

Political Advertising on Facebook During the 2022 Australian Federal Election: A Social Identity Perspective

Stefano Civelli, Pietro Bernardelle, Frank Mols, Gianluca Demartini

The University of Queensland, Brisbane, Australia
 {s.civelli,p.bernardelle,f.mols,g.demartini}@uq.edu.au

Abstract

Political advertising on social media has become an important component of contemporary election campaigns, yet most research has focused on voluntary voting systems in which mobilizing supporters is a central goal. This study examines political advertising on Meta platforms (Facebook and Instagram) during the 2022 Australian federal election, a context shaped by compulsory voting and therefore distinctive campaign incentives. Using data from Meta’s Ad Library, we analyze temporal, demographic, geographic, and textual patterns in advertisements published by political actors during the campaign. The Australian Labor Party and the Liberal Coalition dominated the digital advertising landscape in both spending and impressions, while smaller parties and independents adopted more selective strategies. Ad delivery varied across demographic groups, with major parties, the Greens, and the United Australia Party showing different age and gender skews, though the clearest form of targeting was geographic. Across content, major parties emphasized party branding and criticism of opponents, whereas smaller parties more often relied on issue-based messaging. Interpreted through political psychology perspectives, and particularly Social Identity Theory, these patterns are consistent with different forms of identity appeal: major parties appear to reinforce established partisan identities, while smaller actors seek to connect with voters through issue-based identities to capture the support of disaffected voters who are obligated to participate.

Code — <https://github.com/Stefano-Civelli/Political-Advertising-2022-AU-Federal-Election>

Datasets — <https://doi.org/10.5281/zenodo.13661563>

1 Introduction

Research on social media’s impact on politics has evolved over time. Early studies focused on Twitter’s role in reflecting public opinion and predicting election outcomes (O’Connor et al. 2010; Tumasjan et al. 2011). However, subsequent studies identified several shortcomings in these initial approaches, pointing to issues with methodology and representative sampling across demographic groups (Gayo-Avello 2012a,b; Metaxas, Mustafaraj, and Gayo-Avello 2011). As the research landscape matured, researchers began

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

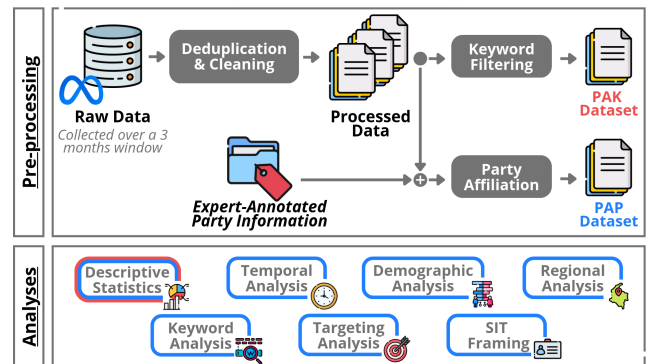


Figure 1: Pipeline for collecting, cleaning, and annotating Meta Ad Library data, followed by analyses.

examining the dual impact of social media on political campaigns—its ability to enable precise audience targeting and efficient campaign spending (Fowler et al. 2021) through precise audience segmentation and dynamic campaign adjustments (Boerman, Kruijemeier, and Zuiderveen Borgeusius 2017; Kreiss 2016), alongside growing concerns about disinformation and manipulation of public opinion (Entous, Timberg, and Dvoskin 2017). Facebook’s emergence as a critical platform for political advertising marked a significant shift in this research trajectory. In 2018, Meta (formerly Facebook) launched its Ad Library in response to transparency concerns, providing unprecedented access to political advertisement data. This development has sparked new research directions examining Facebook’s role in political campaigning across various contexts (Calvo, Cano-Orón, and Baviera 2021; Capozzi et al. 2021; Mejova and Kalimeri 2020; Pierri 2023).

However, most research focuses on voluntary voting systems, where mobilizing a party’s base to turn out is a central objective. By contrast, Australia’s system of compulsory voting ensures consistently high participation¹, shifting the strategic focus from mobilization—or “Get Out The Vote” (GOTV) efforts—toward persuasion (Jackman 2001; Birch 2009). Rather than expending resources to secure turnout,

¹<https://results.aec.gov.au/27966/Website/HouseTurnoutByState-27966.htm>

campaigns concentrate on convincing undecided voters and reinforcing party loyalty, incentivizing broad-based appeals over the micro-targeted mobilization common in voluntary systems (Hillygus and Shields 2008). This institutional feature makes Australia a distinctive setting to study digital advertising tactics and often encourages parties to moderate their messaging as they compete for the median voter (Jackman 2001; Bruns and Moon 2018; Skoric et al. 2012; Oprea, Martin, and Brennan 2024).

Our study addresses this empirical knowledge gap by analyzing political advertisements on Meta platforms (Facebook and Instagram) during the 2022 Australian federal election, and by evaluating the findings using three well-established political psychology perspectives on persuasion, namely (a) appealing to rational self-interest, (b) eliciting automatic pilot behavior, and (c) social identity mobilization. By curating a comprehensive dataset, annotating party affiliations, and examining spending, targeting, and content strategies (Figure 1), we provide insights into digital campaign dynamics in a unique electoral system, building on prior work while addressing platform-specific and context-specific nuances (Jungherr 2016; Skoric et al. 2012). Accordingly, our research seeks to answer the following questions:

- **RQ1.** *What was the scale and reach of political advertisements on Meta platforms during the election campaign?*
- **RQ2.** *How did the advertising strategies of major Australian political parties and groups differ in terms of temporal distribution of ads, spending patterns, targeting and the use of themes throughout the campaign period?*
- **RQ3.** *How can the observed patterns in advertising content and targeting strategies be interpreted through theoretical frameworks of voter motivation?*

Contributions. Our work contributes to the understanding of digital political advertising in the following ways: (1) we curated a dataset of political advertisements from the election period; (2) we had a political science expert annotate a list of parties to categorize ads by political affiliation; (3) provide a general analysis of the campaign, comparable to studies conducted in other countries; and (4) offer insights into major parties’ online campaign strategies, examining both quantitative aspects (spending, timing, targeting) and the theoretical frameworks that may explain their persuasion approaches.

Findings. The Australian Labor Party (ALP)² and the Liberal Coalition³ dominated the digital advertising landscape, collectively accounting for over 70% of total ad impressions and spending. The ALP invested A\$5M (40.2% of total political ad spending) across 13,047 ads, while the Liberal Coalition spent A\$3.9M (31.3% of total political ad spending) across 9,364 ads (**RQ1**). These two major parties secured 68.27% of the electors’ first preference votes on election day.

²We use the terms “ALP” and “Labor” interchangeably.

³The Coalition comprises the Liberal Party, National Party, Liberal National Party QLD, and Country Liberal NT.

We observed distinct patterns in advertising strategies throughout the election timeline (**RQ2**). After the April 10th election call, ad activity surged. Labor consistently outspent the Liberal Coalition, though both intensified efforts toward the campaign’s end. The Greens and UAP followed different timelines, with UAP ramping up late. Labor targeted more women across age groups; the Greens skewed younger; UAP focused on male audiences. Geographically, major parties followed population centers, while smaller parties used more targeted regional strategies. Major parties used ads to promote themselves and attack opponents, while smaller parties emphasized specific issues and policy themes.

When viewed through Social Identity Theory (SIT) (**RQ3**), these findings suggest strategic differentiation in identity appeals. Established major parties worked to strengthen and mobilize existing partisan identifications, whereas smaller political entities sought to foster and nurture identity connections based on specific policy concerns and issue salience. This theoretical interpretation helps explain the observed variations in targeting strategies, with parties tailoring their approaches to activate different identity-based motivations among potential voters.

