Failure to Launch:
The impact of major party candidate dropout *

Christopher Burnitt†

November 10, 2026
Click here for the latest version

Abstract

Prominent partisan divides may inflate support for major party candidates, reinforcing their dominance. To test this, I exploit quasi-exogenous variation in the presence of one of two major political parties, arising from an unanticipated failure by the party to register candidates in some local government elections in New South Wales, Australia. I combine a difference-in-differences approach with detailed ballot-level ranking data. With this, I create a benchmark that separates out the mechanical substitution effect that arises from candidate absence even when voters rank candidates the same way. I find that the remaining major party's vote share falls significantly short of this benchmark, by 7.1 percentage points of enrolled voters. This shortfall is equivalent to 14.2 seats' worth of votes across all affected contests in which the party stood. Analysis of a corpus of candidate statements taken from online news articles is consistent with the framing of the contest changing in the absence of one major party — away from references to candidates' own parties, and towards communal moral language. These results suggest that major parties benefit from the partisan framing that their rivalry creates.

Keywords: electoral competition, political campaigns, voter preferences

*I would like to thank my supervisors Sascha O. Becker, Helios Herrera and Mateusz Stalinski for their guidance. I would also like to thank Margot Belguise, Menna Bishop, David Boll, Francesco Ferlenga, Simon Heß, Adam di Lizia, Jaime Marques-Pereira, Hyungmin Park, Devesh Rustagi, Lily Shevchenko, Stephane Wolton, David Yanagizawa-Drott, and numerous seminar participants for their helpful comments.

†University of Warwick, christopher.burnitt@warwick.ac.uk.

1 Introduction

Elections are often framed around a partisan divide, emphasising major parties at the expense of the remaining alternatives. A substantial share of citizens also feel that no political party represents their views — a median of 42% of adults across 24 countries (Pew Research Center, 2024).[1] Together, these facts raise concerns that the prominence of partisan divides is not just a result of underlying voter priorities and preferences. Instead, it is possible that the presence of a prominent divide distorts how voters rank candidates, in a way that weakens voters' ability to discipline candidates. Specifically, the framing may lead voters to treat their vote as a decision on the partisan divide, and thus discount minor parties and independent candidates outside that divide. In this case, we might expect voters to support major parties more when both are present than when the partisan divide is not there.

In this paper, I ask whether a prominent partisan divide shapes an electoral contest in a way that inflates support for major party candidates. To study this, I exploit quasi-exogenous variation in the presence of a major political party in the 2024 local government elections in New South Wales (NSW), Australia, a state roughly equal in population to mid-sized European countries.[2] The centre-right Liberal Party — one of Australia's two dominant political parties — was unable to field 138 of its intended candidates due to an error by party officials, who missed the statutory registration deadline. As a result, these candidates did not appear on ballots in the elections held one month later. In effect, one side of the partisan divide was missing from affected contests, with the candidate field fixed and the campaign yet to run.

This unique setting combines two features that are ideal for testing how a partisan divide shapes voting decisions. The first feature is the variation in major party presence in these contests — similar circumstances may arise in the case of politician deaths,[3] or bans and scandals.[4]

The second feature is rare: the presence of a preferential voting system, in which voters rank candidates. The challenge with variation in the presence of a party is that removing one party from a contest mechanically forces orphaned voters to choose differently. Any resulting

[1] The study also found that a median of 74% of respondents feel elected officials do not care what people like them think and 59% are dissatisfied with how their democracy is functioning.

[2] As of 2024, NSW had a population of 8.5 million, making it slightly smaller than Switzerland (9 million) and Austria (9.2 million), but larger than Serbia (6.6 million), Bulgaria (6.4 million), Denmark (6 million), Finland (5.6 million), and Norway (5.6 million). https://www.abs.gov.au/statistics/people/population/national-state-and-territory-population/latest-release, accessed 2026-02-20; https://data.worldbank.org/indicator/SP.POP.TOTL, accessed 2026-02-20.

[3] See, e.g., Spenkuch (2014); Dinas et al. (2016); Das et al. (2024); Di Leo et al. (2025).

[4] For example, Ecuador suspended its largest opposition party, the Citizens' Revolution (https://cepr.net/newsroom/ecuador-suspends-the-countrys-largest-opposition-party/, accessed 2026-03-31), and there have been calls to ban the AfD, the second-largest party in the German parliament (https://www.bbc.co.uk/news/articles/cwy6zk9wkrdo, accessed 2026-02-21). Scandals can have similar effects: Marine Le Pen narrowly avoided a ban from public office over embezzlement charges ahead of the 2027 French presidential election (https://www.bbc.co.uk/news/articles/cly85qjg45no, accessed 2026-09-17), and Italy's long-dominant Democrazia Cristiana collapsed following the corruption investigations of the early 1990s.

change in electoral support therefore combines substitution under fixed rankings with changes in rankings themselves. To identify changes in rankings, it is necessary to know how voters would substitute under fixed rankings — ranked ballot data provides exactly this information.

The empirical design draws on these two features. The registration error is the basis for the first component: a difference-in-differences (DiD) design, comparing changes in electoral outcomes in contests where the Liberal Party could not run due to the error, with contests where it could run, across elections between 2012 and 2024. The nature of the error makes it well suited for this design — it was the fault of state-level party officials (not local organisational issues) and those responsible realised too late to request help, delegate the remaining paperwork, or prioritise which candidates were registered.[5] These factors suggest that the set of affected contests was unrelated to local conditions, providing a priori support for the identifying assumption of parallel trends. Consistent with this, affected and unaffected contests are balanced across a wide range of local economic, demographic and political characteristics that might be associated with changes in electoral outcomes over time (joint $p = 0.437$).

The second component of the empirical design uses the ranked ballot data to account for the mechanical substitution effect, which would also be captured in the treatment effect on raw vote shares. I do this by removing Liberal candidates from *all* ballot papers on which they appear — unaffected contests and earlier elections — and observing which remaining candidate is ranked highest. From this, I construct the implied first-preference vote shares that would have been realised if the Liberals had been absent and rankings unchanged.[6] For affected contests in 2024, this construction is simply the raw vote shares, since Liberals were genuinely absent.

The main analysis thus compares vote shares in contests where the Liberals were actually absent with vote shares where they were not, while accounting for the mechanical effect of Liberal absence. Under parallel trends, the counterfactual for treated contests acts as a benchmark: the vote shares expected in the absence of the Liberals with rankings unchanged. The estimand then represents the excess — surplus or shortfall — of vote shares in treated units compared to that benchmark, with the benchmark and the actual change in vote shares separately identified. A non-zero excess means that voter rankings departed from the benchmark — the partisan divide shaped how voters ranked candidates rather than merely describing their preferences.

I find that the centre-left Australian Labor Party, the remaining dominant party, performs significantly worse than the fixed-rankings benchmark predicted. The benchmark predicted an increase in Labor's vote share of 8.5 pp, but the actual vote share increased by only 1.4 pp,

[5] See Section 2.3 for further details.

[6] In doing so, I adjust for ballot exhaustion. Valid votes might only rank Liberal candidates, and these votes would be excluded from this construction — ignoring this would understate the transferred votes. In Section 4.1 I explain the construction in more detail. In Section 5.2 I show that the results hold under different approaches to ballot exhaustion.

revealing a net shortfall of 7.1 pp (p=0.003) or 29% of Labor's mean pre-treatment vote share. This shortfall is meaningful — candidates in these multi-member proportional representation electorates are elected if they achieve a quota, substantially less than 50%. Labor missed out on an average of 0.49 seats' worth of votes per affected contest, with local councils having an average of 12.6 councillors each. Across all treated contests in which Labor stood, this is a loss of 14.2 seats' worth of votes. The results indicate that the presence of the Liberal Party in contests alongside Labor does have a substantive impact on how voters rank candidates, with Labor support inflated by the presence of its Liberal rivals.

This result survives a variety of robustness checks. The estimate remains statistically significant under different assumptions about the structure and correlation of errors. The estimate is not sensitive to controlling for general local sentiment towards the Liberal Party as proxied by state-level voting patterns. It is also not sensitive to flexibly controlling for differential trends based on baseline political, social and economic variables through double machine learning (Chernozhukov et al., 2018), which selects how a large set of controls enters while preserving valid inference on the treatment effect. The result is robust to alternative methods of dealing with ballot exhaustion.

Spillovers are unlikely to drive the result, which is largely insulated by construction. Liberal candidates are removed from rankings, so spillover concerns that act through how voters rank Liberal candidates in control contests do not affect it. This includes Liberals investing more resources in campaigns in unaffected contests to compensate for missing candidates elsewhere. Empirically, unaffected contests in the same councils as affected ones, where spillovers are most likely, show no differential change relative to other controls. There is also no evidence that voters changed their registered addresses to vote for Liberals elsewhere, either within the same council or between councils. It is also unlikely that the results are driven by voters blaming Labor for the error or for manipulating the electoral system. Although Labor controlled the state government, the deadline was statutory and enforced by the independent NSW Electoral Commission. Additionally, out of almost 250 online news articles mentioning the error, none assign blame to Labor or any other party apart from the Liberals.

In contrast to the Labor shortfall, the remaining candidates — independents, smaller parties and local groups — gained at least the predicted share of votes. Their combined vote share increased by 18.4 pp against a benchmark increase of 14.5 pp, outperforming the fixed-rankings prediction by 4.0 pp (p=0.080).

I also document a 2.3 pp increase in invalid votes as a share of enrolled voters (p<0.001). About 75% of the increase in invalid votes is due to blank ballot papers. Accordingly, I interpret this as a protest vote, with a substantial share of the orphaned Liberal voters discarding their

votes rather than transferring them to an alternative candidate. I find a small decrease in turnout of 1.1 pp which is marginally significant in the main specification (p=0.087) but does not survive several robustness checks. Together, these estimates indicate that about 10–15% of Liberal voters are unwilling to vote for an alternative candidate. This is in line with the 10% estimate of Fernández-Villaverde and Sanz (2025) for a radical-right party in Spain, and well below the 50% estimate of Pons and Tricaud (2018) for third-placed candidates in French run-off elections.

I then investigate the mechanisms that could account for these results. The shock occurred one month before the election, allowing campaigns to respond while leaving the set of candidates fixed. I first consider whether campaign rhetoric and the framing of the local political debate changed in the absence of the Liberal Party, in ways that may reduce support for Labor. The presence of two major national parties may dominate the debate, framing the contest as between Labor and Liberal, and the values they represent. In the Liberal Party's absence, the contest cannot be framed around this divide. To study this, I use a large language model (LLM) to extract text that can be directly attributed to specific local election candidates from a corpus of 9,357 online news articles which mention at least one candidate by name. Consistent with this narrative, I report that candidates, especially Labor candidates, make fewer references to their own party, and that a larger share of discussion addresses the topic of party politics, its role in local government and the balance of party representation on the council. The moral content of candidate-attributed text also shifts, away from universalist moral values and towards communal moral values, a dimension that Enke (2020) shows predicts voting behaviour more strongly than typical demographic indicators. This is accompanied by a small but statistically insignificant increase in the use of terms associated with localism. Together, these suggest that contests are reframed away from the partisan divide and towards local matters, giving candidates outside the major parties the prominence that the divide otherwise denies them.

I then address whether changes in campaign spending patterns can account for Labor's shortfall. If Labor chose to spend less in treated areas (due to the reduced competition) or other candidates chose to spend more (to seize the opportunity presented by Liberal absence), this could explain the deviation from expectation. Analysing compulsory financial disclosures of all campaign spending by candidates and parties, I do not find any significant evidence of changes in spending behaviour, suggesting that expenditure is inelastic with respect to local conditions. However, the estimates are noisy, so I cannot exclude potentially meaningful effects.

I explore whether policy platforms, and changes in policy platforms, can account for the results. Local groups and independents, with no clear party brand to maintain, may adapt their positions more readily than Labor. I use an LLM to position candidates and groups on a

left-right scale, and a pro-/anti-development scale, based on the body of news text that relates to them. I do not find any evidence of an overall shift in policy positions, nor do I find evidence of a convergence on the usual policy positions of Liberal candidates. I also do not find evidence that right-leaning candidates gain a significant share of orphaned Liberal votes. However, in the absence of comprehensive manifesto data, the classification of candidates into right- and left-leaning is necessarily noisy and incomplete, so the conclusion is not definitive.

Lastly, I discuss the possibility of voter-side explanations for the results. The relative gain of small groups and independents compared to Labor is unlikely to be driven by an increased incentive to learn about these minor candidates when a voter's first preference choice is unavailable: ranking additional preferences is optional when voting for a party ticket, which most Liberal voters do, and therefore ranking Labor second reveals an active choice rather than a passive one that could be driven by lack of information acquisition. Two voter-side explanations remain that could produce the Labor shortfall independently of the observed change in framing, and that the available data cannot rule out. First, voters may hold preferences over the composition of the council (e.g., a balance of different views). Second, the choice of a second preference may be a materially different question from the choice of the first, so that rankings lack Independence of Irrelevant Alternatives for reasons unrelated to how the contest is framed.

1.1 Related Literature

This paper relates to several bodies of literature. It contributes to the literature studying how the set of candidates affects voting outcomes. Closest to this paper are those that study the presence or absence of particular parties and candidates. Fernández-Villaverde and Sanz (2025) study a similar registration failure by a radical-right party in Spain, finding that voters substitute mainly to the traditional right-wing party, or cast a protest vote. Arenas (2016, 2021) studies the effects of a party ban in Spain on protest voting and future support for the banned party. Pons and Tricaud (2018) show that the presence of a narrowly qualifying third candidate in the second round of French run-off elections boosts turnout and reduces votes for the top two candidates. A related strand studies the effect of deaths of politicians, some belonging to major parties, on political behaviour (Das et al., 2024; Di Leo et al., 2025; Dinas et al., 2016). These differ in that the affected party is not absent from the ballot. Absence itself is often an endogenous decision. Granzier et al. (2023) investigate the strategic coordination problem posed by vote splitting between candidates from allied parties, with Nibourel (2025) finding evidence of coordinated dropout by left-wing candidates in response to the radical right. A key factor in who runs appears to be past electoral performance (Granzier et al., 2023; Anagol and Fujiwara, 2016), though this can be mitigated by major party affiliation (Khalil et al., 2024). Others

have used variation in barriers to entry to understand who runs in elections and the effect this has on electoral outcomes (Faravelli et al., 2025; Kapoor and Magesan, 2018; Drometer and Rincke, 2008). I contribute to this literature in two ways. First, I study the absence of one of the two dominant parties, where existing work studies radical, regional or third-placed parties and candidates, or strategic withdrawals. Second, by combining ranked ballot data with a difference-in-differences design, I separate the mechanical effect of substitution from changes in how voters rank the remaining candidates. This allows me to answer a different question to the existing literature, which has studied cases without ranked ballots — rather than asking where voters turn in the absence of their most preferred option, I ask whether the absence changes how they rank the candidates who remain.

The methods in this paper, and the construction of the fixed-rankings benchmark, build on other work that uses ranked ballots as data. Without ranked ballot data, voters' rankings are unobserved — studies instead use structural assumptions (Kawai and Watanabe, 2013) or patterns in aggregate results (Muller and Page, 2015). Previous work using ranked ballot data has sought to describe the rankings themselves (Mitchell, 2014), and how they relate to partisan directions on how to rank secondary alternatives (Reilly, 2021), which parties can negotiate and trade (Sharman et al., 2002). McCune and Wilson (2023) consider hypothetical contests using real ranked ballot data, but have no real-world analogue with which to compare. The literature suggests that these rankings are unlikely to be distorted by strategic voting. In preferential voting settings, strategic voting is an NP-hard problem, meaning it is computationally demanding and potentially intractable (Bartholdi and Orlin, 1991). It also generally offers lower expected returns compared to plurality systems (Eggers and Nowacki, 2024), though the opportunities still exist and may be as likely to arise as in plurality systems (Gibbard, 1973; Baltz, 2025). As such, voters typically follow simple heuristics like sincere voting — Van der Straeten et al. (2010) find 91% of orderings on ballots in a lab setting under a single transferable vote system are consistent with sincere rankings. To my knowledge, this paper is the first to compare actual voting behaviour in an election where a party is absent with implied voting behaviour from real ranked ballot data.

