on algorithmic fairness shows how proxy optimization can reproduce group-level disparities even without explicit intent (Barocas & Selbst, 2016). Because party coalitions differ in age, geography, and engagement, party-asymmetric reach is an empirical question:

Research Question 1 (party asymmetry): Is the association between partisan ad exposure and partisan identification symmetric in magnitude, or is one party's ad exposure more strongly associated with identification than the other's?

Users add demand signals through follows, prior engagement, and inferred interests. Selective-exposure research shows that people gravitate toward congenial content and that social cues and networks reinforce those patterns (Stroud, 2011; Messing & Westwood, 2014; Bode & Dalrymple, 2016; Cotter et al., 2021). Large-scale observational work also finds heavy-tailed consumption of political or misinformation content (Bakshy et al., 2015; Grinberg et al., 2019). When engagement signals shape optimization, political ads can remain rare for most respondents while concentrating among the respondents most likely to engage, donate, or sign up. That distribution motivates a direct load question:

Research Question 2 (explicitly partisan political-ad load): What proportion of respondents' total ad impressions are explicitly partisan political ads, and how does that vary across platforms and respondents?

Ad libraries are useful but incomplete measures of delivery. They often omit key fields, limit access, or change policies in ways that impede longitudinal analysis (Edelson et al., 2018; Mehta & Erickson, 2022; Merrill & Tobin, 2019; DeLong, 2021). Independent audits show that optimization can differentially reach demographic and political groups, raise cross-group reach costs, and reproduce social skews (Ali et al., 2021; Datta et al., 2015; Lambrecht & Tucker, 2019). Passive metering connects delivered impressions to respondent characteristics in ways spend- or creative-based data cannot.

The literature predicts a delivery pattern in which composition matters more than count. If campaigns emphasize reachable supporters, auctions tilt toward low-cost engagement, and user interest signals are strong, then the partisan mix of political ads should track respondent characteristics more closely than total political-ad volume (Allcott et al., 2026; Ridout et al., 2018, 2021; Ali et al., 2021; Bär et al., 2024; Cotter et al., 2021). H1 tests whether delivered mix aligns with identity, RQ1 tests whether that link differs by party, and RQ2 describes the distribution of explicitly partisan political-ad share across platforms and respondents.

Plain-language definitions. We use partisan ad balance to mean the relative mix of explicitly Democratic versus explicitly Republican ad impressions a respondent received. We use partisan volume to mean the total number of explicitly Democratic plus explicitly Republican impressions, regardless of which party they favored. A composition-first result means the party mix of the partisan ads matters more than the sheer number of partisan ads.

Method

Data Collection

YouGov recruited 379 U.S. adults who installed a RealityMine passive metering application on smartphones and tablets. Each respondent contributed the 30 days before their party-identification survey; pooled impressions span October 4 through November 5, 2024 (UTC). The analysis file contains 435,006 classified ad impressions: 330,798 on Facebook, 63,430 on X, 34,448 on TikTok, and 6,330 on Instagram. Unweighted panel descriptors are 56.7% female and 43.3% male; 66.5% White, 16.6% Black, and 12.4% Hispanic by the separate ethnicity item; mean age 48.7 years as of 2025; and 26.9%

with a four-year degree or higher. These statistics contextualize the metered panel; they do not establish national representativeness.

Sample construction and analytic units. The analysis file contains 435,006 classified ad impressions linked to 379 YouGov/RealityMine case IDs. These records produce 689 observed respondent-platform cells, defined as respondent-platform combinations with at least one observed ad impression. For H1 and RQ1, impressions were aggregated to one respondent-level analytic record because partisan identification is measured once per respondent; seven respondents who answered 'Not sure' on the seven-point PID item were excluded from PID models, yielding $N = 372$. RQ2 uses respondent-platform cells because its estimand is platform-specific explicitly partisan political-ad load.

External validity and scope. We compare unweighted panel descriptors with 2024 American Community Survey one-year profile benchmarks only as a scope check (U.S. Census Bureau, 2025). The metered panel is more female (56.7% vs. 50.5%), more Black (16.6% vs. 12.1%), more White (66.5% vs. 59.8%), less Hispanic by the separate ethnicity item (12.4% vs. 20.0%), and less likely to report a four-year degree or higher (26.9% vs. 36.8%). The sample mean age is also higher than the ACS median age (48.7 vs. 39.2), a rough comparison because the metrics differ. Claims are therefore bounded to delivered mobile/tablet social-media ad exposure in a YouGov/RealityMine U.S. metered panel. Platform scope is also uneven: Facebook accounts for 76.0% of observed impressions, X for 14.6%, TikTok for 7.9%, and Instagram for 1.5%.