2 Background and Related Work

2.1 Theoretical Frameworks of Voter Motivation

Political scientists have long explained voting behavior through (egoistic or sociotropic) rational self-interest, a view rooted in early-20th-century economic thought. This began to shift in the 1950s when researchers like Simon (1957) introduced the concept of ‘bounded rationality’, showing that individuals often rely on heuristics under cognitive strain. This opened the door for insights from cognitive psychology—such as ‘groupthink’, nudging, and behavioral public policy—though these have yet to significantly influence psephology (election studies). This is surprising, given that many campaign messages clearly aim to trigger automatic emotional responses rather than deliberative reasoning, as seen in the Liberal party’s Frankencredit⁴ and Labor’s Mediscare⁵ campaigns. One psychological theory that has been incorporated in election studies, albeit belatedly, is SIT (Tajfel 1978; Tajfel and Turner 1979), which emphasizes partisan identity as a key driver of vote choice. This aligns with Campbell et al. (1960), who found that most voters are firmly loyal to their chosen party—what would be described in Australia as ‘rusted-on’ supporters. Recent interest in SIT has grown alongside rising affective polarization (Garrett and Bankert 2020) and negative partisanship (Abramowitz and Webster 2018; Bankert 2020), with strategists increasingly using ‘us-them’ narratives to build partisan identities (Huddy and Bankert 2017; West and Iyengar 2022). However, this literature is still largely absent from digital campaigning research.⁶

⁴<https://independentaustralia.net/politics/politics-display/franking-credits-the-coalitions-scare-campaign,12365>

⁵<https://theconversation.com/why-scare-campaigns-like-mediscare-work-even-if-voters-hate-them-62279>

⁶Two major models in this area are ‘Expressive Partisanship’—viewing partisanship as emotional group attachment

This identity-based perspective is particularly salient in Australia's compulsory voting system. Because turnout is high, parties focus less on mobilization and more on persuasion (Jackman 2001; Birch 2009). This environment encourages appeals aimed at the median voter, pushing parties to focus on broadly popular messages while avoiding highly salient or polarizing issues (Oprea, Martin, and Brennan 2024). Over time, such strategies can reinforce partisan loyalties. Cross-national research suggests that compulsory voting further strengthens habitual partisan identities (Singh and Thornton 2013). As partisan identity becomes more central to vote choice (Huddy, Mason, and Aarøe 2015), parties can either positively reinforce the in-group or rely on negative partisanship—defining themselves against an out-group—to strengthen these attachments (Abramowitz and Webster 2018; Bankert 2020).

In sum, there is significant opportunity for interdisciplinary research into why particular campaign messages resonate. We suggest that insights from three psychological models—(a) material self-interest, (b) flawed cognition, and (c) partisan social identity—could help explain why particular tactics hit the mark in digital political campaigns.

2.2 Social Media Political Advertising

Digital channels have fundamentally altered modern political messaging (Fowler et al. 2021). These new platforms offer distinct advantages for political campaigns, including unprecedented audience targeting precision and cost-effectiveness compared to traditional media (Boerman, Kruijemeier, and Zuiderveen Borgesius 2017; Kreiss 2016). Building on this foundation, Dommett and Power (2019) explored the political economy of Facebook advertising, documenting its growing prominence in election spending and strategic campaign planning.

The introduction of the Meta Ad Library in 2018 has enabled researchers to examine Facebook's role in political advertising across diverse contexts (Calvo, Cano-Orón, and Baviera 2021; Capozzi et al. 2021; Leerssen et al. 2021; Mejova and Kalimeri 2020; Pierri 2023). Recent studies have also introduced weakly and minimally supervised machine learning methods to classify ad stances, themes, and moral foundations in political and issue-based campaigns, such as those on elections, climate change, and COVID-19 vaccines, revealing targeting patterns and narrative strategies (Islam, Roy, and Goldwasser 2023; Islam, Zhang, and Goldwasser 2023; Islam and Goldwasser 2022).

These studies have uncovered both new campaign opportunities and significant transparency concerns related to platform-based political messaging (Entous, Timberg, and Dwoskin 2017), including imprecise moderation with high false-positive and false-negative rates that allow undeclared coordinated campaigns to evade detection (Bouchaud

(Green, Palmquist, and Schickler 2004; Huddy and Bankert 2017)—and 'Instrumental Partisanship', which sees it as a rational attachment based on party performance (Fiorina 1981; Orr and Huber 2020). In practice, these processes likely interact, making it more useful to consider how different persuasion strategies draw on self-interest, cognition, or identity.

and Liénard 2024). This recent work extends earlier research that examined Twitter's potential for reflecting public opinion and predicting electoral outcomes (O'Connor et al. 2010; Tumasjan et al. 2011), though these approaches faced substantial methodological challenges and critiques (Gayo-Avello 2012a,b; Metaxas, Mustafaraj, and Gayo-Avello 2011).

The evolution of digital political advertising has increasingly centered on the strategic deployment of micro-targeting techniques, which leverage the granular demographic and behavioral data available on social media platforms. This shift represents a significant departure from traditional broadcast approaches to political messaging. Ridout, Fowler, and Franz (2021) investigated how political campaigns deploy ads on Facebook, focusing on the influence of campaign goals and timing. Their work provides valuable insights into the strategic considerations behind digital ad placement. Brodnax and Sapiezynski (2022) studied the evolution of digital advertising strategies during the 2020 US Presidential Primary, highlighting the shift from broad to more targeted approaches as campaigns progressed. Recent analyses have further revealed systematic discrepancies between targeted and actual ad deliveries, with algorithmic biases favoring certain parties and skewing audiences by age and gender (Bär et al. 2024).

While our study focuses on a different political context, these established methodological approaches inform our analysis of partisan advertising strategies in Australia.

2.3 Australian Context

While numerous studies have examined political advertising on social media in various countries, research specific to the Australian context is relatively limited. Bruns and Moon (2018) provided an analysis of Twitter usage during the 2013 and 2016 Australian federal elections, finding significant differences in how major parties utilized the platform. They noted that Labor candidates were consistently more active on Twitter than their Liberal Coalition counterparts across both elections, a trend that aligns with our findings. More recently, Decker et al. (2024) explored topic diversity in social media campaigning during the 2022 Australian federal election. Their study found that while there were some differences between organic and paid content, both types predominantly focused on core political topics aligned with party ideologies. They also noted the emergence of government critique as a distinct topic in both organic and paid content, signaling the use of negative campaigning.

While Arya (2022) and Meguellati et al. (2025) also examined political advertising during the 2022 Australian Federal election, our work extends theirs in several ways. We release a manually curated party affiliation list to support reproducibility, analyze an extra month of data with more advanced analysis methods (including keyword filtering and targeting) and frame the findings through the lens of SIT.

2.4 The 2022 Australian Federal Election

The 2022 Australian federal election was called on April 10th, with election day set for May 21st. Approximately 17.2 million eligible voters were registered to participate.

The election resulted in a change of government, with the Australian Labor Party (ALP) securing 77 seats in the 151-seat House of Representatives, allowing them to form a majority government. The Liberal-National Coalition, which had been in power since 2013, won 58 seats. Minor parties and independents claimed the remaining 16 seats, 12 more as compared to the previous election, reflecting a growing trend towards diversification in Australian politics. The ALP received 32.6% of the primary votes, while the Coalition garnered 35.7%. However, after preference distribution under Australia’s two-party-preferred system, these figures adjusted to 52.1% and 47.9% respectively. The Greens achieved their best-ever result with 12.2% of the primary votes, followed by independent candidates with 5.3%. This election was notable for the significant swing away from the two major parties, with their combined primary vote falling below 70% for the first time in Australian history.

3 Data Preparation

3.1 Data Collection

We utilized the Meta Ad Library API⁷ to collect online advertisements published by parties, political candidates, and other entities related to the 2022 Australian federal election. Our data collection process captured actively running advertisement campaigns at six-hour intervals from March 1st, 2022, to June 1st, 2022. This resulted in a dataset of 56,224 ads from 1,549 sponsors.

Each ad includes attributes such as the creation time, target demographic distributions, ad text, page name, sponsor, ad URL, and metrics of ad impression and spending. For these last two, the Meta Ad Library API provides upper and lower bounds. In our analysis, we calculated ad spend and impressions as the mean of these bounds.

We determined impressions by age, region, and gender by multiplying the mean impressions for each ad by the proportion of ad reach attributed to each category. Ads lacking distribution data for age, gender, or region were excluded from calculations for those specific impression categories. The 13–17 age group and users with an “unknown” gender were excluded from the demographic analysis due to their minimal presence, accounting for only 0.1% and 1.1% of ad impressions, respectively.

It is important to note that our initial dataset, PA (Political Advertisements), comprised only ads explicitly flagged as political, as required by Australian law⁸. This approach may introduce potential biases, as some political ads might not have been flagged and thus excluded, while others could have been mis-flagged as political when they were not, a phenomenon studied by Sosnovik and Goga (2021). Reinforcing this concern, Bouchaud and Liénard (2024) studied Meta’s EU ad moderation and estimate that 92.3% of undeclared political ads (those not flagged as political by the publisher) are not detected by Meta’s systems, highlighting a critical loophole in its reliance on voluntary self-declaration by advertisers

⁷<https://www.facebook.com/ads/library/>

⁸In Australia the posting entity is legally required to flag political ads with a “paid by” disclaimer.

