Allocating Influence: Government Advertising in Mexico's News Market

Camila Blanes*

August 20, 2026

Abstract

Governments, whether autocracies or democracies, seek to influence the media. Autocracies use strategies such as propaganda and censorship, but these tools are limited in democracies. In this paper I show how governments can still influence the media through political favoritism. By reallocating advertising toward aligned outlets, a government enables them to invest in their newsrooms and raise their news production. Combining the universe of federal advertising contracts in Mexico with the full text of national print media, I use the 2018 transition from a right-wing to a left-wing administration and estimate a difference-in-differences specification. Outlets that share the incumbent's ideology become more likely to receive a contract. These allocations reflect favoritism and not the audiences a government would need to reach, whether the demographic a campaign names or the incumbent's own voters. Aligned outlets then hire more journalists, publish more articles, cover more news events, and report more facts in each. Their circulation also rises relative to misaligned outlets. Favoritism toward competing firms therefore reallocates productive capacity across a market instead of transferring rents within it, a mechanism of political control that operates through market competition.

Keywords: Government advertising; Media capture; Political favoritism; Allocation efficiency; Ideological alignment
JEL Codes: D72, L82, M37, D73

*Ph.D. student in Political Economics, Stanford Graduate School of Business. Email: cblanes@stanford.edu. Address: 655 Knight Way Stanford, CA 94305. I thank participants at the XXIII RIDGE/LACEA Workshop on Political Economy (2026) and the 9th Economics of Media Bias Workshop (2026) for helpful comments and suggestions.

1 Introduction

Government advertising is a major source of revenue for media outlets in many developing and developed countries, particularly where private advertising markets are either limited or concentrated.[1] In such contexts, outlets may become financially dependent on public sector advertising. Much attention has focused on how this dependence might undermine editorial independence through quid pro quo arrangements, where outlets trade favorable coverage for advertising revenue (Di Tella and Franceschelli, 2011; Szeidl and Szucs, 2021).

However, in a competitive media market, such explicit arrangements are more costly and harder to sustain. When outlets compete for readers, any outlet that suppresses newsworthy information in exchange for government advertising risks losing its audience to rivals that cover it. This competitive discipline limits the scope for capture, because silencing information would require the government to negotiate with every outlet simultaneously, which is prohibitively expensive (Besley and Prat, 2006). I study a different mechanism through which the government can still shape the information environment in competitive markets. Rather than conditioning advertising on favorable coverage, the government uses its role as the dominant advertiser to selectively subsidize aligned outlets, allowing market competition to amplify the resulting asymmetry.

I focus on Mexico, where the federal government is one of the largest advertisers in a competitive newspaper market. While the stated purpose of government advertising is to promote public programs, highlight achievements, or disseminate prevention messages, its discretionary nature allows incumbents to reward supportive outlets or penalize critical ones. Such favoritism distorts media competition by directing resources toward ideologically aligned outlets. As a result, ideologically misaligned outlets are under-resourced, limiting their ability to expand their newsrooms or maintain coverage.

I first show that the government allocates advertising toward ideologically aligned outlets on the extensive and intensive margins. I implement a difference-in-differences specification that leverages the 2018 federal transition from a right-wing to a left-wing administration to compare changes in advertising spending between left- and right-wing outlets. I measure outlet ideology from text, using a penalized classifier trained on the editorials of outlets with established ideological positions to place all outlets on a left-right dimension. On the extensive margin, left-wing outlets become 18.3 percentage points more likely to receive a contract from a given ministry after the transition, a 33 percent increase over their pre-transition rate of 0.56. On the intensive margin, both groups lost revenue in absolute terms, because the incoming administration cut federal advertising across these ministries by 48 percent, but left-wing outlets receive $105,362 differentially per ministry-year compared to right-wing outlets. Spending on aligned outlets fell by 9 percent and spending on misaligned outlets by 72 percent, so the estimate measures how far alignment shielded an outlet from a contraction that fell almost entirely on its competitors.

[1] Countries where the government is a major advertiser in the market include Bulgaria, Greece, Hungary, Poland, Romania, Slovakia, Slovenia, Albania, Macedonia, Serbia, Turkey, Argentina, Brazil, Colombia, Mexico, Peru, India, Pakistan, Myanmar.

Next, I show that governments reallocate advertising toward ideologically aligned outlets, which reflects favoritism toward the outlet's ideology instead of the government's efforts to reach different demographic groups, like their voter base. One alternative explanation is that the reallocation reflects changes in the types of programs being advertised instead of ideology, a left-wing government may promote more social programs whose intended beneficiaries overlap with the readership of left-wing outlets, making the reallocation appear ideologically motivated even in the absence of favoritism. A second explanation is that the government prioritizes reaching its own voter base, which is disproportionately represented among readers of aligned outlets.

I use two complementary exercises to rule out both explanations. First, I restrict the analysis to advertising campaigns that are constant across federal administrations and whose content and demographic focus remain fixed. If allocations were driven by the desire to reach specific groups, outlet choice should have remained stable. Yet even within these bipartisan campaigns, aligned outlets are 11.7 percentage points more likely to receive advertising and receive higher spending, suggesting that favoritism also drives reallocations.

Second, I construct two counterfactual benchmarks. Each solves an advertiser's problem. It spreads a campaign's observed budget across newspapers at their market advertising prices and reaches as many distinct readers as it can. The first counts readers in the campaign's stated demographic, which the legal cost-efficiency standard obliges ministries to buy. The second counts the incumbent's predicted supporters, the objective an office-seeking politician would hold (Strömberg, 2004). The benchmarks also price the distortion in two units. Dollars measure how much of a campaign's budget alignment moves, and readers measure what the public forgoes, since these campaigns exist to inform citizens about public programs.

Aligned outlets receive $20,876 more per campaign than the legal cost-efficiency standard prescribes and $22,198 more than voter targeting prescribes, and under either benchmark they are roughly 14 percentage points more likely to receive any funding at all. Aligned outlets receive 21 percent of the average campaign budget above the cost-efficiency benchmark and 23 percent above the voter-targeting benchmark. Outlets closer to the incumbent's position receive larger amounts above either benchmark, so the gap widens with the alignment score. Readers bear the cost of this allocation, because campaigns inform between 23 and 36 percent of the audience the same budget could have reached.

Advertising revenue moves toward aligned outlets and away from misaligned ones, so each outlet can afford to produce more news or less. I measure production capacity by the size of the newsroom and by the number of articles an outlet publishes, and I estimate both with the difference-in-differences design of equation 1. After the transition aligned outlets employ 60 percent more journalists than misaligned ones, roughly 84 at their pre-transition newsroom of 140, and they publish 55 percent more articles.

I then turn to three measures of quality, which record different dimensions of how well an outlet reports instead of how much it publishes. The first separates hard news from soft news. Aligned outlets raise hard news output by 66 percent and soft news by 54 percent, so the additional resources reach substantive

reporting instead of filler. The second counts the articles an outlet takes from news agencies. That share does not move, so outlets write the additional articles themselves, which suggests original reporting from the newspaper. The third measure captures outlets that covered the same news event. I group articles into events by textual similarity, and because every event appears in the panel for all 18 outlets, event fixed effects absorb the news itself and I compare differently aligned outlets that reported the same story. Aligned outlets cover 7.8 percentage points more of the events that occur under their own government, a 29 percent increase over their pre-transition rate of 0.27, and they contribute 37 percent more new information to the average event.

Finally, I examine whether the expansion in content production and quality translates into reader demand. Print circulation declined for both groups of outlets between 2014 and 2022, so the difference-in-differences identifies the divergence between their trajectories instead of an absolute increase for either. By the end of the sample, circulation at aligned outlets stands roughly 40 percent above misaligned outlets relative to their pre-transition difference, so alignment slowed the decline instead of producing growth, though the estimate carries a standard error of 0.236 and does not separate from zero at conventional levels. Two forces could generate that divergence. The first is a supply-side effect, in which more content of higher quality attracts new readers. The second is a demand-side effect, in which the 2018 election shifted readers' political preferences toward left-wing media independently of any change in content. To separate them I exploit cross-state variation in the leftward shift in Morena's vote share between 2015 and 2018 and estimate a triple difference-in-differences specification. The identifying assumption is that a state's change in vote share measures the change in its readers' political preferences. An outlet prints one national edition, so a content-driven gain should appear in every state it reaches, whereas a preference-driven gain should concentrate where the electoral shift was largest. The effect in states whose vote share moved least therefore isolates the content channel, and it accounts for roughly 77 percent of the average differential. The interaction that identifies the demand-side component is positive but imprecisely estimated, so the split should be read as suggestive and not as a precise decomposition.

This paper relates to three literatures. The first studies media capture. Besley and Prat (2006) characterize the conditions under which a government can suppress unfavorable information, and show that the cost of suppression rises with the number of independent outlets, because the government must secure the silence of every outlet able to break the story. Later work identifies the conditions that lower that cost of suppression. Suppression becomes cheaper when wealth concentrates sufficiently for a narrow elite to pay for silence (Petrova, 2008), when the government prefers to own an outlet instead of paying it (Gehlbach and Sonin, 2014), or when audience attention concentrates in a few outlets however many exist (Prat, 2018). Empirical work documents the instruments governments use to influence the informational environment, namely bribery (McMillan and Zoido, 2002), state ownership (Qin et al., 2018), and the exchange of advertising for favorable coverage (Di Tella and Franceschelli, 2011; Szeidl and Szucs, 2021). Private firms can also exert comparable pressure through their own advertising spending (DellaVigna et al., 2016; Beattie et al., 2021).

Each of these mechanisms pays an outlet to distort its coverage. I document a different mechanism that does not influence the bias directly, but affects the overall market conditions. When readers can switch outlets, distortion costs an outlet the audience that makes it worth subsidizing, so an aligned outlet earns the highest return by spending the transfer on journalism. Capture therefore operates through market competition instead of against it, and the government buys quality where it cannot buy bias.

The second literature studies the demand for aligned media or government-controlled news. Simonov and Rao (2022) show that Russian consumers read state-controlled outlets because those outlets carry information not available elsewhere, so demand follows the content supplied instead of the absence of alternatives. Bowles et al. (2026) find in a field experiment in Türkiye that pro-government outlets persuade durably while opposition outlets do not, and attribute the asymmetry to the resource constraints that keep opposition newsrooms from producing content able to hold an audience. Knight and Tribin (2019) and Durante and Knight (2012) show that audiences abandon programming that turns transparently partisan, so a government that wants to be read has to fund something readers value. This work establishes that in some occasions government-aligned outlets can be better resourced and more persuasive, and takes that advantage as given. In this paper, the government endogenizes that decision, estimating which newsrooms the government finances and what it buys by doing so.

The third literature studies political favoritism in the allocation of public resources. Padró i Miquel (2007) shows that a leader can direct resources toward their own group and retain its support, because that group fears the alternative leadership more than it resents the leader's own predation. Empirically, favoritism is more pronounced where political institutions constrain leaders least, whether that constraint is measured by regime type (Burgess et al., 2015) or by the quality of governance (Hodler and Raschky, 2014). It also operates in both directions. Brollo and Nannicini (2012) compare Brazilian municipalities in which an opposition mayor won narrowly with those in which an aligned mayor did, and find that the former receive smaller discretionary transfers, so incumbents withhold resources from opponents as well as direct them toward allies. In each of these settings the recipients do not compete with one another, favoritism redistributes between recipients and its cost is borne by whoever was passed over. In this paper, I show that favoritism toward competing firms reallocates productive capacity across a market instead of transferring rents within it.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting for federal advertising in Mexico. Section 3 presents the data and the measurement of outlet ideology. Section 4 documents ideological favoritism in allocation. Section 5 constructs the cost-efficiency and voter-targeting benchmarks. Section 6 traces the downstream effects on newsroom capacity, content production, and circulation. Section 7 concludes.

2 Institutional Background

The Mexican government has long held economic leverage over the press and has changed only the instrument it uses. From 1935 to 1990 it supplied the press with newsprint through PIPSA, a state-owned company that produced and imported most of the paper newspapers required, and it withheld or degraded deliveries to outlets that criticized it. Cash payments to journalists ran alongside PIPSA throughout the PRI era. The government privatized PIPSA in 1998, and the democratic transition of 2000 raised the political cost of both instruments, so advertising contracts replaced them. Contracts gave outlets a revenue stream they could forecast, and they gave the government a decision it could make each year over which outlets to fund.

Federal ministries allocate advertising contracts through annual campaign programs. For each campaign a ministry states a communication objective, such as promoting a social program or informing citizens about a public service, and a target demographic defined by age, income, and education. Mexican law then obliges the ministry to buy the placement that reaches that demographic at least cost, a rule I refer to throughout as the legal cost-efficiency standard. Outlets seeking federal contracts must register with the National Print Media Registry, which records each outlet's circulation, readership demographics, and advertising prices, so the registry supplies every input the standard requires. No independent body audits allocation decisions, and a ministry that departs from the standard faces no penalty. Ministries submit annual advertising budgets to Congress for approval. These rules govern communication campaigns, which account for 95 percent of federal advertising expenditure. General announcements, which cover routine operational notices, name no audience and face no oversight.

Mexico's print media market expanded and fragmented alongside the democratic opening of the 1990s. As single-party dominance eroded, new newspapers entered national circulation, and by the 2000 transition more than a dozen daily titles competed for the same readers and advertisers. Private advertising grew alongside circulation but never grew enough for commercial revenue to substitute for a federal contract. Concentration stays low across the sample period. In every year, half of total circulation comes from at least 39 percent of the outlets.[2]

3 Data and Measurement

3.1 Data Sources

I link three datasets, government advertising contracts from federal administrative records, outlet-level variables from the National Print Media Registry, and newspaper content collected from printed editions.

Government advertising contracts. I use administrative records from the Ministry of the Interior's Social Communication Strategy program, which cover every federal advertising contract signed between

[2]The Herfindahl-Hirschman Index averages 679, below the 1,500 mark at which US merger guidelines classify a market as unconcentrated.

2014 and 2022. Each contract identifies the ministry, the recipient outlet, the date, the amount spent, the number of advertisements by size, and the price paid per advertisement. I complement these records by collecting a novel dataset on campaign characteristics, which reports the communication objective, the target demographic, and whether the campaign promotes government policies and achievements, social programs, prevention and awareness, or civic and cultural obligations.[3]

Outlet-level variables. I construct a novel panel dataset on time-varying outlet characteristics from the National Print Media Registry, which every outlet seeking federal advertising must update each year. For each outlet and year the panel reports average daily circulation by state, the share of readers by socioeconomic level, gender, education, and age group, the legal name of the owner, and list prices per cm^2 for advertising space across sections and page locations, which I average into a single market price per outlet. I restrict the sample to national outlets that publish news daily (or at least five days per week), dropping magazines and outlets focused only on entertainment or sports. This leaves approximately 25 national outlets per year. I further limit the analysis to 18 outlets for which complete data on daily articles and historical registry information are available. On average, these 18 outlets account for approximately 82% of national print circulation each year.

News content. I collect the text of all print issues for the 18 national outlets from 2014 to 2022, totaling approximately 6 million articles.[4] From the text I measure the number of articles an outlet publishes, its split between hard and soft news, and the news events on which I compare coverage across outlets. From bylines I measure newsroom size, and I separate wire copy (Reuters, Associated Press, Notimex) from house staff and named journalists so that the measure counts an outlet's own reporters. I classify every printed advertisement as public or private. The classification validates the advertising spending ministries report and measures the share of an outlet's advertising space that the government occupies against the share private advertisers buy.

Table 1 reports summary statistics for the full sample, separately for the right-wing (PRI) and left-wing (Morena) administrations. Government advertising falls from \$1.7 million per outlet under PRI to \$467,000 under Morena, a decline that, if distributed unevenly across outlet types, translates into revenue shocks that differ by outlet ideology. Even at these levels, government advertising accounts for roughly one-third of all advertising space in the average outlet, with a standard deviation of 17 percentage points indicating that some outlets depend far more heavily on public contracts than others. The same discretion that governs outlet selection shapes the prices ministries pay, the government pays approximately twice the market rate, \$21.93 per cm^2 under PRI against a market price of \$9.00, and \$13.41 under Morena against \$8.28. The average outlet holds a 5.6% market share, consistent with a market where no outlet dominates. Together, these patterns, a government advertiser that accounts for one-third of ad space, prices that exceed market rates, and allocation decisions that no independent body audits, describe a setting where government choices reshape the market.

[3] I obtain these characteristics through information requests to federal ministries under the National Transparency Platform.
[4] See Table B1 for per-outlet totals and yearly averages.

Table 1: **Summary Statistics**

	Right-Wing Government (PRI) (2012–2018)		Left-Wing Government (Morena) (2018–2024)	
	Mean	SD	Mean	SD
Panel A: Government Advertising				
Gov. advertising per outlet (USD, annual)	1,709,158	1,769,201	466,711	826,105
Price paid by government (USD per cm^2)	21.93	16.90	13.41	12.10
Share of public advertising (%)	34.1	17.0	33.8	16.9
Number of ministries contracting	11.0	5.9	6.0	3.8
Panel B: Outlet Characteristics				
Daily circulation (thousands)	115.1	56.6	115.7	57.8
Average market share (%)	5.7	2.8	5.6	2.7
Monthly articles	2,497	1,625	2,275	1,273
Newsroom size	198	160	204	139
Market advertising price (USD per cm^2)	7.83	4.16	7.52	3.42

Outlet-year observations. The sample covers 18 national newspapers from 2014 to 2022. USD values in 2022 constant dollars. Gov. advertising per outlet aggregates all federal ministry contracts per outlet and year. Share of public advertising is the fraction of observed advertisements classified as government-sourced from newspaper content. Market advertising price is the average list price per cm^2.

3.2 Measuring Ideological Alignment

Identifying ideological favoritism requires a measure of where each outlet stands on the left-right dimension. I construct one from text, following the logic of Gentzkow and Shapiro (2010) that the differential use of partisan language identifies ideological position, and using the penalized estimation to recover a sparse and interpretable set of partisan phrases (Gentzkow et al. (2019)). I estimate the measure using editorial board pieces since they reflect the outlet's institutional stance.

Fifteen of the eighteen outlets publish a dedicated editorial section, and I use the 8,938 editorials they published between 2014 and 2017. I end the estimation window before the transition for two reasons. The measure never sees vocabulary specific to the incoming administration, so the recovered dimension separates left from right and not support for the 2018 winner from opposition to it. Outlet ideology also becomes a pre-treatment characteristic, fixed before any of the advertising reallocation I study.[5]

Two outlets serve as left-wing anchors, La Jornada and La Prensa, and three as right-wing anchors, Reforma, El Universal, and El Economista. I select these five based on their editorial histories and established positions in the Mexican media landscape.[6] The estimation vocabulary is built from the 3,836 editorials these five outlets published in the window, subject to three restrictions, a phrase must appear in at least three of the five anchor outlets, so that no single outlet's recurring vocabulary can define the dimension; it must appear in at least five anchor editorials; and it must not contain an outlet name, a column title, or a formatting artifact.

I represent each anchor editorial as a vector of phrase proportions that sums to one, and I fit an L_1-penalized logistic regression of an indicator for right-wing anchor authorship on that vector. The penalty is selected by scanning sixty values of the regularization parameter and retaining the solution closest to

[5]Phrases are the unigrams and bigrams formed from the stemmed text. I lowercase each editorial, strip accents, digits, and punctuation, remove Spanish stopwords, lemmatize, and stem, dropping tokens shorter than three characters.

[6]Lawson and Media Ownership Monitor

a target of five hundred non-zero phrases, following the sparsity target in Gentzkow et al. (2019). The selected model retains 514 phrases, 224 of which load right-wing and 290 left-wing. Table 2 reports a subset chosen for interpretability. Phrases loading left-wing include pension, Cuba, promise, and countryman, while phrases loading right-wing include investment, institution, economic growth, and competence. The two sets divide social provision and national identity from markets and technocratic competence, which is the axis on which Mexican parties have divided since the 1990s. The two sets divide social provision and national identity from markets and technocratic competence, which is the axis on which Mexican parties have divided since the 1990s.

Table 2: **Top Partisan Phrases**

Right-wing Phrases			
American (estadounidens)	Capacity (capac)	Official (funcionari)	Local (local)
Competence (competent)	Union (sindicat)	Investment (inversion)	Structure (estructur)
Professor (profesor)	Institution (institucion)	Org. crime (crim organiz)	Internet (internet)
Transformation (transform)	Corruption (corrupcion)	Citizenship (ciudadani)	Process (proces)
State-run (estatal)	Service (servici)	Opportunity (oportun)	Infrastructure (infraestructur)
Freedom (libert)	Risk (riesg)	Institutional (institucional)	Poll (encuest)
Criminal (criminal)	Growth (crecimient)	Fiscal (fiscal)	
Left-wing Phrases			
Tax auth. (sat)	Injustice (injust)	Environmental (ambiental)	Hunger (hambr)
War (guerr)	Ministry (secretari)	Gas hike (gasolinaz)	Common sense (sent comun)
People (puebl)	Att. General (procurador)	Pollution (contamin)	Teacher (maestr)
Electoral roll (padron)	Promise (promes)	Countryman (paisan)	Disarmament (desarm)
Government (gobern)	Gov. official (mandatari)	Odebrecht (odebrecht)	Misery (miseri)
Ayotzinapa (ayotzinap)	Femicide (feminicidi)	Health (salud)	Electoral process (proces electoral)
Pension (pension)	Cuba (cub)	North American (norteamerican)	

The L_1-penalized logistic regression retains 514 phrases with non-zero loadings, 224 loading right-wing and 290 loading left-wing. This table reports 55 of them, selected for interpretability and not by the rank of the loading. Phrases are shown in English with the Spanish stem in parentheses.