Human participants, consent, and risk reduction. This study was reviewed by the authors' Institutional Review Board; protocol details are reported on the separate title page. Participants entered the YouGov/RealityMine panel and installed the passive metering application after consenting to the panel survey and metering procedures. The analysis uses de-identified secondary records linked by case IDs. No author-initiated contact, experimental manipulation, or attempt to identify individual feeds occurred. Risk was minimized by using case IDs rather than direct identifiers, analyzing records locally, reporting aggregate patterns, and excluding respondent-level browsing or feed histories from publication.

Table A1. Sample construction and analytic units.

Step	N	Unit	Note
Raw classified ad impressions	435006	impression	All rows in master_ad_user_merged.csv
Unique respondents in ad file	379	respondent	RealityMine/YouGov case IDs
Observed respondent-platform cells	689	respondent-platform	Cells with at least one observed ad impression on a platform
Respondent-level analytic records before PID exclusions	379	respondent	Ads aggregated to respondent-level counts
Excluded: PID not sure / unmapped	7	respondent	Cannot order on seven-point Democrat-Republican PID scale
Respondent-level model records after PID exclusions	372	respondent	Primary H1/RQ1 analysis file

Political and civic characteristics were heterogeneous. On the seven-category party identification measure, unweighted shares are 21.4% strong Democrat, 11.1% not very

strong Democrat, 9.0% lean Democrat, 20.3% Independent, 6.9% lean Republican, 10.8% not very strong Republican, 18.7% strong Republican, and 1.8% not sure. Ideology spans the spectrum, with 25.6% liberal or very liberal, 35.1% moderate, 29.3% conservative or very conservative, and 10.0% not sure. Civic and information engagement also varied: 87.6% report being registered to vote, and news interest ranges from most of the time (37.5%) to hardly at all (16.9%). We do not use panel weights for these descriptors because the analysis does not claim national representativeness.

Topic Modeling

Raw advertiser and metadata fields were concatenated into a combined_text field for each advertisement. This field preserved advertiser identity and ad context for embedding, topic modeling, and later classification.

Ad text was embedded with all-mpnet-base-v2 from Sentence Transformers, a Sentence-BERT model family (Reimers & Gurevych, 2019). Embeddings were clustered with BERTopic, which combines transformer document embeddings with class-based TF-IDF topic representations (Grootendorst, 2022). Cluster quality was assessed with the silhouette score using cosine distance and excluding noise entries (Rousseeuw, 1987). The final model used $k = 11,012$ topics and achieved a silhouette score of 0.49.

Topic Metadata Generation

Qwen2.5-72B-Instruct-AWQ summarized topic content and assigned topics to the Interactive Advertising Bureau taxonomy (Qwen et al., 2024). The pipeline used a fixed prompt template, fixed category definitions, and guided-choice category selection. No temperature field is retained, so validation carries the reliability claim; the analysis does not assume deterministic decoding. Responses were parsed into brand lists, summaries, and category assignments.

IAB-to-politics bridge. The IAB taxonomy supplied the first topic screen. Topics assigned to the IAB Politics category were treated as candidate political topics; all other categories were retained as nonpolitical contrast material unless the later audit indicated ambiguity. We then narrowed the Politics topics to a deliberately conservative party-count set: the largest, clearly U.S.-partisan topics whose advertiser, topic summary, and example ads consistently pointed to Democratic or Republican electoral actors, committees, or partisan issue appeals. This produced ten selected partisan topics for the main Democratic and Republican exposure measures. Politics topics that were issue-oriented, mixed, non-U.S., unclear, or party-valenced but not in the pre-specified party-count set remained in the broader Politics audit but were excluded from party-specific counts.

Partisan Political Ad Topics

We validated the topic labels in two steps. First, the retained 100-topic category-agreement file shows agreement for 85 topic-category assignments and disagreement for 15. That file supports percent agreement; it does not contain the normalized second-coder categories needed for Cohen's kappa. Second, the author team audited 70 topics: the ten selected partisan topics, the 40 highest-impression Politics topics outside that

party-count set, and the 20 highest-impression non-Politics contrast topics. The Politics screen performed strongly: 50/50 predicted Politics topics were author-coded as political, electoral, or issue-political, while 19/20 high-impression non-Politics contrast topics were nonpolitical and one was mixed/unclear. The estimable quantities are predicted-Politics precision of 100%, high-impression non-Politics confirmation of 95% when the mixed/unclear topic is counted as non-confirmed, and selected partisan-direction precision of 100%. Recall and F1 are not estimable because the audit was enriched for high-impression and selected topics rather than drawn as a probability sample of all true political and nonpolitical topics. The ten selected partisan topics were all classified as Politics and confirmed in the expected party direction: five Democratic topics (10, 28, 173, 208, and 224) and five Republican topics (47, 50, 72, 77, and 90). These topics account for 2,623 explicitly partisan impressions, including 1,400 Democratic and 1,223 Republican impressions. The 40 high-impression nonselected Politics topics included 23 party-valenced topics outside the pre-specified party-count set (11 Democratic and 12 Republican) and 17 issue, news, electoral-information, non-U.S., mixed, or unclear topics. Selected partisan topics account for 13.3% of Politics-category impressions on Facebook, 18.7% on Instagram, 2.6% on TikTok, and 11.5% on X. The remaining Politics-category impressions were excluded from party counts rather than forced into Democratic or Republican labels, reducing partisan-direction error while understating political and party-valenced volume.