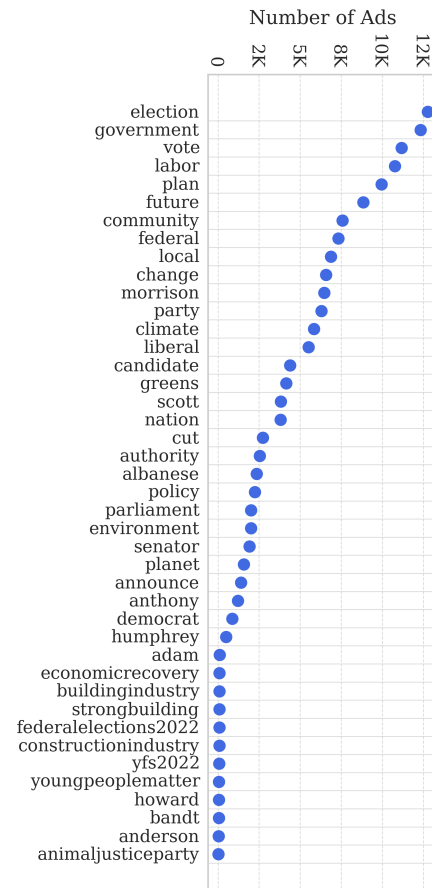


Figure 2: Frequency distribution of selected keywords used in political advertisements during the considered period. Keywords are ranked by the number of ads they appear in.

3.2 Data Preprocessing

While not much can be done for ads that were not flagged as political, we tried to exclude non-political content that was incorrectly flagged as such. To do so, we (1) further refined the collected data by employing a *keyword filtering* mechanism. Additionally, to gain insights into the political campaign strategies, we (2) filtered ads based on their *party affiliation*. The overall pipeline is summarized in Figure 1.

Keyword Filtering. We used a snowball sampling technique, as employed in previous studies (Conover et al. 2012; Yang, Hui, and Menczer 2019; Aiyappa et al. 2023), to develop a list of relevant keywords. Starting with seed terms “government” and “election”, we iteratively expanded our keyword list by identifying frequently co-occurring terms. Two authors independently reviewed and agreed on each new keyword addition. This process resulted in a final list of 43 keywords (see Figure 2). The observed exponential decay in keyword popularity suggests that the inclusion of additional keywords would not substantially alter the corpus structure. Subsequently, we used a stemmed version of these keywords as a filtering mechanism to the initial dataset (PA) to extract a subset of ads more aligned to our research

Name	# Entries	% of total	Description
PA	56,224	100%	All ads
PAK	47,503	84.5%	Ads containing ≥ 1 keyword
PAP	30,767	54.7%	Ads associated to a party

Table 1: Summary of datasets derived from political advertisements on Meta platforms during the 2022 Australian federal election campaign. PA represents the full dataset, while PAK and PAP are subsets based on keyword presence and party association respectively.

focus. This process led to the creation of a second dataset, named PAK (Political Advertisements Keywords), containing 47,503 ads, and excluding 111 sponsors. While some of these keywords are of somewhat general nature and can therefore be the cause of false positives, it is important to note that they are just an additional layer of filtering on top of what is already a dataset comprising political ads (according to Meta).

Party Affiliation. To gain insights into campaign strategies of major Australian political parties, a political science expert annotated a list comprising of all potential sponsors and page names, identifying each with the appropriate party and electorate where applicable. Despite expecting this to be a fairly objective annotation task, to assess reliability, we had a second political science expert independently annotate 10% of the dataset, resulting in a Cohen’s κ inter-annotator agreement of 0.906. We then joined this annotated list with our original data. The resulting dataset, PAP (Political Advertisements Parties), consists of 30,767 advertisements and exclusively contains ads from candidates contesting the 2022 federal election. Importantly, PAP only includes advertisements sponsored by recognized political parties, ensuring a focused analysis of official party strategies.

We want to clarify that the two preprocessing strategies are orthogonal and provide unique insights. This approach allows us to conduct two levels of analysis: a broad examination of the entire political advertising ecosystem, and a targeted study of official party campaign strategies. A summary of the 3 datasets⁹ with their characteristics can be found in Table 1.

4 Descriptive Statistics

This section presents a general overview of the statistics of the PAK dataset obtained with keyword filtering which in turn contains nearly all of the PAP dataset ($PAP \cup PAK \simeq PAK$).

In the analyzed months leading up to the election, 47,503 ads across 1,438 different sponsors were published. These ads generated over 1.11B impressions Australia-wide, with a total ad spend of A\$18.4M. Each impression represents

⁹Code and datasets are available online at <https://doi.org/10.5281/zenodo.13661563> (code repository link is available in Zenodo’s dataset description).

Metric	Mean	Median	Max
Impressions	23,452	4,500	1,000,000
Spend (A\$)	387	49.5	137,500
Duration (days)	8.1	6	89
Efficiency (impr./A\$)	67.5	50	4,319

Table 2: Descriptive statistics of individual ads in the PAK dataset.

Pairs	R (R^2)	ρ	Strength
Spend – Impressions	0.72 (0.52)	0.82	Strong
Spend – Duration	0.10 (0.009)	0.26	Weak
Impressions – Duration	0.17 (0.029)	0.30	Weak

Table 3: Correlation coefficients between ad-level metrics. Pearson’s R measures linear correlation, with R^2 denoting the proportion of explained variance. Spearman’s ρ measures rank-based (monotonic) correlation. All results remain statistically significant after Bonferroni correction ($\alpha = 0.0042$).

a single ad view, with multiple impressions possible per viewer, i.e., a single user might see the same advertisement many times.

Individual Ads Statistics. Table 2 summarizes the distribution of ad-level performance. Across impressions, spending, duration, and efficiency, averages were consistently higher than medians, reflecting highly skewed distributions driven by a small number of extreme outliers. For example, while the median spend was under A\$50, one ad exceeded A\$137,500, and the top-performing ad reached 1 million impressions. These disparities highlight the uneven scale of advertising on the platform, with most ads remaining relatively modest while a few campaigns dominated exposure.

Correlation Insights. Our analysis revealed strong correlations between key metrics (Table 3). Ad expenses strongly correlated with impression numbers, suggesting that higher spending generally led to greater reach. This large effect indicates that approximately 52% of the variance in impressions can be explained by ad spending alone. Ad duration also showed positive, albeit weaker, correlations with both expenses and impressions. All correlations reported in this paragraph and in §5.1 remain statistically significant even after applying a Bonferroni correction for multiple hypothesis testing (adjusted $\alpha = \frac{0.05}{12} = 0.0042$), confirming these relationships are robust (more details in Appendix A.4).

The consistently higher Spearman coefficients relative to Pearson’s suggest that the underlying relationships are monotonic rather than strictly linear, likely reflecting the non-normal nature of the data. While all correlations are statistically significant due to our large sample size ($n=47,503$), the effect sizes for duration-related correlations are small, with duration explaining less than 3% of the variance in impressions and less than 1% of the variance in expenses. This suggests that while ad duration has a statis-

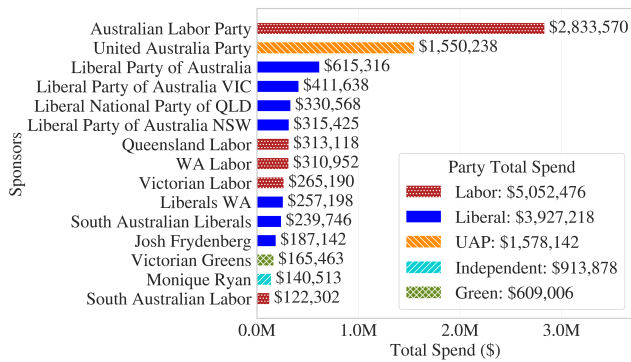


Figure 3: Total spend on political advertising by top sponsors during the 2022 election. Bars represent individual sponsors, color-coded by party affiliation. The legend shows the aggregated total spend for each major party.

tically detectable relationship with both spending and impressions, its practical influence on these metrics is limited compared to other factors.

Sponsors Strategies. On average, sponsors created 33 ads each, though the median of 6 ads per sponsor indicates that most were less prolific, with the most active sponsor producing 3,239 ads. Average sponsor expenditure reached A\$12,782, but the median of A\$645 suggests that most sponsors spent relatively modest amounts, while a few big spenders significantly raised the average. The highest-spending entity, ALP, invested over A\$2.8M (Figure 3). Similarly, while sponsors averaged 774k impressions, the median of 49,747 impressions again points to a skewed distribution, with the top performer achieving over 114.8 million views. Efficiency at the sponsor level achieved an average of 85.05 impressions per dollar spent (median: 70.7).