Another body of literature studies how campaigns respond to the set of opponents they face. In the Downsian tradition, competition shapes policy platforms (Downs, 1957). Di Tella et al. (2025) show that policy platforms converge towards the centre between first and second round elections in French run-offs and US primary and general elections. Tähtinen (2025) and Abou-Chadi and Krause (2018) show that mainstream party platforms respond to the success of populist and radical right parties. Policy platforms may be costly to move — politicians may instead emphasise certain issues, either to broaden their appeal (Le Pennec, 2023) or to emphasise

their differences (Ash et al., 2017; Glaeser et al., 2005). The structure of electoral competition has important implications for the incentives to engage in certain types of campaigning, such as negative advertising (Gandhi et al., 2015; Galasso et al., 2021). Several papers have studied the incentives for and returns to campaign messaging, advertising and spending (Erikson and Palfrey, 2000; Iaryczower and Mattozzi, 2013; Levitt, 1994; Spenkuch and Toniatti, 2018; Bekkouche et al., 2022; Kendall et al., 2015; Le Pennec and Pons, 2023). Enke (2020) and Grillo et al. (2026) emphasise the importance of moral language in particular in explaining voting behaviour. I study campaigns in which a major party is missing. Campaigns shift in rhetoric away from partisan identity and towards communal moral values, while I find no evidence of a shift in policy platforms. Together with Labor's shortfall relative to the fixed-rankings benchmark, these results suggest that major parties may benefit from a contest framed by their rivalry.

My results on invalid votes and turnout speak to the literature on expressive and instrumental motivations for voting. Expressive motivations have long been considered essential in explaining observed political behaviours, such as turnout and voting patterns (Fiorina, 1976; Brennan and Lomasky, 1993). Spenkuch (2018) studies these motivations when voter preferences are unobserved, finding that salient features of voting data remain unexplained by both motivations, and that voters may exhibit both motivations. The motivation and extent of strategic voting can vary with the electoral system (Fujiwara, 2011). The motivation to turn out and vote expressively is responsive to the closeness of the election (Bursztyn et al., 2023; Faravelli et al., 2020; Morton et al., 2015), who is on the ballot (Pons and Tricaud, 2018; Fernández-Villaverde and Sanz, 2025; Bol and Ivandic, 2022), and the electoral system (Eggers, 2014b; Faravelli and Sanchez-Pages, 2014). Voters may also turn out due to valuing collective actions or overweighting small probabilities (Faravelli et al., 2025). Where voting is compulsory, as in Australia, voters may instead use invalid voting as a means of expression. A considerable share of invalid votes in Australian elections are likely generated by mistakes (Hill and Young, 2007), while Hill and Rutledge-Prior (2016) argue that blank and defaced ballot papers may reflect disaffection. Gonzales et al. (2022) show that larger penalties for abstention do little to change electoral outcomes, as they shift abstention to invalid ballots. Ujhelyi et al. (2020) study the "None of the above" option on Indian ballots, which reduces abstention. I contribute to this broad literature by providing an estimate of the share of voters from a major party that are unwilling to vote for any of the remaining alternatives. I find that 10–15% of these voters choose not to vote for a remaining candidate, in line with about 10% for a radical right party in Spain (Fernández-Villaverde and Sanz, 2025) and much lower than the roughly 50% estimate for third-placed candidates in French run-off elections (Pons and Tricaud, 2018), suggesting that the status of the candidate and the presence of compulsory voting are unimportant

in determining this share, but that anchoring to a candidate may occur in the first round of multi-stage elections.

This paper also relates to the literature on how salient partisan divides shape beliefs and preferences. Affective polarisation — the dislike and distrust of people from the other party — has grown in several countries, most notably the US (Boxell et al., 2024). Media and journalists face incentives to emphasise division (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2010; Perego and Yuksel, 2022; Balbuzanov et al., 2026). In this context, partisan signals can distort economic decisions and trust in regulators (Burnitt et al., 2026), as well as shape policy preferences (Kashner and Stalinski, 2024). I provide evidence from a real election that a prominent partisan divide also distorts how voters rank candidates: when one major party is absent, the other receives fewer votes than the fixed-rankings benchmark implies. The changes in campaign rhetoric are consistent with a shift in the salience of party identity, which could explain the shortfall of the remaining major party. This further connects to theories of salience in decision-making, especially in politics (Ansolabehere and Puy, 2018; Bordalo et al., 2013; Nunnari and Zápal, 2025; Bonomi et al., 2021), and to empirical evidence that salience shocks can shift votes (Yousaf, 2021).

This paper contributes to the literature on why two parties may dominate electoral contests. Duverger's Law predicts that two-party dominance will arise under plurality rule (Duverger, 1951). Fujiwara (2011) finds empirical support for this. Established parties may deter or defeat new entrants by strategically selecting policy positions (Palfrey, 1984; Callander, 2005), even in non majoritarian run-off elections (Buisseret, 2017). These explanations cannot explain why major party support is inflated in my setting: the electoral system is proportional representation with ranked ballots rather than plurality rule, and the set of candidates is fixed at the time of the shock. I also do not find any evidence that candidates shifted policy positions. I contribute by showing a further source of two-party dominance, which requires neither plurality rule nor entry deterrence: the presence of its rival raises a major party's support. In this sense, two-party dominance is self-reinforcing.

Finally, the results have implications for the literature on electoral competition and accountability. The politician-voter relationship can be seen as that of principal and agent, with elections providing the disciplining tool (Ferejohn, 1986). Greater electoral competition can strengthen this accountability effect, leading to better policies, lower political rents and less corruption (Besley et al., 2010; Svaleryd and Vlachos, 2009; Ferraz and Finan, 2011). Competition from more than just the top two ranked candidates can be important: Chatterjee et al. (2024) show that the presence of a relevant third-ranked candidate, whose vote share exceeds the winning margin and thus poses a coalition threat for the winner, improves development in

Indian constituencies as measured by nightlights. Partisanship or other strong group identities can weaken accountability (Eggers, 2014a; Padró i Miquel, 2007), since political power can be used to favour specific groups (Cheng and Swee, 2024). Information about candidates can combat this weakened accountability (Casey, 2015). My results suggest that the competition provided by minor parties and independent candidates can be limited by the presence of a prominent partisan divide, which inflates support for a major party and shapes the narrative and rhetoric of campaigns.

2 Background

2.1 Local government in Australia

Local governments in Australia are responsible for managing local services and infrastructure. They also play a key role in one of Australia's most significant policy issues — housing — being responsible for approving developments up to AU$30 million (US$21 million) in cost.[7] Local governments are run by a council of 6–15 councillors, led by a Mayor.[8]

Elections for councillors are held every 4 years concurrently across the state. Local government areas (LGAs) may be divided into wards to elect their councillors.[9] These wards are multi-member, proportional representation electorates with optional preferential (ranked-choice) voting, also known as a single transferable vote system. Voters may choose to vote for either individual candidates or group tickets and must rank a minimum number of alternatives for the vote to be valid, but may choose to rank more than the required number.[10] Voting is compulsory — NSW citizens over the age of 18 may be fined AU$55 (US$39) for failing to turn up to vote without a valid reason. Consequently, total turnout at local elections has ranged from 82.1–85.2% since 2004. In comparison, the 2024 US presidential election had a turnout of 64.1% and the 2024 UK general election had 59.7%.

2.2 Political parties in NSW local elections

The Liberal Party of Australia is one of two major political parties in Australia, the other being the Australian Labor Party. The Liberal Party emphasises small government, low taxes,

[7] According to a 2024 Ipsos survey, housing ranked 2nd as a top issue facing Australia, behind cost of living. It was rated as a top-3 issue among 42% of Australians, and around 50% among those from NSW. https://www.ipsos.com/en-au/issuesmonitor, accessed 2026-02-13.

[8] Mayors may be selected by the council from among the councillors, or directly elected at the Local Government Area level.

[9] I will use the term ward to refer to undivided LGAs. The division of LGAs into wards is at the discretion of councils, meaning that there is substantial variation in the population of a ward and number of candidates representing a ward. To account for this imbalance, I weight observations by population in the main regression specifications.

[10] If voting for groups, voters must rank at least 1 group. If voting for individual candidates, voters must rank at least $\frac{n}{2}$ candidates, where n is the number of councillors to be elected in the ward.

encouraging private enterprise, and often more conservative social values. The Liberal Party was founded in 1944, and in the following 80 years held the prime ministership at the national level for more than 50 years through its coalition with the smaller National Party. Centre-left Labor held the prime ministership for the remainder of those 80 years. In local government elections, these major parties primarily contest urban LGAs. The *Australian Greens*, another national party, also contest local elections in many places. The majority of remaining candidates in local elections run as independents, or for small or local political parties.

Figure 1 shows the 30 LGAs, consisting of 93 wards, in which the Liberal Party contested the most recent local election before 2024. This area covers much of urban NSW, and accounts for approximately 50% of the state's population.

2.3 Failure to register for 2024 elections

In the lead up to the local government elections to be held on 14 September 2024, the Liberal Party failed to register a large number of its candidates with the NSW Electoral Commission by the statutory deadline at midday on 14 August 2024. After investigation, it was found to be the fault of state-level party officials who were tasked with undertaking the registration process on behalf of their local government candidates. This process was historically organised at the local level, but had been newly centralised for the 2024 elections. A large scale registration error like this had never occurred previously, and so was not expected.

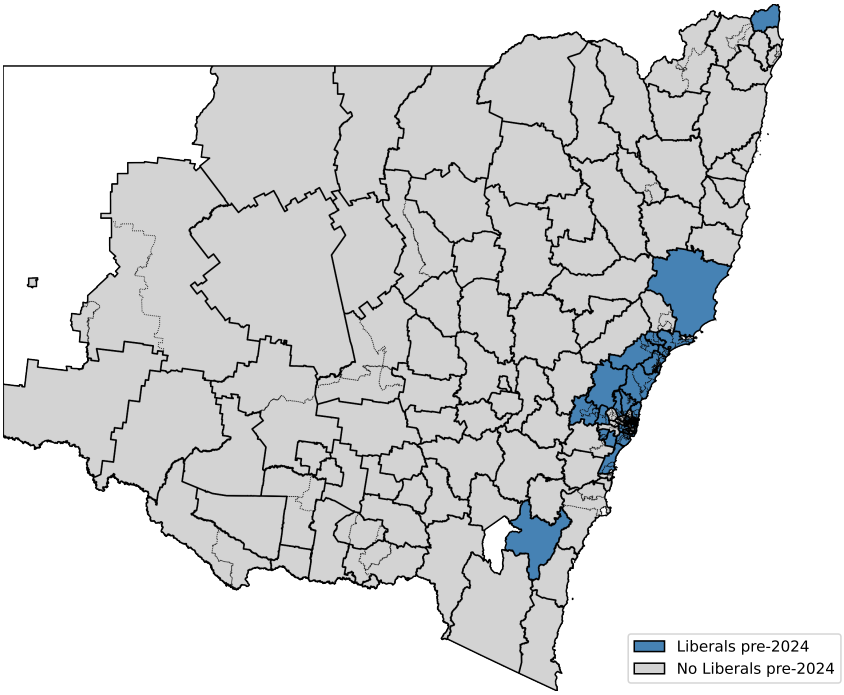
Accounts revealed that the party officials had greatly underestimated the resources required to centrally register the candidates, and had realised too late to request help or additional resources.[11] On the morning of the midday deadline, party officials were still endorsing new prospective candidates (the first step in the registration process).[12] In these final hours, they were also providing assurances to candidates that the paperwork would be completed in time, despite the fact that candidates could alternatively register themselves.[13] Some candidates did not know until the following day, when the official ballot lists were revealed, whether their paperwork had been submitted by the deadline. The NSW Liberal Party director Richard Shields, who was in charge of the registration process, was removed from his director position.

As a result of the registration failure, 138 Liberal candidates that were expected to participate in their respective electoral races were left off the ballot, accounting for about one third of all planned Liberal Party candidates. Among the 138 missing candidates were 39 non-retiring

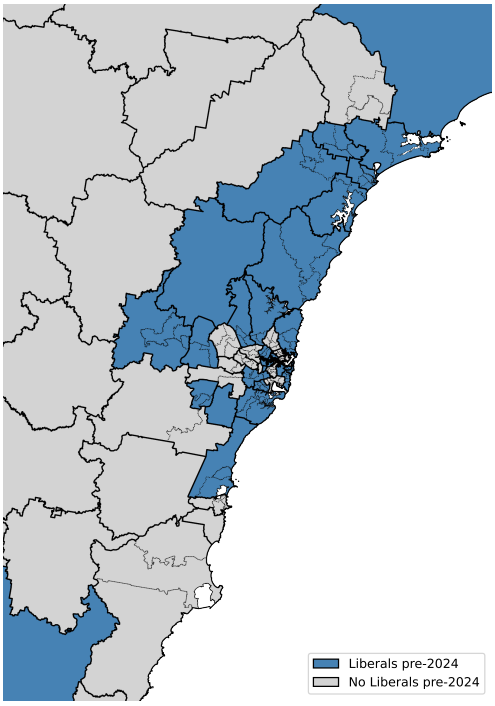
[11] The party official responsible for the registration process stated "I was not made aware that the deadline would not be met until it was too late to rectify the deficiencies." https://www.abc.net.au/news/2024-08-15/nsw-liberal-director-richard-shields-sacked/104231340, accessed 2025-07-27.

[12] https://www.smh.com.au/politics/nsw/a-monumental-stuff-up-liberal-leaders-call-for-director-s-head-20240815-p5k2kc.html, accessed 2025-07-27.

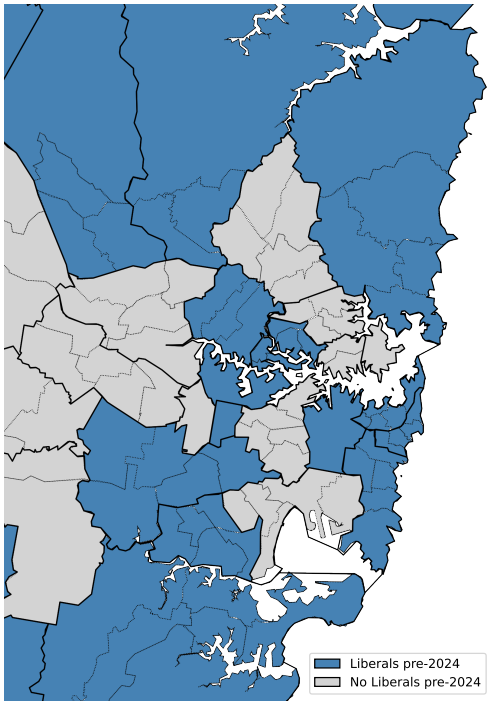
[13] https://www.abc.net.au/listen/programs/radionational-drive/nsw-liberals-local-election-debacle-/104231334, accessed 2026-02-13.



(A) NSW



(B) Coastal NSW



(C) Coastal Sydney

Figure 1: Liberal Party Contests pre-2024

Notes: Panel A displays a map of NSW, with LGA boundaries marked with solid lines and ward boundaries with dotted lines, as of 2024. LGAs which did not have a Liberal Party candidate in the most recent election prior to 2024 are shaded in light grey (the western unincorporated areas and Australian Capital Territory in white are excluded). LGAs which did have a Liberal Party candidate in the most recent pre-2024 election are shaded blue. **Panel B** focuses on coastal NSW which covers the major urban areas of Sydney, Newcastle, Wollongong and the Central Coast. **Panel C** further focuses on the central and coastal areas of Sydney.

incumbents. The issue impacted a variety of LGAs including those considered "conservative strongholds," while some candidates were successfully registered in wards where the Liberal Party was reportedly unlikely to win a seat; within some LGAs the Liberal Party were only able to register candidates for their least-promising ward.[14] The registration failure affected 40 wards in 17 LGAs, with 37 wards having no Liberal candidates at all.[15] Four mayoral races were also affected. Figure 2 shows the wards in which the Liberal Party were unable to register any of their intended council candidates.