Each article receives a score equal to the inner product of the estimated loadings with its own vector of phrase proportions, and each outlet's score is the mean of that quantity across its articles. I subtract the midpoint between the left- and right-wing anchor group means, so that zero falls equidistant between the two clusters, and I classify an outlet as left-wing when its recentered score is negative and right-wing otherwise. The five anchors are scored by the same rule as every other outlet and are not held at their training labels.

Diario de México, El Heraldo de México, and Esto publish no editorial section, so I score them by applying the estimated loadings to their news articles under the same preprocessing, scoring rule, and midpoint. Appendix C.1.1 resamples each of these three corpora down to the size of a typical editorial corpus and reports the resulting distribution of scores, which is stable in every draw.

The right panel of Figure 1 plots each outlet's recentered score, and this classification remains fixed throughout the analysis. Nine outlets fall on each side. Several score close to the midpoint, among them the four with the thinnest editorial corpora. Ovaciones, Publimetro, El Sol de México, and La Razón each contribute fewer than forty editorials, which motivates the robustness checks using a continuous score and

a three-category measure reported in Appendix C.1. The left panel applies the same phrase loadings to editorials published in each year from 2014 to 2022, for the fifteen outlets with editorial sections. The two groups remain clearly separated across the full period and their ordering never reverses, indicating that the classification estimated on the pre-transition window is not specific to that window.

Figure 1: **Outlet Ideology Scores and Stability (2014–2022)**

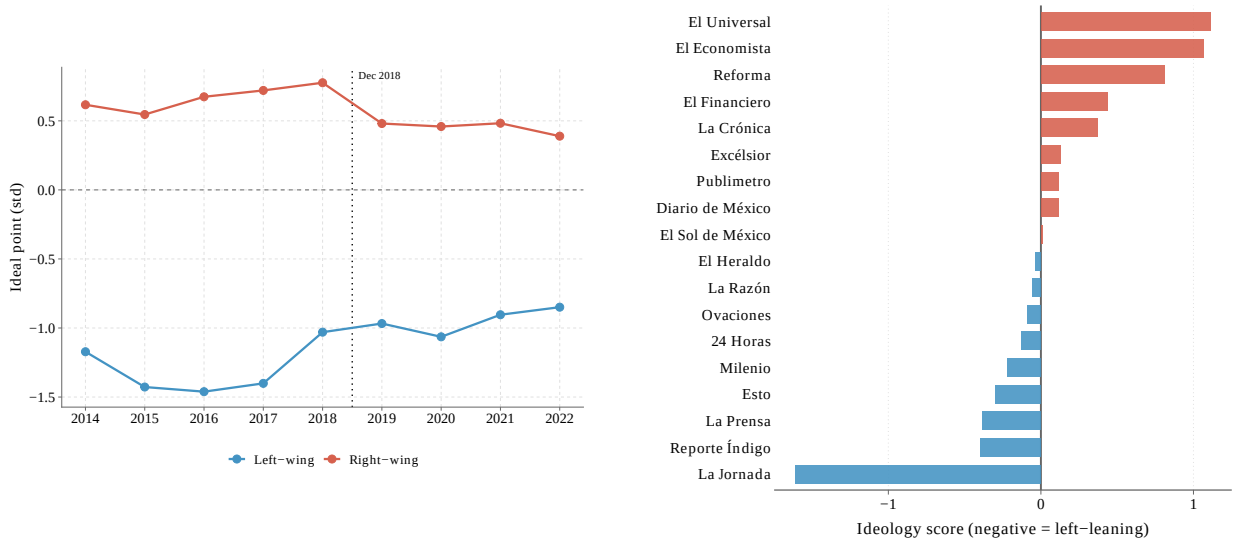

Phrase loadings are estimated by L_1-penalized logistic regression on the editorials of the five anchor outlets. The right panel plots the score of each of the 18 outlets, recentered so that zero lies equidistant between the left- and right-wing anchor group means. Negative values indicate left-leaning positions and positive values indicate right-leaning positions. The left panel applies the same loadings to editorials published in each year and plots the article-weighted group mean.

3.3 Motivating Patterns

Before turning to the causal analysis, I document descriptive patterns in government advertising allocation. Figure 2 plots the time series of government advertising expenditures across left- and right-wing newspapers among the 18 national circulating newspapers from the sample.

Figure 2: **Government Advertising Spending by Outlet Ideology (2014–2022)**

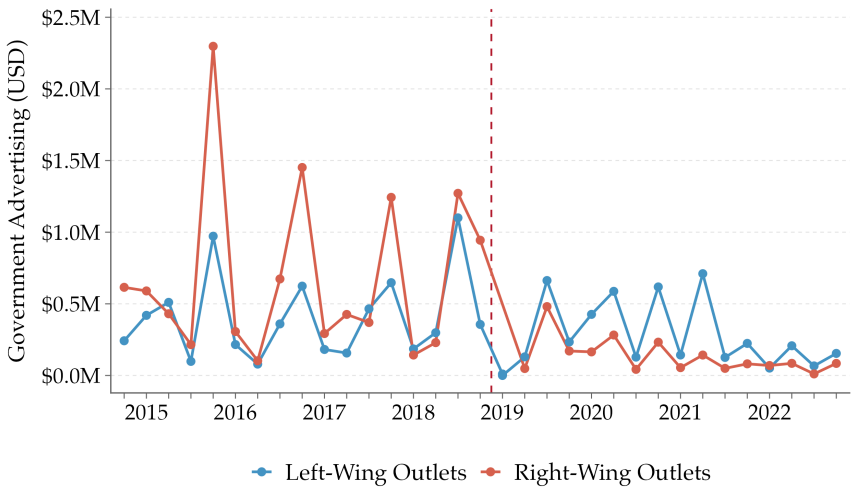

Average quarterly spending on ministry advertising campaigns per outlet, by outlet ideology (9 left-wing, 9 right-wing). Spending is deflated to 2022 pesos and converted to USD at the 2022 average exchange rate. The red dashed line marks the December 2018 presidential transition from PRI to Morena.

Under the right-wing PRI administration, right-wing outlets received more government advertising per quarter than left-wing outlets. Following the December 2018 transition to Morena, this pattern reverses. Spending shifts toward left-wing outlets while right-wing outlets see a decline. Immediately after the transition, total spending drops sharply across all outlets, consistent with the austerity the incoming government announced, before stabilizing at the lower levels shown in Table 1. Spending also peaks in the fourth quarter of every year. Congress appropriates advertising budgets annually and ministries forfeit whatever they have not committed by December, so they place the unspent balance before the fiscal year closes.

The raw patterns are consistent with ideological favoritism but do not establish it. Three explanations would produce the same reversal without any preference for aligned outlets. The incoming government has different policy priorities and therefore advertises different programs. Those programs target different audiences, and if left-wing outlets reach the audiences a left-wing government cares about, the reallocation reflects whom each campaign set out to inform. The government may also buy access to its own voters, who may concentrate among the readers of aligned outlets. Ministries themselves also changed, since the incoming administration advertised through a different set of them and ministries differ in which outlets they use. Sections 4 and 5 identify the effect of ideological alignment on allocation and address each of these explanations in turn.

4 Favoritism in Advertising Allocation

Governments may prefer to advertise in outlets that already provide favorable coverage, either because they want to reach supportive audiences or because of shared ideology. To test whether such affinity shapes allocation decisions, I examine whether changes in government lead to shifts in spending toward ideologically aligned outlets.

I employ a difference-in-differences specification that leverages variation in government transitions and the ideological alignment of newspapers. I focus on the presidential transition in 2018, where the left-wing party Morena came to power, replacing the right-wing party PRI. Newspaper ideological positions are fixed over time and inferred from the text-based measures described in Section 3.2. I estimate the following regression:

$$Y_{ngt} = \beta \text{ Left-Wing Gov}_t \times \text{ Left-Wing Outlet}_n + \mu_t + \delta_{ng} + \epsilon_{ngt}, \tag{1}$$

where Y_{ngt} represents one of two outcome variables of government advertising to outlet n from ministry g in year t. First, an indicator for whether the outlet receives any advertising contract (extensive margin). Second, the level of total advertising spending in 2022 USD, including zero values for outlet–ministry pairs with no contracts in a given year (intensive margin). The price paid, conditional on a contract, is a third margin the government could in principle use, but it is observed only where a contract was placed and alignment moves that margin, so it is not point identified by this design; Appendix E.1 reports it with

bounds. The variable Left-Wing Gov_t is an indicator equal to 1 when the presidency is held by a left-wing party, and Left-Wing Outlet_n equals 1 if the outlet is classified as left-leaning and 0 if right-leaning.

Without ministry fixed effects, the estimates could reflect cross-ministry variation in budget trajectories, a ministry that structurally favors left-wing outlets for non-ideological reasons (e.g., because its target audience skews left) may also have expanded its budget under the left-wing government, generating a spurious correlation. δ_{ng} removes this confound by identifying β from within-pair changes. Hence, the panel is structured at the outlet × ministry × year level. The specification also includes year fixed effects (μ_t) to absorb shocks common to all outlets in a given year, including the aggregate path of the federal communication budget. Standard errors are clustered at the newspaper level, as the treatment variable is defined at the newspaper level and errors across all ministry relationships of the same newspaper are correlated by construction. The identification strategy relies on the common trends assumption, in the absence of a shift in the party in power, advertising allocations to left- and right-leaning outlets would have followed common trends.

Table 3 presents the results. Columns (1)–(3) use the full sample of all advertising campaigns, restricted to the 10 ministries that ran campaigns across both administrations.[7] These columns report the extensive margin, the intensive margin in levels, and the inverse hyperbolic sine of spending.

Ideological alignment favors aligned outlets on both the extensive and the intensive margin. On the extensive margin, a left-wing outlet is 18.3 percentage points more likely to receive a contract from a given ministry once the left-wing government takes office, against a pre-transition rate of 0.56 for the same outlets, a 33 percent increase. The intensive margin requires more care, because the incoming administration cut advertising across these ten ministries by 48 percent and both groups of outlets lost revenue. Spending on aligned outlets fell by \$8,946 per ministry-year, or 9 percent of their pre-transition average, while spending on misaligned outlets fell by \$115,995, or 72 percent. The coefficient of \$105,362 is the difference between those two declines, approximately 0.36 standard deviations of the outcome. Alignment therefore protected an outlet from a contraction that fell almost entirely on its competitors, and did not raise what any outlet received. The levels estimate in Column (2) averages over pairs that begin receiving advertising after the transition and pairs that received advertising under both administrations, so it captures both margins at once. Column (3) reports the same outcome under the inverse hyperbolic sine, which the prior literature estimates. That coefficient is not a percentage change and varies with the units the outcome is measured in (Chen and Roth, 2024), so I report it for comparability and read Column (2) as the effect.

Higher spending on aligned outlets could reflect ideology, or it could be rationalized by audience composition, if the government allocates more to left-wing outlets because those outlets reach its supporting audiences. To disentangle these explanations, I restrict Columns (4)–(6) to bipartisan campaigns, programs that ran under both administrations with identical stated target demographics, so that any differential allocation toward aligned outlets cannot be attributed to differences in whom the government was trying

[7] The 10 ministries are Culture, Defense, Education, Energy, Foreign Affairs, Health, Interior, Navy, Presidency, and Tourism.

to reach. All six columns are estimated at the outlet–ministry–year level and have identical numbers of observations, making the two panels directly comparable.

Table 3: **Differential Effects of Government Partisanship on Advertising Spending Across Left- and Right-Wing Outlets**

	All Campaigns			Bipartisan Campaigns		
	Prob(Ads > 0)	**Gov Ads (Levels, USD)**	**Gov Ads IHS**	**Prob(Ads > 0)**	**Gov Ads (Levels, USD)**	**Gov Ads IHS**
	(1)	(2)	(3)	(4)	(5)	(6)
Left-Wing Government × Left-Wing Outlet	0.183**	105,362*	2.365**	0.117*	15,913*	1.316**
	(0.082)	(53,615)	(1.019)	(0.056)	(9,017)	(0.592)
Mean (pre-2018, left)	0.56	107,562	6.49	0.26	7,563	2.66
Mean (pre-2018, right)	0.78	160,559	9.05	0.36	16,484	3.83
Year FE	✓	✓	✓	✓	✓	✓
Outlet × Ministry FE	✓	✓	✓	✓	✓	✓
Observations	1,410	1,410	1,410	1,410	1,410	1,410
R^2	0.58	0.58	0.63	0.41	0.16	0.41

Note: *p<0.1; **p<0.05; ***p<0.01. SEs clustered at the outlet level. Cols. (1) and (4) use an indicator for positive advertising spending (extensive margin). Cols. (2) and (5) use the level of advertising spending in 2022 USD, including zeros. Cols. (3) and (6) use the inverse hyperbolic sine of spending and are reported for comparability with prior work, where the coefficient is not interpretable as a percentage change and is sensitive to rescaling of the outcome (Chen and Roth, 2024). All columns restrict to the 10 ministries with bipartisan campaigns so that observations are comparable across panels (N = 1,410). The price paid per unit of space is reported in Appendix Table E3. The Bipartisan Campaigns panel restricts to the 17 campaigns implemented across both administrations with identical target demographics. Means computed over the pre-2018 period.

Columns (4)–(6) disentangle ideological favoritism from audience-driven allocations by restricting the sample to bipartisan campaigns. Aligned outlets are still 11.7 percentage points more likely to receive advertising, a 45% increase relative to their pre-transition baseline of 0.26, and receive $15,913 more per ministry per year, confirming that favoritism drives the reallocation independently of audience composition.

Favoritism therefore operates through whom the government buys from and how much it spends, not through what it pays per unit of space. Holding the advertisement's size and the campaign fixed, the price paid to an aligned outlet is statistically indistinguishable from the price paid to a misaligned one, and the bounds that account for selection into a contract contain zero, as Appendix E.1 reports. Appendix D.1 provides two further exercises following Chen and Roth (2024), namely a Poisson regression whose coefficient is scale-invariant and reads as a semi-elasticity, and Lee bounds (Lee, 2009) for the conditional intensive margin.

To assess the validity of the parallel trends assumption, Figure 3 presents event study estimates corresponding to Column (1) of Table 3. Event time is defined as 12-month periods relative to the December 2018 transition, with the period immediately preceding the transition (December 2017–November 2018) normalized to zero. Following the transition, the coefficients become positive and statistically significant, indicating a sharp increase in advertising spending to left-wing outlets relative to right-wing outlets, and the effect grows persistently through 2022.

Figure 3: **Event Study of the Extensive Margin of Government Advertising by Outlet Ideology**

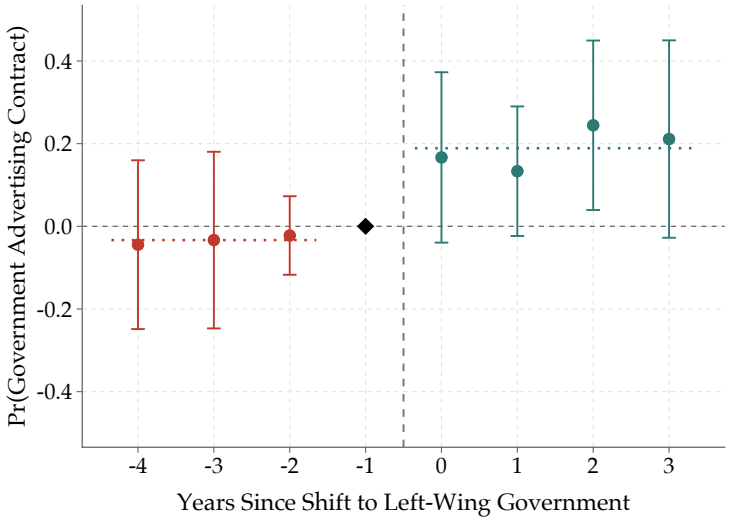

Coefficients on the interaction between event time and a left-wing outlet indicator, from Column (1) of Table 3. Reference group: right-wing outlets. Omitted period: December 2017–November 2018. Error bars show 95% confidence intervals; standard errors clustered at the outlet level. Outlet–ministry fixed effects included.

The flat pre-transition coefficients and the increasing post-transition estimates are consistent with a discontinuous change in advertising allocations driven by the party transition and not by pre-existing differences in trajectory. Event study plots for other specifications are presented in Appendix Figure D1.

5 Benchmarking the Allocation

Section 4 shows that ministries direct advertising toward aligned outlets. A government pursuing an audience would produce the same pattern, so those estimates do not separate favoritism from targeting. I therefore construct the allocation a ministry would have chosen had it pursued an audience and nothing else, one for each of the two audiences a ministry might pursue, and each spends the campaign's observed budget to reach its audience as widely as the budget allows. The gap between the observed and the benchmark allocation measures favoritism in two units, dollars per campaign and readers a campaign did not inform.

5.1 The Advertiser's Problem

A ministry runs campaign j in year t with budget B_{jt}. It allocates that budget across outlets, and outlet i sells one insertion at price p_{it} and delivers a_{ijt} readers in the campaign's target audience, so spending s_{ijt} buys $x_{ijt} = s_{ijt}/p_{it}$ insertions. The ministry chooses spending to inform as many distinct readers in that audience as its budget allows. In particular, the ministry maximizes

$$\max_{\{s_{ijt} \geq 0\}} \sum_i a_{ijt} \left[1 - (1-\phi)^{x_{ijt}}\right] \qquad \text{s.t.} \qquad \sum_i s_{ijt} \leq B_{jt} \tag{2}$$

where ϕ denotes the probability that a reader of an outlet observes any one of its insertions. The bracketed term is the probability that a reader of outlet i observes the campaign at least once, so the objective counts readers informed and not advertisements placed.[8]

Equation 2 restricts how advertising informs readers in three ways. An insertion is one advertisement a ministry places in one issue of one outlet, and an exposure is one reader observing one insertion.

First, exposures are independent across insertions. A reader observes each insertion with probability ϕ whether or not that reader observed the earlier ones, so a reader misses all x insertions with probability $(1-\phi)^x$. The assumption fixes the rate at which reach accumulates. Habitual readers would violate it, because the same readers would observe every insertion while others observed none, and reach would saturate faster than equation 2 implies.

Second, only the first exposure carries value, so returns to insertions diminish. The x-th insertion adds only the readers who missed the previous $x-1$, a share $\phi(1-\phi)^{x-1}$ of the audience, and each insertion adds fewer readers than the one before it. The assumption matches the campaigns the government runs, which inform and do not persuade. A campaign announces that a social program has opened for applications or where a vaccine is available, and a reader who holds that fact has no use for a second insertion. Persuasive advertising would violate it, because repetition there builds recall and shifts beliefs. A ministry would then value total exposures instead of distinct readers, and it would concentrate its budget in the outlet that sells exposures most cheaply.

Third, readers do not multi-home, so no reader takes two of the sample newspapers. The assumption lets me add readers across outlets instead of tracking who reads which combination, and a campaign informs a_{ijt} readers in outlet i and a disjoint $a_{i'jt}$ readers in outlet i'. A reader who takes two newspapers enters equation 2 twice. The National Print Media Registry reports circulation and audience composition outlet by outlet and records nothing about joint readership. The error nonetheless carries a known sign. Benchmark campaigns fund more outlets than observed campaigns do, so overlap would subtract more from benchmark reach than from observed reach, and the reach gap I report below is an upper bound.

Reach saturates within an outlet, and that ceiling shapes the solution. A ministry that places its entire budget with one newspaper informs $a_{ijt}[1-(1-\phi)^{B_{jt}/p_{it}}]$ readers and never more than a_{ijt}, however large the budget. A ministry that maximizes reach therefore spreads spending across outlets.

5.2 Constructing the Benchmarks

The problem admits a closed-form solution that reduces to finding one number per campaign. The objective is additively separable and strictly concave in each s_{ijt}, and the constraint set is convex, so the Kuhn-Tucker conditions pin down a unique optimum. Marginal reach per dollar stays positive everywhere, so the budget binds. Writing λ_{jt} for its multiplier and $\kappa \equiv -\ln(1-\phi)$,

[8] Equation 2 is the reach-only case of Gentzkow et al. (2024), who value an advertiser reaching a viewer M times at $a \cdot u(M)$ for a submodular $u(\cdot)$ and set $u(M) = \mathbf{1}[M > 0]$. It also inherits the constant-reach, independent-readership technology of Grossman and Shapiro (1984), who write the reach of a campaign placed in m vehicles each covering a fraction r of the population as $1-(1-r)^m$.

$$s^{*}_{ijt}(\lambda_{jt}) \;=\; p_{it} \cdot \max\left\{0,\; \frac{1}{\kappa}\ln\!\left(\frac{a_{ijt}\,\kappa}{\lambda_{jt}\,p_{it}}\right)\right\} \tag{3}$$

λ_{jt} is the multiplier on the budget constraint, equal at the optimum to the marginal reach a ministry obtains from an additional dollar. Total spending under equation 3 falls continuously and strictly as λ_{jt} rises, from more than any budget when λ_{jt} approaches zero to zero once λ_{jt} exceeds every outlet's return on its first dollar. Exactly one value of λ_{jt} therefore exhausts a given budget. I find that value numerically for each of the 383 campaigns, starting from an interval that brackets it and halving the interval until benchmark spending equals the observed budget.