Table A2. Validation summary for automated and author-adjudicated labels.

Validation check	Unit	N	Result
Retained 100-topic category agreement check	topic assignment	100	85/100 topic-category assignments marked agreement (85%); the retained sheet supports percent agreement but not the normalized second-coder category structure required for kappa
Stratified topic audit: predicted Politics	topic	50	50/50 audited Politics topics were author-confirmed as political/electoral/issue-political
Stratified topic audit: non-Politics contrast	topic	20	19/20 high-impression contrast topics were confirmed nonpolitical; 1 large catch-all topic was mixed/unclear
Topic-level agreement metrics	topic	70	Predicted-Politics precision = 50/50; high-impression non-Politics specificity proxy = 19/20 when the mixed/unclear catch-all is counted as non-confirmed; selected partisan direction precision = 10/10; recall/F1 are not estimable from the enriched topic sample

Validation check	Unit	N	Result
Confusion-style Politics/non-Politics screen audit	topic	70	Predicted Politics: 50 author-coded political, 0 nonpolitical/mixed. High-impression predicted non-Politics contrast: 19 nonpolitical, 1 mixed/unclear, 0 author-coded political.
Selected partisan topic category cross-check	topic	10	10/10 selected partisan topics are also classified as Politics
Selected partisan topic direction audit	topic	10	5/5 Democratic topics and 5/5 Republican topics were confirmed with no issue/unclear recodes among selected party topics
Selected partisan topic impression coverage	impression	2623	1400 Democratic and 1223 Republican explicitly partisan impressions
High-impression nonselected Politics topic audit	topic	40	23/40 appeared party-valenced but outside the pre-specified party-count set (11 Democratic, 12 Republican); 17 were issue/ news/electoral/mixed/unclear
Issue/unclear/nonselected Politics topics	topic	626	Politics-category topics outside the ten pre-specified partisan topics were excluded from party counts
Issue/unclear/nonselected Politics impressions	impression	18218	These impressions define the likely conservative undercount of political and party-valenced volume
Platform-specific label coverage diagnostic	platform	4	Selected partisan topics account for 13.3% of Politics-category impressions on Facebook, 18.7% on Instagram, 2.6% on TikTok, and 11.5% on X; remaining Politics-category impressions are excluded from party counts rather than forced into Democratic/Republican labels.

Semantic, Platform, and Ad-Type Networks

Semantic networks used political-ad metadata and balanced platform ties. Advertisement text was tokenized after removing English stopwords and common ad terms (e.g., "click," "shop," "sponsored"), retaining unique words with at least three characters per document. Word-platform associations were ranked by conditional probability, with raw co-occurrence counts breaking ties. Candidate words were expanded from the top 50 terms until each platform contributed 25 associations. In Appendix A, platform nodes identify services, word nodes are colored blue for Democratic and red for Republican, device type nodes are orange, ad-type nodes are purple, and edge thickness reflects raw co-occurrence frequency.

Variables

Political advertising counts (party-specific exposure variables)

Party-specific exposure variables. Each qualified ad impression was assigned to one party only. Ads were included when sponsor/paid-for fields and topic labels explicitly indicated Democratic or Republican candidates, committees, causes, or attacks on the other party. Issue-only political ads, non-U.S. political ads, and ads whose partisan alignment could not be verified were excluded from both party counts. Respondent-level counts sum qualified impressions across Facebook, Instagram, X, and TikTok; repeated exposures to the same creative count as distinct impressions because exposure, not unique creative reach, is the estimand.

total_dem_ads - Democratic-aligned impressions: ads explicitly advocating for Democratic candidates, committees, or causes, or attacking Republican candidates.

total_rep_ads - Republican-aligned impressions: ads explicitly advocating for Republican candidates, committees, or causes, or attacking Democratic candidates.

Total ad exposure. In this analysis, total_impressions_z is the respondent's total number of observed ad impressions across platforms, standardized as a z-score. APOL is a survey item about attention to American politics and is not used as the ad-opportunity control.

Covariates. The primary models adjust for political-news interest, voter registration, age, gender, education, race/ethnicity indicators, family income, urbanicity, homeownership, and platform coverage indicators. Missing or don't-know responses for income, urbanicity, news interest, and voter registration are represented with separate indicators rather than silently dropped.

Table A3. Variable dictionary and coding choices.