These accounts of the error suggest that it is an ideal situation to study as a quasi-exogenous shock to the pool of candidates. First, various elements of the reported accounts indicate that prioritisation in the registration process was very unlikely. If there was insufficient time to ask for help or even to delegate submission back to candidates, actions which would have increased the number of successfully registered candidates, it is unlikely that there was opportunity to assess the prospects of candidates and selectively complete the paperwork for stronger candidates. The large number of incumbents that planned to re-contest but were not registered, the number of favourable contests that were affected and unfavourable contests that were not affected all reinforce the conclusion that prioritisation did not occur. Second, the shock affected multiple contests at once, allowing for the use of a difference-in-differences design, relying on parallel trends for causal interpretation. Third, the shock affected a single, major party allowing for clean isolation of the impact of major party absence, being one side of the main partisan divide in Australian politics. Lastly, the pool of candidates was fixed at the deadline, with no ability for new candidates to enter the contests, meaning that the level of electoral competition did not endogenously respond to the shock.

3 Data

Information on which wards were affected by the registration failure is taken from election analyst Ben Raue, who provides the only comprehensive list of affected wards.[16] The list provides information about the number of successfully registered Liberal candidates compared to the intended number for affected contests.

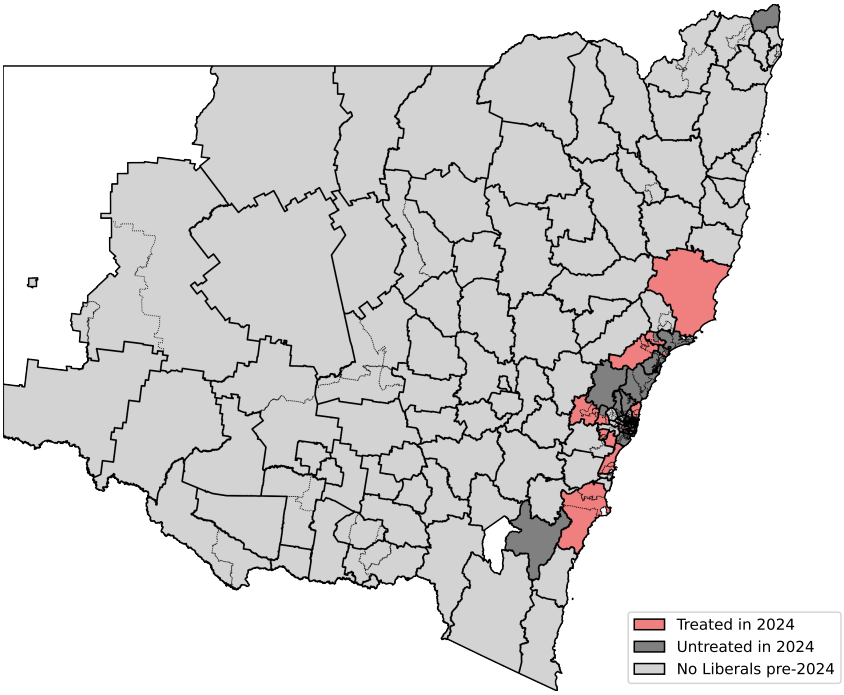
Election outcomes for local elections from 2012–2024 are collected from the NSW Electoral Commission (NSWEC) website.[17] The data provides the complete ranking of candidates or

[14] https://www.abc.net.au/news/2024-08-15/nsw-liberal-party-shields-resign-council-election-nominations/104227072, accessed 2025-07-27.

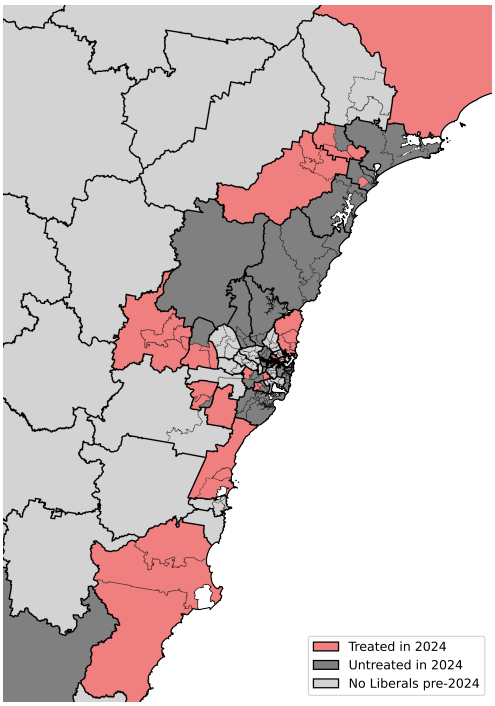
[15] The remaining three affected wards were able to register some, but not all, of their intended candidates. These wards are excluded from the analysis, since they are neither treated in the sense of having no Liberal candidates, nor are truly control, since the partially missing Liberal list means that they could not qualify for a group ballot option — voters were required to vote for these Liberal candidates individually.

[16] The list is available on his blog at https://www.tallyroom.com.au/56329, accessed 2025-06-23.

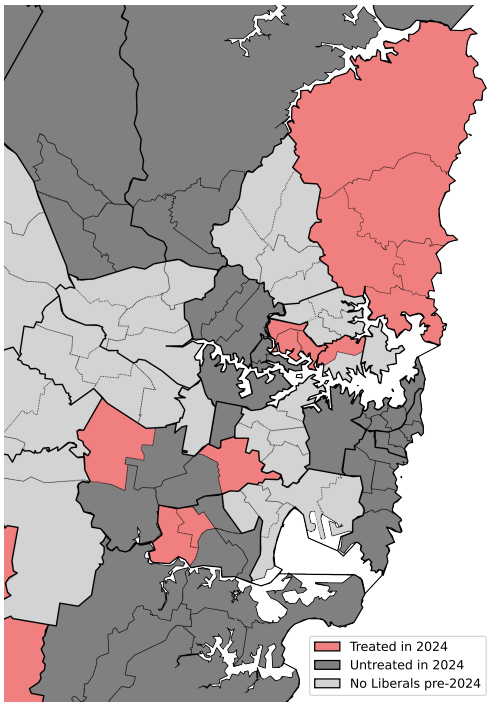
[17] https://elections.nsw.gov.au/.



(A) NSW



(B) Coastal NSW



(C) Coastal Sydney

Figure 2: Missing Liberals in 2024 (Treatment)

Notes: Panel A displays a map of NSW, with LGA boundaries marked with solid lines and ward boundaries with dotted lines, as of 2024. Wards with no Liberal candidate in 2024 due to the registration failure are shaded in red. LGAs which did not have a Liberal Party candidate in the most recent election prior to 2024 are shaded in light grey (the western unincorporated areas and Australian Capital Territory in white are excluded). The remaining wards in LGAs which did have a Liberal Party candidate in the most recent pre-2024 election are shaded dark grey. **Panel B** focuses on coastal NSW which covers the major urban areas of Sydney, Newcastle, Wollongong and the Central Coast. **Panel C** further focuses on the central and coastal areas of Sydney.

candidate groups recorded on each ballot paper, the number of invalid ballots at polling place level including the markings on each which lead to the spoilage, and the number of enrolled voters at ward level. The ballot-level data might show, for example, that a voter at polling place X ranked Party A candidates first, Party B candidates second, and did not rank Party C candidates at all.

I also scrape polling place locations and ward boundaries for local elections from the NSWEC website. Since these geographies shift over time, I construct population-based geographic crosswalks in the spirit of Ferrara et al. (2024) to harmonise outcomes to 2024 geographies. For polling place level outcomes, I first use locations of polling places to construct Voronoi partitions of wards by polling places for each election. Further information about the process is provided in Appendix C.

To generate these population weights, I take exact granular population counts at the statistical mesh block level from the 2021 Australian Census (Australian Bureau of Statistics, 2021). NSW contains 112,000 mesh blocks with an average of 71.6 people in each. I also collect information on demographic and economic variables at the SA1 level (SA1 areas contain about 400 people on average), which I harmonise to 2024 electoral geographies.

To control for area-specific changes in political sentiment over time, specifically sentiment towards the Liberal Party, I collect data on state election outcomes from the NSW Electoral Commission and similarly harmonise to 2024 polling places.

To analyse the financial response of political candidates, I again scrape the NSWEC website for mandatory financial disclosure reports. Candidates, parties and groups of candidates, elected councillors, and associated entities must disclose all electoral expenditure incurred, on an annual basis. The data is itemised and dated, with expenditure categorised among advertising, production and distribution of election material, staff costs, office costs, travel, research, and fundraising. I match expenditure records to ward-level ballot groups and parties (or individuals for ungrouped candidates) by name and geography, successfully matching almost 90% of candidates, resulting in about 95% of all candidates belonging to a ballot group or individual ticket with at least one matched expenditure record. I focus on expenditure in the 31 days prior to the election (for 2024, this coincides with the period after the registration failure).[18] I express all expenditure in 2024 Australian dollar terms.

To investigate the effect of this shock on campaign rhetoric and the broader political debate, I collect data on online newspaper articles. I take as my universe of articles all ".au" domain articles from the GDELT Global Knowledge Graph (GKG) that mention a relevant candidate by name in the 60 days prior to a local election.[19] Since GDELT coverage begins in 2013,

[18] Most campaign spending occurs in this period, see Appendix Figure 7.

[19] Additionally, I manually exclude a number of out-of-state newspapers that do not plausibly serve local

this does not cover the 2012 local elections.[20] Full article text is recovered from the article's URL where possible, else the Internet Archive's Wayback Machine. This process successfully recovered full text for more than 90% of targeted URLs, resulting in a corpus of 9,357 articles.

4 Empirical Strategy

I conduct the primary electoral analysis at the 2024 ward level, weighting observations by population.[21] I use a difference-in-differences (DiD) design, comparing wards affected by the registration failure in 2024 (i.e., those who were expecting Liberal Party candidates but due to the failure did not have any) with those that did have at least one Liberal Party candidate, before and after the incident. I do so by estimating the two-way fixed effects model given by:

$$y_{wt} = \alpha_w + \alpha_t + \beta D_{wt} + u_{wt}, \tag{1}$$

where y_{wt} is the outcome of interest in ward w at election t,[22] α_w are the ward fixed effects, α_t are the election fixed effects, and D_{wt} is the treatment dummy equal to 1 for affected wards for $t = 2024$ and 0 otherwise. I cluster standard errors at the ward level, which is the level of treatment assignment since candidates compete at the ward level.[23] Since treatment occurred for all treated units simultaneously and in a single period, no concerns regarding negative weights of the kind contemplated by de Chaisemartin and D'Haultfœuille (2020) or Goodman-Bacon (2021) arise. The identifying assumption for this DiD design is parallel trends: that areas affected by the registration failure (where Liberal candidates were expected but could not run) would have evolved similarly to other unaffected areas (where Liberals did run) in the absence of the failure.

I limit the sample to LGAs in which the Liberals contested the most recent election before 2024, and for the pre-treatment period I include only ward-year contests where the Liberals had a candidate on the ballot. Excluding contests in which the Liberals did not compete is essential for maintaining the plausibility of parallel trends. Additionally, I exclude partly treated wards wherein the failure led to only a subset of intended candidates being registered — as such, these wards are neither unaffected (because at least one candidate is missing) nor truly treated (since

news in NSW. Candidate name mentions are based on GDELT's extracted person names, or from full text for articles that GDELT classifies as having a local government election-related theme (e.g., 'ELECTION' or 'GOV_LOCALGOV').

[20] Coverage in 2021 is also poor, consistent with changes in the news landscape in Australia (consolidation of local papers) and issues with the GDELT crawler.

[21] Thus, treatment effects can be interpreted as if at the voter level. In the appendix, I show that the main results are robust to analysis at the polling place level, and persist without population weights.

[22] Note that local elections took place in 2016 for some LGAs and 2017 for others, with no overlap. Thus, I pool the 2016 and 2017 elections together as one cohort of elections.

[23] In the appendix, I show that the results are substantially robust to allowing for broader correlations at the LGA level.

the Liberals do still compete there). The final estimation sample includes 31 treated wards and 53 control wards in 2024.

One concern with this design is that, despite anecdotal evidence to the contrary, treated wards may differ systematically from control wards. To investigate this concern, I analyse the balance of treatment assignment on a range of political, social and economic characteristics on the pre-treatment (2021) cross-section. Table 1 provides the associated results. Out of 22 characteristics, none display statistically significant differences at the 5% significance level. Three are statistically significant at the 10% level, which is not unexpected among 22 variables. Indeed, a permutation test for balance yields a p-value of 0.437, suggesting there is insufficient evidence to conclude the treatment is unbalanced.

The minor imbalances may nonetheless raise some concern for identification. Additionally, the power of these tests is limited by the relatively small sample in this cross-section, meaning that we cannot exclude sizeable imbalances on the available evidence. Even among variables with no statistically significant imbalances, the 95% confidence intervals can include differences slightly less than 1 standard deviation. The ward fixed effect term α_w in (1) captures any time-invariant ward-level political, social and economic characteristics, including general favourability towards the Liberal Party, which might correlate with both political outcomes and the treatment assignment. Therefore these potential imbalances do not directly threaten identification.

As identification relies on parallel trends, the potential identification concern comes if these imbalances predict *changes* in the political outcome variables. Motivated by this concern, I account for the possibility that treated wards may differ systematically from control wards in terms of their trends, as a robustness check. The first specification that addresses these concerns controls for the share of votes cast for the Liberal-National Coalition in the most recent state election in the Legislative Assembly (lower house), harmonised to 2024 ward geographies. I estimate the following equation:

$$y_{wt} = \alpha_w + \alpha_t + \beta D_{wt} + \gamma \text{State VS}_{wt} + u_{wt}, \tag{2}$$

where State VS_{wt} is the state Legislative Assembly election vote share for the Liberal-National Coalition,[24] in ward w for the most recent state election before local election year t. The vote share for the Coalition at the state election captures general local sentiment towards the Liberal Party, and related factors that may change over time. The benefits of relying on state election results are (1) they are available at the same frequency as local elections (every 4 years), (2) the timing of these elections is such that the previous state election closely

[24] I control for the total vote share for Liberals plus Nationals because (1) voters can express their support for a Liberal-National government by voting for the Nationals, and (2) since Liberals and Nationals often do not compete in the same electorates, this provides greater coverage than controlling for the Liberal vote share alone.