5 Campaigning Strategies

This analysis focuses on the five political parties in Australia,¹⁰ examining their advertising strategies on Meta platforms during the 2022 federal election campaign. We utilized the PAP dataset, which exclusively contains ads directly sponsored by political parties, ensuring a focused analysis of official campaign strategies.

Spending Patterns. The ALP led in spending, investing over A\$5M (40.2% of total political ad spending) across 13,047 ads in the analysed period. This surpassed the Liberal Coalition’s expenditure by 28%, with the latter spending over A\$3.9M on 9,364 ads. The UAP secured the third position, allocating A\$1.578M to their ad campaigns (RQ1).

Figure 3 (legend) shows the total spend for each of the main parties. Among the top 15 party sponsors by ad spend, the ALP and Liberal Party accounted for five and seven pages respectively. The highest-ranking Australian Greens page was the Victorian Greens, placing 13th overall. Rounding out the top 15 were one page each from the UAP and an

¹⁰ Australian Labor Party, Liberal Coalition, United Australia Party, Independents and Green.

Party	Spend	Ads	Impressions	CPI	Vote (%)
ALP	5.0 A\$M	13,047	292 M	17.3	32.58
Liberal Coalition	3.9 A\$M	9,364	261 M	15	35.70
UAP	1.6 A\$M	416	37.3 M	42.2	4.12
Independents	0.9 A\$M	2,844	51 M	17.9	5.29
Greens	0.6 A\$M	2,088	31.8 M	19.1	12.25

Table 4: Party-level spending, impressions, and vote share. Vote percentages do not sum to 100% due to other minor parties not shown. ALP: Australian Labor Party; UAP: United Australia Party; CPI: Cost Per Impression (A\$).

independent candidate, Monique Ryan, who contested the Kooyong electorate in Victoria taking the seat held since 2010 by Josh Frydenberg (Liberal), Treasurer and Liberal Party Deputy Leader.

Reach and Impressions. Overall, party-sponsored ads generated over 714M impressions across Australia during the campaign. As shown in Table 4, ALP accounted for 40.9% of all impressions, with the Liberal Coalition following at just over 36%. The distribution of spending broadly mirrors the eventual vote shares, with the Liberal Coalition converting slightly fewer impressions into a higher vote share (35.7%) compared to ALP (32.6%). The Greens and the UAP attracted significantly smaller audiences. Despite outspending the Greens by more than double, the UAP generated a similar number of impressions, leading to the highest cost per impression (A\$42.2) among the major parties and translating into just 4.1% of the vote. In contrast, the Greens achieved a lower cost per impression and converted their reach into a considerably higher vote share (12.3%), suggesting a more engaged or better-targeted audience.

5.1 Temporal Analysis

To understand the evolution of political advertising strategies over time (RQ2), we conducted a time series (TS) analysis of daily ad spend, impressions, and ad count for major Australian parties (Figure 4).

A visual inspection reveals a notable increase in activity beginning April 10th—the day the election was called—especially for the Labor and Liberal parties. This trend intensifies as election day approaches, with a sharp peak on May 18th, just before the voting blackout period.

To statistically validate this trend, we used the Mann-Kendall test, which confirmed a significant upward trend in both spend and impressions over time ($\tau=0.925$, $p=8.86 \times 10^{-35}$).

However, strong temporal trends can produce spurious correlations due to autocorrelation in non-stationary data. Therefore, we tested for stationarity using the KPSS test, which confirmed that all three series—spend, impressions, and ad count—were non-stationary (p-values = 0.01).

To remove the trend and allow valid correlation analysis, we applied first-order differencing. A second KPSS test showed that the differenced series were now stationary (p-values = 0.10), validating the transformation (see Appendix A for details).

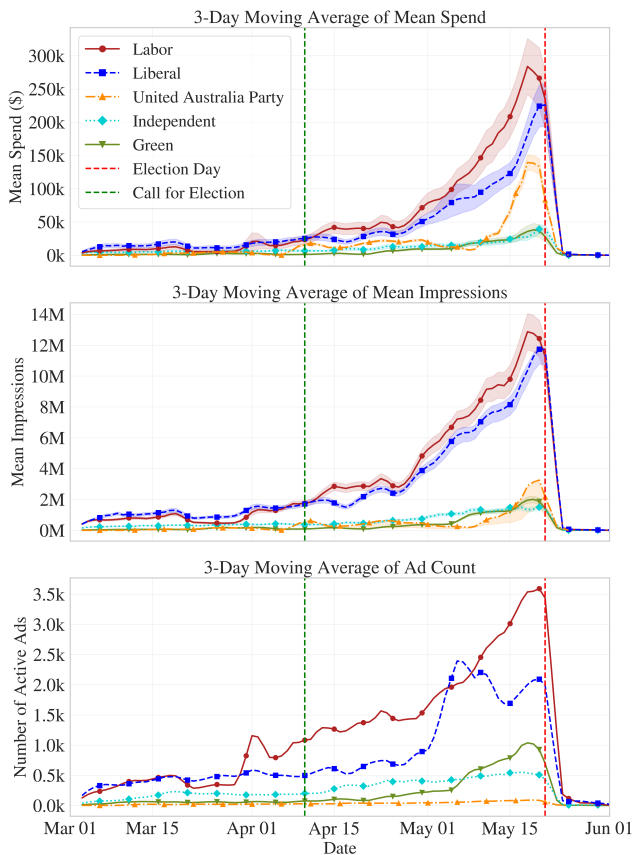


Figure 4: Time series of the total AUD spend, the number of impressions, and the number of unique ads for each day and party in the PAP dataset. In the two upper panels the dashed area represents lower and upper bounds, with the solid line corresponding to the mean value. In all plots we use a 3-day moving average.

With stationary data, we computed correlations between the differenced series (Table 5). We found a robust relationship between changes in spend and impressions, indicating that 89% of the variance in day-to-day changes in impressions can be explained by changes in spend. Additional strong correlations were found between differenced ad count and impressions, and between spend and ad count. More details in Appendix A.4. The exceptionally high correlation between spend and impressions is not only expected in digital advertising, but also likely amplified by Meta’s pricing model, where costs are often directly tied to reach metrics. This systematic relationship meant that parties could reliably predict the scale of their audience based on investment levels, enabling deliberate calibration of ad reach. Different parties exploited this dynamic in varied ways: Labor and Liberal used high ad volume, while the UAP focused on fewer but more impactful ads. Online advertising peaked just before the election, capitalizing on digital platforms’ exemption from Australia’s broadcast blackout.¹¹

¹¹www.acma.gov.au/election-and-referendum-blackout-periods

Pairs	R (R^2)	ρ	Strength
Spend – Impressions	0.94 (0.89)	0.93	Very strong
Ad Count – Impressions	0.83 (0.69)	0.69	Strong
Spend – Ad Count	0.76 (0.58)	0.59	Strong

Table 5: Correlations between differenced time-series of daily ad metrics. Pearson’s R measures linear correlation, with R^2 denoting the proportion of explained variance. Spearman’s ρ measures rank-based (monotonic) correlation. All results remain statistically significant after Bonferroni correction ($\alpha = 0.0042$).

5.2 Demographic Analysis of Ad Impressions

Figure 5 presents how Facebook ad impressions are distributed by age and gender across major Australian political parties (RQ2).

Most parties target the 18–24 and 25–34 age groups, reflecting the dominant demographics on Facebook in Australia.¹² The Greens particularly focus on the 18–24 group, aligning with their youth-centric mission. This emphasis on younger users could be interpreted as an attempt to engage with a demographic traditionally less involved in political processes or more susceptible to changing voting preferences. This targeting approach aligns with SIT (Tajfel 1978; Tajfel and Turner 1979), which suggests that political identities are still forming in younger demographics, making them potentially more responsive to partisan messaging that establishes “us-them” narratives as discussed in §2.1 (RQ3). In contrast, the UAP distributes impressions more evenly across age groups, emphasizing older users and showing minimal engagement with the youngest demographic.

Gender-wise, the ALP tends to reach more women, especially those over 65, a trend seen in most parties. The Greens heavily target young women (18–24), while the Liberal Coalition and Independents maintain a balanced gender reach. The UAP stands out with a male-skewed audience in the 18–54 range. These gendered targeting patterns may reflect parties’ understanding of how gender intersects with political identity formation, with the UAP’s male-skewed approach potentially appealing to a specific social identity that aligns with their right-wing populist messaging (RQ3).