Variable	Role	Definition
pid7_num_z	Outcome	Seven-point PID coded 1 = strong Democrat through 7 = strong Republican, then z-scored.
total_dem_ads	Exposure count	Respondent-level count of impressions from explicitly Democratic candidate, committee, or anti-Republican ads.
total_rep_ads	Exposure count	Respondent-level count of impressions from explicitly Republican candidate, committee, or anti-Democratic ads.
Balance	Composition	$\ln((\text{total_dem_ads} + 0.5) / (\text{total_rep_ads} + 0.5))$; positive values indicate a more Democratic ad mix.
balance_z	Composition	Balance z-scored across respondents.
partisan_vol	Volume/intensity	$\log(\text{total_dem_ads} + \text{total_rep_ads} + 1)$.
partisan_vol_z	Volume/intensity	partisan_vol z-scored across respondents.
total_impressions_z	Ad opportunity	Total ad impressions across all observed platforms, z-scored; this is not APOL.
dem_share	Companion descriptive	$(\text{total_dem_ads} + 0.5) / (\text{total_dem_ads} + \text{total_rep_ads} + 1)$, defined for respondents with any partisan exposure.
fb_any / ig_any / x_any / tt_any	Platform coverage	Binary indicators that the respondent had at least one observed impression on each service.
political_ad_load	Tail metric	explicitly partisan political impressions / total observed impressions at the respondent or respondent-platform level

Variable	Role	Definition
age_z	Covariate	2025 - birth year, z-scored.
female	Covariate	1 = female, 0 = male.
education_6_z	Covariate	Six-category education scale, z-scored.
race_white / race_black / race_hispanic	Covariate	Race/ethnicity indicators.
faminc_ord_z	Covariate	Ordinal income bracket, z-scored; prefer-not-to-say marked separately.
urbanicity_ord_z	Covariate	Rural-to-big-city ordinal measure, z-scored; missing urbanicity marked separately.
ownhome_bin	Covariate	1 = owns home, 0 = rents/other.
newsint_num_z	Covariate	Political-news interest from hardly at all to most of the time, z-scored; don't-know marked separately.
votereg_bin	Covariate	1 = registered to vote, 0 = not registered/don't know; don't-know marked separately.

Analytical Strategy

Estimand and temporal design. The primary estimand is the association between pre-survey partisan ad mix and party identification. Each respondent's exposure measures use impressions delivered in the 30 days before their PID survey date.

Exposure construction. The composition metric is $\text{Balance}_i = \ln((\text{total_dem_ads}_i + 0.5) / (\text{total_rep_ads}_i + 0.5))$, z-scored across respondents. Positive Balance values indicate a more Democratic ad mix; because pid7_num_z runs from Democratic to Republican, negative balance coefficients mean that a more Democratic ad mix is associated with more Democratic identification. The intensity metric is $\text{PartisanVol}_i = \log(\text{total_dem_ads}_i + \text{total_rep_ads}_i + 1)$, also z-scored. For example, 8 Democratic and 2 Republican partisan impressions yield $\text{Balance} = \ln(8.5/2.5) = 1.22$ before standardization; 2 Democratic and 8 Republican impressions yield $\text{Balance} = -1.22$.

Primary model and units. The H1 model is respondent-level: $\text{pid7_num_z}_i = \alpha + f(\text{balance_z}_i) + \beta_1 \text{partisan_vol_z}_i + \beta_2 \text{total_impressions_z}_i + \gamma \text{'X}_i + \pi \text{'Platform}_i + \epsilon_i$, where $f(\cdot)$ is a natural cubic spline. Platform variables are respondent-level coverage indicators, not repeated platform rows. RQ1 uses the same respondent-level sample and replaces the spline with separate standardized Democratic and Republican exposure counts. RQ2 is descriptive and uses respondent-platform cells, with explicitly partisan political-ad load defined as explicitly partisan political impressions divided by all observed impressions for that respondent on that platform.

Inference uses HC3 heteroskedasticity-consistent standard errors. The nonlinear balance term is summarized with a joint Wald test rather than separate interpretations of each spline basis coefficient. The RQ1 model replaces Balance with standardized Democratic and Republican exposure counts.

Respondents with no explicitly partisan impressions remain in the models with zeros in the party-specific counts. The 0.5 pseudocount avoids undefined log ratios when one party count is zero. Because H1 and RQ1 aggregate impressions to one respondent-level record, respondent-clustered standard errors are not required; clustering would matter only for repeated respondent-level observations.

Results

H1a & H1b

The respondent-level spline model (Table 1) explained modest variance in standardized party identification, $R^2 = .238$, adjusted $R^2 = .185$, $F(24, 347) = 7.43$, $p < .001$; AIC = 1004.50; BIC = 1102.47. The balance spline was jointly significant under HC3 inference, Wald chi-square(4) = 62.14, $p < .001$. Partisan volume was not statistically distinguishable from zero in the same model ($b = -0.099$, $SE = 0.121$, $p = .413$), supporting the composition-first interpretation.

Table 1. Respondent-level spline model predicting standardized PID.