Table 1: Balance of Treatment

	Full Sample			Control	Treated	|Δ|	p-value
Variable	mean	sd	N	mean	mean	|Control − Treatment|	Control=Treatment
Population	46,087.920	33,767.227	91	47,480.815	43,519.769	3,961.046	0.596
				(4,419.209)	(5,979.431)	(7,435.254)	
Age	38.879	3.185	89	38.950	38.749	0.201	0.843
				(0.476)	(0.890)	(1.009)	
Proportion Female	0.508	0.010	89	0.508	0.509	0.000	0.906
				(0.003)	(0.001)	(0.003)	
Proportion Born Overseas	0.284	0.145	89	0.297	0.261	0.036	0.321
				(0.027)	(0.025)	(0.036)	
Proportion Finished High School	0.488	0.128	89	0.509	0.450	0.059	0.077
				(0.026)	(0.020)	(0.033)	
Proportion Secular	0.361	0.092	89	0.369	0.346	0.023	0.432
				(0.017)	(0.023)	(0.029)	
Proportion Married	0.439	0.037	89	0.440	0.437	0.004	0.719
				(0.007)	(0.007)	(0.010)	
Household Size	2.660	0.315	89	2.625	2.724	0.099	0.277
				(0.066)	(0.062)	(0.091)	
Proportion Renting	0.119	0.056	89	0.127	0.105	0.022	0.176
				(0.015)	(0.005)	(0.016)	
Weekly Rental Cost	432.092	96.315	89	444.237	409.699	34.538	0.187
				(13.729)	(22.074)	(25.996)	
Proportion of Outright Home Owners	0.113	0.031	89	0.112	0.115	0.003	0.747
				(0.005)	(0.009)	(0.010)	
Monthly Mortgage Cost	1,685.347	225.148	89	1,705.634	1,647.938	57.696	0.317
				(28.729)	(49.647)	(57.360)	
Unemployment Rate	0.047	0.013	89	0.046	0.048	0.002	0.646
				(0.002)	(0.003)	(0.004)	
Individual Income	866.818	205.984	89	900.833	804.094	96.739	0.083
				(37.802)	(40.303)	(55.257)	
Liberal Vote Share	0.252	0.107	85	0.268	0.224	0.044	0.079
				(0.020)	(0.014)	(0.025)	
Labor Vote Share	0.251	0.110	79	0.238	0.275	0.037	0.267
				(0.019)	(0.026)	(0.033)	
Greens Vote Share	0.095	0.066	71	0.094	0.097	0.003	0.868
				(0.010)	(0.016)	(0.019)	
Other (excl. Greens) Vote Share	0.222	0.131	82	0.211	0.243	0.032	0.424
				(0.023)	(0.033)	(0.040)	
Invalid Ballot Share	0.042	0.014	85	0.041	0.043	0.002	0.702
				(0.003)	(0.003)	(0.004)	
Turnout	0.836	0.049	85	0.830	0.847	0.017	0.217
				(0.013)	(0.004)	(0.014)	
State Liberal Vote Share (2023)	0.306	0.142	84	0.327	0.269	0.058	0.120
				(0.025)	(0.028)	(0.037)	
State Coalition Vote Share (2023)	0.325	0.125	86	0.340	0.299	0.041	0.224
				(0.022)	(0.026)	(0.034)	
Joint permutation test							0.437

Notes: The table presents balance on covariates between treatment and control wards, and descriptive statistics for the full estimation sample at the ward level. For each variable I report the mean, standard deviation and number of observations, and the mean for the treatment and control groups with standard errors in parentheses. I also regress each variable on a pair of dummies equal to one if the ward was in the control or treatment group, without a constant term, and report the p-value of an F-test for equality between the two coefficients, alongside the absolute difference in the means. For all covariate-level F-tests, p-values were computed using standard errors clustered at the ward level. I report the following variables: (1) population count, (2) average age, (3) proportion of female residents, (4) proportion of residents born overseas, (5) proportion of residents that have completed year 12 of school or equivalent, (6) proportion of residents that identify as secular/non-religious, (7) proportion of residents that are married, (8) average household size, (9) proportion of residents renting their housing, (10) average weekly cost of renting among renters, (11) proportion of residents that own their housing outright, (12) average monthly mortgage among mortgage holders, (13) proportion of labour force participants seeking work, (14) average weekly income per resident, (15) first preference votes for Liberal Party at 2021 local election, as a share of enrolled voters, (16) first preference votes for Australian Labor Party at 2021 local election, as a share of enrolled voters, (17) first preference votes for Australian Greens at 2021 local election, as a share of enrolled voters, (18) first preference votes for all other candidates at 2021 local election, as a share of enrolled voters, (19) invalid ballots cast at 2021 local election, as a share of enrolled voters, (20) turnout, (21) share of valid first preference votes for Liberal Party in the Legislative Assembly at 2023 state election, (22) share of valid first preference votes for Coalition (Liberal Party or Nationals) in the Legislative Assembly at 2023 state election. In the final row, I report the p-value of a joint permutation test of balance, by reassigning treatment at random across wards 999 times and comparing the sum of the squared t-statistics of the rows above with its permutation distribution. In all cases except for population, observations were weighted by population. All variables are as of 2021, except where otherwise indicated. Standard errors are clustered at the ward level.

precedes the corresponding local elections by one or two years, and (3) the data is available at similarly granular geographic levels (polling places), allowing for meaningful harmonisation to local election geographies. Under this specification, threats to the identifying assumption would need to come in the form of differential trends in political factors that are not captured by the Liberal Party's local success at the state elections, or which developed between the 2023 state election and the 2024 local elections; in other words, conditional parallel trends is required.

In addition to conditional parallel trends, Sant'Anna and Zhao (2020) show that this specification identifies the ATT only if treatment effects are homogeneous across strata of changes in State VS. Otherwise, it overweights observations with uncommon changes in State VS. Therefore, I additionally use the doubly robust DiD estimator which instead relies on correct specification of either the control-group outcome evolution or treatment propensity, not homogeneity, for consistency.

As an alternative, I also provide estimates based on double machine-learning in a partially linear model (Chernozhukov et al., 2018). Here, I include all baseline variables contained in the balance table,[25] with time-invariant values interacted with time fixed effects.[26] That is, I allow for changes over time to differ conditional on baseline characteristics. The relative benefit of double machine-learning is the ability to flexibly control for a larger set of covariates. It is also not limited to one pre-period, as doubly robust DiD is.

Additionally, I present results from unconditional event study analyses based on estimating the following equation:

$$y_{wt} = \alpha_w + \alpha_t + \sum_{k \in \mathcal{K} \setminus \{2021\}} \beta_k \times \mathbb{1}_{t=k} \times \text{Ever-treated}_w + u_{wt}, \tag{3}$$

for $\mathcal{K} = \{2012, 2016/2017, 2021, 2024\}$, and where $\text{Ever-treated}_w = 1$ in all years for wards that are treated eventually in 2024, 0 otherwise. Due to the limited data available online from the NSW Electoral Commission, the pre-treatment period consists of relatively few local elections. Still, the event study analysis may be somewhat informative regarding underlying trends in the outcome variables in the period before treatment, which could otherwise cast doubt on the parallel trends assumption.

4.1 Outcomes

The main analysis leverages the richness of preference ranking data available in NSW local elections. The results use three sets of dependent variables. The first set is the *actual first preference vote shares* for the Labor Party, and for all remaining candidates combined as 'Other'

[25] With missing vote share values zero filled and accompanied with a missing indicator variable.
[26] For lasso and ridge learners, I additionally include the second-order polynomial expansion of these variables.

(the Greens, smaller and local parties, and independents), as a share of enrolled voters.[27] That is, the share of enrolled voters who rank candidates from these respective groups as number 1 on their ballot paper. For each of these variables, I treat values as missing in contests where they did not run.[28] In this set, I also include invalid votes as a share of enrolled voters,[29] turnout, and first preference votes for the Liberal Party as a share of enrolled voters, which is mechanically zero for treated wards in 2024. Therefore, the absolute value of the estimated effect on the Liberal vote share is the estimated share of orphaned Liberal voters in treated wards.

The second set of dependent variables is the *expected transferred vote shares*, constructed from the preference ranking data. I construct these variables by taking all ballot papers which rank Liberal as their first preference (in all ward-years) and observing which group, if any, is ranked next. Some of these ballots rank no candidate outside the Liberal Party and so exhaust. Since voting is compulsory, most of these voters are still expected to turn out when no Liberal stands, so dropping these votes would understate transfers.[30] Therefore, I multiply these counts by a scale factor so that the combined number of implied second preference votes equals the number of actual Liberal first preference votes per ward, which assumes that exhausted ballots would transfer in the same ratios as non-exhausted ballots. I then divide by the number of enrolled voters to express them as shares. These values reflect the expected increase in vote share of each group if Liberals were removed from rankings — i.e., they provide an account of substitution flows from Liberal voters. Again, this is mechanically zero in treated wards in 2024, meaning the corresponding treatment effect is (in absolute value) the expected vote share to be transferred to these groups in treated wards under parallel trends. There is no equivalent of this variable for invalid votes, Liberal vote share or turnout. Among voters who ranked a Liberal candidate first in treated wards in 2021, and who also ranked at least one non-Liberal candidate, 38% ranked Labor next with the remainder going to Other.

The third set of dependent variables is the *implied vote shares* (if Liberals are removed from all rankings), equal to the sum of the actual vote shares and the expected transfer of votes. For control and pre-treatment observations, these are the vote shares the two groups would have received had the Liberals been absent and rankings unchanged; for treated wards in 2024,

[27] This ensures the vote shares do not vary mechanically with any turnout effect.

[28] As a consequence, the sum of changes to raw vote shares does not equal the estimated share of orphaned Liberal votes. This identity can be returned by treating missing values as zero, or by restricting the sample to the set of ward-years for which there is full coverage.

[29] Ballots may be invalid either on purpose or due to a mistake. For instance, ranking two different candidates as preference 1, or leaving the ballot paper empty.

[30] Empirically, only about 5% of orphaned Liberal voters (1.1 pp of enrolled voters) did not turn out, though the estimate is not robust across specifications. In comparison, ballot exhaustion would have affected about three-quarters of Liberal ballots, so dropping exhausted ballots would remove most Liberal voters from the benchmark.

they are simply the actual vote shares. In this sense, the control units provide a Liberal-free comparison that already accounts for the mechanical reallocation of Liberal first preferences to the next-ranked alternative. Under parallel trends, the DiD counterfactual for treated wards built from this variable is the *fixed-rankings benchmark* — the vote shares expected in treated wards in the absence of the Liberals had rankings not changed.

Together, these three sets of variables have a useful property, specifically in terms of Labor and Other vote shares (for which all three exist). Estimates from specifications with the actual raw vote shares provide the actual change in votes from Liberal Party absence, combining both the mechanical substitution effect under fixed rankings with any changes from the rankings themselves. Estimates from specifications with the expected transferred votes give the additional share of votes these groups would be expected to receive under fixed rankings, given trends in control wards.[31] Specifications with the implied vote shares provide an estimate of how much actual votes deviate from expected votes, mechanically the difference of the two prior estimates. In effect, this decomposes the actual change in vote shares into a part expected under fixed rankings and an *excess*, whether a surplus or a shortfall, that fixed rankings cannot explain. Under the null that rankings are unaffected by the absence of the Liberals, the excess is zero; a non-zero excess means that rankings themselves changed.

5 Results

In this section, I study the impact of Liberal Party absence on first preference vote shares. In Section 5.1 I decompose the change in actual vote shares into what is expected from mechanical substitution and what is not. Section 5.2 investigates the robustness of the main results to alternative specifications, and assesses a number of identification concerns. I decompose invalid votes by reason of invalidity in Section 5.3. Lastly, I discuss what the estimates imply about Liberal voters' loyalty in Section 5.4.

5.1 Main Results

The main results are presented in Table 2. Panel A provides the estimated effect of Liberal Party absence on actual vote shares (including mechanically on Liberal vote share), giving the combined effect of substitution with fixed rankings and changes in rankings for the Labor and Other vote shares. Panel B provides the estimated effect on the magnitude of Liberal second preference flows, giving the expected change in vote shares under substitution with fixed rankings. Panel C provides the estimated effect on the implied vote shares, isolating the excess change in vote shares that is not explained by substitution with fixed rankings. This

[31] By construction, this value is negative — I report the absolute magnitude of effects for ease of interpretation.

decomposition, and the effects on the invalid vote share and turnout, are additionally presented in Figure 3.

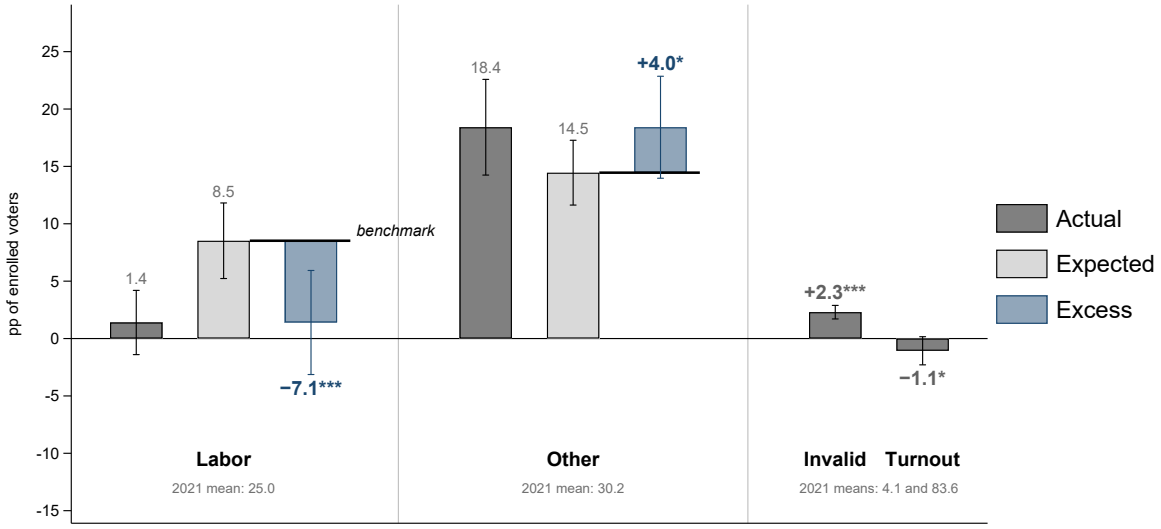


FIGURE 3: MAIN RESULTS

Notes: The figure presents the estimates of the effect of Liberal Party absence from estimating Equation 1, as reported in Table 2, and the associated 95% confidence intervals. Actual refers to estimates from regressions using the actual first preference vote shares as the outcome variables. Expected refers to (the absolute value of) estimates from regressions using the expected transferred vote shares as outcome variables. Excess refers to estimates from regressions using the implied vote shares (in the absence of Liberals, holding rankings fixed) as outcome variables. The excess is positioned relative to the fixed-rankings benchmark. These are reported for the Labor vote share, the Other vote share, the invalid vote share (where possible) and turnout (where possible). Pre-treatment means in 2021 are also provided. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. Confidence intervals are based on standard errors clustered at the ward level. * significant at 10%; ** significant at 5%; *** significant at 1%.

In Column (1), I report that the benchmark predicted the Labor vote share to increase by 8.5 pp in the absence of the Liberal Party, or about 38% of the orphaned Liberal vote reported in Column (5). The Labor vote share actually increased by 1.4 pp, capturing only about 6% of orphaned Liberal votes. Consequently, Labor experienced a large and statistically significant decrease in vote share of 7.1 pp (p=0.003) relative to the fixed-rankings benchmark. Given this represents 7.1% of the entire electorate of enrolled voters (including the 15–20% who did not turn out), this is extremely meaningful in size.

To further understand the magnitude of this shortfall, we can compare to the approximate share of votes required for a single seat on the council. In 2024, the average number of seats available to be filled per ward was 4.71. Average turnout was 82.94%, so a typical quota (the number of votes required to be elected) was equivalent to 14.53% of enrolled voters.[32] Therefore, the shortfall relative to the fixed-rankings transfer benchmark was roughly equivalent to 0.49 council seats — almost half of an elected councillor. Given that councils contain on average 12.6 members, this represents a potentially substantial impact on who controls local government. Across all affected contests in which Labor competed, the shortfall was equivalent to 14.2

[32] Quota share $\approx$ Turnout $\times \frac{1}{\text{N positions}+1}$.

Table 2: Decomposition of Effect on Vote Shares

	(1) Labor Share	(2) Other Share	(3) Invalid Share	(4) Turnout	(5) Liberal Share
Panel A: Actual change					
Treatment	0.0140	0.184***	0.0230***	-0.0107*	-0.227***
	(0.0141)	(0.0210)	(0.00297)	(0.00617)	(0.0153)
p-value	0.3230	< 0.0001	< 0.0001	0.0873	< 0.0001
Wild-bootstrap p-value	0.3318	< 0.0001	< 0.0001	0.0909	< 0.0001
Mean dep. var. (2021)	0.2497	0.3018	0.0414	0.8360	0.2540
SD dep. var. (2021)	0.1094	0.1255	0.0140	0.0488	0.1055
Panel B: Expected transferred votes					
Treatment	0.0852***	0.145***			
	(0.0166)	(0.0142)			
p-value	< 0.0001	< 0.0001			
Wild-bootstrap p-value	< 0.0001	< 0.0001			
Mean dep. var. (2021)	0.0920	0.1662			
SD dep. var. (2021)	0.0755	0.0889			
Panel C: Excess					
Treatment	-0.0712***	0.0396*			
	(0.0228)	(0.0224)			
p-value	0.0025	0.0800			
Wild-bootstrap p-value	0.0038	0.0831			
Mean dep. var. (2021)	0.3416	0.4680			
SD dep. var. (2021)	0.1522	0.1567			
Observations	277	299	301	301	301
N ward clusters	85	88	88	88	88

Notes: The table reports regression estimates of the effect of Liberal Party absence from estimating Equation (1). **Panel A** provides the estimates using the actual first preference vote shares as the outcome variables. Column (1) uses the actual first preference vote share for Labor as the outcome variable. Column (2) uses the actual first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. Column (3) uses the share of invalid votes as the outcome variable. Column (4) uses the share of voters that turned out as the outcome variable. Column (5) uses the actual first preference vote share for Liberals as the outcome variable. **Panel B** uses the same specification as Panel A but with the expected transferred vote shares as the outcome variables (only available for Labor and Other). The estimate reported is the absolute value of the estimate from Equation (1). **Panel C** uses the same specification as Panel A but with the implied vote shares as the outcome variables (only available for Labor and Other). All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

council seats.