5.3 Regional Distribution

The distribution of ad impressions across Australian regions highlights distinct strategies among political parties, as shown in Figure 6 (RQ2). All parties show a strong focus on eastern Australia—Queensland, New South Wales, and Victoria—home to around two-thirds of the nation’s population. However, the degree of this focus varies between parties.

Liberal & Labor. The major parties, Liberal Coalition and ALP, distribute their ad impressions in line with national population trends, showing a relatively balanced focus

¹²<https://www.statista.com/statistics/680607/australia-facebook-user-distribution-by-age-and-gender/>




Figure 5: Distribution of ad impressions across demographic categories for major Australian political parties during the 2022 campaign. Each box plot represents the variation in the percentage of ad impressions reaching a specific age group and gender combination. The whiskers show the full range of values, boxes indicate the interquartile range, and triangles mark the mean.

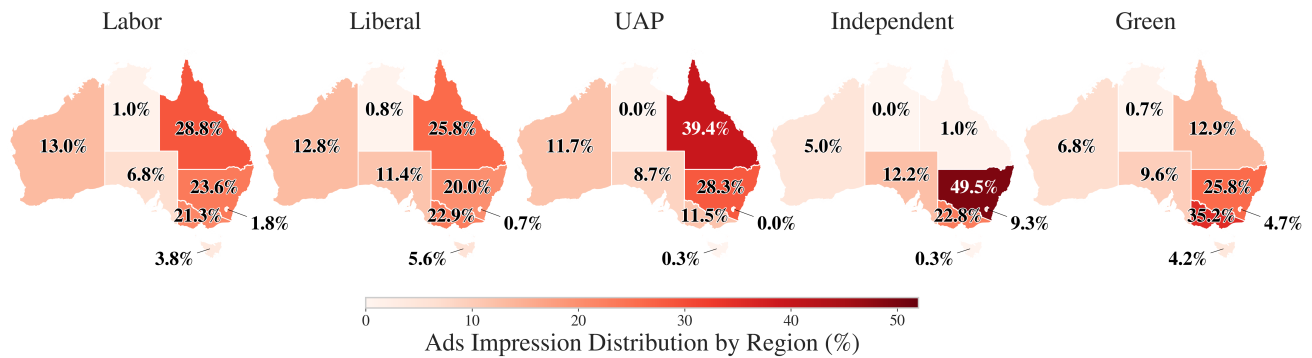


Figure 6: Regional distribution of ad impressions for major Australian political parties. Color intensity indicates the relative concentration of ad impressions in each region. We determined impressions by region multiplying the mean impressions for each ad by the proportion of ad reach attributed to each of them. The top-spend state for the Labor, Liberal, and UAP is QLD, NSW for the independents, and VIC for the Greens.

across the eastern states, with a slight tilt toward Queensland. This proportionality suggests their broader reach stems from greater financial resources, enabling high ad volume and nationwide coverage.

Greens & Independents. The Greens and Independents adopt more targeted approaches. The Greens concentrate on Victoria, likely due to limited resources and stronger urban support in Melbourne. However, unlike their geographic concentration, the Greens’ messaging is thematically consistent across the country: a very large share of their advertising spend (50.03%) is devoted to anti-mining messaging (refer to Appendix B), reaching very high proportions in every state, peaking at 82% in Western Australia and 81% in Tasmania. This indicates that opposition to mining is a core focus of their campaign and is promoted nationwide rather than concentrated in one or two regions. Independents, on the other hand, cluster ad impressions in New South Wales, as illustrated in Figure 7, reflecting their local focus in specific electorates. Key examples include Kylea Tink (North Sydney), Sophie Scamps (Mackellar), and Allegra Spender (Wentworth), each concentrating ads in their areas.

United Australia Party. UAP places heavy emphasis on Queensland, accounting for 39.4% of their impressions—the highest regional share for any party. This reflects the party’s

roots with Clive Palmer. UAP also targets New South Wales (28.3%) and Victoria (11.5%), while showing relatively high activity in Western Australia (11.7%). In terms of campaign themes, UAP is the party most focused on positive, mining-related advertising, devoting 4.36% of its total spend to this topic—still a relatively small proportion overall. These mining-related ads are geographically concentrated, with New South Wales receiving the largest share relative to the party’s total spend.

5.4 Keywords Analysis

The analysis of keyword ad spending shows how Australian political parties used digital advertising to reinforce their ideologies and target specific voter groups (RQ2). As shown in Figure 8, the Greens prioritized climate-related keywords (e.g., “climate,” “change,” “future”), aligning with their environmental agenda and appeal to younger voters, as also observed in §5.2. Major parties spent heavily on promoting their own names and used contrast advertising by targeting rival party names (e.g., ALP targeting “scott_morrison”; UAP targeting “liberal” and “labor”), reflecting a negative campaigning strategy designed to reinforce in-group identity by attacking an out-group (Abramowitz and Webster 2018). This tendency toward negative campaigning on so-

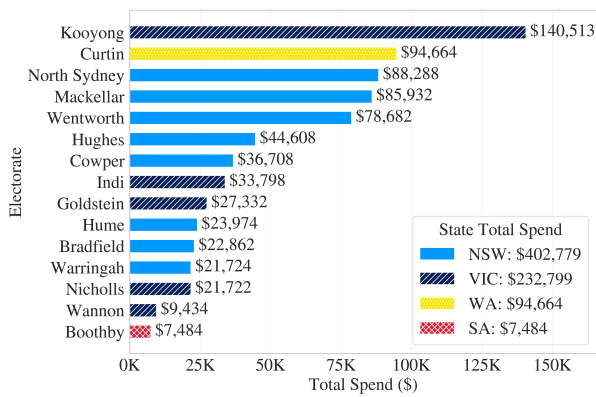


Figure 7: Total campaign spending by Independent candidates in the top 15 electorates where they ran. Bars are colored by state. The legend shows the total spend grouped by state, ranked from highest (NSW) to lowest (SA).

cial media platforms aligns with findings from (Decker et al. 2024). Independents also focused on climate-related keywords, leveraging recent climate crises to attract voters dissatisfied with major parties. This tactic of leveraging salient events to differentiate from established parties is often employed by smaller political entities to gain traction, as noted in Meguid (2005). Overall, the keyword strategies reveal attempts to boost party visibility, attack opponents, and connect with voters on salient issues. This is especially true for smaller parties, who used issue-based appeals to form identity-based voter blocs, linking to broader political theory on tribal identity and issue alignment (RQ3).

5.5 Targeting Strategies

Our initial analysis indicates that parties are possibly using targeted political advertising, focusing their campaigns in specific regions and demographics with tailored messaging. To verify these, we conduct in this section a more detailed investigation of targeting practices, drawing on methodologies established by Capozzi et al. (2021) (RQ2).

Geographical Targeting. Geographical targeting is the most common approach, with 84.7% of ads aimed at a single Australian region (Figure 9 bottom-left panel). This pattern holds across all parties, with roughly 80% of ads state-specific and 20% broadly distributed. Surprisingly, regional and national ads have similar CPI (A\$21.4 per thousand impressions), possibly explaining the popularity of this strategy. UAP stands out with significantly higher CPI (1.9x for targeted and 2.9x for broad ads), likely due to its reliance on fewer high-spending campaigns (see §5.1).

Age-based Targeting. The top-left panel in Figure 9 reveals that the vast majority of ads (85.6%) target all six adult age intervals (18 and above). This approach suggests a broad-reach strategy across adult age groups. We observe that these ads generate an even larger proportion of impressions (87.8%). This indicates that parties are not extensively using age-based targeting.

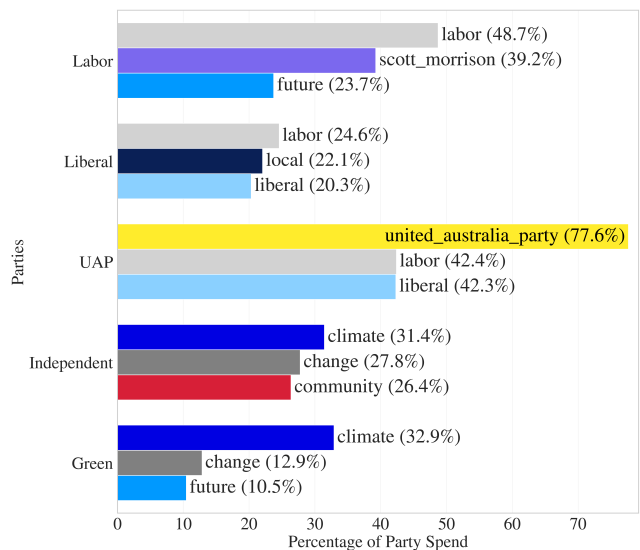


Figure 8: Most prominent keywords in each major Australian party’s ads, shown as a share of their total ad spend (common terms like “vote,” “election,” “party,” and “plan” excluded).