Variable	Coefficient	HC3 SE	t	p	95% CI
Balance spline 1	1.328	0.383	3.472	< .001	[0.578, 2.078]
Balance spline 2	0.377	0.265	1.421	.155	[-0.143, 0.897]
Balance spline 3	-0.631	0.237	-2.662	.008	[-1.095, -0.166]
Balance spline 4	-0.740	0.423	-1.748	.081	[-1.569, 0.090]
Partisan volume	-0.099	0.121	-0.819	.413	[-0.337, 0.138]
Total impressions	-0.032	0.071	-0.450	.653	[-0.170, 0.107]
Facebook coverage	0.035	0.185	0.188	.851	[-0.328, 0.397]
Instagram coverage	-0.107	0.105	-1.016	.310	[-0.314, 0.100]
X coverage	0.181	0.129	1.402	.161	[-0.072, 0.435]
TikTok coverage	0.111	0.111	0.999	.318	[-0.107, 0.329]
Age	0.136	0.054	2.500	.012	[0.029, 0.242]
Female	-0.012	0.110	-0.109	.913	[-0.229, 0.204]
Education	-0.045	0.054	-0.841	.401	[-0.151, 0.060]
White	0.003	0.170	0.020	.984	[-0.330, 0.337]
Black	-0.699	0.196	-3.563	< .001	[-1.084, -0.315]
Hispanic	-0.272	0.190	-1.431	.153	[-0.645, 0.101]
Income	0.037	0.059	0.627	.530	[-0.079, 0.153]
Urbanicity	-0.035	0.055	-0.641	.522	[-0.143, 0.073]
Home owner	0.080	0.114	0.709	.479	[-0.142, 0.303]
News interest	-0.058	0.052	-1.116	.264	[-0.159, 0.044]
Registered voter	-0.283	0.136	-2.082	.037	[-0.549, -0.017]

Table A4. Model fit and core tests.

Model	N	df	R2	Adj. R2	F/Wald	p
H1 spline	372	24, 347	0.238	0.185	F = 7.426; balance Wald = 62.136	< .001
RQ1 separate party exposure	372	22, 349	0.182	0.131	F = 4.141; symmetry Wald = 0.073	.788

The individual spline coefficients are basis-function terms, not separate linear slopes. The joint test carries the substantive result: partisan ad balance is associated with party identification after adjustment for total ad impressions, partisan volume, platform coverage, and covariates. The model is respondent-level throughout, so the degrees of freedom correspond to respondents rather than respondent-platform cells.

Sensitivity checks tested whether the separate-party exposure pattern depended on a small set of labels. A Monte Carlo label-error check flipped or dropped 5%, 10%, 15%, and 20% of explicitly partisan labels and re-estimated the separate-party exposure model. Republican exposure stayed positive in all scenarios; Democratic exposure

stayed negative but less precise. The sign pattern is not an artifact of a small number of labels, but Democratic exposure should be interpreted cautiously.

The main measurement sensitivity is the narrow party-count definition. Party-specific counts use only the ten pre-specified, author-verified partisan topics. The label audit found 18,218 Politics-category impressions outside those topics, and the high-impression audit showed that the excluded set includes issue/news/electoral material and additional likely party-valenced topics. H1/RQ1 therefore measure explicitly partisan exposure from a narrow verified set; total political and party-valenced volume are conservative lower bounds.

RQ1

RQ1 used the same respondent-level analytic sample ($N = 372$). The separate-party exposure model was significant, $R^2 = .182$, adjusted $R^2 = .131$, $F(22, 349) = 4.14$, $p < .001$. Democratic exposure was negative but not statistically reliable ($b = -0.146$, $SE = 0.114$, $p = .197$); Republican exposure was positive and statistically reliable ($b = 0.113$, $SE = 0.053$, $p = .033$). The symmetry test did not reject equal-and-opposite magnitudes, Wald chi-square(1) = 0.073, $p = .788$.

RQ1 shows expected-direction point estimates but no evidence that one party's association is larger. Only the Republican coefficient is individually statistically reliable, and the formal symmetry test provides no evidence of asymmetric magnitude.

RQ2

Across respondent-platform cells, explicitly partisan political-ad load was low for most respondents and right-skewed. At the 10% threshold, three Facebook cells, five Instagram cells, one X cell, and no TikTok cells exceeded the threshold. Instagram mean-share intervals are wide because there are fewer Instagram cells and a few high-share cases. Threshold checks at 5%, 10%, and 15% show the same pattern, but denominator restrictions matter: with at least 50 total ads, Instagram falls from five high-load cells to one and X falls to zero; with at least 100 total ads, only the three Facebook cells remain.

Table 2. Platform summary for explicitly partisan political-ad load.

Platform	N cells	Mean share	Bootstrap 95%			
			CI	Cells > 10%	P95	Max
Facebook	333	0.005	[0.003, 0.007]	3	0.026	0.204
Instagram	133	0.013	[0.005, 0.024]	5	0.058	0.382
X	108	0.009	[0.005, 0.014]	1	0.044	0.167
TikTok	115	0.000	[0.000, 0.000]	0	0.000	0.003

Table 3. Denominator-threshold sensitivity for cells above 10% explicitly partisan political-ad load.