The remaining candidates outperformed the benchmark: these candidates were expected to gain 14.5 pp of vote share, but gained 18.4 pp, leading to a surplus of 4.0 pp. This surplus is marginally significant (p = 0.080). This indicates that the increase in actual votes for Other candidates was only marginally distinguishable from what was expected by substitution with fixed rankings.

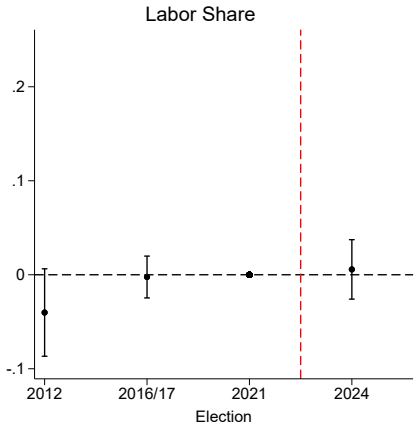
I also report that the share of invalid votes increased by 2.3 pp, which is statistically significant at the 1% level (p<0.001). This is extremely large, representing a 56% increase over the pre-treatment mean, or 1.64 standard deviations. This result is consistent with a protest vote, wherein some voters may seek to express dissatisfaction with the available pool of candidates. It could also be consistent with an increase in accidental errors arising from voters ranking candidates (as opposed to parties) at a higher rate, which has more stringent requirements on how many options need to be ranked. In Section 5.3, I provide support for the protest vote interpretation by breaking down the reasons for votes being invalid.

Turnout also marginally decreases by 1.1 pp, suggesting a mild demobilisation effect. However, this result is only statistically significant at the 10% level, and as Section 5.2 will expand on, it is sensitive to the precise specification used. It is not surprising that the effect on turnout is small and statistically weak, given Australia's compulsory voting laws.

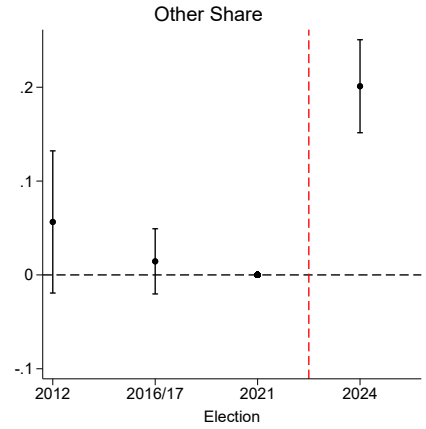
The results of the event studies for the raw electoral outcomes are presented in Figure 4. There is no obvious evidence of systematic pre-trends that might undermine confidence in the identification assumption. In particular, it is reassuring that the Liberal vote share does not display any evidence of pre-trends since this would have been most likely to drive any prioritisation if the state party officials were focused on maximising Liberal Party wins in a restricted set of wards. Figure 5 presents the event studies estimating the excess vote share, which also do not arouse any specific concerns. Still, the limited number of pre-treatment periods (3), their low frequency (roughly every 4 years) and wide confidence intervals mean that any conclusions are necessarily limited.

5.2 Robustness

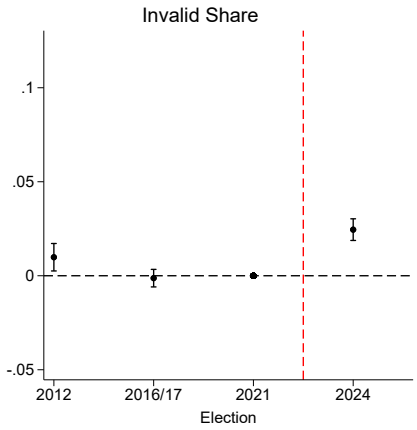
In this section, I assess the robustness of the main results, focusing on the estimated excess vote share for Labor and Other candidates, and the actual change in the invalid vote share and turnout.



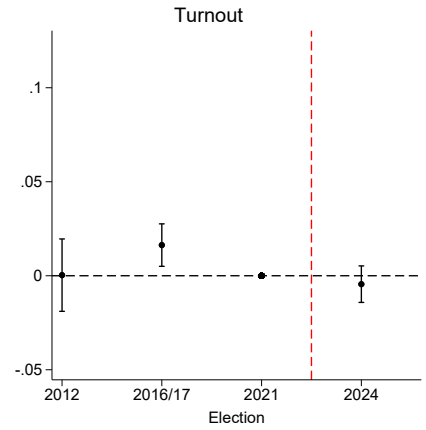
(A) Labor (Actual Vote Share)



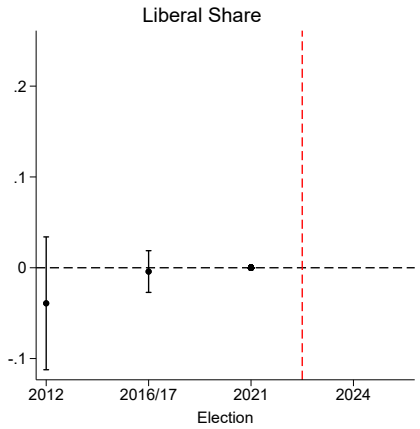
(B) Other (Actual Vote Share)



(C) Invalid Vote Share



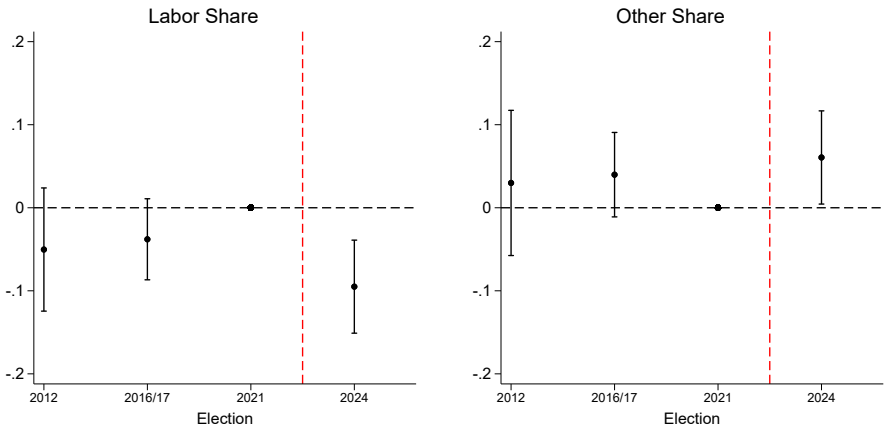
(D) Turnout



(E) Liberal (Actual Vote Share)

Figure 4: Event Studies (Raw electoral outcomes)

Notes: The figure presents the event study estimates based on Equation 3 and the associated 95% confidence intervals, with the following dependent variables: (1) actual first preference vote share for Labor, (2) actual first preference vote share for other candidates, excluding Labor and Liberal, (3) invalid vote share, (4) turnout, and (5) actual first preference vote share for the Liberal Party. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. The 2021 election is the omitted reference period. The red dashed line indicates when treatment took place. Confidence intervals were calculated based on standard errors clustered at the ward level.



(A) Labor (Excess) (B) Other (Excess)

Figure 5: Event Studies (Implied vote shares)

Notes: The figure presents the event study estimates based on Equation 3 and the associated 95% confidence intervals, with the following dependent variables: (1) implied first preference vote share for Labor, and (2) implied first preference vote share for other candidates, excluding Labor and Liberal. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. The 2021 election is the omitted reference period. The red dashed line indicates when treatment took place. Confidence intervals were calculated based on standard errors clustered at the ward level.

5.2.1 Clustering Approach

First, I consider the appropriateness of the cluster level. While p-values are generally valid for more than 50 clusters, Carter et al. (2017) suggest that cluster heterogeneity can lead to problems similar to insufficient cluster units. Structurally, some LGAs are not split into wards, while others may be split into 4 or 5, drastically affecting the ward's population. Reassuringly, Table 2 reports the wild cluster bootstrapped p-values which perform well so long as the number of treated units is not small, as in these circumstances (Cameron et al., 2008; MacKinnon and Webb, 2016), and do not lead to any different conclusions on statistical significance at the 1%, 5% and 10% levels.

It might be possible that errors are correlated within LGAs, which is the actual political unit of local government. It is significantly less likely for errors to correlate across LGA bounds, and allowing for spatial correlations would largely coincide with LGA clusters without having a clear theoretical motivation. To address these broader concerns, I report the main specifications with LGA level clusters in Table 6 in the Appendix. There, the number of clusters is substantially smaller (30), so wild cluster bootstrapped p-values are preferred.

I report that the primary finding, that Labor did worse than the fixed-rankings benchmark would have predicted, remains statistically significant at the 5% significance level (bootstrapped $p = 0.043$). The excess votes for Other candidates (compared to the benchmark) is not statistically significant with a bootstrapped p-value of 0.223, reinforcing the conclusion that the fixed-rankings expectations largely accounts for the increase in raw votes for this group of

candidates. The effect on turnout similarly loses statistical significance at conventional levels (bootstrapped p=0.150), in line with the muted (de)mobilisation effect under compulsory voting. The rise in invalid votes remains significant at the 1% significance level (bootstrapped p<0.001).

5.2.2 Specifications

Next, I consider the potential concern that unconditional parallel trends does not hold. The event studies presented above in Section 5.1 do not generate this concern, but due to the noise of the estimates, cannot convincingly exclude the possibility of pre-existing differential trends. In Appendix Table 8, I provide estimates for the series of specifications mentioned in Section 4. I report that across specifications which control for state level Liberal Coalition vote shares, doubly robust DiD controlling for state vote shares, and double machine-learning specifications which flexibly control for differential changes related to all covariates listed in the balance table,[33] the negative excess of Labor votes relative to the fixed rankings benchmark remains statistically significant at the 5% significance level. The estimated excesses are relatively consistent in magnitude, ranging from −5.4 pp to −8.3 pp (unconditional OLS = −7.1 pp). The excess for Other candidates has p-values ranging from 0.066 to 0.125, meaning it is either marginally significant or marginally insignificant at the 10% significance level. The invalid vote share effect remains significant at the 1% level across all specifications, with the point estimate varying by no more than 0.5 pp. The turnout effect remains small and negative, but inconsistently significant. Thus, neither of the clear results in the unconditional specifications (Labor excess and increased invalid vote share) is sensitive to specification choice, and only the weaker marginal results are inconsistent across specifications, reinforcing the initial doubt placed in them.

5.2.3 Representativeness of Second Preferences

The construction of the implied vote shares in the absence of the Liberals carries with it an important assumption in order for it to be considered a reasonable benchmark. This assumption is representativeness — that Liberal voters who also rank non-Liberal candidates on their ballots have similar preferences to Liberal voters who only rank Liberal candidates. The need for this assumption arises because a considerable share (about three quarters) of Liberal voters only rank the Liberal Party group, which is the minimum required for a valid vote. However, as demonstrated above by the relatively low share of Liberal votes that go to invalid votes or non-turnout, this is likely due to the effort cost of ranking more alternatives, especially

[33] Except for state Liberal vote share which is essentially Coalition vote share but with poorer coverage.

since only a fraction of the single transferable vote is likely to transfer beyond a major party. This is not the case when Liberals are missing altogether. Neglecting this fact would lead to drastic underestimation of the transferred votes. Likewise, if the representativeness assumption fails because Liberal-only-ranking voters have systematically different preferences to those that express a ranking over additional candidates, the benchmark values will also be significantly less informative.

It is therefore important to probe this representativeness assumption. I do this by taking different assumptions over the behaviour of the latent rankings of Liberal-only-ranking voters, and different sample restrictions. First, I construct the implied vote shares in the absence of the Liberals without making the representativeness assumption at all — this only assumes that voters who reveal a secondary ranking represent a similar share of future voters, though this underestimates the expected transferred votes as discussed above. Nonetheless it provides a useful but extreme lower bound to the effect.

Second, I restrict counted ballots to those that rank at least two distinct groups, eliminating Liberal-only-ranking voters altogether. However, in treated units, voters who typically rank Liberals and exactly one other alternative are observationally equivalent to those voters who normally only vote Liberals but vote for someone else under the treatment — assuming the former does not wish to change their rankings, both groups are observed as voting for exactly one option in 2024 treated units. The solution is to employ this new assumption to restore parity: add the share of implied votes from 2021 for Liberal voters who express exactly one other preference back to the 2024 shares. The implied vote shares are then valid as long as the same share of voters express a consistent preference across the elections, which is precisely how the benchmark should behave. The downside of this approach is that it focuses only on the subset of voters that express further preferences.

Finally, I further restrict counted ballots to ballots which rank two or more distinct non-Liberal groups. This avoids the observational equivalence problem of the previous specification, at the cost of further reducing the pool of voters.

Table 9 in the Appendix presents these different specifications alongside the primary specification, for the implied vote share outcomes.[34] Across the four specifications, the point estimates for the implied Labor vote share are consistently negative, with all but the most conservative specification statistically significant at at least the 5% significance level. Columns (3) and (4) indicate that even among the multi-ranking voter population, the Labor shortfall relative to the benchmark remains. The point estimates are -1.65 (p=0.011) and -1.34 pp (p=0.047) respectively. Importantly, Column (2) shows that even when every voter that only ranks a Lib-

[34] Apart from the predicted transferred vote shares, these are the only outcomes which depend on representativeness — raw vote shares, and in particular invalid vote shares and turnout, do not.

eral option is assumed to not vote Labor in their absence, the point estimate remains negative, though statistically insignificant. That is, the effect is only removed if representativeness fails in the more extreme sense. The estimates for Other candidates' implied vote shares are also quite consistent across columns (1), (3) and (4), with only the conservative column differing substantially. Predictably, column (2) suggests that Other candidates did significantly better than just receiving transferred votes from those that would have been expected to rank them second if the Liberals were still available. This reinforces the assertion that the conservative approach grossly underestimates vote transfers. Together, these results do not provide substantive evidence that the reliance on representativeness is problematic.

5.2.4 Spillovers

The remaining discussion relates to the potential for spillovers, which might impact the interpretation of the results.

One obvious concern is that the Liberal Party now competes in fewer contests than planned and thus might have more resources to distribute over the fewer contests. If this substantially increases the Liberal vote share in control locations, vote shares for the other candidates could be suppressed, leading to a violation of the Stable Unit Treatment Value Assumption (SUTVA). Although it pushes in the other direction, one my also worry that the error caused reputational damage for the Liberals in control wards, reducing their support in those places.

By construction, however, this concern is largely excluded for the implied vote share outcomes. These outcomes remove Liberal candidates from voter rankings altogether, and in particular for control units in 2024 where this concern would live. An increase (or decrease) in Liberal rankings has no effect on these outcomes. This means the most striking result — that Labor fell significantly short of its expected vote share in the absence of the Liberals — is insulated from these concerns unless Liberal spending has second-order effects on the relative ranking of non-Liberal candidates. SUTVA threats must therefore act by changing the relative ranking of non-Liberal candidates.

A number of other factors suggest SUTVA is unlikely to be a concern, even for outcomes not insulated in this way. First, financial disclosures data indicates that the Liberal Party spent substantially less on campaigns across the 2024 local elections compared to 2021.[35] This suggests that even if surplus resources did flow to control units, they did not do so in full.

Second, the structure of local government limits exposure to spillovers in a predictable way. If voters desire a particular composition of parties in local government, some voters may switch to voting Liberal to balance the effect of the registration failure. This concern is limited to the

[35] In 2024 terms, the Liberal Party spent $899,000 in the 31 days leading up to the 2021 election, and $703,000 in the same period for the 2024 election.

level of the polity — i.e., the council level. In Appendix Table 10, I show that the results are robust to excluding these higher-risk control units from the control group, with each statistically significant coefficient in Panels A and C being at least as large in magnitude as their main specification counterpart. These control units in treated LGAs do not experience a statistically significant differential change compared with the unexposed control group, though power is necessarily limited by the small number of this category of ward.