The max operator in equation 3 produces a corner solution and with it an extensive margin. Marginal reach per dollar at outlet i equals $a_{ijt}\kappa/p_{it}$ at zero spending and falls thereafter, so an outlet whose marginal reach at zero spending lies below λ_{jt} receives nothing at any budget. The benchmark therefore determines which outlets a ministry funds as well as how much each one receives, so I can compare the outlets a ministry chose against the outlets an audience objective implies. Panel A of Table 4 reports that comparison.

I set $\phi = 0.10$, so one insertion informs one reader in ten and the two insertions a funded outlet receives at the median inform one in five. The advertising data do not identify ϕ, so Appendix E.3 re-estimates every deviation over a range of values, starting at the $\phi \to 0$ corner. The intensive-margin deviation stays positive and significant throughout, so no conclusion in this section turns on the value. The number of outlets a benchmark campaign funds does move with ϕ, so Figure 4 supports a qualitative comparison and not a point estimate.

Dividing κ by the multiplier turns it into a reservation price, the highest cost per target reader at which a ministry still buys an outlet. Write $c_{ijt} = p_{it}/a_{ijt}$ for what an outlet costs per target reader and set $\bar{c}_{jt} \equiv \kappa/\lambda_{jt}$. A ministry buys outlet i if and only if $c_{ijt} < \bar{c}_{jt}$, and buys $x^{*}_{ijt} = \kappa^{-1}\ln(\bar{c}_{jt}/c_{ijt})$ insertions, so a campaign buys more insertions from the outlets that cost less per target reader. The median reservation price reaches \$115 per thousand target readers, while outlets charge a median of \$89 with an interquartile range of \$58 to \$163. The reservation price falls inside the price distribution, so the benchmark funds some outlets and not others. The outlets that clear it most easily are the high-circulation popular and sports titles, so Appendix E.6 re-solves both benchmarks with those titles removed.

The two benchmarks differ only in a_{ijt}, the audience each one counts, and I use two because the two audiences answer different questions. The *cost-efficiency benchmark* sets a_{ijt} to an outlet's readers in the campaign's stated demographic, which is the audience the legal cost-efficiency standard obliges ministries to buy, so it prices the departure from the rule the law imposes.[9] The *voter-targeting benchmark* sets a_{ijt} to the outlet's readers who support the incumbent, which I predict by weighting official state-level vote shares by where the outlet circulates, so it tests whether a government pursuing its own voters would have bought

[9]Campaigns state a demographic on two dimensions, income and age, and the registry reports each outlet's readership shares on the same categories. I compute the share of an outlet's readers falling in the stated groups separately on each dimension and multiply the two, which treats income and age as independent within an outlet. The registry reports only marginal distributions, so no other combination is available. Three campaigns of 386 state a demographic that no national newspaper reaches, which leaves the benchmark undefined, so I drop them. The estimation sample holds 383 campaigns and 6,711 outlet–campaign cells.

what ministries bought. Appendix E.2 checks the measure against an independent estimate built from individual polling microdata, with which it correlates at 0.80.

Each benchmark yields an allocation s^*_{ijt} for every outlet in every campaign. I regress the gap between what a ministry spent and what the benchmark prescribes on alignment,

$$Y_{ijt} = \alpha \cdot \mathbf{1}[\text{Aligned}]_{it} + \gamma_j + \delta_i + \varepsilon_{ijt} \tag{4}$$

where γ_j and δ_i denote campaign and outlet fixed effects and $\mathbf{1}[\text{Aligned}]_{it}$ equals one when the party running the campaign and the outlet share an ideology. The outcome Y_{ijt} takes three forms, the extensive margin $\mathbf{1}[s^{\text{obs}}_{ijt} > 0] - \mathbf{1}[s^*_{ijt} > 0]$, the intensive margin $s^{\text{obs}}_{ijt} - s^*_{ijt}$ in dollars, and that dollar gap divided by B_{jt}, since campaign budgets span three orders of magnitude. Campaign fixed effects absorb the budget, the stated demographic and the party in power, so α compares outlets within a campaign, and outlet fixed effects absorb every time-invariant outlet characteristic. I cluster standard errors by outlet and re-estimate each specification with the continuous ideology score in place of the binary indicator.

5.3 Deviations from the Benchmarks

I first present results on the extensive margin, which records whether an outlet receives government advertising at all. Figure 4 plots how many outlets ministries funded per campaign and how many the benchmark prescribes. The observed distribution peaks at a single outlet and reaches a median of six, against a benchmark median of eleven with mass at fifteen outlets and above. Ministries concentrate their campaigns beyond what informing the stated demographic would justify.

Figure 4: **Outlets Used per Campaign, Observed and Benchmark**

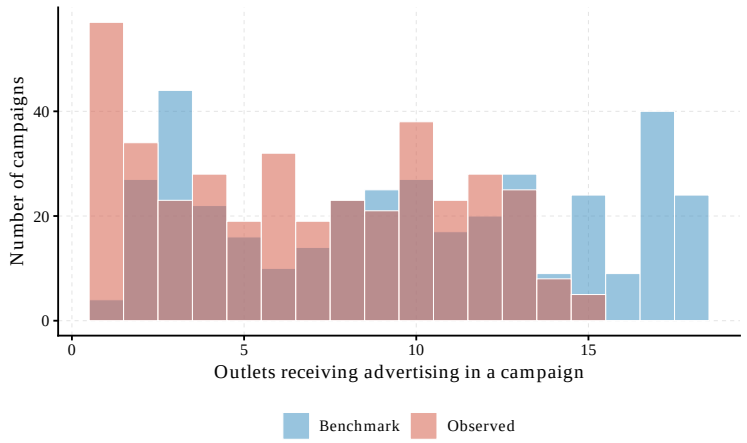

The benchmark solves equation 2 at $\phi = 0.10$ for each campaign, holding the observed budget fixed. Bars overlap where the distributions coincide.

Table 4 puts that margin in regression form, with each panel reporting one outcome from equation 4. Panel A asks how often a ministry funds an outlet the benchmark would not. Aligned outlets receive money the cost-efficiency benchmark withholds 13.9 percentage points more often than misaligned outlets, and

13.1 points when the continuous ideology score replaces the binary split. Columns (3) and (4) repeat the comparison against the voter-targeting benchmark, and both estimates rise. Alignment therefore predicts which outlets a ministry funds beyond what either benchmark prescribes.

I then present results on the intensive margin, which records how much a funded outlet receives. Panel B prices the same gap in dollars. Aligned outlets receive $20,876 more per campaign than the cost-efficiency benchmark prescribes, and $28,749 more under the continuous score. Panel C divides each deviation by the campaign budget, since budgets run from $3,642 to $7.7 million, and returns 2.45 and 3.50 percentage points, of which only the continuous estimate separates from zero at conventional levels. A benchmark campaign funds 10.2 outlets, so a funded outlet holds 9.8 percent of a budget on average, and the two deviations amount to a quarter and a third of that share. Neither the ideological threshold nor the audience a ministry might have pursued changes the result, because ministries directed money toward aligned outlets under every specification.

Table 4: **Deviations from the Benchmark Allocations by Outlet Alignment**

	Cost-efficiency benchmark		**Voter-targeting benchmark**	
	(1)	(2)	(3)	(4)
Panel A. Extensive margin				
Alignment	0.139***	0.131***	0.141***	0.153***
	(0.039)	(0.018)	(0.048)	(0.022)
Panel B. Intensive margin, 2022 dollars				
Alignment	20,876**	28,749***	22,198**	31,038***
	(8,912)	(6,008)	(9,865)	(6,863)
Panel C. Intensive margin, share of campaign budget (pp)				
Alignment	2.45	3.50**	3.01*	4.08**
	(1.43)	(1.51)	(1.71)	(1.65)
Alignment measure	**Binary**	**Continuous**	**Binary**	**Continuous**
Ad Campaign FE	✓	✓	✓	✓
Outlet FE	✓	✓	✓	✓
Observations	6,711	6,711	6,711	6,711

Estimates of equation 4 at $\phi = 0.10$. Panel A takes $\mathbf{1}[\text{observed} > 0] - \mathbf{1}[\text{benchmark} > 0]$ as the outcome, Panel B takes the difference in 2022 dollars, and Panel C divides that difference by the campaign's total budget. Each column reports a separate regression in which the stated alignment measure is the only regressor. Binary columns split outlets at the ideological midpoint. Continuous columns replace the split with the ideology score of Section 3.2, signed so that higher values indicate closer alignment with the incumbent. Binary and continuous coefficients are measured on different scales and are not comparable in magnitude. Standard errors clustered by outlet across 18 clusters. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.10$.

The copartisan measure weights state vote shares by where an outlet circulates, so it registers no sorting within a state. Appendix E.4 rebuilds the voter-targeting benchmark under assumed sorting and reports how much of the deviation survives. Measured as a share of the campaign budget, the deviation holds as long as the odds of an incumbent supporter reading an aligned newspaper instead of a misaligned one stay within roughly 1.3 times what I measure, which would put 24 percent of an aligned outlet's readers on the incumbent's side against 21 percent of a misaligned outlet's.Measured in dollars it holds to 2.7 times. Two further results speak against sorting of that magnitude. Neither of the two independently constructed copartisan measures separates left-wing from right-wing outlets by even half a percentage

point, as Appendix E.2 reports. Appendix E.5 also shows that every weighting of reader demographics predicts within-campaign allocation with the wrong sign, while a single parameter loading circulation on outlet ideology predicts it with the right one.

Finally, I quantify what the departure costs readers, who learn what a campaign announces only if it reaches them. Campaigns inform between 178,000 and 275,000 target readers against the benchmark's 765,000, or 23 to 36 percent.[10] Figure 5 plots that shortfall campaign by campaign, and the distribution lies entirely above zero with a mean of 586,000 readers, so not one of the 383 campaigns informs as many readers as its own budget could have. Outlet choice produces the shortfall and price does not. A ministry's own choices, priced at the market rate, deliver 275,000 target readers, while the benchmark spends the same money to deliver 765,000, and both allocations convert spending into insertions at the same price. Appendix E.1 shows that pricing those choices at market rates recovers 4,422 readers per campaign in cells recording at least three advertisements, against a shortfall of 490,000.

Figure 5: **Difference in Unique Target Readers Reached, Benchmark Minus Observed**

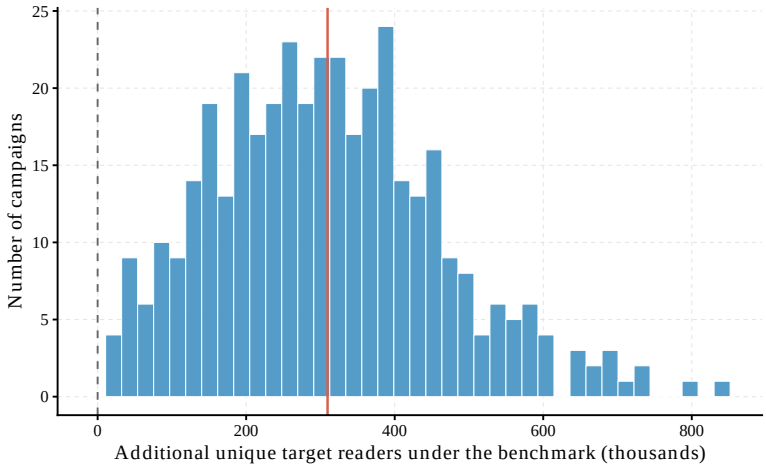

One observation per campaign. Reach counts readers in the campaign's named groups who see the advertisement at least once, computed at $\phi = 0.10$. Observed reach converts spending into insertions using the advertisement counts ministries recorded, and benchmark reach converts it at the private market price. The dashed line marks zero and the solid line marks the mean difference.

Ministries therefore departed from every allocation an audience objective would justify, and aligned outlets absorbed that departure. The deviation holds against the standard the law imposes and against the one an office-seeking politician would hold, and it grows with an outlet's alignment. Section 6 estimates the effect of changes in revenue on three market outcomes, the capacity to produce news, journalistic quality, and circulation.

[10] Ministries record both the number of advertisements they placed and the amount they paid, and the two do not always agree. The lower count uses the recorded advertisements. The higher converts spending into advertisements at the private market price, and it exceeds the first because ministries pay above that price and because the recorded counts appear to miss insertions in the cells where the discrepancy is largest, as Appendix E.1 documents.

6 Downstream Effects in the Media Market

Sections 4 and 5 show that ministries directed advertising toward aligned outlets. Government advertising fell for every outlet after the 2018 transition, from $1.7 million per outlet to $467,000, and it fell by less at aligned outlets. Government contracts fund roughly a third of the advertising space in the average outlet, so that differential decline left aligned outlets with more resources than their competitors. I now estimate the downstream consequences in the media market, with a focus on the supply of news and on the demand from readers.

6.1 Newsroom Capacity and Content Production

Outlets pay journalists out of revenue, so an outlet that loses advertising employs fewer of them and publishes less. I measure three outcomes in sequence, the number of journalists an outlet employs, the articles they produce, and the quality of that reporting. I use the difference-in-differences design of Section 4 throughout and compare left-wing and right-wing outlets before and after the 2018 transition.

I first examine the raw patterns of newsroom size across left-wing and right-wing outlets. Figure 6 plots the average number of journalists publishing in each month from 2014 to 2022.[11] Right-wing outlets ran the larger newsrooms before 2019, the ordering reverses after the transition, and left-wing outlets stay larger through 2022. The gap narrows over 2018, before the December handover. Readers alone could account for that narrowing if they turned toward left-wing outlets once the July election returned a left-wing president. Ministries reallocated their advertising only after the new administration took office, so the break at the end of 2018 dates the revenue shift and the earlier narrowing does not. The raw series cannot separate that shift from other shocks reaching both groups at the same time, so I estimate the difference-in-differences, which removes shocks common to all outlets and measures how much more left-wing newsrooms grew than right-wing ones.

Table 5 estimates equation 1 on the outlet–month panel. Column (1) takes newsroom size, the number of journalists who publish under their own byline in a given month. Columns (2) to (4) take what those journalists produce, the total article count and its split into hard and soft news, which I assign from a topic model fitted on the article text and described in Appendix F.1. Column (5) takes the share of articles an outlet buys from news wires, which measures how much of its output it does not report itself. Appendix F.3.2 reports the classification procedure and its robustness checks.

[11] I count a journalist as active in a month when an article carries their byline, excluding agency copy, and I retain those who publish at least 16 articles across the study period.

Figure 6: **Newsroom Size in Time across Right- and Left-Wing Outlets (2014–2022)**

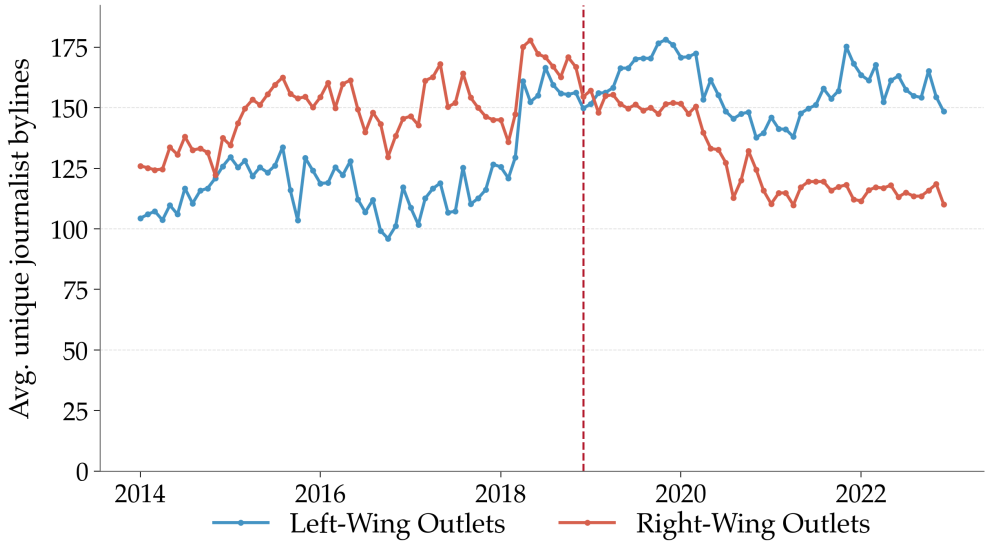

Newsroom size represents the number of classified journalists active per outlet-month. A journalist is classified as active in a given month if they have published at least 16 articles across the entire study period. Journalists are identified through byline analysis, excluding agency content.

Table 5: **Differential Effects of Government Partisanship on Content Production Across Left- and Right-Wing Outlets**

	Dependent variable:				
	Newsroom	Total Articles	Hard News	Soft News	Wire Copy (%)
	(1)	(2)	(3)	(4)	(5)
Left-Wing Gov × Left-Wing Outlet	0.600***	0.549***	0.657***	0.540***	−1.001
	(0.133)	(0.185)	(0.212)	(0.195)	(2.252)
Mean pre-2018 (Left)	139.73	1841	800.55	174.08	12.59
Mean pre-2018 (Right)	251.54	3105.23	1399.36	310.16	15.06
SD pre-2018	160.96	1667.79	835.23	187.34	7.1
Year FE	✓	✓	✓	✓	✓
Outlet FE	✓	✓	✓	✓	✓
Observations	1,908	1,908	1,908	1,908	1,908
R^2	0.842	0.569	0.780	0.755	0.609

Note: *p<0.1; **p<0.05; ***p<0.01. SEs clustered at the outlet level. Outlet–month panel, 2014–2022. Seventeen outlets span all 108 months and El Heraldo de México, which began publishing in 2017, contributes 72, giving N = 1,908. Cols. (1)–(4) use the inverse hyperbolic sine, so coefficients approximate percentage changes. Col. (5) is a percentage in levels, weighted by the monthly article count, and the nine outlet–months with no articles carry zero weight. Means and standard deviations cover the pre-2018 period and are reported in levels, so each coefficient can be read against the base it applies to.

Left-wing outlets expand their newsrooms after the transition. Column (1) reports 60 percent more journalists relative to right-wing outlets, or roughly 84 journalists at the left-wing average of 140 before the transition, and the added capacity reaches the page. Total article production rises 55 percent in Column (2), hard news 66 percent in Column (3), and soft news 54 percent in Column (4). Column (5) finds no change in the share of articles bought from news wires, so left-wing outlets published more of their own reporting and not more purchased copy. Figure 7 plots the event study behind Columns (1) to (4), and the pre-transition coefficients sit at zero in all four.

Figure 7: **Event Studies for Newsroom Size and Content Production**

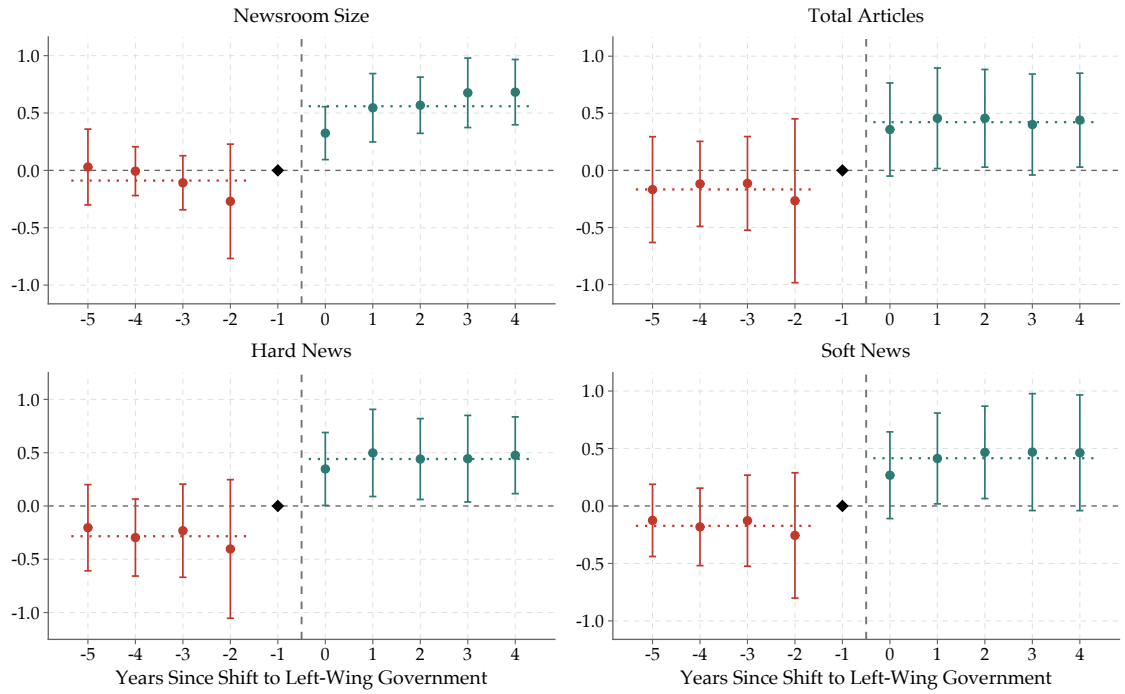

Coefficients on the interaction of a left-wing outlet indicator with each year relative to the December 2018 transition, for Columns (1) to (4) of Table 5. Event years run from December to November, and the year before the transition is the omitted category, marked by the diamond. Outlet fixed effects absorb every time-invariant outlet characteristic. Bars are 95 percent confidence intervals from standard errors clustered by outlet across 18 clusters. Dotted lines mark the average coefficient before and after the transition.