Minimum total ads	Facebook > 10%	Instagram > 10%	X > 10%	TikTok > 10%	Rank by mean share
0	3	5	1	0	Instagram=1, X=2, Facebook=3, TikTok=4

Minimum total ads	Facebook > 10%	Instagram > 10%	X > 10%	TikTok > 10%	Rank by mean share
50	3	1	0	0	X=1, Instagram=2, Facebook=3, TikTok=4
100	3	0	0	0	X=1, Instagram=2, Facebook=3, TikTok=4

Note. Explicitly partisan political-ad load is the share of explicitly partisan political impressions among all observed ad impressions for a respondent on that platform. Counts reflect respondent-platform cells whose explicitly partisan political-ad load exceeded 0.10 (10%) on that platform.

Few respondent-platform cells exceeded a 10% explicitly partisan political-ad load. Denominator checks make platform rankings descriptive, not evidence of platform mechanisms.

Discussion

The results describe explicitly partisan political-ad exposure in everyday mobile social media use. Each respondent contributed the 30 days before their party-identification survey; pooled impressions spanned October 4 through November 5, 2024 (UTC). Explicitly partisan political-ad shares were small for most respondents and concentrated in a thin upper tail. Party identification was associated with partisan exposure composition more than total partisan volume. These are delivered-exposure patterns, not persuasion effects or causal delivery mechanisms.

The findings support a composition-first description of digital political-ad delivery. The observed associations do not identify whether campaigns, algorithms, or users produced particular feeds. They do show that partisan mix is more informative than the sheer number of explicitly partisan political impressions.

The balance-identification association is nonlinear. The spline pattern is consistent with stronger alignment in some portions of the PID scale than others, but the data do not identify the mechanism. Selective exposure, auction dynamics, and campaign segmentation remain plausible mechanisms for future work to isolate. These are delivery-side interpretations, not persuasion claims; temporal ordering helps but does not make the models causal.

The formal test provides no evidence that Democratic and Republican exposure associations differ in magnitude. Point estimates run in the expected directions: Democratic exposure is negative on the Republican-coded PID scale and Republican exposure is positive, although only the Republican coefficient is individually statistically reliable. Cost asymmetries in prior audits can coexist with this respondent-level symmetry because differential prices need not produce asymmetric associations with identification once delivery and demand are observed as realized impressions (Ali et al., 2021; Bär et al., 2024).

Explicitly partisan political-ad load is sparse but uneven. Median shares are essentially zero, yet a small number of respondent-platform cells exceed ten percent. Because those

cells are sparse and denominator-sensitive, the results use bootstrap intervals, threshold checks, and denominator restrictions instead of a single high-load cutoff.

The null result for overall partisan volume separates opportunity from partisan alignment. Overall ad activity and platform coverage adjust for opportunity; the partisan mix of explicitly partisan impressions carries more information about party identification than the count of those impressions. Oversight should therefore report composition, not only volume. Advertiser libraries and platform spending aggregates capture supply; metered delivery captures realized audience mix (Edelson et al., 2018; Mehta & Erickson, 2022; Fowler et al., 2021, 2023).

Three implications follow. First, microtargeting theory should treat delivered composition as a core parameter. Second, evaluation should report tail-aware statistics, such as the 95th percentile of per-respondent political share by platform and week, alongside means. Third, transparency regimes that disclose only spend or creative miss the realized mix of exposures. None of these implications presumes persuasion effects; they concern the structure of exposure.

Limitations

The limits are clear. The data are mobile and feed-oriented, topic labels remain imperfect despite validation, and the models describe associations rather than causal effects. Stronger designs should instrument auction or policy changes, meter across devices and placements, and pair delivery data with advertiser-side objectives (Harrell, 2015; MacKinnon & White, 1985; Ali et al., 2021; Bär et al., 2024). Within these limits, the core pattern remains: explicitly partisan political advertising is sparse for most respondents, concentrated for a few, and more politically sorted than saturated.

Generalizability is bounded. Delivery is observed for one U.S. general-election season, on smartphones and tablets, among a YouGov-recruited metered panel. Facebook dominates observed impressions, with fewer impressions on Instagram and TikTok. The estimates should not be projected to desktop use, other countries, off-cycle periods, or platforms not instrumented here. The 10% ad-load cases are few and denominator-sensitive. Broad scarcity and narrow intensity can coexist, but their exact frequencies will vary with context and coverage.

The measures capture impressions, not persuasion or unique reach. Party-specific measures are conservative: Politics-category topics outside the ten pre-specified Democratic and Republican topics are excluded even when some high-impression nonselected topics appear party-valenced. This reduces partisan-direction error but likely understates total political and party-valenced volume.