Alternatively, voters in treated wards might seek to vote in control wards instead, in order to support Liberal candidates. Legally, this can only occur if the voter changes their registration to a new address. Since the electoral roll closed on 5 August 2024,[36] 9 days before the candidate registration deadline, voters responding to the Liberal failure would have to vote via declaration vote in order to change their registered address.[37] These votes are identified in the ballot data (directly in 2021 and 2024, pooled with declared institutions for 2012–2017), allowing us to test for evidence that voters moved ward. Appendix Table 11 shows that there is no statistically significant evidence of a flow of voters from treated wards to untreated wards overall, with 95% confidence intervals excluding effects larger than about 0.5 pp for the 2021/2024 sample and 0.28 pp for the full sample. Focusing on flows within councils that have at least one treated and at least one control ward, where such flows are most likely, we see that the estimates are again statistically insignificant across both the 2021/2024 sample (bootstrapped p=0.293) and the full sample (bootstrapped p=0.329), although power is severely limited in this subsample.[38] Still, 95% confidence intervals exclude effects as large as 0.71 pp and 0.59 pp respectively, an order of magnitude smaller than the Labor shortfall. As such, any movement of voters is minimal relative to the magnitude of the main effects.

Lastly, inspecting Appendix Figure 8 shows that there was no large increase in the Liberal vote share for control wards between 2021 and 2024 on average (the increase was smaller than from 2016/17 to 2021), while the differential changes in the Other candidates' vote shares and invalid vote shares correspond to large changes in treated unit outcomes.

Redistribution of expenditure by non-Liberal groups should itself be considered a competing mechanism for explaining why electoral outcomes in places missing Liberal candidates differ from places which do have Liberal candidates. This is studied in Section 6.3 below. This has a direct analogue to situations where a candidate is banned from contesting a single electorate, or a regional party only contests a subset of electorates. Extrapolating to a general absence of a major party from all electorates requires additional assumptions which may not hold. This is a necessary limitation of the DiD framework which requires the co-existence of a comparable

[36] https://elections.nsw.gov.au/bulletins/2024-nsw-local-government-elections-election-notice-no-1, accessed 2026-09-12.

[37] https://elections.nsw.gov.au/voters/voting-options/declaration-voting, accessed 2026-09-12.

[38] As the number of clusters is small (20), bootstrapped p-values are preferred.

control group.

5.3 Protest Vote

The results above demonstrate a clear increase in the share of invalid votes when the Liberal Party is absent. This is consistent with an expressive protest vote. To reinforce this interpretation, I disaggregate the invalid ballots by their reason for invalidity, which can be inferred from the ballot level data beginning from the 2016 elections.[39] In particular, I separate out forms of spoiled votes which are likely to be intentional (and thus validly be interpreted as a protest) and forms which are more likely to be accidental. Blank ballots are likely to represent an intentional expressive protest, since there has been no attempt to fill in the ballot at all. If the rise in invalid votes is driven by blank ballots, then we may conclude there was a protest vote.

Alternatively, it is possible that the rise in invalid voting is driven by an increase in mistakes. This is especially the case if voters attempt to vote for candidates that do not form a group ticket, which may be more likely in the absence of the Liberals given the large transfer of votes towards small parties and independents. This is because voting for individuals (termed voting 'below the line') requires filling in more boxes on the ballot paper compared to voting for group tickets (which only require one group to be ranked, compared with multiple individuals), increasing both the odds of mistakenly listing the same number twice, and failing to fill in the required number of boxes. Invalid votes which attempt to fill in boxes below the line are easily observed in the data.

The results are presented in Figure 6, with the companion table included in the Appendix (Table 12). I report that blank invalid votes rise significantly as a share of enrolled voters, by 2.0 pp, a change that is statistically significant at all conventional levels ($p<0.001$). This represents more than 75% of the increase in the total invalid vote share on the same sample (2016–2024). This equals a 77% increase over the pre-treatment mean, or 2.5 standard deviations.

Below the line invalid votes also rise by a statistically significant amount ($p<0.001$), but substantially smaller in magnitude at just 0.58 pp. The change is also smaller in relative magnitudes — it is a 44% increase over the pre-treatment mean, or 0.8 standard deviations. This makes up the majority of the remaining 25% of the increase in invalid votes, with only a small share attributable to ballots that were invalid for other reasons.

To conclude, the detailed invalid vote share results strongly favour an interpretation that this is a protest vote, being driven by a sizeable influx of blank ballots.

[39] As such, detailed invalid ballot data is not available for 2012.

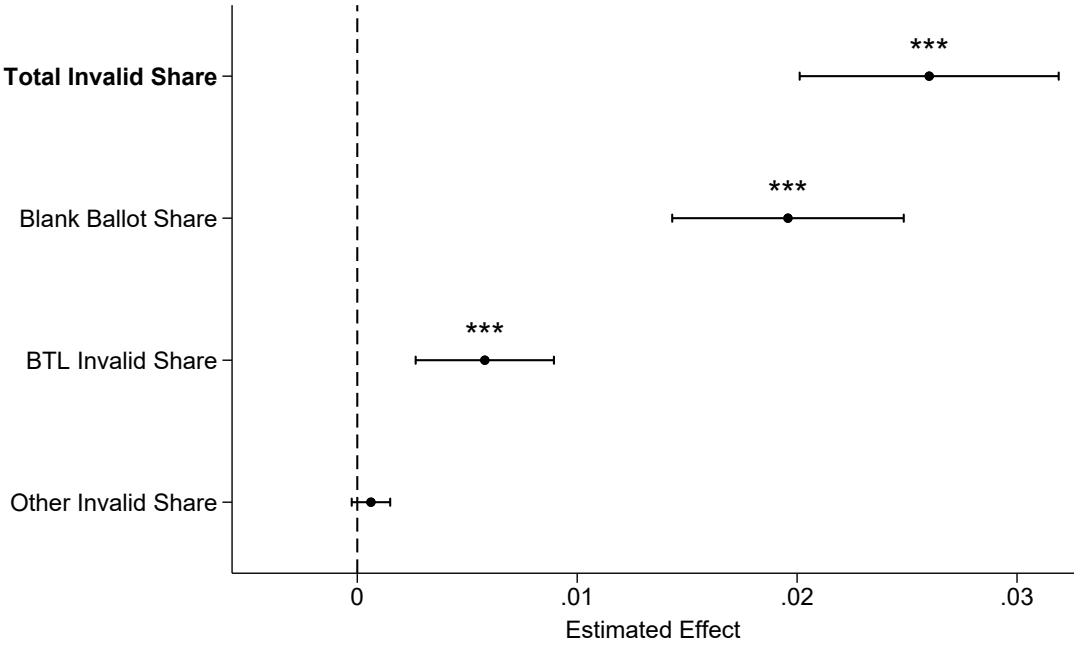


FIGURE 6: INVALID VOTING COMPONENTS

Notes: The figure presents estimates of the effect of Liberal Party absence from estimating Equation 1 and the associated 95% confidence intervals, with the following dependent variables: (1) invalid vote share, (2) blank invalid vote share, counting ballots on which no preference was marked at all, (3) below the line invalid vote share, counting ballots on which at least one individual candidate was marked, labelled BTL, and (4) the share of all other remaining invalid votes. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. The sample excludes the 2012 election, for which detailed data on invalid ballots is not available. Confidence intervals are calculated based on standard errors clustered at the ward level. * significant at 10%; ** significant at 5%; *** significant at 1%. The regressions that are the basis of this figure are reported in Table 12.

5.4 Voter Loyalty

In this section, I consider what the estimated treatment effects imply about voter loyalty and the willingness to support a candidate other than their most-preferred, through the framework of Pons and Tricaud (2018). Taking the estimates at face value, blank ballots account for around 9% of the orphaned Liberal votes. As a protest vote, this suggests that a small but appreciable minority of Liberal voters would rather discard their vote to express support for the missing Liberal Party, or their distaste for the remaining options, than express a preference ranking over those options. Likewise, we might include voters who decide not to turn out at all, around 5% of the orphaned Liberal voters, in this group, though the estimate is inconsistently significant across specifications. Combined, they give an estimate of about 13%. Even taking the high end of the 95% confidence intervals for these estimates as an upper bound, these account for less than 25% of the orphaned Liberal votes combined. This does not account for the possibility that at least some of the turnout effect (minimal though it may be) could be attributed to demobilisation of non-Liberal supporters due to reduced competition. This is in sharp contrast with Pons and Tricaud (2018), who estimate an equivalent share of more than 50% in their analysis of French run-off elections. They argue that this is indeed driven by 'loyals' — voters

who usually vote for the third candidate and do not vote for a remaining top-two candidate when the third is marginally absent due to the inclusion rule for the second round. The point estimates, which give about 15% of Liberal voters as loyal, are instead more in line with that of Fernández-Villaverde and Sanz (2025), who estimate that 10% of voters for a radical-right party in Spain exhibit loyalty through turnout and invalid votes.

Comparing the results in these three settings suggests that multi-round elections may play a role in generating loyalty, at least between rounds, since this is the most obvious distinction of French contests to Spanish and Australian ones. Voters may feel committed to a single candidate due to their actions in supporting them in the recent past, thus creating an anchoring effect. On the other hand, there are several factors that do not appear to explain the gap — whether the candidate is from a major party or not, and whether voters habitually support other candidates through ranked-choice voting.

6 Mechanisms

In this section, I address a number of potential mechanisms which could account for the observed results, in particular the shortfall of Labor votes relative to the fixed-rankings expectation. In Section 6.1, I use textual data from online news articles to study how the rhetoric around campaigns is affected by Liberal absence. I use the same data to classify candidate and party positions, the implications of which I study in Section 6.2. In Section 6.3, I investigate evidence of differential spending patterns in treated and control competitions. In Section 6.4, I discuss potential voter-side mechanisms that could explain the results.

6.1 Campaign Rhetoric

In this section, I leverage the corpus of online news articles to understand how campaign rhetoric changed in the absence of the Liberals. I focus on text which is attributed to candidates using GPT 5.4. In Appendix B, I provide the specific LLM prompts used at each stage of the text analysis. The analysis excludes 2012, which is not covered by the GDELT GKG dataset, and has only limited coverage for 2021.

Political campaigns have the potential to shift voting behaviour, as a considerable share of voters only make their mind up in the two months before an election (Le Pennec and Pons, 2023). In this case, which issues are discussed, and how they are discussed, also have the potential to shift dramatically due to the absence of one of two major parties. In particular, it would not be reasonable for the campaign to focus on the differences between the Labor and Liberal candidates (who usually get a large share of votes) when the Liberals are missing.

One dimension that has shown the ability to predict and move voting behaviour is the

moral content of speech (Enke, 2020; Grillo et al., 2026). Panel A of Table 3 reports the main specification (Equation (1)) for outcomes based on the moral foundations dictionary used by Enke (2020). The relative universalism index summarises the extent to which a body of text features moral language in relation to universalist values like fairness, as compared to communalist values like loyalty. Text is pooled to the ward-year level, and the outcome is standardised.

Column (1) reports that there is a statistically significant decrease in the relative universalism index (p=0.093), in the order of 1.4 standard deviations. This reflects a shift away from universalist language and towards communalist language. Columns (2)–(5) break down some important components of this shift. Column (2) shows the effect on relative universalism on text excluding Liberal candidates. The point estimate is −0.9 standard deviations (of the original relative universalism index used in Column (1)), and is only marginally insignificant (p=0.140). This suggests that a substantial amount of the effect comes from the compositional effect of Liberal absence — in the order of 35% of the effect.[40] The remainder is due to changes in the moral content of speech by non-Liberal candidates. Columns (3) and (4) further break this down into the universalist and communal components, the difference of which forms the index. The majority of the change in relative universalism among non-Liberal candidates is due to an increase in communal moral language (86%), with the remainder coming from a decrease in universalist moral language. The rise in communal language is itself only marginally insignificant (p=0.113). Lastly, I report that the overall share of moral language increases but not in a statistically meaningful way (p=0.146).

It is possible that the shift towards communal moral language is in part due to an increase in discussion of local community issues and related topics in the absence of a major party which otherwise might frame the debate as abstract Liberal values against abstract Labor values. To explore this possibility, I study the share of 'localist' language of candidate-attributed text using a domain-specific dictionary generated by GPT 5.4. Panel B of Table 3 reports the estimated effects on the share of candidate-attributed words which appear in the localism dictionary, for all candidates (Column (1)) and non-Liberal candidates (Column (2)). The point estimates are consistent with a shift towards community discussion, both due to composition and behavioural shifts, however they are very noisy and not statistically significant (p=0.507, 0.572). As such, the shift in moral language is likely more than just a shift towards focusing on local community.

I examine the way parties are discussed in candidate-attributed speech. Panel C of Table 3 reports the estimated effects on the rate of own-party mentions per candidate, for all candidates, non-Liberal candidates, and Labor and Greens candidates (the next-largest national

[40] Due to coverage issues, the sample changes slightly between Column (1) and the remaining columns, meaning this decomposition of the main result in Column (1) is not precise.

Table 3: Campaign Rhetoric

	(1) Relative Universalism	(2) Relative Universalism (excl. Lib.)	(3) Universalism (excl. Lib.)	(4) Communalism (excl. Lib.)	(5) Share Moral (excl. Lib.)
Panel A: Moral Language					
Treatment	-1.425*	-0.923	-0.128	0.795	0.667
	(0.833)	(0.613)	(0.338)	(0.492)	(0.450)
p-value	0.0931	0.1394	0.7066	0.1130	0.1459
Wild-bootstrap p-value	0.0715	0.1295	0.7131	0.1037	0.1565
Mean dep. var. (2016/17)	-0.3836	-0.3259	0.8528	1.1787	1.6246
SD dep. var. (2016/17)	0.9533	0.8065	0.6659	0.8724	1.0169
Observations	119	104	104	104	104
N ward clusters	51	46	46	46	46

	Share Local	Share Local (excl. Lib.)
Panel B: Local Focus		
Treatment	0.0166	0.0109
	(0.0248)	(0.0191)
p-value	0.5071	0.5721
Wild-bootstrap p-value	0.7337	0.7184
Mean dep. var. (2016/17)	0.0417	0.0403
SD dep. var. (2016/17)	0.0280	0.0260
Observations	119	104
N ward clusters	51	46

	All	Non-Lib.	Labor	Greens
Panel C: Own-party References				
Treatment	-0.0743*	-0.0490	-0.144**	0.00734
	(0.0416)	(0.0348)	(0.0571)	(0.122)
p-value	0.0788	0.1641	0.0146	0.9523
Wild-bootstrap p-value	0.0930	0.1896	0.0262	0.9521
Mean dep. var. (2016/17)	0.0863	0.0664	0.1089	0.2169
SD dep. var. (2016/17)	0.1552	0.1185	0.1926	0.3049
Observations	156	156	137	105
N ward clusters	63	63	56	44

Notes: The table reports estimates of the effect of Liberal Party absence from estimating Equation 1. **Panel A** presents the effects on the moral content of candidate-attributed text, pooled to the ward-year level. Column (1) uses the relative universalism index of Enke (2020), based on the Moral Foundations Dictionary, as the dependent variable, standardised. Column (2) uses the same outcome but excludes text attributed to Liberal candidates, and normalises by the standard deviation uses in Column (1) for comparability, as the dependent variable. Columns (3) and (4) use the universalism and communalism components of the Column (2) outcome respectively as the dependent variables. Note that the outcome in Column (2) is the difference between the outcomes in Columns (3) and (4). Column (5) uses the share of words from candidate-attributed text for non-Liberal candidates that belong to the universalism or communalism arms of the Moral Foundations Dictionary, standardised by its own standard deviation, as the dependent variable. **Panel B** presents the effects on the localism content of candidate-attributed text, pooled to the ward-year level. Column (1) uses the share of non-stop words belonging to the LLM-generated localism dictionary as the dependent variable. Column (2) uses the same outcome as Column (1) but excludes text attributed to Liberal candidates. **Panel C** presents the effects on the number of own-party mentions per candidate in available text. Column (1) uses the number of own-party mentions per candidate for all candidates as the dependent variable. Column (2) excludes Liberal candidates from the construction of the dependent variable. Column (3) limits the construction to Labor candidates. Column (4) limits the construction to Greens candidates, Greens being the largest national party that is included in Other for election outcomes. All dependent variables in this table are constructed using text from the 31 days prior to election day. Pre-treatment descriptive statistics are presented for the 2016/17 elections, due to limited GDELT GKG coverage in 2021. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

party that contest local elections). I report that in the absence of the Liberals, own-party mentions decrease by 0.07 per candidate, almost 0.5 standard deviations. Part of this effect is likely the composition effect — without Liberals to talk about their own party, treated contests have a greater share of independents that have no party identity to invoke. Disaggregating reveals that Labor candidates greatly reduce the rate at which they mention their own party relative to the control wards — by 0.14 mentions per candidate or 0.75 standard deviations. This cannot be driven by a composition effect. Relative to the pre-treatment mean of 0.109 party mentions per candidate, this reflects a substantial decrease. This effect is statistically significant at the 5% level (p=0.015). Interesting, there is no evidence that the Greens (the largest national party that is included in Other for election outcomes) reduce their rate of mentioning their own party, though the estimates are much less precise than they are for Labor, possibly due to the slightly weaker coverage.