The expansion is concentrated in substantive reporting instead of spread evenly across the newspaper. Appendix Table F3 decomposes the article count into finer content categories and shows that the gains are largest in economics and social policy, followed by politics, while the effect on sports is small and insignificant.

I also examine how private advertising changed, to assess whether the transition moved the whole advertising market or only the government's share of it. If private advertisers maximize profits and place advertisements in the outlets with the highest commercial value, then any movement toward left-wing outlets means the transition changed what those outlets were worth to any advertiser, and not only to the government. Appendix Figure F1 reports the event study for the number of private advertisements an outlet carries. Left-wing outlets gain on right-wing outlets in every year after the transition, by 0.20 asinh points in the transition year and 0.38 in the year that follows. Both revenue sources therefore moved toward the same outlets, and the estimates above measure the combined change and not the government contract alone.

Newsroom size and article counts record how much an outlet can produce and how much it publishes. Two features of those estimates already bear on quality. The added output runs mostly through hard news, which starts from a base more than four times larger than soft news, so roughly 85 percent of the additional articles are hard reporting. The share bought from news wires does not move, so the outlets filled those pages with their own reporting. Neither estimate measures how well an outlet reports a given story, so I

now estimate measures of quality that capture original reporting and the ability to break new stories.

I measure quality around news events, which separates the number of stories an outlet covers from the amount of each story it reports. I build the events with an unsupervised clustering classifier that groups articles on the similarity of their text and their timing, following Cagé et al. (2020) and Durante et al. (2026), so an event collects every article the sample published about the same underlying occurrence. Within an event I compute measures that capture different dimensions of how an outlet reports a story, the detail it carries and the information it brings that others do not. All three rest on the distinct factual elements of an article, meaning a named actor, an organization, a figure, a place, or a date, and I weight each element by the share of other outlets that reported it, leaving the outlet's own reporting out of that share.

In particular I construct three measures. First, completeness records the share of the event's weighted elements the outlet carried, which measures how much of an event an outlet reported conditional on covering it. Second, new information counts the elements the outlet carried that few other outlets reported, which measures original reporting. Third, story elements counts the factual elements the article holds, which measures how much verifiable material it puts in front of a reader. Together the three record how an outlet reported a story against how the rest of the market reported the same story. Appendix F.2 reports the clustering procedure, the element extraction, and their validation.

Table 6 reports the estimates. Column (1) is an indicator for whether the outlet published anything about the event, which measures the breadth of what a newsroom covers, and Columns (2) to (4) are the three within-event measures. Event fixed effects absorb all event-level variation, so the coefficient is identified from differently aligned outlets covering the same story.

Table 6: **Effect of Left-wing Government on Journalistic Quality**

	Dependent variable:			
	Reports Event	Completeness	New Information	Story Elements
	(1)	(2)	(3)	(4)
Left-Wing Government × Left-Wing Outlet	0.0775**	0.0286**	2.8435**	3.7234**
	(0.0319)	(0.0134)	(1.1129)	(1.4349)
Event FE	✓	✓	✓	✓
Outlet × Category FE	✓	✓	✓	✓
Weighted by story size		✓		
Pre-period mean	0.2652	0.0700	7.6981	10.19
Wild bootstrap p	0.03	0.05	0.02	0.02
Observations	3,404,704	2,700,888	2,700,888	2,700,888
Events	193,066	153,109	153,109	153,109

Note: *p<0.1; **p<0.05; ***p<0.01. Each column is a separate regression of the outcome on Left-Wing Gov × Left-Wing Outlet with event and outlet × category fixed effects. Left-Wing Government equals one for events dated on or after 1 December 2018, and the pre-period mean is the left-wing-outlet mean before that date. Events enter only if at least three outlets covered them. All columns run on 2014–2022. Col. (1) uses every event in the coverage panel, 193,066 in total. Cols. (2)–(4) require the event to yield at least one extracted factual element, which drops the 39,957 events where no covering article produced one and leaves 153,109. Element extraction excludes wire copy, so the elements come from outlets' own reporting. Element prevalence is computed leaving out the outlet's own reporting, so an element no other outlet carries scores as fully exclusive; the share outcomes in cols. (2) and (3) are weighted by the number of prevalence-weighted elements in the event. Standard errors clustered on outlet are in parentheses and shown for comparability; with 18 clusters they over-reject, so stars and the reported p-value come from a wild cluster bootstrap with the null imposed and Rademacher weights, clustered on outlet.

Aligned outlets cover 7.8 percentage points more of the events in the panel after the transition, against a pre-transition baseline of 0.27, an increase of 29 percent (Column 1). The three within-event measures rise by comparable proportions. Completeness rises 0.029 on a baseline of 0.070, or 41 percent (Column 2). New information rises by 2.84 elements on a baseline of 7.70, or 37 percent (Column 3). Story elements rise by 3.72 on a baseline of 10.20, or 37 percent (Column 4). Every estimate clears the five percent level under the wild cluster bootstrap. Aligned outlets therefore covered more events and reported more of each event they covered, and the two margins moved by similar amounts.

I now address the interpretation of these measures, and whether they capture quality or only volume. The count of events an outlet reports may record newsroom capacity instead of newsworthiness, and the count of factual elements may record detail that adds nothing a reader wants. These measures capture original reporting and the capacity to produce new information, which is one dimension of journalistic quality (Turkel et al., 2021; Armona et al., 2024). Additionally, Durante et al. (2026) show that the number of facts an article carries and its distance from the article that broke the story predict which outlet a news aggregator places on its front page, which suggests these measures capture something readers value.

6.2 Circulation

Government advertising reshapes the information environment only if it moves readers, because an outlet that publishes more without reaching more readers leaves the ideological composition of the market unchanged. I therefore examine circulation at left-wing and right-wing outlets.

Figure 8 presents the event study of equation 1 on the balanced regional circulation panel, with the inverse hyperbolic sine of circulation as the outcome. Left-wing outlets gain circulation relative to right-wing outlets after the transition, and the point estimate reaches 0.34 by 2022. The pre-transition coefficients sit near −0.3, and with 18 outlet clusters every confidence interval covers zero, so the event study matches the pattern the decomposition reports below without establishing it on its own.

Figure 8: **Event Study of Circulation by Outlet Ideology**

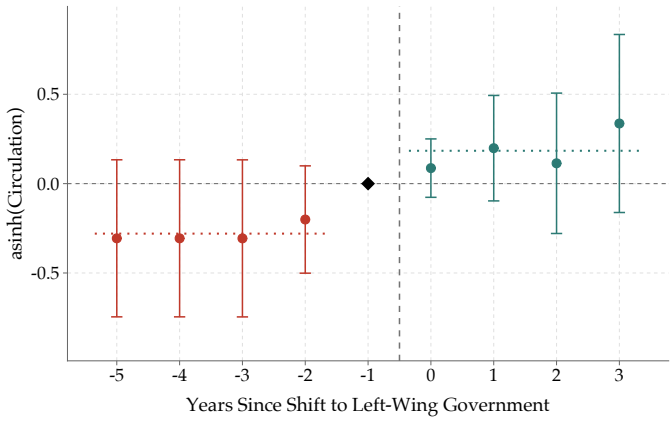

Interaction of a left-wing outlet indicator with each year relative to the December 2018 transition, from a regression of asinh(Circulation) on outlet-state and year fixed effects. The omitted year is marked by the diamond, dotted lines mark the pre- and post-transition averages, and bars are 95 percent confidence intervals clustered by outlet across 18 clusters.

Two channels could raise circulation at left-wing outlets after the transition. On the supply side, the revenue these outlets gained paid for journalists and articles, and readers may buy a newspaper that covers more of what they want to read. On the demand side, the 2018 election may have moved readers toward left-wing outlets on its own, if the voters who turned to Morena also turned to the newspapers aligned with it. The two channels carry different implications, because only the first traces circulation back to what the government spent.

I separate the two channels by exploiting cross-state variation in the leftward shift in Morena's vote share between 2015 and 2018. The intuition is that content improvements at left-wing outlets are national, because the same newspaper is distributed across all Mexican states, while preference-driven gains should be larger where more voters shifted toward Morena. States with little or no leftward vote shift therefore serve as a counterfactual for the supply-side component; with preferences largely unchanged in those states, any differential circulation gains for left-wing outlets reflect content improvements alone. The additional gains in high-shift states then identify the preference-driven component. This design relies on two assumptions. First, content effects are uniform across states. Second, demand-side shifts are proportional to the local electoral shift. In particular, I estimate a triple difference-in-differences specification:

$$\begin{aligned} Y_{nst} &= \beta_1 \left(\text{Left-Wing Gov}_t \times \text{Left-Wing Outlet}_n\right) \\ &\quad + \beta_2 \left(\text{Left-Wing Gov}_t \times \text{Left-Wing Outlet}_n \times D_s\right) \\ &\quad + \beta_3 \left(\text{Left-Wing Gov}_t \times D_s\right) + \alpha_{ns} + \gamma_t + \varepsilon_{nst} \end{aligned} \tag{5}$$

where $D_s = \mathbf{1}[\Delta\text{Vote Share Left-Wing Party}_s > \text{median}]$ is an indicator for states with an above-median increase in left-wing vote share between 2015 and 2018, α_{ns} are outlet-state fixed effects, and γ_t are year fixed effects. Appendix Figure F2 maps the state-level variation in Morena's vote share shift that defines D_s; southern and southeastern states experienced shifts of 50–65 percentage points, while northern states shifted by 20–35 percentage points. Under this design, β_1 captures the content-driven component of circulation gains, common across all states regardless of political shift. β_2 identifies the additional preference-driven component, and β_3 tests whether the overall political shift drives circulation independently of outlet ideology.

Table 7 presents the results. In Column (1) the coefficient on Left-Wing Gov × Left-Wing Outlet is 0.411 with a standard error of 0.236. Column (2) splits that effect by the size of the local electoral shift. The coefficient on Left-Wing Gov × Left-Wing Outlet falls to 0.317, and the triple interaction with the high-shift indicator is 0.187 with a standard error of 0.136, positive but not distinguishable from zero. The coefficient on Left-Wing Gov × $\mathbf{1}\{\Delta$ Left-Wing Vote Share > median$\}$ is −0.002, so the leftward political shift does not by itself move circulation across all outlet types. Weighting the triple interaction by the share of high-shift states, one half by construction of the median split, the component common to all states accounts for approximately 77 percent of the average circulation gain and the remainder is concentrated in states with

above-median realignment. Because the triple interaction is imprecisely estimated, that split should be read as suggestive and not as a sharp decomposition.

Table 7: **Decomposing Circulation Gains**

	asinh(Circulation)		
	(1)	(2)	(3)
Left-Wing Gov × Left-Wing Outlet	0.411*	0.317*	0.411*
	(0.236)	(0.183)	(0.236)
Left-Wing Gov × Left-Wing Outlet × **1**{Δ Left-Wing Vote Share > median}		0.187	
		(0.136)	
Left-Wing Gov × **1**{Δ Left-Wing Vote Share > median}		−0.002	
		(0.066)	
Left-Wing Gov × Left-Wing Outlet × Δ Left-Wing Vote Share (SD)			0.107
			(0.072)
Left-Wing Gov × Δ Left-Wing Vote Share (SD)			−0.008
			(0.035)
Outlet × State FE	✓	✓	✓
Year FE	✓	✓	✓
Wild bootstrap p	0.051	0.062	0.051
Observations	5,194	5,194	5,194
R^2	0.958	0.959	0.959

Note: *p<0.1; **p<0.05; ***p<0.01. The outcome is the inverse hyperbolic sine of outlet–state circulation, 2014–2022. Standard errors are clustered on outlet, because outlet ideology is fixed over time and the interaction therefore varies only across the 18 outlets. **1**{Δ Left-Wing Vote Share > median} marks states whose increase in left-wing vote share between 2015 and 2018 exceeded the median.

The event study and the decomposition point the same way on circulation without settling it. Left-wing outlets gain relative to right-wing outlets after the transition, the gain does not concentrate in the states where Morena advanced most, and both estimates are too imprecise to rule out the alternative. What the evidence does establish is the step before circulation, that ministries directed advertising toward aligned outlets and those outlets expanded what they produced.

7 Conclusion

Governments that buy advertising choose which outlets to buy it from, and this paper shows that in Mexico they chose outlets that shared their ideology. Aligned outlets received $20,876 more per campaign than the legal cost-efficiency standard prescribes and $22,198 more than voter targeting prescribes, and they were roughly 14 percentage points more likely to receive any funding at all. Neither the audience the law obliges ministries to buy nor the audience an office-seeking politician would pursue accounts for where the money went.

The money changed what those outlets produced. Left-wing outlets employed 60 percent more journalists than right-wing outlets after the 2018 transition, published 55 percent more articles, and bought no more of their pages from news wires, so they filled the additional space with their own reporting. They also covered more of the events in the market and carried more of the facts in each event they covered. Readers bear the cost of the allocation that produced this, because campaigns informed between 23 and 36 percent

of the audience the same budget could have reached.

This is media capture without censorship. A government that wants friendly coverage need not suppress a story or buy a favorable one. It can move a routine procurement budget toward the outlets it prefers and let the market do the rest, because an outlet that loses a third of its advertising revenue employs fewer journalists and covers fewer stories whatever its editors intend. The instrument leaves no trace in any single article, which is what makes it hard to observe and hard to regulate.

The conditions that make this work are common. They are competitive media markets in which outlets compete for readers while depending on a large public advertiser, which describes much of Latin America and Eastern Europe alongside Mexico. Where private advertising is thin and the state is the largest client, the discretionary allocation of public advertising deserves attention as an instrument of influence over the press.

References

Angelucci, Charles and Julia Cagé, "Newspapers in times of low advertising revenues," *American Economic Journal: Microeconomics*, 2019, *11* (3), 319–364.

Armona, Luis, Matthew Gentzkow, Emir Kamenica, and Jesse M. Shapiro, "What is Newsworthy? Theory and Evidence," *American Economic Review: Insights*, 2024. Forthcoming.

Beattie, Graham, Ruben Durante, Brian Knight, and Ananya Sen, "Advertising spending and media bias: Evidence from news coverage of car safety recalls," *Management Science*, 2021, *67* (2), 698–719.

Besley, Timothy and Andrea Prat, "Handcuffs for the grabbing hand? Media capture and government accountability," *American Economic Review*, 2006, *96* (3), 720–736.

Bowles, Jeremy, Horacio Larreguy, Shelley Liu, and Ahmet Akbiyik, "Polarization and Exposure to Cross-Partisan Media in an Electoral Autocracy," 2026. Working paper.

Brollo, Fernanda and Tommaso Nannicini, "Tying Your Enemy's Hands in Close Races: The Politics of Federal Transfers in Brazil," *American Political Science Review*, 2012, *106* (4), 742–761.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel, "The value of democracy: Evidence from road building in Kenya," *American Economic Review*, 2015, *105* (6), 1817–1851.

Cagé, Julia, "Media competition, information provision and political participation: Evidence from French local newspapers and elections, 1944–2014," *Journal of Public Economics*, 2020, *185*, 104077.

_ , Nicolas Hervé, and Marie-Luce Viaud, "The Production of Information in an Online World," *The Review of Economic Studies*, 2020, *87* (5), 2126–2164.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *Quarterly Journal of Economics*, 2024, *139* (2), 891–936.

DellaVigna, Stefano, Ruben Durante, Brian Knight, and Eliana La Ferrara, "Market-based lobbying: Evidence from advertising spending in italy," *American Economic Journal: Applied Economics*, 2016, *8* (1), 224–256.

Durante, Ruben and Brian Knight, "Partisan control, media bias, and viewer responses: Evidence from Berlusconi's Italy," *Journal of the European Economic Association*, 2012, *10* (3), 451–481.

_ , Ruben Enikolopov, Daniil Mikhailov, and Andrey Simonov, "Platform Power in News Aggregation," 2026. Working paper.

Gehlbach, Scott and Konstantin Sonin, "Government control of the media," *Journal of public Economics*, 2014, *118*, 163–171.

Gentzkow, Matthew and Jesse M Shapiro, "What drives media slant? Evidence from US daily newspapers," *Econometrica*, 2010, *78* (1), 35–71.

_ , _ , and Matt Taddy, "Measuring group differences in high-dimensional choices: method and application to congressional speech," *Econometrica*, 2019, *87* (4), 1307–1340.

_ , Jesse M. Shapiro, Frank Yang, and Ali Yurukoglu, "Pricing Power in Advertising Markets: Theory and Evidence," *American Economic Review*, 2024, *114* (2), 500–533.

Grossman, Gene M. and Carl Shapiro, "Informative Advertising with Differentiated Products," *The Review of Economic Studies*, 1984, *51* (1), 63–81.

Hodler, Roland and Paul A. Raschky, "Regional Favoritism," *The Quarterly Journal of Economics*, 2014, *129* (2), 995–1033.

Knight, Brian and Ana Tribin, "The limits of propaganda: Evidence from Chavez's Venezuela," *Journal of the European Economic Association*, 2019, *17* (2), 567–605.

Lee, David S., "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects," *Review of Economic Studies*, 2009, *76* (3), 1071–1102.

McMillan, John and Pablo Zoido, "How to subvert democracy: Montesinos in Peru," *Journal of Economic perspectives*, 2002, *18* (4), 69–92.

Padró i Miquel, Gerard, "The Control of Politicians in Divided Societies: The Politics of Fear," *The Review of Economic Studies*, 2007, *74* (4), 1259–1274.

Petrova, Maria, "Inequality and media capture," *Journal of public Economics*, 2008, *92* (1-2), 183–212.

Prat, Andrea, "Media power," *Journal of Political Economy*, 2018, *126* (4), 1747–1783.

Qin, Bei, David Strömberg, and Yanhui Wu, "Media bias in China," *American Economic Review*, 2018, *108* (9), 2442–2476.

Roberts, Margaret E., Brandon M. Stewart, and Dustin Tingley, "stm: An R Package for Structural Topic Models," *Journal of Statistical Software*, 2019, *91* (2), 1–40.

_ , _ , _ , Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson, and David G. Rand, "Structural Topic Models for Open-Ended Survey Responses," *American Journal of Political Science*, 2014, *58* (4), 1064–1082.

Simonov, Andrey and Justin Rao, "Demand for online news under government control: Evidence from Russia," *Journal of Political Economy*, 2022, *130* (2), 259–309.

Strömberg, David, "Radio's impact on public spending," *Quarterly Journal of Economics*, 2004, *119* (1), 189–221.

Szeidl, Adam and Ferenc Szucs, "Media capture through favor exchange," *Econometrica*, 2021, *89* (1), 281–310.

Tella, Rafael Di and Ignacio Franceschelli, "Government advertising and media coverage of corruption scandals," *American Economic Journal: Applied Economics*, 2011, *3* (4), 119–151.

Turkel, Eray, Anish Saha, Rhett Carson Owen, Gregory J Martin, and Shoshana Vasserman, "A method for measuring investigative journalism in local newspapers," *Proceedings of the National Academy of Sciences*, 2021, *118* (30), e2105155118.

A Setting

A.1 Regulatory Framework for Government Advertising

Two obligations govern federal advertising over the sample period, and both were in force throughout it. A 2007 reform to Article 134 of the Constitution requires that official advertising be institutional in character and bars it from carrying the name, image, voice, or symbol of any public servant. From 2003 onward, an annual *Acuerdo* issued each December by the Secretaría de Gobernación required every federal entity to file an annual communication program, obtain prior authorization for each campaign, exhaust official airtime before buying commercial space, register the outlets it contracts with, and report what it spent.

The *Ley General de Comunicación Social*, published on 11 May 2018 and in force from 1 January 2019, gave that apparatus statutory standing. It extended the rules to state and municipal government (Art. 3), fixed the annual planning cycle (Arts. 22–25), placed the media registry in law (Arts. 35–37), and set the disclosure requirements (Arts. 40–43). Two 2022 reforms, in April and December, changed language and state-level spending limits without touching the allocation rules. The substantive obligations a ministry faces therefore

do not change at the 2018 transition, which matters for the design in Section 4: the legal environment is constant across the comparison, so nothing in the reallocation I document can be attributed to a change in what the law permitted.

Article 8 of the law requires that ministries select outlets by cost efficiency, defined as the cost of reaching the campaign's stated audience. This is the legal standard the cost-efficiency benchmark in Section 5 operationalizes. The law names no auditor for that requirement and attaches no penalty to departing from it, which is why the deviations I measure are best read as discretion rather than error.