Gender-based Targeting. Gender-specific targeting is even less common: only 4.8% of ads are shown exclusively to one gender, with most of those aimed at females (62.7%). When using a looser threshold ($\geq 60\%$ of impressions to one gender), 14.2% of ads target females and 12.8% target males, indicating a mild preference for female-oriented targeting. This contrasts with findings in Italy (Pierri 2023), where male-focused targeting is more prevalent, despite similar user demographics.

While targeting strategies—especially geographical—are evident, they reflect relatively simple use of the platform’s capabilities. This may be due to limitations in the Meta API, which could obscure more advanced targeting tactics. However, this preference for broad-reach strategies is also consistent with the strategic logic of a compulsory voting system. With no need to micro-target specific segments for get-out-the-vote efforts, campaign resources are logically reallocated toward broad persuasive messaging aimed at a universal electorate (Hillygus and Shields 2008).

6 Discussion

6.1 Summary of Results

Our analysis of political advertising on Meta platforms during the 2022 Australian federal election reveals that:

- **Advertising Dominance:** The ALP and Liberal Coalition dominated the digital landscape, collectively accounting for over 70% of total ad impressions and spending.
- **Temporal Strategies:** Advertising activity intensified following the election call, with parties escalating their efforts as election day approached.




Figure 9: Distribution of ads and impressions by age buckets (top) and regions (bottom) targeted. Left panels display the number of ads, while right panels show the number of impressions. Colors represent different political parties. Percentages indicate the proportion of ads or impressions targeting a specific number of age buckets or regions.

- **Ads Content:** Major parties focused on self-promotion and opponent criticism, while smaller parties emphasized issue-specific messaging. Greens and Independents strategically differentiated themselves by focusing on distinct topics such as climate change and other issues neglected by major parties, appealing to voters dissatisfied with the mainstream political discourse.
- **Targeting Techniques:** Political entities primarily employed geographical targeting, with 84.7% of ads targeted to a single Australian region. Age-based targeting was minimal, with 85.6% of ads targeting all adult age groups. Gender-based targeting was limited, with only 4.8% of ads exclusively targeting one gender. This suggests a preference for broad-reach strategies over sophisticated micro-targeting

6.2 Rethinking On-Line Targeting in Light of Political Psychology

Research into online campaign messaging has largely focused on who is being targeted, when, and how. While these are important empirical questions for understanding the shift from traditional to digital campaigning, less attention has been given to what persuades voters to change their minds. To address this, we can draw on longstanding research into persuasion from social and political psychology.

Humans shift between personal (“I”) and group (“We”) identities (Durkheim 1912), a duality foundational to SIT and Self-Categorization Theory (SCT) (Tajfel 1978; Tajfel and Turner 1979; Turner 1987b,a). These theories show that persuasion is more effective (i.e., lead to enduring behaviour change) when it engages social identities (Mols 2012; Mols et al. 2015, 2023).

Emerging as a critique of rational choice models, SIT, SCT, and social constructivism (Berger Peter and Luckmann

1966) argue that self-interest is shaped by context and identity. This lens helps explain why voters often prioritize group success over personal gain.

Applying these theories to Facebook ads from the 2022 Australian Federal Election reveals strong alignment with a “tribal” partisan social identity model (Abramowitz and Webster 2018; Campbell et al. 1960; Maggiotto and Pierson 1977; Green, Palmquist, and Schickler 2004; Huddy and Bankert 2017).

Major parties, like ALP and the Liberal Coalition, focused on self-promotion and attacking rivals—strategies that reinforce existing identities and deepen polarization. Their investment in ads naming themselves and opponents (Figure 8) illustrates this identity-reinforcing tactic (Huddy, Mason, and Aarøe 2015; Abramowitz and Webster 2018). In contrast, Independents and Greens emphasized issue-based appeals, consistent with niche party theory (Meguid 2005). Their keyword spending (31.4% and 32.9% respectively) suggests targeting identity groups focused on environmental concerns. This strategy creates or reinforces niche partisan affiliations beyond the major party divide. Parties likely combine identity- and issue-based appeals, reinforcing partisan identities while targeting demographics—for example, the ALP’s focus on women may reinforce partisanship while appealing to gender-specific concerns.

This reflects a divide between party loyalists and issue-driven voters. “Rusted-on” voters maintain deep psychological attachments to parties, often inherited and resilient to policy change.¹³

Overall, we find that appeals to social identities are central to campaign strategies. However, our data cannot confirm whether such appeals change minds. Future research should go beyond correlations and use experimental designs to test causal effects. System-level factors may also shape messaging: compulsory voting in Australia might explain scare campaigns aimed at nudging the undecided, while voluntary systems may require stronger identity mobilization (e.g., MAGA in the US). These are empirical questions best addressed through established theories in social psychology and communication science.

7 Limitations

We acknowledge that the reliance on Meta’s Ad Library data may not capture the full scope of political advertising during the election period. Some political ads may not have been properly flagged as such, leading to potential omissions in our dataset. Furthermore, our analysis does not extend to other digital platforms that may have played a role in campaign strategies. Additionally, it is crucial to consider that the apparent simplicity in targeting strategies observed may be a result of limitations in the data provided by the Meta API. Some “invisible” targeting methods could interact with social identities in nuanced ways, potentially amplifying identity-based appeals in ways analysis cannot

¹³Campbell et al. (1960) defined partisanship in *The American Voter* as both a set of beliefs and feelings that culminate in a sense of psychological attachment to a political party.

capture. Political advertisers might be employing more advanced targeting strategies that are not visible due to restricted information disclosure. Our analysis also focused solely on ad text, excluding images, videos, and audio components. As visual and audiovisual content often carry significant persuasive power, future work incorporating multi-modal analysis could reveal additional patterns and deepen understanding of campaign strategies.

Acknowledgments

This work is partially supported by the Swiss National Science Foundation (SNSF) under contract number CRSII5_205975 and by an Australian Research Council (ARC) Future Fellowship Project (Grant No. FT240100022).