The exposure-identification association remains descriptive. Reverse causation and unobserved confounding remain plausible: Democratic and Republican respondents may seek, attract, or be delivered congenial ads for reasons not captured by the covariates.

Instrumentation is uneven across ad surfaces. The meter records in-feed mobile ads, but capture may vary across formats and placements. The data exclude organic political content, cross-device exposures, and off-platform channels such as connected TV. Feed-only measures may therefore understate total campaign pressure and miss substitution across channels.

Further Research

Two research programs follow. First, establish causality and calibrate salience. High-frequency panels that meter exposure across devices while repeatedly surveying identification, attitudes, donations, and signup behavior could model within-person change as ad mix fluctuates. Field or quasi-experimental designs using policy or auction changes could instrument shifts in partisan mix. Perception studies should test when a 10% political-ad share becomes noticeable or memorable.

Second, identify mechanisms and tail dynamics. Audits that pair metered delivery with advertiser-side metadata could separate campaign choices, auction dynamics, and engagement optimization. Tail metrics such as the 95th percentile of per-respondent political share by platform and week should accompany averages. Researchers should model entry, persistence, and exit from high-frequency political tracks.

References

- Aggarwal, M., Allen, J., Coppock, A., Frankowski, D., Messing, S., Zhang, K., Barnes, J., Beasley, A., Hantman, H., & Zheng, S. (2023). A 2 million-person, campaign-wide field experiment shows how digital advertising affects voter turnout. *Nature Human Behaviour*, 7(3), 332-341. <https://doi.org/10.1038/s41562-022-01487-4>
- Ali, M., Sapiezynski, P., Korolova, A., Mislove, A., & Rieke, A. (2021). Ad delivery algorithms: The hidden arbiters of political messaging. *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 13-21. <https://doi.org/10.1145/3437963.3441801>
- Allcott, H., Gentzkow, M., Levy, R., Crespo-Tenorio, A., Dumas, N., Mason, W., Moehler, D., Barbera, P., Brown, T. W., Cisneros, J. C., Dimmery, D., Freelon, D., González-Bailón, S., Guess, A. M., Kim, Y. M., Lazer, D., Malhotra, N., Nair-Desai, S., Nyhan, B., ... Tucker, J. A. (2026). The effects of political advertising on Facebook and Instagram before the 2020 US election. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-025-02328-w>
- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130-1132. <https://doi.org/10.1126/science.aaa1160>
- Bär, D., Pierri, F., De Francisci Morales, G., & Feuerriegel, S. (2024). Systematic discrepancies in the delivery of political ads on Facebook and Instagram. *PNAS Nexus*, 3(7), pgae247. <https://doi.org/10.1093/pnasnexus/pgae247>
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732. <https://doi.org/10.15779/Z38BG31>
- Barrett, B., & McGregor, S. C. (2025). No better than soup? Comparing null experimental effects of political Facebook ads across persuasive and instrumental measures of effectiveness. *Social Media + Society*, 11(1), 20563051251316117. <https://doi.org/10.1177/20563051251316117>
- Bode, L., & Dalrymple, K. E. (2016). Politics in 140 characters or less: Campaign communication, network interaction, and political participation on Twitter. *Journal of Political Marketing*, 15(4), 311-332. <https://doi.org/10.1080/15377857.2014.959686>
- Coppock, A., Hill, S. J., & Vavreck, L. (2020). The small effects of political advertising are small regardless of context, message, sender, or receiver: Evidence

from 59 real-time randomized experiments. *Science Advances*, 6(36), eabc4046. <https://doi.org/10.1126/sciadv.abc4046>

Cotter, K., Medeiros, M., Pak, C., & Thorson, K. (2021). "Reach the right people": The politics of "interests" in Facebook's classification system for ad targeting. *Big Data & Society*, 8(1), 1-14. <https://doi.org/10.1177/2053951721996046>

Datta, A., Tschantz, M. C., & Datta, A. (2015). Automated experiments on ad privacy settings. *Proceedings on Privacy Enhancing Technologies*, 2015(1), 92-112. <https://doi.org/10.1515/popets-2015-0007>

DeLong, L. A. (2021, August 21). Facebook disables Ad Observatory; academicians and journalists fire back. NYU Center for Cyber Security. <https://cyber.nyu.edu/2021/08/21/facebook-disables-ad-observatory-academicians-and-journalists-fire-back/>

Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1), 242-259. <https://doi.org/10.1257/aer.97.1.242>

Edelson, L., Sakhuja, S., & McCoy, D. (2018). An analysis of Facebook's archive of ads with political content. New York University. <https://online-pol-ads.github.io/FBPoliticalAds/>

Fowler, G. A. (2020, October 27). How politicians target you: 3,000 data points on every voter, including your phone number. *The Washington Post*. <https://www.washingtonpost.com/technology/2020/10/27/political-campaign-data-targeting/>

Fowler, E. F., Franz, M. M., & Ridout, T. N. (2018). *Political advertising in the United States*. Routledge.