Lastly, I explore the prevalence of topics being discussed in candidate-attributed text. Table 4 reports the estimated effects on the eight most commonly-discussed topics, as coded by GPT 5.4 against a manually constructed list of domain-specific topics. Consistent with the results above, there is an increase in the discussion of 'party politics' (p=0.040), a topic that includes discussion of the role of national parties in local politics, the use of community-first or non-partisan framing, and discussion over the party balance of council. The economy and social topics experience a statistically significant reduction in how much they are discussed, at the 10% significance level. However, this exploratory analysis involves a large number of hypothesis tests and so should be considered suggestive only. It is reassuring, however, that the party politics effect is thematically related to the previous results.

Taken together, these results suggest that the absence of the Liberals may transform the framing of the local political debate in a meaningful way. They are consistent with a political debate that in their presence is framed as a contest between two major national parties and their general ideologies, and which in their absence is framed as grass-roots local candidates campaigning on local issues against a national party. If voting behaviour follows the cues given by this framing, this could explain why support for Labor falls when the Liberal Party is absent — they would otherwise benefit from the contest being framed as Labor vs Liberal, weakening the platform of alternative candidates.

Table 4: Topic Analysis

	(1) Development	(2) Integrity	(3) Party politics	(4) Party politics (excl. Lib error)	(5) Economy	(6) Services	(7) Environment	(8) Fiscal	(9) Social
Treatment	-0.00504	0.0698	0.188**	0.0931*	-0.0628*	-0.00266	-0.0163	-0.0830	-0.0501*
	(0.0410)	(0.0549)	(0.0893)	(0.0516)	(0.0371)	(0.0400)	(0.0261)	(0.0719)	(0.0263)
p-value	0.9027	0.2097	0.0403	0.0769	0.0971	0.9473	0.5352	0.2539	0.0622
Wild-bootstrap p-value	0.9103	0.2221	0.0131	0.0561	0.1202	0.9627	0.5393	0.3038	0.0595
Mean dep. var. (2016/17)	0.1490	0.1016	0.0720	0.0720	0.0536	0.0414	0.0383	0.0315	0.0132
SD dep. var. (2016/17)	0.1735	0.1833	0.1264	0.1264	0.0630	0.0641	0.0700	0.0427	0.0272
Observations	120	120	120	120	120	120	120	120	120
N ward clusters	51	51	51	51	51	51	51	51	51

Notes: The table reports estimates of the effect of Liberal Party absence from estimating Equation 1. Each column uses the share of words in candidate-attributed text, including text attributed to Liberal candidates, that discuss the topic named in the column heading as the dependent variable, pooled to the ward-year level. The non-mutually exclusive topics are coded by GPT 5.4 against a manually constructed list of fourteen domain-specific topics and a residual 'other' policy category, with words that discuss more than one topic split equally between those topics. The prompt, including the description of each topic, is provided in Appendix B.2. The table reports the eight most commonly discussed of the fourteen topics across the full sample, ordered by their 2016/17 mean. Column (4) uses the same outcome as Column (3) but excludes text in articles that report the registration failure, as coded by GPT 5.4 (Appendix B.5), from the numerator only. Since no such articles appear before 2024, the descriptive statistics in Columns (3) and (4) are identical. All dependent variables in this table are constructed using text from the 31 days prior to election day. Pre-treatment descriptive statistics are presented for the 2016/17 elections, due to limited GDELT GKG coverage in 2021. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

6.2 Policy Positions

In this section I discuss the extent to which the data supports a policy response by candidates driving the observed Labor shortfall. It is possible that local candidates sought to fill the vacated policy position of the missing Liberals, something that national party candidates may not be able to do without risking damage to their party's brand. Using text identified as being about a candidate, I have GPT 5.4 code candidates on a left-right policy scale on 5 local policy issues, and use that as the basis for an overall left-right position between -1 and 1 in increments of 0.25, with left corresponding to negative values. In Table 13 in the Appendix, I report the average position of candidates on this scale by grouping. The LLM measures are on average in line with expectation: Liberals are ranked to the right of centre, and specifically to the right of Labor, who are in turn ranked to the right of the Greens.

I do the same on the issue of development (e.g., approval of housing developments) which does not sit cleanly on the left-right spectrum, but which is an important issue in Australian local politics. I then ask whether candidates in treated wards present with different policy positions, and in particular whether they shift towards the pre-treatment Liberal average position. Here I take candidates as the unit of observation, but otherwise retain the features of the main specification (in particular, ward-level fixed effects). I use the online news articles corpus as it is the best available information about candidate positions, since no manifesto-type data is available for local government elections.

Appendix Table 14 presents the results of these regressions. Panel A shows that there is no statistically significant shift by candidates either in the left-right dimension, or the pro-/anti-development dimension. However, the policy position data is only available for the restricted set of candidates for which there is sufficient text to place them on each of the two spectrums.[41] As such, the estimates are noisy and cannot exclude the possibility of potentially meaningful changes in policy position. In particular, the 95% confidence interval for the left-right position of non-Liberal candidates includes leftward shifts as large as 0.16 on a -1 to 1 scale. Similarly, Panel B of Table 14 does not show a statistically significant convergence of policy towards the gap left by the absent Liberal candidates, with the 95% confidence interval including convergent shifts as large as 0.11 and divergent shifts (away from the average pre-treatment Liberal position) of almost equal magnitude. While the results do not provide support for a shift in policy positions driving the results, available data also cannot exclude the possibility.

Another way to test this channel is to see if disproportionately more of the increase in the Other vote share went to right-leaning candidates. In Appendix Table 15, I disaggregate Other votes into right-leaning (with a left-right score >0), not right-leaning (≤ 0) and uncat-

[41] Coverage of all left-right sub policy areas was not required for an overall left-right classification.

egorised, with the Greens separated out with an established left-wing agenda. Right-leaning candidates account for about 20% of pre-treatment Other votes, with about 70% able to be classified. Despite uncategorised being only marginally the largest of the groupings, I find that it experiences a statistically significant increase accounting for almost 90% of the increase in Other votes (p=0.004). The point estimates for right- and not-right-leaning Other candidates are substantially smaller, negative, and not statistically significantly different from 0 (p=0.853, 0.751). The Greens vote share increases by about 3.8 pp and is only marginally significant (p=0.087). Therefore, available data does not support there being electoral returns to being a close ideological substitute for the Liberal Party.

6.3 Campaign Expenditure

Another possibility is that the absence of the Liberal Party in some contests changes the return to campaign expenditure. On the one hand, there is less competition in treated wards, so candidates might expect to be able to perform just as well but with less expenditure. On the other hand, several of the missing Liberal candidates were incumbents, meaning there is an incumbency advantage to be gained, increasing the incentive to spend on campaigning. If the former channel is stronger for Labor, or the latter is stronger for Other candidates, then campaign spending could account for the observed deviation from the fixed-rankings benchmark. I use financial disclosures data to investigate whether spending patterns could drive the main results, specifically Labor's shortfall relative to the benchmark.

Table 5 provides the results of estimating Equation (1) with real campaign spending in Australian dollars per enrolled voter, a binary equal to 1 if there was any spending (extensive margin), spending per enrolled voter conditional on spending being positive (intensive margin), and spending per candidate. Panel A provides the estimates for Labor spending, Panel B for spending by Other candidates (candidates other than Liberal or Labor), and Panel C for all candidates and groups. I do not look at the effect on Liberal spending, as their ward-specific spending was zero in 2024 — instead, campaign expenditure was incurred at the state level by the party, which means a DiD specification on those outcomes is not possible.

I report that there is no statistically significant effect of Liberal absence on Labor, Other, or Total campaign expenditure, on any of the four reported margins. There is insufficient evidence to conclude that expenditure behaviours drove the Labor shortfall. The estimates are, however, noisy. Focusing on spending per enrolled voter, the point estimate for Labor spending is −2.6 cents per voter but the 95% confidence intervals span −19 cents to +14 cents, or −0.9 to +0.7 in terms of pre-treatment standard deviations. The 95% confidence interval for Other spending spans −0.5 to +0.6 standard deviations. For total spending, it spans −0.6 to

Table 5: Campaign Expenditure

	(1) $/voter	(2) Any spend	(3) $/voter |> 0	(4) $/candidate
Panel A: Labor				
Treatment	-0.0263	-0.0331	0.0448	273.8
	(0.0841)	(0.147)	(0.0866)	(768.7)
p-value	0.7553	0.8228	0.6067	0.7225
Wild-bootstrap p-value	0.7662	0.8234	0.6216	0.7370
Mean dep. var. (2021)	0.1396	0.8625	0.1702	1308.3
SD dep. var. (2021)	0.2107	0.3466	0.2222	2033.8
Observations	277	277	214	277
N ward clusters	85	85	72	85
Panel B: Other				
Treatment	0.0187	0.0836	0.0276	78.05
	(0.0911)	(0.119)	(0.136)	(222.5)
p-value	0.8382	0.4848	0.8404	0.7266
Wild-bootstrap p-value	0.8435	0.4974	0.8502	0.7331
Mean dep. var. (2021)	0.2172	0.8752	0.2583	608.1
SD dep. var. (2021)	0.3126	0.3325	0.3281	689.4
Observations	299	299	217	299
N ward clusters	88	88	74	88
Panel C: Total				
Treatment	0.0103	-0.0235	0.0551	413.2
	(0.123)	(0.106)	(0.141)	(263.4)
p-value	0.9337	0.8258	0.6978	0.1203
Wild-bootstrap p-value	0.9369	0.8264	0.7048	0.1341
Mean dep. var. (2021)	0.3505	0.9169	0.3785	635.0
SD dep. var. (2021)	0.3901	0.2776	0.3931	619.2
Observations	301	309	269	301
N ward clusters	88	90	86	88

Notes: The table reports estimates of the effect of Liberal Party absence on campaign spending from estimating Equation 1. **Panel A** presents the effects on campaign spending by Labor candidates. Column (1) uses real spending per enrolled voter as the dependent variable. Column (2) uses a dummy equal to one if there was any spending as the dependent variable. Column (3) uses the same outcome as Column (1) but limits the sample to ward-years with positive spending. Column (4) uses real spending per candidate as the dependent variable. **Panel B** uses the same specification as Panel A but with spending by all other non-Liberal, non-Labor candidates. **Panel C** uses the same specification as Panel A but with spending by all candidates and groups. Spending is taken from candidate and ballot group financial disclosures, and excludes spending by central party organisations. It is measured in 2024 Australian dollars. All dependent variables in this table are constructed using spending from the 31 days prior to election day. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

+0.7 standard deviations. Therefore, while there is no statistically significant evidence for the campaign expenditure mechanism, I also cannot exclude the possibility of meaningful changes in spending behaviour in either direction.

6.4 Other Mechanisms

Lastly, I consider a number of alternative mechanisms which could account for the Labor shortfall, although data limitations mean I am not able to provide substantive statistical evidence on them.

One possibility is that endorsement by a major party acts as an informative signal about the quality of a candidate. Voters may see little benefit in investing time and effort into investigating the quality of independents and small party candidates when their usual first preference (Liberal) is available. Putting Labor as second might not express full-information preferences, but rather ensure that one's vote goes to a relatively competent candidate even if policy alignment is poor. When the most-preferred option is completely absent, the incentive to acquire information rises, and may overcome the friction. This would mean preference data overstates the rate at which Liberal votes would transfer to Labor. However, to the extent that voters vote above-the-line for a group ticket, expressing a second preference is completely optional. Indeed, Liberal candidates almost always run on a group ticket in the sample, and about 90% of 2021 Liberal votes are for the group ticket. Therefore, it is less likely that the Labor second-preferences are simply an artefact of poor information, since they also reflect an active decision to provide a second preference which is optional.

Voters may have seen Labor as manipulating the electoral system to their advantage, given they controlled the NSW state government. The particular circumstances of the error do not support this — the deadline was published well in advance of the election, was a statutory requirement, and was enforced by the independent NSW Electoral Commission, not the NSW Labor government. The dominant media narrative blames the Liberal Party for the issue: based on 246 articles on the issue from the news corpus, 241 blame the Liberal Party and zero blame Labor. In any case, this channel should be common to treated and control units alike, as many national and state news outlets covered the story.

Another possibility is that some Liberal voters in fact have a preference for balanced representation in the council. They may prefer their residual votes to support Labor as a counterbalance to the Liberal Party, but when Liberals are missing they can only maintain balance by supporting a third party or independent candidate. This would mean the expected transferred votes overstate the share of Liberal votes which would transfer to Labor. The Labor shortfall may also expose a broader violation of the Independence of Irrelevant Alternatives assumption.

If voters view ranking a candidate second as inherently different to ranking them first, then preference rankings are unable to predict voting behaviour when candidate pools change. That is, substitution may not follow the fixed rankings. Available data cannot provide evidence on these mechanisms.

7 Conclusion

In this paper, I analyse how the absence of a major party affects support for the remaining candidates. I rely on a specific instance in which a major party was excluded from a number of contests in local elections in New South Wales, Australia due to an unexpected administrative error. Using ranked ballot data, I construct a benchmark that predicts substitution patterns based on the way voters ranked candidates on actual ballot papers. This allows me to disentangle the substantive effect of major party absence on the way voters rank candidates from the mechanical substitution effect that arises under fixed rankings.

I find that the remaining major party falls significantly short of the fixed-rankings benchmark by 7.1 pp, missing out on the equivalent of 14.2 seats' worth of votes across affected elections. Smaller parties and independents do marginally better than predicted. I document a substantial protest vote, with invalid voting rates rising significantly, predominantly through blank ballots. I also find a small but not robust decrease in turnout.

I observe patterns in local news coverage consistent with a narrative in which the absence of the major party reshaped the political debate away from a contest between the two major parties, towards one in which a communalist, non-partisan focus and discussion over the role of parties in local politics gained prominence. I find no evidence that campaign expenditure or changes in policy platforms contribute, though estimates are imprecise. I also argue that it is unlikely that the results are caused by information acquisition or by voters blaming Labor for the Liberals' error.

Together, these results suggest that the presence of a partisan divide reinforces the dominance of major parties, possibly by framing the debate around the divide rather than around dimensions on which other candidates can compete. This provides motivation for further investigation into the stability of major parties alternating power, despite many voters being dissatisfied with political parties and politicians.

References

Abou-Chadi, T. and W. Krause (2018). The causal effect of radical right success on mainstream parties' policy positions: A regression discontinuity approach. *British Journal of Political*

Science 50(3), 829–847.

Anagol, S. and T. Fujiwara (2016). The runner-up effect. *Journal of Political Economy 124*(4), 927–991.

Ansolabehere, S. and M. S. Puy (2018). Measuring issue-salience in voters' preferences. *Electoral Studies 51*, 103–114.

Arenas, A. (2016). Sticky votes. *Journal of Economic Behavior & Organization 132*, 12–25.

Arenas, A. (2021). Party bans: Deterrence or backlash? Evidence from the Basque Country. *Quarterly Journal of Political Science 16*(3), 325–358.

Ash, E., M. Morelli, and R. Van Weelden (2017). Elections and divisiveness: Theory and evidence. *The Journal of Politics 79*(4), 1268–1285.

Australian Bureau of Statistics (2021). Australian Census.

Balbuzanov, I., J. Gars, M. Stalinski, and E. Tjernström (2026). Incentivizing engagement: Experimental evidence on journalist performance pay. Technical report.

Baltz, S. (2025). The probability of casting a pivotal vote in an instant runoff voting election. *Political Analysis 33*(4), 447–455.

Bartholdi, J. J. and J. B. Orlin (1991). Single transferable vote resists strategic voting. *Social Choice and Welfare 8*(4), 341–354.