The full text of each enacted version is available in `source/lgcs/` of the project repository.

A.2 Government Advertising Descriptives

Table A1: **Government Advertising Budgets by Ministry and Party (Millions MXN)**

Ministry	**PRI**		**Morena**	
	Millions	%	Millions	%
Health	361.82	21.0	366.14	46.0
Finance	235.09	14.0	90.99	11.0
Education	228.66	14.0	1.36	0.0
Energy	121.22	7.0	91.60	11.0
Social Development	180.11	11.0	29.30	4.0
Environment	141.00	8.0	48.56	6.0
Communications and Transportation	104.98	6.0	21.91	3.0
Labor	83.58	5.0	15.26	2.0
Interior	35.96	2.0	31.21	4.0
Navy	34.32	2.0	30.68	4.0
Agrarian, Territorial and Urban Development	59.37	4.0	0.00	0.0
National Defense	10.85	1.0	37.36	5.0
Foreign Affairs	29.87	2.0	5.08	1.0
Presidency	20.41	1.0	8.07	1.0
Culture	12.28	1.0	11.22	1.0
Economy	15.18	1.0	0.00	0.0
Agriculture	7.50	0.0	6.16	1.0
Tourism	8.65	1.0	1.15	0.0
Public Administration	0.00	0.0	1.07	0.0

The table shows government advertising expenditures during the two administrations across federal ministries. Total spending is reported in millions of Mexican pesos, and the percentage represents the proportion of each ministry's budget for a given administration.

Figure A1: **Share of Advertising Campaigns by Content Type and Party**

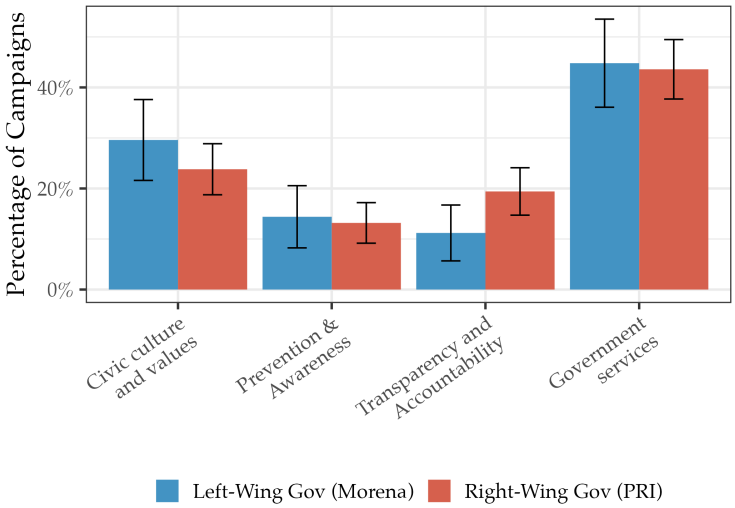

The figure shows the distribution of government advertising campaigns across campaign types, as reported by the contracting government ministry. Campaigns are classified by the party in power at the time of placement. The sample includes 273 campaigns during the PRI administration and 125 campaigns during the Morena administration, covering the period 2014–2022.

Figure A2: **Share of Advertising Campaigns Targeting Each Demographic Group by Party**

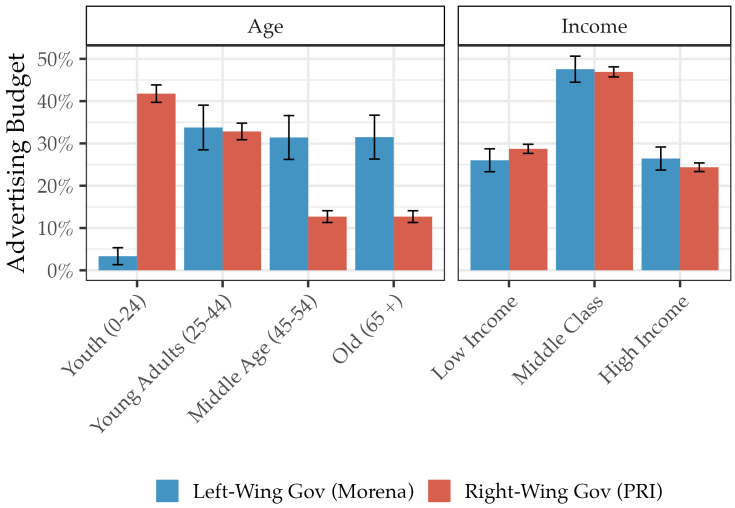

The figure shows the demographic groups targeted by government advertising campaigns conditional on the campaign specifying a target audience, as reported by the contracting government entity. Campaigns are classified by the party in power at the time of placement. Campaigns that report no specific targeting are excluded. The full sample consists of 398 government communication campaigns between 2014 and 2022.

Table A2: **Bipartisan Government Advertising Campaigns**

Ministry	Type	Title	Total Spending (USD)	
			PRI	Morena
Health	Prevention	Obesity Prevention	1,946,087	2,786,190
Interior	Civic Culture	Human Rights	66,200	1,895,410
Navy	Government Services	National Maritime Authority	1,325,999	157,456
Navy	Civic Culture	The Navy Close to You	1,181,894	131,892
Defense	Civic Culture	Military Commemorations	891,358	276,259
Foreign Affairs	Government Services	Protection of Mexicans Abroad	1,089,583	35,021
Education	Civic Culture	Sports in Mexico	1,087,621	25,974
Defense	Civic Culture	Commitment to Mexico	432,652	379,743
Defense	Civic Culture	Army and Family Outreach	384,237	58,197
Presidency	Civic Culture	In Search of Equality	153,061	209,533
Tourism	Prevention	Green Angels Road Service	302,958	11,356
Culture	Civic Culture	Linguistic Diversity	164,744	145,441
Defense	Civic Culture	Military Education	214,854	44,258
Tourism	Transparency	Local Tourism Promotion	214,352	23,920
Energy	Civic Culture	Daylight Saving Time	109,721	4,525
Education	Civic Culture	National Sports Award	2,710	18,701

The table shows government advertising campaigns that did not change their target audience during the sample period. Spending during the PRI administration covers 2014 to November 30, 2018; spending during the Morena administration covers December 1, 2018 to 2022.

Figure A3: **Example of a Cross-Administration Campaign at the Ministry of Defense**

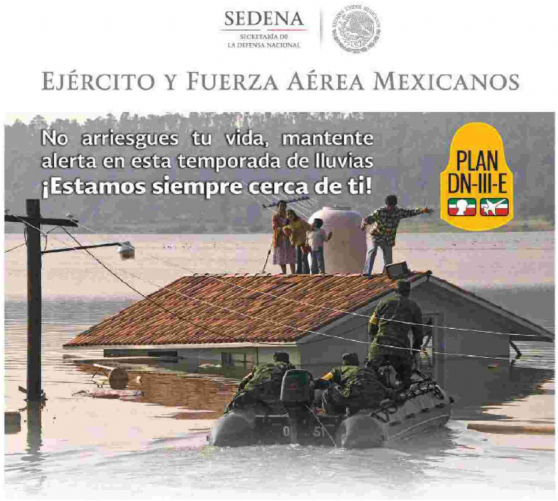

(a) **PRI Administration (2014–2018)**

(b) **Morena Administration (2018–2022)**

B Data

B.1 Media Market

Table B1: **Summary Statistics for News Articles**

Newspaper	N Total Articles	Articles per Year	
		Mean	SD
La Prensa	198,133	22,014.78	5,148.32
Esto	129,351	14,372.33	8,254.49
Ovaciones	245,451	27,272.33	7,839.42
Reporte Indigo	81,175	9,019.44	1,916.97
Publimetro	72,942	8,104.67	2,795.21
El Universal	399,328	44,369.78	7,629.34
Reforma	622,705	69,189.44	9,873.62
La Jornada	397,990	44,221.11	3,524.11
Excelsior	545,266	60,585.11	9,675.98
Milenio	339,008	37,667.56	8,352.47
La Razon	59,330	17,703.33	2,460.09
El Financiero	203,748	22,638.67	5,182.94
24 Horas	157,708	17,523.11	2,282.09
La Crónica de Hoy	253,724	28,191.56	8,771.31
El Sol de México	373,994	41,554.89	13,712.01
El Heraldo	143,123	17,890.38	11,760.17
Diario de México	83,352	9,261.33	3,337.80
El Economista	261,679	29,075.44	3,260.44

This table reports summary statistics of news articles for 18 national newspapers from 2014 to 2022. N Total Articles shows the total number of articles published by each outlet over the entire period. Articles per Year Mean and SD report the average number of articles per year and the corresponding standard deviation.

Figure B1: **Reader Demographics Distribution Across Newspaper Outlets (2022)**

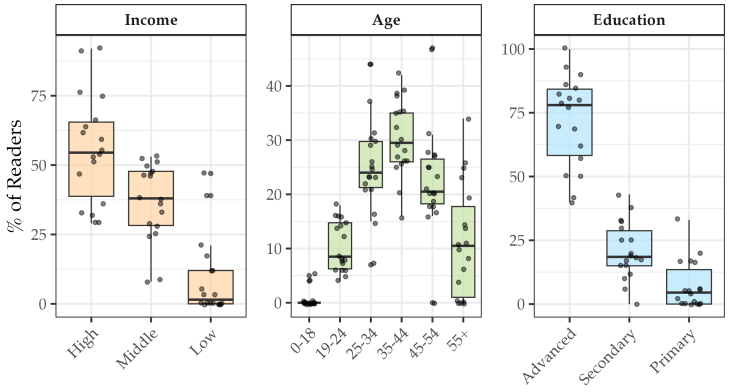

This figure displays the distribution of reader demographics across 18 newspaper outlets in 2022. Each panel shows boxplots representing the variation in the share of readers within each demographic category. Income levels are classified as High (AB/C+), Middle (C/D+), and Low (D/E). Education categories include Advanced, Secondary, and Primary education.

Figure B2: **Cumulative Market Share (2014–2022)**

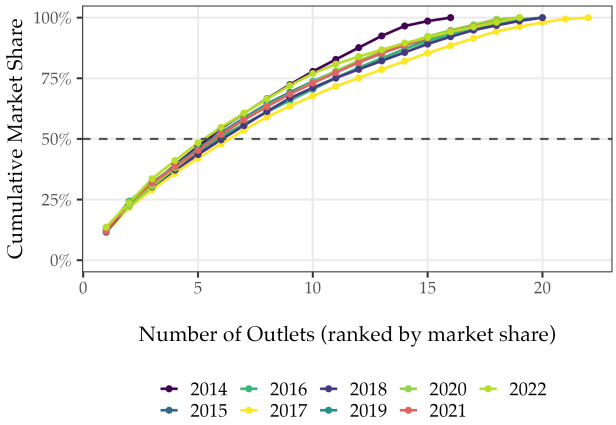

This figure shows the cumulative distribution of market share across the universe of national print media outlets for each year from 2014 to 2022. Outlets are ranked within each year by their circulation, from largest to smallest.

Table B2: **Ownership Structure of National Newspapers**

Newspaper	In sample	Owner	Conglomerate	**Market share**	
				Mean	SD
El Grafico		Ealy Ortíz	Grupo Universal	12.30	0.75
La Prensa	1	Vázquez Raña	OEM	11.50	1.13
Esto	1	Vázquez Raña	OEM	7.57	0.80
Ovaciones	1	Vázquez Raña	OEM	6.81	0.75
Reporte Indigo	1	Garza García	Mac Multimedia	6.69	1.29
Publimetro	1	Rivas Garza	Publicaciones Metropolitanas	6.66	0.83
El Universal	1	Ealy Ortíz	Grupo Universal	6.16	1.59
Reforma	1	Junco de la Vega	Grupo Reforma	6.09	0.66
La Jornada	1	–	Demos	5.23	1.81
Mas por Mas		–	–	4.85	2.52
Uno Mas Uno		Libien Kaui	–	4.77	1.35
Excelsior	1	Vázquez Aldir	Grupo Imagen Multimedia	4.34	1.03
Milenio	1	González Family	Multimedios	4.32	0.47
La Razon	1	Garza Cantú	LRHG Informativo	4.19	0.65
El Financiero	1	Arroyo Rodríguez	Grupo Lauman	3.83	0.89
24 Horas	1	Zabludovsky	Información Integral 24/7	3.67	1.13
La Crónica de Hoy	1	Kahwagi Gastine	La Crónica	3.16	0.50
El Sol de México	1	Vázquez Raña	OEM	2.91	0.63
El Heraldo	1	Mieres Zimmermann	Grupo Andrade	2.09	0.18
Diario de México	1	–	Editorial DDM	1.85	0.21
El Economista	1	Nacer Gobera	–	1.73	0.37
Eje Central		Oliver Fernández	OFEM Media Group	0.63	0.09

Market share is computed as the outlet's share of total national print circulation in a given year, averaged over 2014–2022. Ownership and conglomerate information is obtained from the National Print Media Registry (PNMI) and the Media Ownership Monitor Mexico database. When ownership information is not reported, this reflects either the absence of a clearly identifiable controlling shareholder or a cooperative ownership structure.

B.2 Constructing the Market Advertising Price

The market advertising price p_{it} is the price a private buyer pays outlet i for one square centimetre of space in year t. It enters the benchmark of Section 5 as the cost at which a ministry converts spending into insertions, so it has to be observed independently of what the government paid.

Table B3: **Classification of Journalists in Database (2014–2022)**

Employment Pattern	Category	N	%
Exclusive Affiliation	*Total*	2,560	**43.10**
	Left-Wing Outlets	1,163	12.97
	Right-Wing Outlets	1,397	15.58
Primary Affiliation	*Total*	2,214	**37.27**
	Left-Wing Outlets	869	9.69
	Right-Wing Outlets	1,345	15.00
Outlet Movers	*Total*	380	**6.39**
	Within-Ideology Moves	214	2.38
	Left → Right	78	0.87
	Right → Left	75	0.84
	Multiple Switches	14	0.16
Freelance/Multi-Outlet		785	**13.21**
Total		**5,939**	**100.00**

This table summarizes the classification of 5,939 journalists in the database. *Exclusive Affiliation* journalists publish exclusively at one outlet. *Primary Affiliation* journalists publish predominantly at one outlet but occasionally contribute to others. *Outlet Movers* switched between outlets during the period, either within the same ideological category or across ideologies. *Freelance/Multi-Outlet* exhibit diverse publication patterns across multiple outlets.

I take it from the *tarifarios* published quarterly by mediavyasa, a commercial media directory that reproduces the published price list of each outlet. The source covers editions 86 to 121, one per quarter from 2014Q1 to 2022Q4 with no gaps. From each edition I extract the Mexico City daily-newspaper block of the *Prensa* section, which is where the national titles appear.[12]

Each quoted price is a full-page rate. I divide it by the outlet's page area to obtain a price per square centimetre, deflate to 2022 pesos, and convert at the 2022 average exchange rate of 20.11 pesos per dollar. Averaging within outlet and year gives a panel of 157 outlet-years across the 18 outlets in the sample. Coverage is not complete: an outlet is missing from a quarter when it does not appear in that edition's CDMX daily block, and I audit every missing cell against the source PDF rather than imputing it.

Two features of the series matter for how it is used. First, it is a published list price, quoted before any negotiated discount, so it is an upper bound on what a private buyer transacts at. Because the benchmark converts a ministry's budget into insertions by dividing by this price, an upper bound on the price is a lower bound on the insertions a budget could have bought, and therefore a conservative estimate of the reach the observed allocation forgoes. Second, it is quoted for a generic full page. Appendix E.1 reports what position and colour add to that price and shows the range is too narrow to account for the gap between what ministries paid and the market rate.

Averaged over the sample the market advertising price is \$7.83 per square centimetre under the PRI and \$7.52 under Morena, against \$21.93 and \$13.41 paid by the government, as Table 1 reports.

[12]The block is headed *México, D.F.* in editions 86–97, *CDMX* in 98–108, and *Ciudad de México* in 109–121. Blocks covering the state press and the CDMX weeklies use a different column set and are not extracted. The edition number maps onto the calendar quarter as year = $2014 + \lfloor (e-86)/4 \rfloor$ and quarter = $(e-86) \bmod 4 + 1$, verified against the PDF creation dates at editions 86, 100, 102 and 121.

C Measurement of Ideology

C.1 Alternative Ideology Measures

Table C1 replicates the main analysis using the continuous ideology score rather than the binary classification. The three outcomes are the same as in Table 3. The interaction term is Left-Wing Gov × Outlet's Ideal Score, where higher scores indicate more right-leaning positions. A negative coefficient indicates that more right-wing outlets receive less advertising under the left-wing government, consistent with ideological favoritism.

Table C1: **Robustness to a Continuous Ideology Score**

	All Campaigns			Bipartisan Campaigns	
	Prob(Gov Ad > 0) (1)	**IHS(Gov Ad) (2)**	**Price Premium (3)**	**Prob(Gov Ad > 0) (4)**	**IHS(Gov Ad) (5)**
Left-Wing Gov × Outlet's Ideal Score	−0.161*	−2.193*	−0.212***	−0.116*	−1.357**
	(0.061)	(0.838)	(0.027)	(0.049)	(0.474)
Mean (Pre-2018, Left)	0.46	—	—	0.26	—
Mean (Pre-2018, Right)	0.66	—	—	0.36	—
Year FE	✓	✓	✓	✓	✓
Outlet × Ministry FE	✓	✓		✓	✓
Campaign FE			✓		
Observations	2,880	2,880	2,617	1,440	1,440

All columns use the full outlet–ministry–year panel across all ministries, so the number of observations exceeds that in Table 3, which restricts to the ten ministries running bipartisan campaigns. Outlet's Ideal Score is the continuous GST score; higher values indicate more right-leaning positions. A negative coefficient indicates that more right-wing outlets receive less advertising under the left-wing government. SEs clustered at the newspaper level. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$.

Table C2 replicates the analysis using a three-category ideology classification. Outlets are divided into terciles of the continuous score (6 outlets each): left-wing, centrist, and right-wing. The reference category is right-wing outlets. Both left-wing and centrist outlets gain advertising relative to right-wing outlets under the left-wing government, with left-wing outlets gaining the most.

Table C2: **Robustness to a Three-Category Ideology Classification**

	All Campaigns			Bipartisan Campaigns	
	Prob(Gov Ad > 0) (1)	IHS(Gov Ad) (2)	Price Premium (3)	Prob(Gov Ad > 0) (4)	IHS(Gov Ad) (5)
Left-Wing Gov × Left-Wing Outlet	0.288***	3.890***	0.283**	0.171*	1.924*
	(0.063)	(0.774)	(0.089)	(0.066)	(0.688)
Left-Wing Gov × Centrist Outlet	0.229**	3.119***	0.008	0.133*	1.457*
	(0.060)	(0.764)	(0.139)	(0.062)	(0.670)
Mean (Pre-2018, Left)	0.46	—	—	0.26	—
Mean (Pre-2018, Right)	0.66	—	—	0.36	—
Year FE	✓	✓	✓	✓	✓
Outlet × Ministry FE	✓	✓		✓	✓
Campaign FE			✓		
Observations	2,880	2,880	2,617	1,440	1,440

Reference category: right-wing outlets (6 most right-wing by GST score). Left, centrist, and right-wing categories are terciles of 18 outlets ranked by continuous GST score (6 outlets each). SEs clustered at the newspaper level. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$.

C.1.1 Outlets Scored Without Editorials

Diario de México, El Heraldo de México, and Esto publish no editorial section and are scored by applying the estimated phrase loadings to their news articles, as described in Section 3.2. Two features of these corpora differ from the editorial corpora used for the other fifteen outlets. They are far larger, with 26,493 articles for Esto against a median of 469 editorials per outlet, and they contain agency copy, which carries no editorial stance. Table C3 addresses both. I first restrict each corpus to journalist-authored content, dropping wire and agency material, and then draw 500 bootstrap samples at the median and mean editorial corpus size, rescoring each draw.

The journalist restriction moves each score by less than 0.02, and the bootstrap confirms that corpus size is not what drives the classification of Diario de México or Esto, whose signs are stable in 99.8 and 100 percent of draws respectively. El Heraldo de México is genuinely ambiguous. Its score sits within 0.03 of the anchor midpoint and its sign is stable in only 78 to 83 percent of draws. It is coded as left-wing in the main analysis. The outlet enters a single ministry–outlet series and the main estimates are unchanged when it is coded as right-wing.

Table C3: **Bootstrap Validation for Outlets Scored From News Articles**

	Articles (all)	Journalist articles	Score (all)	Score (journalist)	95% CI (bootstrap)	Sign stable
Diario de México	5,102	4,539	+0.114	+0.102	[+0.025, +0.174]	99.8%
El Heraldo de México	3,081	2,903	−0.034	−0.027	[−0.097, +0.051]	77.8%
Esto	26,493	21,502	−0.296	−0.296	[−0.347, −0.245]	100.0%

Scores are recentered at the anchor midpoint, so negative values indicate left-leaning language. Journalist articles exclude wire and agency copy. Bootstrap draws 500 samples of 469 articles, the median editorial corpus size, from the journalist-authored corpus and rescores each draw. Sign stable reports the share of draws whose sign matches the full-corpus classification. At the mean editorial corpus size of 596 articles the corresponding shares are 99.8, 82.6, and 100.0 percent.