References

- Abramowitz, A. I.; and Webster, S. W. 2018. Negative partisanship: Why Americans dislike parties but behave like rabid partisans. *Political Psychology*, 39.
- Aiyappa, R.; DeVerna, M. R.; Pote, M.; Truong, B. T.; Zhao, W.; Axelrod, D.; Pessianzadeh, A.; Kachwala, Z.; Kim, M.; Seekin, O. C.; et al. 2023. A multi-platform collection of social media posts about the 2022 US midterm elections. In *Proceedings of the international AAAI conference on web and social media*, volume 17.
- Arya, P. 2022. *Political advertising on social media platforms during the 2022 federal election*. Australia Institute.
- Bankert, A. 2020. The origins and effect of negative partisanship. In *Research handbook on political partisanship*. Edward Elgar Publishing.
- Bär, D.; Pierri, F.; De Francisci Morales, G.; and Feuerriegel, S. 2024. Systematic discrepancies in the delivery of political ads on Facebook and Instagram. *PNAS nexus*, 3(7).
- Berger Peter, L.; and Luckmann, T. 1966. The social construction of reality. *A Treatise in the Sociology of Knowledge*.
- Birch, S. 2009. 59Compulsory voting and election campaigns. In *Full Participation: A Comparative Study of Compulsory Voting*. Manchester University Press. ISBN 9780719077623.
- Boerman, S. C.; Kruikemeier, S.; and Zuiderveen Borgesius, F. J. 2017. Online behavioral advertising: A literature review and research agenda. *Journal of Advertising*, 46(3).
- Bouchaud, P.; and Liénard, J. F. 2024. Beyond the Guidelines: Assessing Meta's Political Ad Moderation in the EU. In *Proceedings of the 2024 ACM on Internet Measurement Conference*.
- Brodnax, N. M.; and Sapiezynski, P. 2022. From Home Base to Swing States: The Evolution of Digital Advertising Strategies during the 2020 US Presidential Primary. *Political Research Quarterly*, 75(2).
- Bruns, A.; and Moon, B. 2018. Social media in Australian federal elections: Comparing the 2013 and 2016 campaigns. *Journalism & Mass Communication Quarterly*, 95(2).
- Calvo, D.; Cano-Orón, L.; and Baviera, T. 2021. Global spaces for local politics: An exploratory analysis of Facebook Ads in Spanish election campaigns. *Social Sciences*, 10(7).
- Campbell, A.; Converse, P. E.; Miller, W. E.; and Stokes, D. E. 1960. *The american voter*. John Wiley & Sons.
- Capozzi, A.; De Francisci Morales, G.; Mejova, Y.; Monti, C.; Panisson, A.; and Paolotti, D. 2021. Clandestino or Rifugiato? Anti-immigration Facebook Ad Targeting in Italy. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- Conover, M. D.; Gonçalves, B.; Flammini, A.; and Menczer, F. 2012. Partisan asymmetries in online political activity. *EPJ Data science*, 1.
- Decker, H.; Angus, D.; Bruns, A.; Dehghan, E.; Matich, P.; Tan, J.; and Vodden, L. 2024. Topic Diversity in Social Media Campaigning: A Study of the 2022 Australian Federal Election. *Politics and Governance*, 12(0).
- Dommett, K.; and Power, S. 2019. The political economy of Facebook advertising: Election spending, regulation and targeting online. *The Political Quarterly*, 90(2).
- Durkheim, E. 1912. Le probl'eme religieux et la dualit'e de la nature humaine.
- Entous, A.; Timberg, C.; and Dwoskin, E. 2017. Russian operatives used Facebook ads to exploit America's racial and religious divisions. *Washington Post*, 25.
- Fiorina, M. P. 1981. Retrospective voting in American national elections.
- Fowler, E. F.; Franz, M. M.; Martin, G. J.; Peskowitz, Z.; and Ridout, T. N. 2021. Political advertising online and offline. *American Political Science Review*, 115(1).
- Garrett, K. N.; and Bankert, A. 2020. The moral roots of partisan division: How moral conviction heightens affective polarization. *British Journal of Political Science*, 50(2).
- Gayo-Avello, D. 2012a. "I Wanted to Predict Elections with Twitter and all I got was this Lousy Paper"—A Balanced Survey on Election Prediction using Twitter Data. *arXiv preprint arXiv:1204.6441*.
- Gayo-Avello, D. 2012b. No, you cannot predict elections with Twitter. *IEEE Internet Computing*, 16(6).
- Green, D. P.; Palmquist, B.; and Schickler, E. 2004. *Partisan hearts and minds: Political parties and the social identities of voters*. Yale University Press.
- Hillygus, D. S.; and Shields, T. G. 2008. The Persuadable Voter: Wedge Issues in Presidential Campaigns.
- Huddy, L.; and Bankert, A. 2017. Political partisanship as a social identity. In *Oxford research encyclopedia of politics*.
- Huddy, L.; Mason, L.; and Aarøe, L. 2015. Expressive partisanship: Campaign involvement, political emotion, and partisan identity. *American Political Science Review*, 109(1).
- Islam, T.; and Goldwasser, D. 2022. Understanding covid-19 vaccine campaign on facebook using minimal supervision. In *2022 IEEE International Conference on Big Data (Big Data)*. IEEE.

- Islam, T.; Roy, S.; and Goldwasser, D. 2023. Weakly supervised learning for analyzing political campaigns on facebook. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17.
- Islam, T.; Zhang, R.; and Goldwasser, D. 2023. Analysis of climate campaigns on social media using bayesian model averaging. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*.
- Jackman, S. 2001. Compulsory voting.
- Jungherr, A. 2016. Twitter use in election campaigns: A systematic literature review. *Journal of information technology & politics*, 13(1).
- Kreiss, D. 2016. *Prototype politics: Technology-intensive campaigning and the data of democracy*. Oxford University Press.
- Kwiatkowski, D.; Phillips, P. C.; Schmidt, P.; and Shin, Y. 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of econometrics*, 54(1-3).
- Leerssen, P.; Dobber, T.; Helberger, N.; and de Vreese, C. 2021. News from the ad archive: How journalists use the Facebook Ad Library to hold online advertising accountable. *Information, Communication & Society*.
- Maggiotto, M. A.; and Piereson, J. E. 1977. Partisan identification and electoral choice: The hostility hypothesis. *American Journal of Political Science*.
- Meguellati, E.; Civelli, S.; Bernardelle, P.; Sadiq, S.; King, I.; and Demartini, G. 2025. Towards Detecting Persuasion on Social Media: From Model Development to Insights on Persuasion Strategies. arXiv:2503.13844.
- Meguid, B. M. 2005. Competition between unequals: The role of mainstream party strategy in niche party success. *American political science review*, 99(3).
- Mejova, Y.; and Kalimeri, K. 2020. COVID-19 on Facebook ads: competing agendas around a public health crisis. In *Proceedings of the 3rd ACM SIGCAS Conference on Computing and Sustainable Societies*.
- Metaxas, P. T.; Mustafaraj, E.; and Gayo-Avello, D. 2011. How (not) to predict elections. In *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*. IEEE.
- Mols, F. 2012. What makes a frame persuasive? Lessons from social identity theory. *Evidence & Policy*, 8(3).
- Mols, F.; Haslam, S. A.; Jetten, J.; and Steffens, N. K. 2015. Why a nudge is not enough: A social identity critique of governance by stealth. *European Journal of Political Research*, 54(1).
- Mols, F.; Haslam, S. A.; Platow, M. J.; Reicher, S. D.; and Steffens, N. K. 2023. The social identity approach to political leadership.
- O'Connor, B.; Balasubramanian, R.; Routledge, B.; and Smith, N. 2010. From tweets to polls: Linking text sentiment to public opinion time series. In *Proceedings of the international AAAI conference on web and social media*, volume 4.
- Oprea, A.; Martin, L.; and Brennan, G. H. 2024. Moving toward the median: Compulsory voting and political polarization. *American Political Science Review*, 118(4).
- Orr, L. V.; and Huber, G. A. 2020. The policy basis of measured partisan animosity in the United States. *American Journal of Political Science*, 64(3).
- Pierri, F. 2023. Political advertisement on Facebook and Instagram in the run up to 2022 Italian general election. In *Proceedings of the 15th ACM Web Science Conference 2023, WebSci '23*. New York, NY, USA: Association for Computing Machinery. ISBN 9798400700897.
- Ridout, T. N.; Fowler, E. F.; and Franz, M. M. 2021. The influence of goals and timing: How campaigns deploy Ads on Facebook. *Journal of Information Technology & Politics*, 18(3).
- Simon, H. A. 1957. *Models of man: social and rational; mathematical essays on rational human behavior in society setting*. Wiley.
- Singh, S.; and Thornton, J. 2013. Compulsory voting and the dynamics of partisan identification. *European Journal of Political Research*, 52(2).
- Skoric, M.; Poor, N.; Achananuparp, P.; Lim, E.-P.; and Jiang, J. 2012. Tweets and votes: A study of the 2011 singapore general election. In *2012 45th hawaii international conference on system sciences*. IEEE.
- Sosnovik, V.; and Goga, O. 2021. Understanding the Complexity of Detecting Political Ads. In *Proceedings of the Web Conference 2021, WWW '21*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383127.
- Tajfel, H. 1978. The achievement of inter-group differentiation. Differentiation between social groups.
- Tajfel, H.; and Turner, J. 1979. An Integrative theory of inter-group Conflict (In WG Austin and S. Worchel, The social psychology of inter-group relations, Monterey, CA: Brooks).
- Tumasjan, A.; Sprenger, T. O.; Sandner, P. G.; and Welpe, I. M. 2011. Election forecasts with Twitter: How 140 characters reflect the political landscape. *Social science computer review*, 29(4).
- Turner, J. 1987a. Introducing the problem: Individual and group. Rediscovering the social group: A self-categorization theory (pp. 1–18).
- Turner, J. C. 1987b. The analysis of social influence. *Rediscovering the social group: A self-categorization theory*.
- West, E. A.; and Iyengar, S. 2022. Partisanship as a social identity: Implications for polarization. *Political Behavior*, 44(2).
- Yang, K.-C.; Hui, P.-M.; and Menczer, F. 2019. Bot election-eering volume: Visualizing social bot activity during elections. In *Companion Proceedings of The 2019 World Wide Web Conference*.