Fowler, E. F., Franz, M. M., Martin, G. J., Peskowitz, Z., & Ridout, T. N. (2021). Political advertising online and offline. *American Political Science Review*, 115(1), 130-149. <https://doi.org/10.1017/S0003055420000696>

Fowler, E. F., Franz, M. M., Neumann, M., Ridout, T. N., & Yao, J. (2023). Digital advertising in the 2022 midterms. *The Forum*, 21(1), 53-73. <https://doi.org/10.1515/for-2023-2006>

Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. presidential election. *Science*, 363(6425), 374-378. <https://doi.org/10.1126/science.aau2706>

Grootendorst, M. (2022). BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv*. <https://doi.org/10.48550/arXiv.2203.05794>

Harrell, F. E., Jr. (2015). *Regression modeling strategies: With applications to linear models, logistic and ordinal regression, and survival analysis* (2nd ed.). Springer. <https://doi.org/10.1007/978-3-319-19425-7>

Hersh, E. D. (2015). *Hacking the electorate: How campaigns perceive voters*. Cambridge University Press.

Hillygus, D. S., & Shields, T. G. (2008). *The persuadable voter: Wedge issues, candidate strategy, and campaign effects*. Princeton University Press.

Kalla, J., & Broockman, D. (2018). The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments. *American Political Science Review*, 112(1), 148-166. <https://doi.org/10.1017/S0003055417000363>

- Lambrecht, A., & Tucker, C. (2019). Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Science*, 65(7), 2966-2981. <https://doi.org/10.1287/mnsc.2018.3093>
- MacKinnon, J. G., & White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29(3), 305-325. [https://doi.org/10.1016/0304-4076\(85\)90158-7](https://doi.org/10.1016/0304-4076(85)90158-7)
- Mehta, S., & Erickson, K. (2022). Can online political targeting be rendered transparent? Prospects for campaign oversight using the Facebook Ad Library. *Internet Policy Review*, 11(1), 1-31. <https://doi.org/10.14763/2022.1.1648>
- Merrill, J. B., & Tobin, A. (2019, January 28). Facebook moves to block ad transparency tools-including ours. *ProPublica*. <https://www.propublica.org/article/facebook-blocks-ad-transparency-tools>
- Messing, S., & Westwood, S. J. (2014). Selective exposure in the age of social media: Endorsements reduce news selectivity. *Communication Research*, 41(8), 1042-1063. <https://doi.org/10.1177/0093650212466406>
- Nickerson, D. W., & Rogers, T. (2014). Political campaigns and big data. *Journal of Economic Perspectives*, 28(2), 51-74. <https://doi.org/10.1257/jep.28.2.51>
- Qwen, Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Li, C., Liu, D., Huang, F., Wei, H., Lin, H., Yang, J., Tu, J., Zhang, J., Yang, J., Yang, J., Zhou, J., Lin, J., ... Qiu, Z. (2024). Qwen2.5 technical report. *arXiv*. <https://doi.org/10.48550/arXiv.2412.15115>
- Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3982-3992. <https://doi.org/10.18653/v1/D19-1410>
- Ridout, T. N., Fowler, E. F., & Franz, M. M. (2021). The influence of goals and timing: How campaigns deploy ads on Facebook. *Journal of Information Technology & Politics*, 18(2), 133-149. <https://doi.org/10.1080/19331681.2021.1874585>
- Ridout, T. N., Franklin Fowler, E., Franz, M. M., & Goldstein, K. (2018). The long-term and geographically constrained effects of campaign advertising on political polarization and sorting. *American Politics Research*, 46(1), 3-25. <https://doi.org/10.1177/1532673X17721479>
- Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20, 53-65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7)
- Silva, M., Santos de Oliveira, L., Andreou, A., Vaz de Melo, P. O. S., Goga, O., & Benevenuto, F. (2020). Facebook Ads Monitor: An independent auditing system for political ads on Facebook. *Proceedings of The Web Conference 2020*, 224-234. <https://doi.org/10.1145/3366423.3380109>
- Stroud, N. J. (2011). *Niche news: The politics of news choice*. Oxford University Press.
- U.S. Census Bureau. (2025). American Community Survey 2024 one-year estimates data profiles [Data set]. <https://api.census.gov/data/2024/acs/acs1/profile>
- Varian, H. R. (2007). Position auctions. *International Journal of Industrial Organization*, 25(6), 1163-1178. <https://doi.org/10.1016/j.ijindorg.2006.10.002>

Appendix A. Semantic, Platform, and Ad-Type Networks

Network interpretation: Node size reflects word frequency; edge thickness reflects association strength. Platform nodes identify services; word nodes show characteristic terms. Spatial proximity is layout-driven, not semantic distance.




Figure 1. Relationship between explicitly partisan ad balance and partisan identification.




Figure 2. Explicitly partisan political-ad share by platform, with respondent-platform cells above 10% highlighted.




Figure 3. Semantic, platform, and ad-type networks.