C.2 Ideology and Audience Composition

I compare the ideology score against the audience composition each outlet reports to the National Print Media Registry. The profiles cover three income brackets, six age brackets, three education levels, sex, and an urban–rural split, and are available for all 18 sample outlets in every year from 2014 to 2022. Figure C1 presents the comparison.

Audience similarity is unrelated to ideological similarity. Across the 153 outlet pairs the correlation between audience distance and the gap in ideology scores is −0.025, with a Mantel permutation p-value of 0.87. Pairs on the same side of the ideological divide are no closer in audience space than pairs on opposite sides, with mean distances of 5.46 and 5.53 respectively.

Among individual dimensions, only age is strongly related to the score. Right-leaning outlets draw older readers, with a correlation of +0.71 for the share aged 55 and above and −0.49 for the share aged 25 to 34. Income runs in the expected direction but weakly, at +0.34 for the highest bracket and −0.43 for the lowest, neither significant in a cross-section of 18 outlets. Education correlations are below 0.03 in absolute value.

The audience space is organized by outlet genre rather than politics. The first two principal components account for 28 and 23 percent of the variance, and the first is uncorrelated with ideology at −0.01. It separates the tabloid and sports titles *La Prensa* and *Esto* from the business and broadsheet titles *El Financiero* and *Excélsior*, while *La Jornada* and *Reforma*, the two ideological poles of the sample, sit close together. The second component correlates with the score at +0.34, reflecting the age gradient.

The measure is therefore not a proxy for readership, and the near-independence of alignment and audience composition supports the identification argument in Section 4. A government reallocating advertising toward demographic groups it wishes to reach would not, as a byproduct, reallocate toward ideologically aligned outlets, because the two dimensions do not line up. Age is the exception, and remains a live alternative explanation for any result that operates through the age composition of readers.

Figure C1: **Ideological Similarity and Audience Composition**

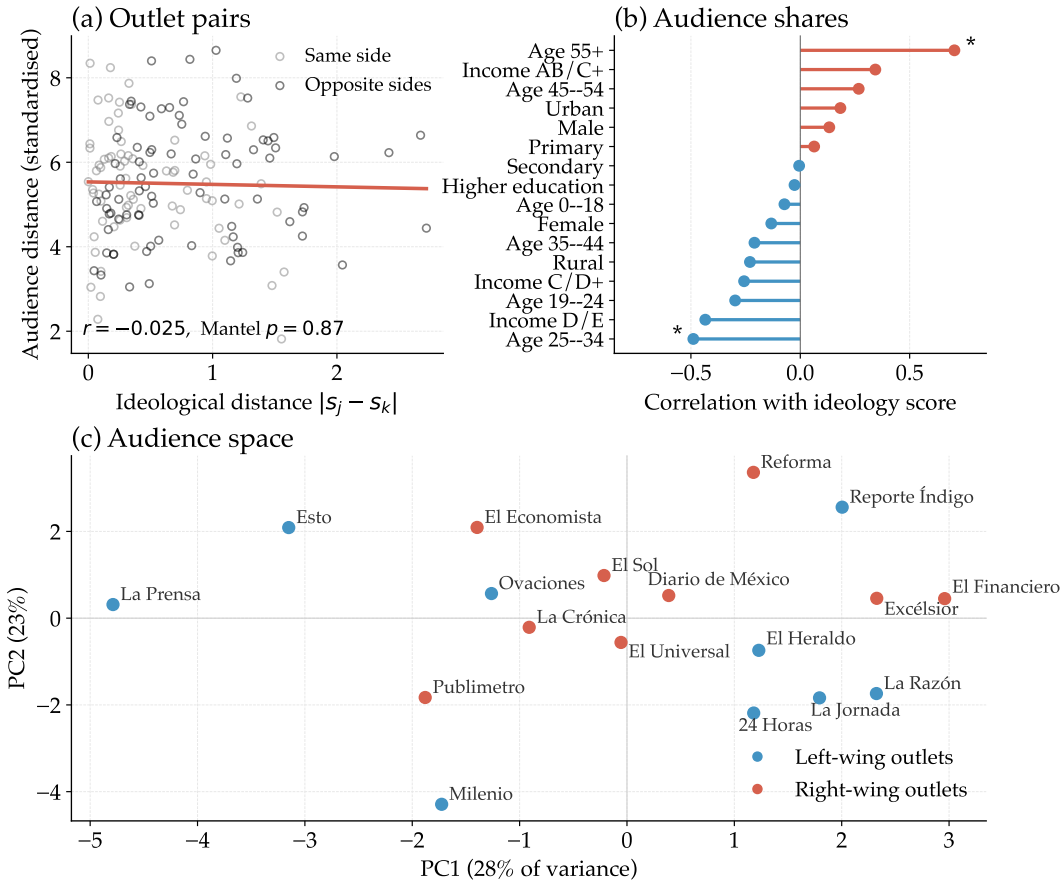

Audience shares reported to the National Print Media Registry, averaged over 2014–2022. Panel (a) plots audience distance against ideological distance for the 153 outlet pairs, with a linear fit and a Mantel permutation p-value. Panel (b) correlates each audience share with the ideology score; asterisks mark significance at 5 percent. Panel (c) plots the first two principal components of the audience profiles. Higher scores are more right-leaning.

D Robustness for the Allocation Estimates

D.1 Alternative Estimators for the Intensive Margin

Following Chen and Roth (2024), Table D1 reports two alternative estimators for the intensive margin. Column (1) presents a Poisson PPML regression, which is scale-invariant; the coefficient can be interpreted as $\exp(\hat{\beta}) - 1 \approx 222\%$, representing the proportional increase in expected spending for aligned outlets. Column (2) reports Lee bounds (Lee, 2009) on the conditional intensive margin effect — the effect for ministry–outlet pairs that receive positive advertising under both administrations — which range from \$11,668 to \$373,764. Both exercises confirm a large positive effect on the intensive margin.

Table D1: **Intensive Margin Robustness Under Poisson and Lee Bounds**

	Poisson PPML (semi-elasticity) (1)	**Lee Bounds** (2022 USD, conditional intensive margin) (2)
Left-Wing Gov × Left-Wing Outlet	1.169**	[+11,668 ; +373,764]
	(0.452)	
Year FE	✓	
Outlet × Ministry FE	✓	
Observations	1,149	

Sample restricted to 10 bipartisan ministries. Poisson drops outlet–ministry cells with all-zero outcomes; $\exp(1.169) - 1 \approx 222\%$. *Lee bounds identify the average treatment effect on the conditional intensive margin under a monotonicity assumption (Lee, 2009). SE clustered at the newspaper level.* $^{**}p<0.05$, $^{*}p<0.1$.

D.1.1 Event Studies

Figure D1: **Event Studies of Differential Government Advertising by Outcome**

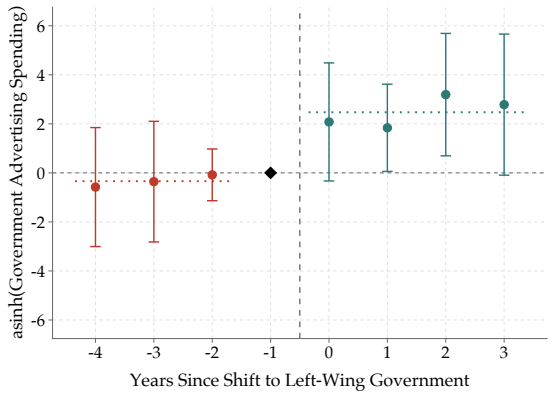

(a) **Outlet–Ministry, IHS(Spending)**

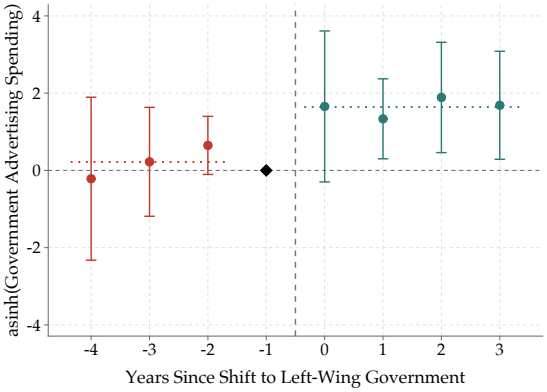

(b) **Outlet–Ministry, Bipartisan Campaigns, IHS(Spending)**

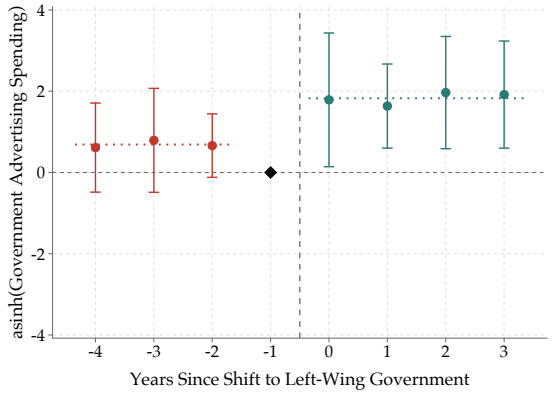

(c) **Campaign–Outlet, IHS(Spending)**

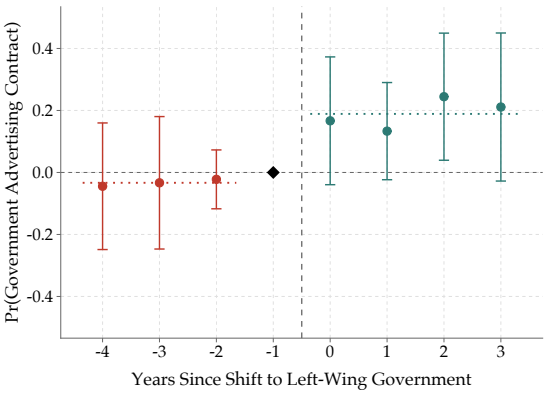

(d) **Outlet–Ministry, Any Spending**

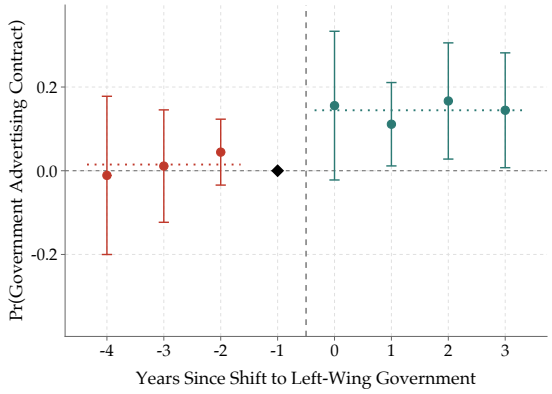

(e) **Outlet–Ministry, Bipartisan Campaigns, Any Spending**

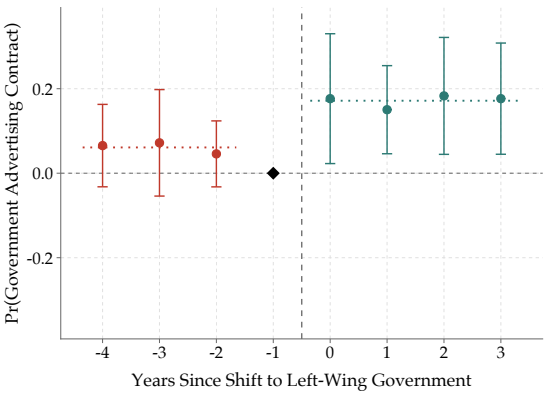

(f) **Campaign–Outlet, Any Spending**

This figure presents event study estimates for all six specifications reported in Table 3. All panels include outlet–ministry or outlet–campaign fixed effects. Coefficients represent the interaction between event time (years relative to the December 2018 government transition) and outlet ideology. Error bars show 95% confidence intervals with standard errors clustered at the outlet level.

E Robustness for the Benchmarks

E.1 Prices, Placement, and the Source of the Reach Gap

Ministries pay more per unit of space than private advertisers do, and this appendix asks what that gap is and whether it drives the reach shortfall in Section 5. Measured against the price private buyers pay for the same space, a ministry pays 1.56 times as much at the median and 1.66 times as much in aggregate. Two explanations compete. Ministries may buy better placements, since a front page in color costs more than an inside page in black and white and the market price averages across both. Or they may pay more for the same product.

Placement cannot account for a gap of that size. Outlets quote their own market advertising prices by position and by color, and Table E1 reports what each combination costs relative to the median quote. The most expensive product available, an odd page in color, costs 1.13 times the median, and the cheapest, an even page in black and white, costs 0.57 times. A ministry that bought nothing but the single most expensive placement on every card would therefore pay 13 percent above the middle of the market, against the 56 percent it actually pays. Premium buying explains at most a quarter of the gap.

Table E1: **Market Advertising Price per Square Centimeter by Position and Color**

	Quotes	Price per cm^2 (median)	Relative to the median quote
No fixed position, colour	213	116.2	1.13×
Odd page, colour	90	114.3	1.12×
Unspecified, colour	1167	108.9	1.06×
Even page, colour	180	94.2	0.92×
No fixed position, black and white	112	84.6	0.83×
Unspecified, black and white	551	75.9	0.74×
Odd page, black and white	40	64.4	0.63×
Even page, black and white	92	58.2	0.57×

Quotes drawn from the outlets' own published tariffs, collected from commercial sites and archived copies, restricted to those reporting a positive price per square centimeter. The final column divides each combination's median by the median across all quotes. Position labels follow the cards: odd and even refer to the page on which an advertisement runs, and no fixed position means the outlet places it at its discretion.

The remaining gap is concentrated where the recorded number of advertisements is least trustworthy. A cell recording one advertisement carries a median price ratio of 4.08 and a ninetieth percentile of 27.2, while a cell recording more than twenty carries a median of 0.96 and a ninetieth percentile of 1.72. Ministries buying at scale therefore pay almost exactly the private market rate. Single-advertisement cells hold 6.6 percent of the advertisements but 34 percent of the spending, which is what a systematic undercount of insertions would produce and what a genuine price premium would not: no plausible placement costs twenty-seven times its market rate.

That distinction determines how much of the reach gap overpayment can explain, and Table E2 shows it explains almost none. Across all cells, pricing a ministry's own outlet choices at market rates would recover 96,588 readers per campaign. Restricting to cells recording at least three advertisements, where the price

ratio approaches one, that figure falls to 4,422, and restricting to six or more it turns slightly negative. The misplacement term moves in the opposite direction and stays close to half a million readers throughout.

Table E2: **Reach Gap Decomposition by Number of Advertisements Recorded**

	Delivered (readers)	Overpayment (readers)	Misplacement (readers)
All cells	178,486	109,817	439,227
Cells with 3 or more	136,565	11,607	469,239
Cells with 6 or more	101,530	-1,839	476,531

Readers per campaign at ϕ = 0.10, averaged across campaigns. Delivered reach uses the recorded number of advertisements. The overpayment column prices a ministry's own outlet choices at the private market rate and reports the additional readers that would follow. The misplacement column moves the same money to the outlets that reach the named audience most cheaply. Rows restrict the sample to cells recording at least the stated number of advertisements, retaining all cells with none.

Two further questions remain about prices, and they are separate from the level of the gap. The first is whether the gap is a gap for the *same* product, since the market benchmark averages across sizes and placements. The second is whether ideological alignment predicts what a ministry pays. Table E3 answers both by comparing two outlets carrying the same campaign at the same advertisement size, so the identity of the product is held fixed rather than averaged over.

Alignment does not move the price. The point estimates are small relative to their standard errors and mixed in sign, and this holds for the price paid, for the price relative to the outlet's published list price, and for the price per reader. With eighteen clusters the normal approximation is not reliable, so the p-values come from a wild cluster bootstrap with the null imposed, and they are larger still. The estimates are also conditional in a way the extensive-margin result is not: a price is observed only where a contract was placed, and alignment raises the probability of placement from 32.7 to 43.9 percent, so the sample the price is measured on is selected on a variable the treatment moves. Trimming the aligned distribution by that excess share gives Lee (2009) bounds, reported in the final row, and the identified interval contains zero. Favoritism operates on whom the government buys from and how much it spends, not on the unit price it pays.

Table E3: **Price Paid to Aligned Outlets, Conditional on a Contract**

	$\log(p_{gov})$		$\log(p_{gov}/p_{mkt})$		$\log(p_{gov}/\text{circ})$	
	(1)	(2)	(3)	(4)	(5)	(6)
Left-Wing Government × Left-Wing Outlet	0.136	0.219	-0.038	0.022	-0.035	0.276
	(0.216)	(0.169)	(0.221)	(0.219)	(0.244)	(0.244)
	[0.544]	[0.358]	[0.835]	[0.929]	[0.978]	[0.404]
Outlet × Size FE	Yes	Yes	Yes	Yes	Yes	Yes
Campaign FE		Yes		Yes		Yes
Observations	8,597	8,597	8,597	8,597	8,597	8,597
Lee (2009) bounds on column (4): [-0.270, 0.278], trimming 25.4% of the aligned distribution.						

Note: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered on outlet in parentheses; wild cluster bootstrap p-values with the null imposed in brackets (2,000 Rademacher draws), which is the relevant inference with 18 clusters. The market price is the outlet-year published list price from the mediavyasa rate cards. A price is observed only where a contract was placed, and alignment raises the probability of placement, so the coefficient is not point identified; the Lee bounds give the sharp identified interval under the assumption that alignment never reduces the chance of placement.

The misplacement term is the one that does not depend on the insertion count. Both quantities it compares, what a ministry's own choices deliver at market prices and what the benchmark delivers, convert spending into insertions by dividing by the same price, so an undercount cancels. Ministries reach 275,073 readers with the outlets they chose and would reach 764,514 with the outlets that serve the campaign's audience, a factor of 2.8. That comparison stands whatever the recorded advertisement counts turn out to be.

E.2 Estimating Outlet-Level Copartisan Shares

I measure each outlet's copartisan reach using geographic variation in circulation. For each outlet j, year t, and state s, I compute the weight w_{js} as outlet j's share of its own total circulation that falls in state s:

$$w_{jst} = \frac{\text{circ}_{jst}}{\sum_{s'} \text{circ}_{js't}} \tag{6}$$

The copartisan share for party p in election year τ is then the circulation-weighted average of state-level vote shares:

$$\widehat{\text{Copartisan}}^{p}_{j\tau} = \sum_{s} w_{jst} \cdot v^{p}_{s\tau} \tag{7}$$

where $v^{p}_{s\tau}$ is the official vote share of party p in state s at election τ. Circulation data come from a balanced panel of regional edition figures for 18 national newspapers.[13]

A key assumption underlying the voter-targeting benchmark is that Morena and PRI governments could, in principle, select outlets whose readership skews toward their own supporters. Figure E1 tests this by plotting each outlet's ideology score, a continuous measure of the partisan valence of its editorial language (Gentzkow et al., 2019), against its estimated Morena copartisan share in 2018. The ideological gap in copartisan reach is negligible. Left-wing and right-wing outlets attract nearly identical partisan audiences, and all 18 outlets fall within a 10-percentage-point band. The average Morena copartisan share is 41.6% among left-wing outlets and 41.8% among right-wing outlets. Voters in Mexico do not sort into newspapers by ideology, which limits the scope for outlet-level partisan targeting.

[13]Circulation is measured as copies per 10,000 inhabitants, obtained from the National Print Media Registry (PNMI). State-level vote shares come from official results published by the National Electoral Institute (INE) for the 2015, 2018, and 2021 midterm and presidential elections. Because advertising campaigns span multiple calendar years, I assign each year to the nearest election cycle: 2014–2016 to the 2015 midterm, 2017–2019 to the 2018 presidential election, and 2020–2022 to the 2021 midterm.

Figure E1: **Outlet Ideology Score and Morena Copartisan Share (2018)**

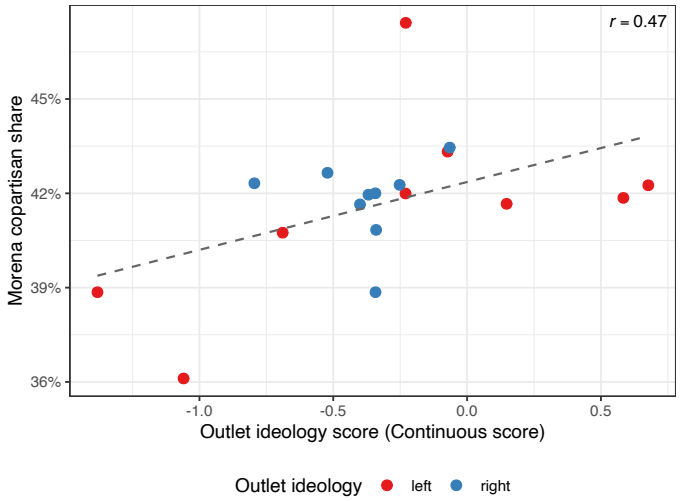

Each point is one of 18 national newspapers. The x-axis is the outlet's continuous ideology score from Gentzkow et al. (2019); more negative values indicate more left-leaning editorial language. The y-axis is the estimated share of the outlet's readers who supported Morena, computed as a circulation-weighted average of state-level Morena vote shares (2018 presidential election). The dashed line is an OLS fit. r denotes the Pearson correlation.