Ethics Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? Yes.
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? Yes.
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? Yes.
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? Yes. See §6.
 - (e) Did you describe the limitations of your work? Yes. See §6.
 - (f) Did you discuss any potential negative societal impacts of your work? N/A
 - (g) Did you discuss any potential misuse of your work? N/A
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? Yes.
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes.
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? N/A
 - (b) Have you provided justifications for all theoretical results? N/A
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? N/A
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? N/A
 - (e) Did you address potential biases or limitations in your theoretical framework? N/A
 - (f) Have you related your theoretical results to the existing literature in social science? N/A
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? N/A
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? N/A
 - (b) Did you include complete proofs of all theoretical results? N/A
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? N/A
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? N/A
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? N/A
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? N/A
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? N/A
 - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? N/A
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
 - (a) If your work uses existing assets, did you cite the creators? N/A
 - (b) Did you mention the license of the assets? N/A
 - (c) Did you include any new assets in the supplemental material or as a URL? Yes. See §3 in which we provide code and datasets.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? N/A
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? N/A
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? Yes. In §3.2, we link to a Zenodo dataset repository, as per FAIR guidelines.
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? Yes. The Zenodo repository includes a detailed Datasheet file where we addressed the required questions.
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
 - (a) Did you include the full text of instructions given to participants and screenshots? N/A
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? N/A
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? N/A
 - (d) Did you discuss how data is stored, shared, and de-identified? N/A

A Statistical Details for Temporal Analysis

A.1 Stationarity Testing

To assess the validity of correlation metrics in the presence of trending data, we applied the KPSS test to our three series: daily ad count, spend, and impressions. All were non-stationary in their original form:

- Ad Count: KPSS = 1.37, $p = 0.01$
- Spend: KPSS = 1.08, $p = 0.01$
- Impressions: KPSS = 1.21, $p = 0.01$

After first-order differencing, all series became stationary:

- Differenced Ad Count: KPSS = 0.12, $p = 0.10$
- Differenced Spend: KPSS = 0.16, $p = 0.10$
- Differenced Impressions: KPSS = 0.16, $p = 0.10$

It is important to note that the p-values reported by the KPSS test are limited to the values 0.01 and 0.10. This is because the implementation in statsmodels interpolates p-values from Table 1 in Kwiatkowski et al. (1992) and returns boundary points when the test statistic falls outside the table of critical values. Specifically, if the p-value would be smaller than 0.01 or larger than 0.10, the function returns these boundary values instead. This explains why our non-stationary series all show $p = 0.01$ (indicating strong evidence against stationarity) and our stationary series show $p = 0.10$ (indicating no evidence against stationarity).

A.2 Time Series Differencing




Figure 10: Original time series data showing daily ad counts, spending, and impressions from March to May 2022. All three metrics show strong upward trends, particularly accelerating in May, with a sharp decline at the end of the period coinciding with the election blackout. The daily ad count reaches approximately 10,000 at its peak, daily spend approaches 800,000 AUD, and daily impressions reach over 40 million.

Figure 10 displays the raw time series data for our three metrics. The consistent upward trend across all three metrics indicates non-stationarity, which was confirmed by our




Figure 11: First-differenced time series showing daily changes in ad counts, spending, and impressions. The differencing process removes the trend component, resulting in stationary series fluctuating around zero.

KPSS test results. We observe that all three metrics show relatively steady growth from March through April, followed by a more dramatic increase in early-to-mid May, and then a sharp drop-off near the end of the observation period corresponding to the election blackout period.

Figure 11 presents the first-differenced versions of our time series, where each point represents the day-to-day change in the respective metric. The differencing procedure effectively removed the trend component, resulting in stationary series that fluctuate around zero, as confirmed by our second round of KPSS tests.


A.3 Correlation Analysis with Differenced Data

Figure 12a visualizes the Pearson correlation coefficients between our differenced (stationary) time series. The correlation matrix reveals strong positive relationships between all three metrics, with the strongest correlation ($r = 0.94$) observed between daily changes in spend and impressions.

Figure 12b presents the Spearman rank correlation coefficients for the same data. This non-parametric measure assesses monotonic relationships without assuming linearity. The pattern of correlations is similar to the Pearson coefficients, though with some notable differences. The Spend- Impressions correlation remains extremely strong ($\rho = 0.93$), essentially identical to the Pearson coefficient, indicating a highly consistent relationship. However, correlations involving Ad Count show more pronounced differences between Pearson and Spearman metrics (particularly Pearson Spend-Ad Count, dropping from $r = 0.76$ to $\rho = 0.59$), suggesting potential non-linear components in these relationships.

A.4 Multiple Hypothesis Testing Correction

When conducting multiple statistical tests, the probability of obtaining false positive results increases with each ad-



(a) Pearson correlation matrix of stationary time series. (b) Spearman correlation matrix of stationary time series.

Figure 12: Correlation matrices of stationary (differenced) time series. (a) Pearson correlation matrix showing strong positive relationships between all metrics, with the strongest correlation (0.94) between spend and impressions. (b) Spearman correlation matrix exhibiting similar patterns but with some differences in magnitude, suggesting potential non-linear relationships particularly involving ad count.

ditional test. To control for this family-wise error rate, we applied the Bonferroni correction method, which adjusts the significance threshold by dividing the standard alpha level (0.05) by the number of tests performed.

Our analysis included a total of 12 correlation tests across the paper:

1. Descriptive Statistics (§4)

- Pearson and Spearman correlations between ad expenses and impression numbers
- Pearson and Spearman correlations between ad duration and expenses
- Pearson and Spearman correlations between ad duration and impressions

2. Temporal Analysis (§5.1)

- Pearson and Spearman correlations between differenced spend and impressions
- Pearson and Spearman correlations between differenced ad count and impressions
- Pearson and Spearman correlations between differenced spend and ad count

The Bonferroni-corrected significance threshold was $\frac{0.05}{12} = 0.0042$. All our reported correlations have p-values substantially lower than this adjusted threshold, with most being below 10^{-9} , indicating that the observed relationships remain statistically significant even after controlling for multiple comparisons.

Additionally, our analysis included diagnostic tests for time series properties:

1. KPSS tests for stationarity (6 tests total: 3 on original series, 3 on differenced series)
2. Mann-Kendall test for trend (1 test)

These diagnostic tests were conducted sequentially as prerequisite checks for our correlation analyses rather than simultaneous hypothesis tests. The Mann-Kendall test ($p =$

8.86×10^{-35}) remains highly significant under any reasonable correction. The KPSS test results for our original series (all $p = 0.01$) and differenced series (all $p = 0.10$) represent boundary values from the statistical implementation rather than exact p-values, and their interpretation remains consistent regardless of correction.

B Mining-Related Advertisement Analysis

B.1 Methodology

To analyze mining-related political advertisements, we employed a keyword-based detection approach combined with regional and party-level aggregation. We identified mining-related content using a comprehensive set of keywords including: *mine, mining, miner, resource, mineral, iron, ore, gas, oil, coal, drill, energy, fifo, export, and dig*. All keywords were lemmatized and matched with ad texts.

Each advertisement was classified by its dominant region (the region receiving the highest percentage of targeting) and analyzed across three dimensions: regional focus, party focus, and party-region intersections. We calculated both advertisement count percentages and spending percentages to capture different aspects of mining-related campaign emphasis.

B.2 Key Findings

Party-Level Results The analysis reveals stark differences in mining-related advertising focus across political parties. The Greens demonstrate the most intensive mining focus, allocating 50.03% of their total advertising spend (\$304,702 out of \$609,006) to mining-related content—an order of magnitude higher than any other major party. The United Australia Party allocated 4.36% of their spend (\$68,887 out of \$1,578,142) to mining topics, while the major parties showed much lower focus: Liberal (2.99%) and Labor (1.59%).

Regional Distribution Mining-related advertising is disproportionately concentrated in smaller jurisdictions rather than major mining states. The Australian Capital Territory leads with 7.74% of advertising spend on mining issues, followed by Tasmania (7.47%) and South Australia (6.79%). Notably, Queensland—despite its significant mining economy—ranks lowest among major states at 3.34%, while Western Australia shows moderate focus at 6.51%.

Strategic Targeting The party-region analysis reveals distinct strategic approaches. The Greens maintain consistently high mining-related spending across all regions (ranging from 50-82% of regional spend), suggesting a national anti-mining campaign strategy. The United Australia Party shows more targeted geographic focus, with highest mining-related spending in New South Wales (5.4%) while maintaining minimal presence in other states.

Limitations of mining-related keyword analysis Our keyword-based approach provides a practical way to identify mining-related ads but may miss content using alternative phrasing, capture unrelated ads, ignore multimodal cues, and oversimplify regional targeting. Consequently, results should be seen as indicative rather than exhaustive.