The geographic measure can be checked against an independent construction. Individual polling microdata registered with the National Electoral Institute records vote intention alongside state, gender, education, age and income; I estimate vote choice on those covariates and apply the fitted model to each outlet's registry audience profile, giving a demographic estimate of copartisan reach that shares no inputs with the geographic measure beyond state of residence. Figure E2 plots the two against each other across the 76 outlet–years where both are available. They correlate at 0.80. On the demographic measure left-wing outlets reach a Morena share of 30.1 percent against 30.5 percent for right-wing outlets, a gap of −0.4 percentage points; on the geographic measure the figures are 33.3 and 33.4 percent, a gap of −0.2 points. The absence of an ideological gap in copartisan reach is therefore not an artifact of building the measure from circulation geography.

Figure E2: **Demographic and Geographic Estimates of Copartisan Reach**

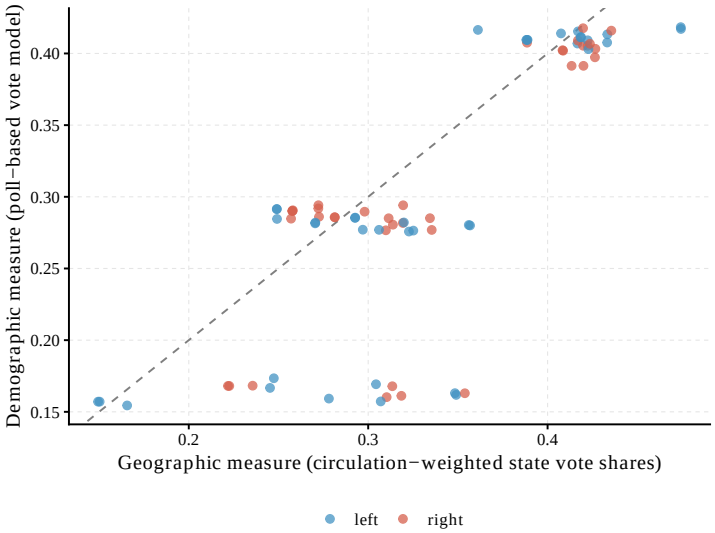

Estimated Morena share of each outlet's readership, 76 outlet–years. The horizontal axis weights state-level official vote shares by the outlet's circulation across states. The vertical axis applies a vote-choice model estimated on polling microdata to the outlet's registry audience profile. The dashed line is the 45-degree line; colors denote the outlet's ideology classification from Section 3.2.

E.3 The Exposure Rate ϕ

The technology of Section 5 has one free parameter, the probability ϕ that a reader sees any one insertion. Rather than defend a value, I report the estimate across the range over which the budget constraint binds. As $\phi \to 0$ reach becomes linear in spending and the optimum is a corner placing the entire budget in the cheapest outlet per target reader, the maximum-concentration allocation. The upper end is degenerate and is therefore not a comparable column. As $\phi \to 1$ a single insertion saturates an outlet, so additional spending buys no further reach and the budget ceases to bind. The corresponding allocation spends only 57 percent of the average campaign budget and underspends in 251 of 383 campaigns, so a dollar deviation measured against it would mechanically contain the unspent remainder rather than a reallocation.

Table E4: **Intensive-Margin Deviation Across the Full Range of ϕ**

	$\phi \to 0$ (corner)	$\phi = 0.05$	$\phi = 0.10$	$\phi = 0.15$
	(1)	(2)	(3)	(4)
Aligned (binary)	38,840	21,763**	20,876**	20,504**
	(26,210)	(8,803)	(8,912)	(8,909)
Aligned (continuous)	38,440**	28,987***	28,749***	28,304***
	(16,124)	(5,560)	(6,008)	(6,141)
Ad Campaign FE	✓	✓	✓	✓
Outlet FE	✓	✓	✓	✓
Observations	6,711	6,711	6,711	6,711

*Deviations from the cost-efficiency benchmark, in 2022 USD. SE clustered by outlet. ***p<0.01, **p<0.05, *p<0.10. Column (1) is the $\phi \to 0$ corner and columns (2) to (4) are interior. Column (3) reproduces Panel B, column (1) of Table 4.*

Table E5 reports the distribution across campaigns of the budget multiplier λ recovered in Section 5, of

the reservation cost per target reader it implies, and of the market advertising prices those are compared against. The reservation cost sits inside the observed price distribution at every quantile, which is what produces interior solutions rather than corner ones.

Table E5: **Recovered Shadow Price and Reservation Cost Across Campaigns**

	p10	p25	Median	p75	p90
Shadow price λ (target readers per USD)	0.15	0.52	1.07	1.93	3.95
Reservation cost $\bar{c}$ (USD per 1,000 target readers)	27	55	99	203	719
Rate-card cost per 1,000 target readers	16	56	88	134	231

One observation per campaign, at $\phi = 0.10$. λ is recovered by bisection on the budget constraint under equation 3. The reservation cost is $\bar{c} = \kappa/\lambda$ with $\kappa = -\ln(1-\phi)$. The last row is the registry market advertising cost per thousand readers in the campaign's stated demographic, pooled across outlet–campaign cells.

The intensive-margin deviation is positive at every point in the bracket, as Table E4 reports, from \$38,608 at the concentration limit to \$20,168 at $\phi = 0.15$, and the continuous-score specification is significant at the five percent level at the corner and at the one percent level throughout the interior. The calibration of ϕ therefore does no work in the conclusion. Two quantities are more sensitive and should be read with that in mind. The number of outlets the efficient allocation uses rises from one at the concentration limit to 10.3 at $\phi = 0.10$ and 15.4 at the coverage limit, so the statement that campaigns are more concentrated than efficiency requires is qualitative rather than a point estimate. And the foregone-audience gap ranges from roughly three to five times delivered reach across the interior of the bracket.

E.4 Reader Sorting and the Voter-Targeting Benchmark

The voter-targeting benchmark in Section 5 builds copartisan reach from where an outlet circulates, so it misses any tendency of incumbent supporters to choose newspapers that share their politics. This appendix asks how strong that sorting would have to be before the result disappears. Writing ρ_{it} for the measured copartisan share, I tilt it by

$$\text{logit}(\tilde{\rho}_{it}) \;=\; \text{logit}(\rho_{it}) + \tau\, z_i, \qquad z_i = \tfrac{1}{2} \text{ if aligned, } -\tfrac{1}{2} \text{ otherwise} \tag{8}$$

so that e^{τ} scales the odds that an incumbent supporter reads an aligned rather than a misaligned newspaper. I apply the tilt in log-odds so that it means the same thing at any base rate, which matters because the incumbent's vote share differs sharply across the two administrations. I then rebuild the voter-targeting benchmark with $\tilde{\rho}_{it}$ and estimate equation 4 again at each value.

Table E6: **Deviations from the Voter-Targeting Benchmark Under Assumed Reader Sorting**

	Assumed sorting, odds ratio e^{τ}					
	1.00	1.28	1.65	2.12	2.72	4.48
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Deviation in 2022 dollars						
Aligned	22,198**	16,035	9,964	4,136	−1,274	−10,686
	(9,865)	(9,567)	(9,285)	(9,032)	(8,835)	(8,609)
Ideology score	31,038***	27,015***	23,003***	19,127**	15,543*	9,408
	(6,863)	(7,045)	(7,349)	(7,731)	(8,175)	(9,120)
Panel B. Deviation as share of campaign budget (pp)						
Aligned	3.01*	1.64	0.26	−1.08	−2.35	−4.65**
	(1.71)	(1.63)	(1.59)	(1.59)	(1.65)	(1.90)
Ideology score	4.08**	3.14*	2.18	1.25	0.39	−1.14
	(1.65)	(1.75)	(1.92)	(2.15)	(2.41)	(2.98)
Implied aligned−misaligned gap (pp)	−0.9	+3.2	+7.4	+11.5	+15.6	+23.8
Ad Campaign FE	✓	✓	✓	✓	✓	✓
Outlet FE	✓	✓	✓	✓	✓	✓
Observations	6,711	6,711	6,711	6,711	6,711	6,711

Intensive-margin estimates of equation 4, with the copartisan share replaced by $\tilde{\rho}_{it}$ from equation 8. Panel A takes the deviation in 2022 dollars and Panel B divides it by the campaign's total budget. Column (1) uses the measured share. The measured baseline odds ratio does not equal one, so the headings report the assumed multiplier applied to it rather than the resulting ratio. The final row reports the implied difference in mean copartisan share between aligned and misaligned outlets, which depends jointly on the multiplier and on the base rate. Standard errors clustered by outlet. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.10$.

How much sorting the result tolerates depends on the scale I measure the deviation on. In Table E6 I raise the assumed sorting from none in column (1) to a multiplier of 4.5 in column (6), rebuilding the voter-targeting benchmark and re-estimating equation 4 at each step. Panel A measures the deviation in dollars and survives the most sorting, since the continuous estimate stays significant at the five percent level out to a multiplier of 2.1 and at ten percent out to 2.7. Panel B divides each deviation by its campaign budget and survives less, since there the continuous estimate holds at ten percent only to 1.3 and the binary estimate loses significance as soon as I impose any sorting at all. The binary estimate turns negative in both panels once the multiplier passes 2.1, and by 4.5 the sign reverses at the one percent level in Panel B, so sorting that strong would make the benchmark want to buy aligned newspapers more heavily than ministries did. Taking the share scale as the binding one, the result survives as long as the odds of an incumbent supporter reading an aligned rather than a misaligned newspaper stay within roughly 1.3 times what I measure, which would put 24 percent of an aligned outlet's readers on the incumbent's side against 21 percent of a misaligned outlet's.

E.5 Reader Demographics and the Allocation

The benchmark treats every reader in a campaign's named groups as equally valuable. Gentzkow et al. (2024) instead allow each viewer an intrinsic value to advertisers, which raises the possibility that the government weights demographic groups differently and that the apparent misallocation reflects those weights rather than ideology. This appendix asks what weights would be required.

Replace the audience with $a_i(w) = N_i \sum_g w_g r_{ig}$, where r_{ig} is outlet i's share of readers in group g across

three income and six age categories, and choose $w \geq 0$ to minimize the squared distance between the resulting optimal allocation and the observed one.

The weights are not identified, for a reason that is exact rather than numerical. Income shares sum to one and age shares sum to one for every outlet, so adding a constant to all three income weights and subtracting the same constant from all six age weights leaves $a_i(w)$ unchanged for every outlet, verified to machine precision at any magnitude. The objective is correspondingly flat: fixing any single group's weight at zero and re-optimizing the rest costs at most 2.7 percent of the objective, and seven of the nine cost under half a percent. A further obstacle is that r_{ig} is time-invariant and outlet-specific, so $a_i(w)$ does not vary across campaigns and cannot in principle reproduce campaign-specific targeting. The estimated weights are therefore not reported.

What the exercise does deliver is a comparison. Table E7 reports how well each specification reproduces the observed allocation. The informative column is the last, the correlation between predicted and observed *within-campaign* shares, since the correlation in levels is largely driven by campaign budget. Splitting each budget equally across outlets, with no model at all, attains a level correlation of 0.49. Every audience-based specification predicts within-campaign allocation negatively. A single parameter loading circulation on outlet ideology reverses that, and outperforms nine free demographic weights.

Table E7: **Fit of Alternative Audience Weightings to the Observed Allocation**

	Objective ($\times 10^{13}$) (1)	**Correlation, levels** (2)	**Correlation, within-campaign shares** (3)
Equal weights, $a_i \propto N_i$	9.604	0.453	−0.103
Benchmark audience weight	9.627	0.454	−0.083
Best demographic weights (9 parameters)	9.273	0.479	−0.056
Circulation × ideology (1 parameter)	9.065	0.505	+0.110
Demographic weights + ideology	8.954	0.508	+0.062

Each row solves equation 2 campaign by campaign with the stated audience weight and compares the result to observed spending. The objective is the summed squared deviation across outlet–campaign cells. The ideology specification sets $a_i = N_i \exp(\theta \cdot score_i)$ *with* θ *estimated freely; the fitted value is* $\theta = 1.28$*. Column (3) computes the correlation of predicted and observed spending shares within each campaign, then averages.*

E.6 Robustness to Low-Cost Titles

The outlets the cost-efficiency benchmark most wants to buy are those with the lowest registry market advertising cost per thousand readers in the campaign's stated demographic, and the three cheapest by that measure are high-circulation popular and sports titles, namely La Prensa at $10.7, Esto at $16.0 and Ovaciones at $22.9, against $53.9 for El Universal and $61.7 for El Sol de México. Two concerns follow. Published list prices for these titles may understate the prices at which the government could actually transact, and registry audience shares may overstate how far a campaign placed in a sports paper reaches the demographic it names. Either would inflate the measured gap between the observed allocation and the benchmark.

Table E8 re-estimates equation 4 dropping these outlets. Column (2) removes the two sports titles

and column (3) removes all three cheapest. Dropping an outlet removes it from the choice set, so both benchmarks are re-solved on the retained outlets and each campaign's budget is reallocated only among them. Every estimate rises rather than falls. On the cost-efficiency benchmark the binary estimate goes from \$20,876 in the full sample to \$24,286 and \$27,183, and the continuous estimate from \$28,749 to \$29,364 and \$30,649; the voter-targeting estimates move in the same direction, from \$22,198 to \$26,239 and \$30,181 on the binary measure. Standard errors on the binary specifications widen as outlets are dropped, which is expected with eighteen clusters, but the continuous specifications remain significant at the one percent level throughout. The result is not produced by the benchmark's preference for cheap high-circulation titles.

Table E8: **Deviations from the Benchmarks Excluding Low-Cost Titles**

	Full sample (1)	Excl. sports titles (2)	Excl. three cheapest (3)
Cost-efficiency, binary alignment	20,876**	24,286**	27,183**
	(8,912)	(10,463)	(11,606)
Cost-efficiency, continuous score	28,749***	29,364***	30,649***
	(6,008)	(5,918)	(5,536)
Voter-targeting, binary alignment	22,198**	26,239**	30,181**
	(9,865)	(12,021)	(13,483)
Voter-targeting, continuous score	31,038***	32,249***	34,231***
	(6,863)	(6,829)	(6,335)
Ad Campaign FE	✓	✓	✓
Outlet FE	✓	✓	✓
Observations	6,711	5,945	5,562

*Intensive-margin estimates of equation 4 in 2022 USD, at $\phi = 0.10$. Column (2) excludes Esto and Ovaciones; column (3) additionally excludes La Prensa. Both benchmarks are re-solved campaign by campaign on the retained outlets. SE clustered by outlet. ***$p<0.01$, **$p<0.05$, *$p<0.10$.*

F Robustness for the Downstream Effects

F.1 Topic Model and the Hard News Classification

The hard news and soft news split of Section 6 comes from a topic model fitted on the newspaper text. I fit the model on a stratified training sample and then apply it to the full corpus. The sample draws up to 1,000 articles from each outlet and year, covering the 18 national outlets and the nine years from 2014 through 2022, which caps the training set at 162,000 documents. An article enters the pool when the archive records it as unpaid editorial content and its body runs to at least 100 characters, and article types covering advertising, paid inserts, obituaries, cartoons, legal notices, tables, and open letters leave the pool before the draw. A fixed random seed governs the draw. I join the headline to the body, transliterate accented characters to their unaccented forms, lowercase the text, and strip punctuation, numbers, and tokens shorter than three characters. Beyond the standard Spanish stopword list I remove 56 terms that recur across Mexican newsrooms, among them the attribution verbs *dijo*, *informó*, and *señaló* and the wire-copy bylines *redacción*, *agencia*, and *corresponsal*, so that the estimated topics track subjects and not house style. A Snowball stemmer

collapses the surviving tokens, and I drop every stem that appears in fewer than 10 documents. These steps run through the `textProcessor` and `prepDocuments` routines of the `stm` package for R (Roberts et al., 2019).

I estimate a structural topic model with $K = 40$ topics (Roberts et al., 2014). Topic prevalence depends on outlet indicators, author type, and a spline in a semester index running from the first half of 2014 to the second half of 2022, so the subject mix differs across outlets and moves over time. Topic content depends on outlet ideology, the left and right classification of Section 3.2, which allows outlets on either side to use different words within a single topic. Estimation starts from a latent Dirichlet allocation solution and runs for 150 expectation-maximization iterations under a fixed seed. I selected K on a pilot sample of up to 300 articles per outlet and year, evaluating $K \in \{10, 20, 30, 40, 50\}$ on held-out likelihood, semantic coherence, exclusivity, and residual dispersion. Every diagnostic moves monotonically over that range. Held-out likelihood rises from -8.42 to -8.14, residual dispersion falls from 13.8 to 9.1, and exclusivity rises from 9.36 to 9.82, while semantic coherence falls from -76.7 to -91.2. Three diagnostics therefore favor the largest model and coherence favors the smallest, so no interior optimum exists and the choice trades one against the others. Improvements in likelihood, residuals, and exclusivity flatten after $K = 30$, and I set $K = 40$, which costs 2.8 points of coherence against the $K = 30$ model.

Scoring proceeds one year at a time. I preprocess that year's articles with the same steps, align the tokens to the training vocabulary, hold the estimated topic-word distributions fixed, and infer the topic proportions of each document out of sample. The procedure returns proportions for 4,595,189 articles between 2014 and 2022, so each article carries a vector of 40 proportions summing to one and no article goes unassigned. Two aggregations of that vector appear in the paper. At the article level I take the topic with the largest proportion as the dominant topic and read the hard or soft assignment off it, which delivers counts of hard and soft articles per outlet and month. Those counts cover the articles that carry a journalist byline and leave out wire copy and unsigned house content, because a topic assigned to an article the newsroom did not write says nothing about what that newsroom produced. Hard and soft news in Table 5 therefore sum to an outlet's bylined output, 2,506,372 articles, and not to its 4,572,411 total articles. Hard news takes 79.1 percent of bylined articles. At the outlet-semester level I sum the proportions across hard topics within each article and average across articles, which retains the second and third topics of a mixed article. The first aggregation places 68.2 percent of articles in hard news and 31.8 percent in soft news; the second places 65.8 percent of content in hard news and 34.2 percent in soft news.

Labeling the 40 topics runs in two steps. For each topic I extract the 10 words with the highest FREX score, which weights frequency against exclusivity, the 10 words with the highest probability, and the 10 words most distinctive to left-wing outlets and to right-wing outlets, the last two available because the model carries an ideology content covariate. I pass those four word lists together with the mean proportion of every topic to a large language model, claude-opus-4-6, in one request under a structured output schema that returns an English label, a Spanish label, a hard or soft assignment, a geographic scope, and a one-sentence rationale per topic. The instructions define hard news as politics, elections, the economy, crime,

security, courts, policy, governance, legislation, protest, public health emergencies, and energy, and they define soft news as sports, entertainment, culture, the arts, lifestyle, celebrity, human interest, tourism, religion, and health coverage outside an emergency. The definitions follow the taxonomy of Cagé (2020) and Angelucci and Cagé (2019) and depart from it on crime, which I place in hard news.

I then reviewed every label against its word list, and the review changed one assignment. The returned file held 39 rows for a 40-topic model, leaving topic 3 without a label, and the leading words of that topic describe personal and family life (*familia*, *hijo*, *amigo*, *abuela*, *marido*). I added it as soft news under the category Human Interest. The other 39 topics keep the assignment they received, so 1 of 40 topics changed and the absence of topic 3 from the original file is the only recorded reason for a change. The classification ends with 26 hard topics and 14 soft topics. Hard news covers presidential politics, congress and legislation, courts and the judiciary, electoral politics, macroeconomics and financial markets, public finance and taxation, crime and public security, and energy and trade policy. Soft news covers Liga MX soccer, European soccer, music and film, arts and literature, tourism and gastronomy, and personal and family life. Six topics describe how an article is written instead of what it covers, and I flag them without recoding them. Topic 2, Development and Innovation Policy, carries the language of official announcements and sits in hard news, which matters in a paper about governments buying coverage; topics 15, 18, and 26 pick up opinion registers; topic 6 picks up event listings; and topic 37 picks up data-heavy writing. The six take 8.6 percent of articles as a dominant topic, and removing them raises the hard news share from 68.2 to 69.8 percent, a move of 1.6 percentage points.

The split can be checked against a label the newspapers produce themselves. Every article carries the section of the page on which it ran, recorded in the print registry and formed without reference to the model. Section names are outlet-specific brands, so I hand-map the 89 most common normalized names onto hard and soft. Names that identify no subject, such as front page and general, and names that identify a register, such as opinion and analysis, stay out of the comparison instead of being forced onto one side. Across 2018 and 2019 the crosswalk covers 784,866 of 985,149 articles, or 79.7 percent, and the two classifications agree on 87.0 percent of them. Agreement reaches 88.0 percent once the six register topics leave the sample and 91.5 percent on articles whose dominant topic takes at least half the document. Table F1 lists the 40 topics with their labels, categories, and assignments.

Table F1: **The Forty Estimated Topics and Their Hard News Assignment**

Topic	Label	Category	Type	Share (%)
1	Aviation & Auto Industry	Infrastructure	Hard	1.5
2	Development & Innovation Policy†	Development	Hard	2.9
3	Personal & Family Life	Human Interest	Soft	2.0
4	Catholic Church & Religion	Culture	Soft	0.9
5	Presidential Politics	National Politics	Hard	2.2
6	Events & Announcements†	Culture	Soft	0.9
7	Public Finance & Taxation	Fiscal	Hard	3.1
8	Venezuela & Cuba Politics	International	Hard	1.0
9	Soccer Match Reporting	Sports	Soft	1.7
10	Brazil Politics	International	Hard	0.3
11	Left-Right Ideology & Populism	International	Hard	2.3
12	US & International Sports	Sports	Soft	3.0
13	Education & Teachers Unions	Education	Hard	2.0
14	Human Rights & Military	National Politics	Hard	2.2
15	Opinion & Informal Commentary†	Opinion	Soft	1.4
16	Macroeconomics & Financial Markets	Economy	Hard	4.4
17	Drug Cartel Leaders & Crime	Crime & Security	Hard	0.9
18	Technology & Science Commentary†	Opinion	Soft	1.3
19	Congress & Legislation	National Politics	Hard	3.4
20	European Soccer (Champions League)	Sports	Soft	2.6
21	State & Municipal Government	Local	Hard	1.0
22	Business & Tech Companies	Firms	Hard	2.1
23	Public Health & Food Policy	Health Policy	Hard	2.3
24	Arts & Literature	Culture	Soft	3.0
25	COVID-19 Pandemic	Health Policy	Hard	3.1
26	Political Opinion & Analysis†	Opinion	Soft	0.6
27	Courts & Judiciary	Justice	Hard	4.1
28	Electoral Politics	Elections	Hard	4.1
29	Crime & Public Security	Crime & Security	Hard	3.3
30	Mexico City Infrastructure	Local	Hard	4.9
31	Telecom & Digital Regulation	Telecom	Hard	2.4
32	Formula 1 & Olympic Sports	Sports	Soft	2.2
33	Energy & Trade Policy	Energy	Hard	2.5
34	Music & Film	Entertainment	Soft	6.1
35	International Conflict & Terrorism	International	Hard	4.0
36	Tourism & Gastronomy	Tourism	Soft	1.8
37	Statistics & Data Reporting†	Economy	Hard	1.5
38	Gender & Women's Rights	Gender	Hard	1.6
39	Liga MX Soccer	Sports	Soft	4.3
40	US Politics & Migration	International	Hard	4.9

Each row is one topic of the structural topic model with $K = 40$ estimated on the stratified training sample of the 18 national outlets, 2014 to 2022. Label and Category record the labeling step described above, and Type records the hard or soft assignment used throughout the paper. Share gives the percentage of the 4,595,189 scored articles for which the topic carries the largest proportion, and the column sums to 100. A dagger marks the six topics whose leading words describe writing register instead of subject. These six keep their assignment, and the robustness check that drops them raises the hard news share from 68.2 to 69.8 percent.

F.2 Constructing News Events and the Within-Event Measures

The quality measures in Section 6 compare outlets on the same story, so I first group articles into news events. The input is the news content of the print corpus, 3,945,226 articles published by the 18 outlets between 1 January 2014 and 31 December 2022, keeping the articles the archive records as unpaid editorial content under a news article type. I join the headline to the body, strip HTML markup, transliterate accented characters to their unaccented forms, lowercase the text, and delete everything that is not a letter. I neither stem nor remove stopwords, because the inverse-document-frequency weight already discounts terms that appear in every article. One term-frequency inverse-document-frequency weighting runs over the pooled nine years, so the vocabulary stays fixed across the sample. It keeps unigrams that appear in at least five articles and in at most 95 percent of them, caps the vocabulary at 100,000 terms, and takes the logarithm of

within-document term counts. Each article becomes a unit-length vector, and the inner product of two such vectors equals the cosine of the angle between them, which is the similarity I use throughout.

Events form inside a moving three-day window that advances two days at a time, following the topic detection design of Cagé et al. (2020) and Durante et al. (2026). Within each window I compute the cosine distance between every pair of articles, build a hierarchical tree under single linkage, and cut that tree at a cosine distance of 0.68, so two articles enter the same cluster when a chain of articles connects them at a cosine similarity of 0.32 or above. Consecutive windows overlap by one day, and a union-find pass merges two window clusters whenever they share an article, which lets an event run past the length of a single window. Windows sit inside a calendar year, so no event crosses 31 December. An article that pairs with no other article stays outside the event set. The procedure places 2,159,555 articles, 54.7 percent of the corpus, into 410,916 events. Table F2 reports the resulting panel.

I selected the linkage rule and the distance threshold against a list of events compiled outside the project. The first specification cut a complete-linkage tree at a distance of 0.60 and failed that check. Complete linkage holds every pair inside a cluster within the threshold, so the median event held two articles at every threshold from 0.40 to 0.70 while the share of clustered articles rose from 14.0 to 54.4 percent, and a story that seventeen outlets headline differently never forms one cluster. The external list takes 38 dated events from the Wikipedia timelines for 2018 and 2019, each selected before any clustering output was inspected and each carrying a proper noun distinctive enough to match on. Headlines matched 37 of the 38. Under the complete-linkage specification the dominant cluster of a matched event spanned at least two outlets for 21 of those 37 events, 56.8 percent, the median event split across six clusters, and the largest cluster held 42.4 percent of an event's matched articles. I then re-ran the clustering over 22 combinations of linkage rule and threshold, three linkage rules and thresholds from 0.50 to 0.90, and scored each against six criteria fixed before the runs. The criteria require a median of at most two clusters per external event, mean concentration of at least 0.70 in the largest cluster, a dominant cluster spanning two outlets for at least 85 percent of external events, a largest cluster holding at most 10 percent of clustered articles, a median duration of at most seven days for those dominant clusters, and a 99th percentile duration of at most 21 days. Two runs clear all six, single linkage at 0.68 and single linkage at 0.70, and I use 0.68, whose largest cluster holds 1,490 articles against 4,957 at 0.70. Under the chosen run 34 of the 37 external events, 91.9 percent, have a dominant cluster spanning at least two outlets, 45.9 percent land in a single cluster, the median event splits across two clusters, the largest cluster holds 82.3 percent of an event's matched articles, the largest cluster in the whole event set holds 0.27 percent of clustered articles, and event duration reaches five days at the 99th percentile.

An event enters the estimation sample when at least three of the 18 outlets published on it, which leaves 193,066 events, 47.0 percent of those formed, holding 1,690,436 articles. The filter applies to the total number of outlets covering an event and not to the outlet whose behavior the regression explains, so an event that three outlets ran and a fourth skipped stays in the sample and the fourth outlet enters it with a zero. Events

covered by one or two outlets hold 467,863 articles and resemble those kept, with hard news taking 67.1 percent of their articles against 70.7 percent among kept events and left-wing outlets writing 41.0 percent against 43.9 percent. Drop rates by subject category run from 40.6 percent for national politics to 60.9 percent for culture, so the filter does not select against political content. Crossing the 193,066 events with the 18 outlets gives 3,475,188 outlet-event observations, and an outlet covered 30.5 percent of them. A separate pass flags same-outlet near-duplicates, defined as two articles from one outlet within three days at a cosine distance below 0.001, and removes the 2,370 it finds.

A factual element is a discrete piece of verifiable content that an article puts on the page. Named entities come from the Spanish news pipeline of spaCy, `es_core_news_sm`, and I keep the four labels it assigns, person, organization, location, and the miscellaneous label that covers events, works, and nationalities. Numbers come from a regular expression over the same text and enter as elements when they carry at least two digits, which keeps death tolls, peso amounts, and vote shares while discarding single digits. I parse the headline and the first 4,000 characters of the body, since the opening paragraphs carry the facts of a print news article. Normalization lowercases each element, strips accents and punctuation, collapses whitespace, and drops elements of two characters or fewer. It performs no entity linking, so Andrés Manuel López Obrador and AMLO stay distinct elements, which lowers the level of every completeness score by the same rule and leaves the comparison across outlets intact. Within an event each element counts once per outlet, so an outlet that repeats a fact across five articles contributes it once. Wire copy leaves the sample before extraction, 329,575 of the 1,690,436 articles, because a dispatch that eight outlets reprint would raise the prevalence of its entities and would credit an outlet that bought the copy with reporting it did not do. Parsing the remaining 1,360,859 articles yields 37,836,207 entity and number mentions. Articles carrying fewer than three elements leave the sample, and an event survives only if three outlets still cover it after both deletions, which leaves 1,303,858 articles, 153,109 events, and 823,169 outlet-event observations with measurable content. The median event carries 81 distinct elements and the median outlet contributes 28 of them. A second extraction reads subject-verb-object triples off the dependency parse as an alternative definition of an element, and every measure below is computed on both definitions. The table in the text uses named entities and numbers.

Elements differ in how central they are to a story, so each carries a prevalence weight. For event e covered by n_e outlets, let k_f count the outlets reporting element f and let F_{je} hold the elements outlet j reported. Raw prevalence is $p_f = k_f/n_e$, and outlet j enters its own weight under that definition. An element only j reports then takes $p_f = 1/n_e$ instead of zero, so the rarity of j's own reporting is capped at $1 - 1/n_e$ and the cap tightens as events get smaller. In the estimation panel raw exclusivity reaches 0.938 against a ceiling of one. Completeness carries j in both its numerator and its denominator. Both distortions move with the number of outlets covering an event, which the treatment itself shifts, so they do not net out across the two

sides of the comparison. I therefore weight each element by its prevalence among the other outlets,

$$p_f^{-j} = \frac{k_f - 1}{n_e - 1},$$

and the event total follows the same rule, $\sum_{f \in E_e} p_f^{-j} = (\sum_f k_f - |F_{je}|)/(n_e - 1)$. An element no other outlet carries now scores exactly one on rarity, leave-one-out exclusivity reaches 1.000 in the panel, and completeness reads as the share of what the rest of the market established that j also carried.

Four outcomes enter the estimation. Coverage is an indicator equal to one when outlet j published at least one article on event e, and it is defined for every outlet-event pair. Completeness is $\sum_{f \in F_{je}} p_f^{-j}$ divided by the event total, and the regression weights it by the event's element mass, because a ratio taken over three elements carries far less information than one taken over five hundred. New information is $\sum_{f \in F_{je}} (1 - p_f^{-j})$, the count of j's elements weighted by how few other outlets carried them. Story elements is $|F_{je}|$, the number of distinct elements j contributed to the event. The three element outcomes run on the balanced panel and take the value zero where an outlet contributed no measurable reporting, which covers 2,581,535 of the 3,404,704 rows that remain once El Heraldo de México's pre-relaunch years leave the sample, 2,344,172 of them pairs where the outlet never covered the event and 237,363 pairs where it covered the event without producing an element. Rows belonging to the 39,957 events where no covering article produced an element leave the sample, which leaves 2,700,888 rows across 153,109 events for Columns (2) to (4) of Table 6. A fifth measure, exclusivity, equals $\sum_{f \in F_{je}} (1 - p_f^{-j})/|F_{je}|$, the mean rarity of what j reported. It is a rate over the outlet's own elements, so it has no value when the outlet published nothing, and coding it as zero would record an outlet that wrote nothing as having reported only ubiquitous facts. The main table omits it, and I report it here conditional on coverage, where it averages 0.713 over the 823,169 outlet-event pairs that carry elements.

Two checks bear on the event set. The first is the external recovery rate reported above, 91.9 percent under the chosen specification against the 92 percent of Europe Media Monitor stories that Cagé et al. (2020) recover. The second is the behavior of the event set across the threshold grid, which shows no knife edge. On the two years used to select the specification, single linkage clusters 24.9 percent of articles at a distance of 0.50, 41.3 percent at 0.60, and 60.0 percent at 0.70, and the count of events covered by at least three outlets moves from 30,037 to 45,893 to 46,850 across those three cuts, while the strict recovery rate rises from 54.1 to 73.0 to 91.9 percent and then holds flat. A hand-coded check of individual article pairs is drawn and not yet annotated. It holds 500 blind pairs, 300 that the algorithm assigns to one event stratified by the number of outlets covering it and 200 same-day pairs it separates, so this appendix reports no human agreement rate for the clustering. The hard and soft classification that splits events by subject is validated separately against newspaper section labels, and the two agree on 87.0 percent of the 784,866 articles the crosswalk covers, as Appendix F.1 reports.

Table F2: **The Event Panel and the Element Panel**

	Value
Panel A. News events, 2014–2022	
Articles in the corpus	3,945,226
Articles assigned to an event	2,159,555
Share of articles assigned to an event	54.7%
Events formed	410,916
Events covered by at least three outlets	193,066
Articles in those events	1,690,436
Articles per event, mean / median	8.76 / 5
Outlets per event, mean / median	5.49 / 4
Events lasting a single day	57.5%
Event duration, 90th / 99th percentile	2 / 7 days
Outlet–event observations	3,475,188
Share covered	30.5%
Panel B. Factual elements	
Wire articles excluded before extraction	329,575
Articles parsed	1,360,859
Entity and number mentions extracted	37,836,207
Articles carrying at least three elements	1,303,858
Events retained	153,109
Outlet–event observations with elements	823,169
Distinct elements per event, median	81
Distinct elements per outlet–event, mean / median	39.1 / 28
Panel C. External event recovery, 2018–2019	
External events tested	38
Events matched by headline	37
Dominant cluster spans at least two outlets	91.9%
Mean concentration in the largest cluster	82.3%
Median clusters per external event	2

Panel A covers the 18 outlets from 1 January 2014 through 31 December 2022. Events come from single-linkage agglomerative clustering of term-frequency inverse-document-frequency vectors, cut at a cosine distance of 0.68 inside a three-day window that advances two days at a time. Panel B covers the same events after wire copy, articles carrying fewer than three elements, and events left with fewer than three covering outlets leave the sample. Elements are named entities from spaCy's `es_core_news_sm` pipeline together with numbers of at least two digits, counted once per outlet within an event. Panel C compares the event set against 38 events dated from the Wikipedia timelines for 2018 and 2019, matched on proper nouns in the headline.

F.3 Revenue, Newsroom, and Circulation

F.3.1 Revenue Shock to Outlets

Government advertising was not the only revenue that could have moved at the transition. Figure F1 runs the event study of Section 6 with the number of private advertisements an outlet carries as the outcome. Left-wing outlets carry more private advertising than right-wing outlets in every year after the transition, by 0.20 asinh points in the transition year and 0.38 in the year that follows, and the difference stays positive through 2022. The pre-transition coefficients scatter below zero without a trend, and none of them separates from zero. Private advertisers therefore moved toward left-wing outlets alongside the government. The panel holds 18 outlets and 178 outlet–years, so the pooled post-transition difference of 0.690 carries a standard error of 0.432, and expressing advertisements per day of coverage shrinks it to 0.323 with a standard error of 0.262. The sign holds across both specifications and the magnitude is not precisely estimated.

Figure F1: **Event Study of Private Advertising by Outlet Ideology**

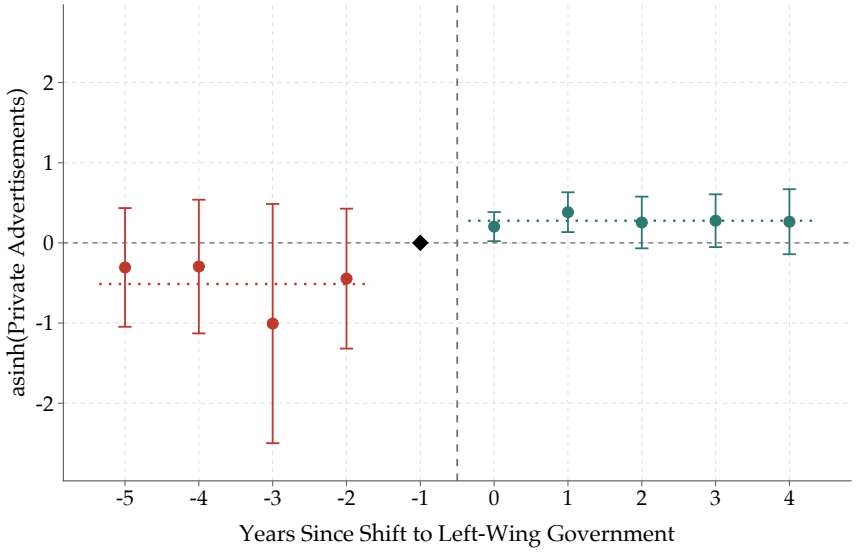

Coefficients on the interaction of a left-wing outlet indicator with each year relative to the 2018 transition, from a regression of the inverse hyperbolic sine of an outlet's private advertisement count on outlet and year fixed effects. The year before the transition is the omitted category, marked by the diamond. Bars are 95 percent confidence intervals from standard errors clustered by outlet across 18 clusters. Dotted lines mark the average coefficient before and after the transition. Outlet ideology follows the classification of Section 3.2.

F.3.2 Newsroom Size

Table F3 decomposes the article count of Column (2) of Table 5 into finer content categories, estimating the same specification on the inverse hyperbolic sine of the monthly article count within each category. Columns (1)–(4) cover hard news topics (total, politics, economics, and social policy), and columns (5)–(6) cover soft news (sports and entertainment).

Table F3: **Effect of Left-Wing Government on News Production by Content Type**

	Total (1)	Politics (2)	Economics (3)	Social Policy (4)	Sports (5)	Entertainment (6)
Left-Wing Government × Left-Wing Outlet	0.416***	0.425***	0.512***	0.474**	0.342	0.356*
	(0.154)	(0.149)	(0.196)	(0.189)	(0.235)	(0.213)
Year FE	✓	✓	✓	✓	✓	✓
Outlet FE	✓	✓	✓	✓	✓	✓
Mean pre-2018 (Left)	10.502	9.559	8.415	8.048	8.319	6.988
Mean pre-2018 (Right)	10.967	9.98	9.297	8.647	8.484	7.528
SD pre-2018	0.787	0.864	1.097	0.955	0.744	0.811
Observations	159	159	159	159	159	159
R^2	0.907	0.923	0.934	0.906	0.854	0.867

Note: *p<0.1; **p<0.05; ***p<0.01. SEs clustered at the outlet level. Outlet–month panel. The outcome is the inverse hyperbolic sine of the monthly count of articles within each content type, so coefficients are interpretable as approximate percentage changes. Cols. (1)–(4) cover hard news, and cols. (5)–(6) cover soft news. All columns include outlet and year fixed effects. Means and standard deviations are computed over the pre-2018 period in the units of the transformed outcome.

Left-wing outlets increase news production across all hard news categories. Total article production rises by approximately 42% (Column 1), corresponding to roughly 776 additional articles per month. The gains are largest in economics (approximately 51%) and social policy (approximately 47%), followed by politics (approximately 43%), consistent with a newsroom expansion concentrated in substantive, resource-intensive

reporting. As a back-of-the-envelope check, the 776 additional monthly articles imply approximately 9 articles per additional journalist — somewhat below the pre-transition baseline of 13 — consistent with new hires ramping up over time. The effect on soft news is positive for entertainment (approximately 36%) and smaller and insignificant for sports, indicating that the extra resources translate primarily into hard news coverage rather than a broad undifferentiated expansion.

F.3.3 Circulation

Figure F2: **Change in Morena Vote Share by State (2015–2018)**

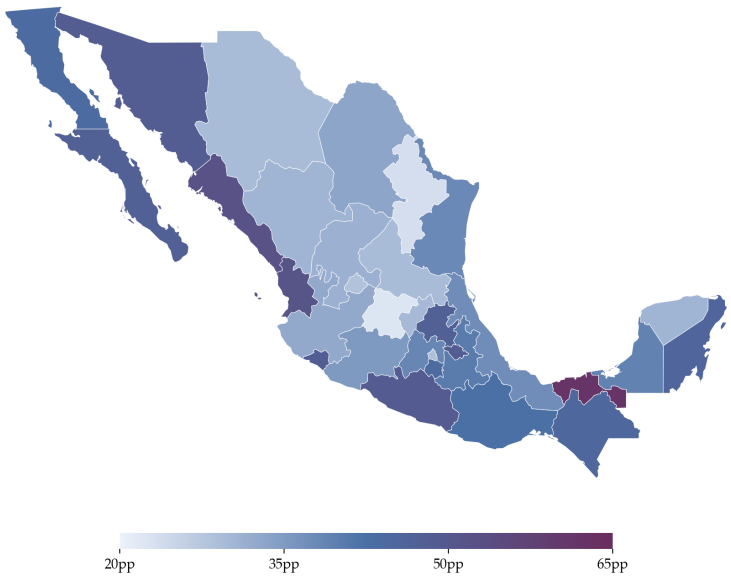

This figure maps the change in Morena's vote share by state between the 2015 midterm and the 2018 presidential election. Darker shades indicate larger shifts toward Morena. The dashed line in the triple difference-in-differences specification classifies states above the sample median shift ($D_s = 1$) as high-shift states. The variation ranges from approximately 20 percentage points in the north and northwest to over 65 percentage points in the south and southeast.