Bekkouche, Y., J. Cagé, and E. Dewitte (2022). The heterogeneous price of a vote: Evidence from multiparty systems, 1993–2017. *Journal of Public Economics 206*, 104559.

Besley, T., T. Persson, and D. M. Sturm (2010). Political competition, policy and growth: Theory and evidence from the US. *Review of Economic Studies 77*(4), 1329–1352.

Bol, D. and R. Ivandic (2022). Does the number of candidates increase turnout? Causal evidence from two-round elections. *Political Behavior 44*(4), 2005–2026.

Bonomi, G., N. Gennaioli, and G. Tabellini (2021). Identity, beliefs, and political conflict. *The Quarterly Journal of Economics 136*(4), 2371–2411.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013). Salience and consumer choice. *Journal of Political Economy 121*(5), 803–843.

Boxell, L., M. Gentzkow, and J. M. Shapiro (2024). Cross-country trends in affective polarization. *Review of Economics and Statistics 106*(2), 557–565.

Brennan, G. and L. Lomasky (1993). *Democracy and Decision: The Pure Theory of Electoral Preference*. Cambridge University Press.

Buisseret, P. (2017). Electoral competition with entry under non-majoritarian run-off rules. *Games and Economic Behavior 104*, 494–506.

Burnitt, C., J. Gars, and M. Stalinski (2026). Politics of food: An experiment on trust in expert regulation and economic costs of political polarisation. *The Economic Journal*.

Bursztyn, L., D. Cantoni, P. Funk, F. Schönenberger, and N. Yuchtman (2023). Identifying the effect of election closeness on voter turnout: Evidence from Swiss referenda. *Journal of the European Economic Association 22*(2), 876–914.

Callander, S. (2005). Electoral competition in heterogeneous districts. *Journal of Political Economy 113*(5), 1116–1145.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics 90*(3), 414–427.

Carter, A. V., K. T. Schnepel, and D. G. Steigerwald (2017). Asymptotic behavior of a t-test robust to cluster heterogeneity. *The Review of Economics and Statistics 99*(4), 698–709.

Casey, K. (2015). Crossing party lines: The effects of information on redistributive politics. *American Economic Review 105*(8), 2410–2448.

Chatterjee, S., P. Maitra, and M. Manchanda (2024). The relevant third: Threat of coalition and economic development. Monash Economics Working Papers 2024-13, Monash University, Department of Economics.

Cheng, H.-P. and E. L. Swee (2024). Farewell president! Political favoritism, economic inequality, and political polarization. *European Journal of Political Economy 81*, 102483.

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal 21*(1), C1–C68.

Das, K., A. Ghosh, and P. Maitra (2024). Exogenous shocks and electoral outcomes. *Public Choice 201*(1-2), 145–179.

de Chaisemartin, C. and X. D'Haultfœuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–2996.

Di Leo, R., E. Dinas, and B. Meiske (2025). Berlinguer, I love you (still): The downstream effects of expressive voting. *Political Behavior*.

Di Tella, R., R. Kotti, C. Le Pennec, and V. Pons (2025). Keep your enemies closer: Strategic platform adjustments during US and French elections. *American Economic Review 115*(8), 2488–2528.

Dinas, E., E. Hartman, and J. van Spanje (2016). Dead man walking: The affective roots of issue proximity between voters and parties. *Political Behavior 38*(3), 659–687.

Downs, A. (1957). An economic theory of political action in a democracy. *Journal of Political Economy 65*(2), 135–150.

Drometer, M. and J. Rincke (2008). The impact of ballot access restrictions on electoral competition: evidence from a natural experiment. *Public Choice 138*(3-4), 461–474.

Duverger, M. (1951). *Les partis politiques*. Librairie Armand Colin.

Eggers, A. C. (2014a). Partisanship and electoral accountability: Evidence from the UK expenses scandal. *Quarterly Journal of Political Science 9*(4), 441–472.

Eggers, A. C. (2014b). Proportionality and turnout: Evidence from French municipalities. *Comparative Political Studies 48*(2), 135–167.

Eggers, A. C. and T. Nowacki (2024). Susceptibility to strategic voting: A comparison of plurality and instant-runoff elections. *The Journal of Politics 86*(2), 521–534.

Enke, B. (2020). Moral values and voting. *Journal of Political Economy 128*(10), 3679–3729.

Erikson, R. S. and T. R. Palfrey (2000). Equilibria in campaign spending games: Theory and data. *American Political Science Review 94*(3), 595–609.

Faravelli, M., B. Gordon, J. Huisman, and V. L. te Velde (2025). What drives voter turnout? Experimental insights on collectivism and probability weighting. *Theory and Decision 100*(2), 427–449.

Faravelli, M., K. Kalayci, and C. Pimienta (2020). Costly voting: a large-scale real effort experiment. *Experimental Economics 23*(2), 468–492.

Faravelli, M., U. Khalil, and S. Ponnusamy (2025). Procedural barriers to political candidacy: Gender, discouragement and candidate persistence. *The Economic Journal 136*(676), 1440–1473.

Faravelli, M. and S. Sanchez-Pages (2014). (Don't) make my vote count. *Journal of Theoretical Politics 27*(4), 544–569.

Ferejohn, J. (1986). Incumbent performance and electoral control. *Public Choice 50*(1-3), 5–25.

Fernández-Villaverde, J. and C. Sanz (2025). How substitutable are the classical and radical right? *Journal of Public Economics 252*, 105521.

Ferrara, A., P. A. Testa, and L. Zhou (2024). New area- and population-based geographic crosswalks for U.S. counties and congressional districts, 1790–2020. *Historical Methods: A Journal of Quantitative and Interdisciplinary History 57*(2), 67–79.

Ferraz, C. and F. Finan (2011). Electoral accountability and corruption: Evidence from the audits of local governments. *American Economic Review 101*(4), 1274–1311.

Fiorina, M. P. (1976). The voting decision: Instrumental and expressive aspects. *The Journal of Politics 38*(2), 390–413.

Fujiwara, T. (2011). A regression discontinuity test of strategic voting and Duverger's law. *Quarterly Journal of Political Science 6*(3-4), 197–233.

Galasso, V., T. Nannicini, and S. Nunnari (2021). Positive spillovers from negative campaigning. *American Journal of Political Science 67*(1), 5–21.

Gandhi, A., D. Iorio, and C. Urban (2015). Negative advertising and political competition. *Journal of Law, Economics, and Organization 32*(3), 433–477.

Gentzkow, M. and J. M. Shapiro (2010). What drives media slant? evidence from U.S. daily newspapers. *Econometrica 78*(1), 35–71.

Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica 41*(4), 587–601.

Glaeser, E. L., G. A. M. Ponzetto, and J. M. Shapiro (2005). Strategic extremism: Why Republicans and Democrats divide on religious values. *The Quarterly Journal of Economics 120*(4), 1283–1330.

Gonzales, M., G. León-Ciliotta, and L. R. Martínez (2022). How effective are monetary incentives to vote? Evidence from a nationwide policy. *American Economic Journal: Applied Economics 14*(1), 293–326.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

Granzier, R., V. Pons, and C. Tricaud (2023). Coordination and bandwagon effects: How past rankings shape the behavior of voters and candidates. *American Economic Journal: Applied Economics 15*(4), 177–217.

Grillo, E., J. S. Morales, and A. Moresi (2026). Moral characters: Social media and elections in the US. Technical report.

Hill, L. and S. Rutledge-Prior (2016). Young people and intentional informal voting in Australia. *Australian Journal of Political Science 51*(3), 400–417.

Hill, L. and S. Young (2007). Protest or error? Informal voting and compulsory voting. *Australian Journal of Political Science 42*(3), 515–521.

Iaryczower, M. and A. Mattozzi (2013). On the nature of competition in alternative electoral systems. *The Journal of Politics 75*(3), 743–756.

Kapoor, S. and A. Magesan (2018). Independent candidates and political representation in India. *American Political Science Review 112*(3), 678–697.

Kashner, D. and M. Stalinski (2024). Preempting polarization: An experiment on opinion formation. *Journal of Public Economics 234*, 105122.

Kawai, K. and Y. Watanabe (2013). Inferring strategic voting. *American Economic Review 103*(2), 624–662.

Kendall, C., T. Nannicini, and F. Trebbi (2015). How do voters respond to information? Evidence from a randomized campaign. *American Economic Review 105*(1), 322–353.

Khalil, U., M. Oak, and S. Ponnusamy (2024). The heterogeneous role of party affiliation in the runner-up effect. *Journal of Applied Econometrics 39*(5), 952–959.

Le Pennec, C. and V. Pons (2023). How do campaigns shape vote choice? Multicountry evidence from 62 elections and 56 TV debates. *The Quarterly Journal of Economics 138*(2), 703–767.

Levitt, S. D. (1994). Using repeat challengers to estimate the effect of campaign spending on election outcomes in the U.S. House. *Journal of Political Economy 102*(4), 777–798.

Le Pennec, C. (2023). Strategic campaign communication: Evidence from 30,000 candidate manifestos. *The Economic Journal 134*(658), 785–810.

MacKinnon, J. G. and M. D. Webb (2016). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics 32*(2), 233–254.

McCune, D. and J. Wilson (2023). Ranked-choice voting and the spoiler effect. *Public Choice 196*(1-2), 19–50.

Mitchell, P. (2014). The single transferable vote and ethnic conflict: The evidence from Northern Ireland. *Electoral Studies 33*, 246–257.

Morton, R. B., D. Muller, L. Page, and B. Torgler (2015). Exit polls, turnout, and bandwagon voting: Evidence from a natural experiment. *European Economic Review 77*, 65–81.

Mullainathan, S. and A. Shleifer (2005). The market for news. *American Economic Review 95*(4), 1031–1053.

Muller, D. and L. Page (2015). A new approach to measure tactical voting: evidence from the British elections. *Applied Economics 47*(36), 3839–3858.

Nibourel, C. (2025). How do voters and parties respond to the radical right? Technical report.

Nunnari, S. and J. Zápal (2025). A model of focusing in political choice. *The Journal of Politics 87*(4), 1465–1481.

Padró i Miquel, G. (2007). The control of politicians in divided societies: The politics of fear. *The Review of Economic Studies 74*(4), 1259–1274.

Palfrey, T. R. (1984). Spatial equilibrium with entry. *The Review of Economic Studies 51*(1), 139–156.

Perego, J. and S. Yuksel (2022). Media competition and social disagreement. *Econometrica 90*(1), 223–265.

Pew Research Center (2024). Representative democracy remains a popular ideal, but people around the world are critical of how it's working. Technical report, Pew Research Center.

Pons, V. and C. Tricaud (2018). Expressive voting and its cost: Evidence from runoffs with two or three candidates. *Econometrica 86*(5), 1621–1649.

Reilly, B. (2021). Ranked choice voting in Australia and America: Do voters follow party cues? *Politics and Governance 9*(2), 271–279.

Sant'Anna, P. H. and J. Zhao (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics 219*(1), 101–122.

Sharman, C., A. M. Sayers, and N. Miragliotta (2002). Trading party preferences: the Australian experience of preferential voting. *Electoral Studies 21*(4), 543–560.

Spenkuch, J. L. (2014). Please don't vote for me: Voting in a natural experiment with perverse incentives. *The Economic Journal 125*(585), 1025–1052.

Spenkuch, J. L. (2018). Expressive vs. strategic voters: An empirical assessment. *Journal of Public Economics 165*, 73–81.

Spenkuch, J. L. and D. Toniatti (2018). Political advertising and election results. *The Quarterly Journal of Economics 133*(4), 1981–2036.

Svaleryd, H. and J. Vlachos (2009). Political rents in a non-corrupt democracy. *Journal of Public Economics 93*(3-4), 355–372.

Tähtinen, T. (2025). Populism and ideological convergence: Evidence from a multiparty system. *Journal of Public Economics 241*, 105271.

Ujhelyi, G., S. Chatterjee, and A. Szabó (2020). None of the above: Protest voting in the world's largest democracy. *Journal of the European Economic Association 19*(3), 1936–1979.

Van der Straeten, K., J.-F. Laslier, N. Sauger, and A. Blais (2010). Strategic, sincere, and heuristic voting under four election rules: an experimental study. *Social Choice and Welfare 35*(3), 435–472.

Yousaf, H. (2021). Sticking to one's guns: Mass shootings and the political economy of gun control in the United States. *Journal of the European Economic Association 19*(5), 2765–2802.

A Appendix: Additional Results and Robustness

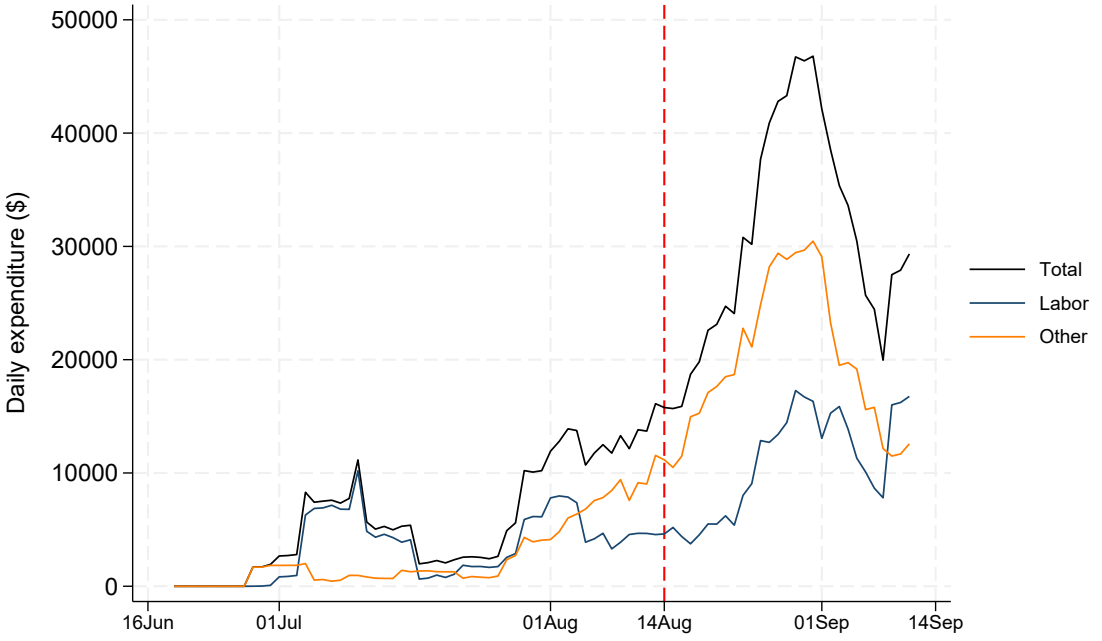


FIGURE 7: DAILY DISCLOSED CAMPAIGN EXPENDITURE, 2024

Notes: The figure presents the 7-day centred moving average of daily campaign spending in 2024, from 16 June to election day on 14 September. The black line is spending by all candidates and groups, the navy line is spending by Labor candidates, and the orange line is spending by all other non-Liberal, non-Labor candidates. Spending is taken from candidate and ballot group financial disclosures, and excludes spending by central party organisations. The sample is restricted to LGAs which did have a Liberal Party candidate in the most recent pre-2024 election. The red dashed line indicates when treatment took place, at the close of candidate nominations on 14 August 2024.

Table 6: Decomposition of Effect on Vote Shares (LGA Clusters)

	(1) Labor Share	(2) Other Share	(3) Invalid Share	(4) Turnout	(5) Liberal Share
Panel A: Actual change					
Treatment	0.0140	0.184***	0.0230***	-0.0107	-0.227***
	(0.0186)	(0.0275)	(0.00379)	(0.00676)	(0.0165)
p-value	0.4580	< 0.0001	< 0.0001	0.1258	< 0.0001
Wild-bootstrap p-value	0.4744	< 0.0001	< 0.0001	0.1502	< 0.0001
Mean dep. var. (2021)	0.2497	0.3018	0.0414	0.8360	0.2540
SD dep. var. (2021)	0.1094	0.1255	0.0140	0.0488	0.1055
Panel B: Expected transferred votes					
Treatment	0.0852***	0.145***			
	(0.0255)	(0.0214)			
p-value	0.0023	< 0.0001			
Wild-bootstrap p-value	0.0053	< 0.0001			
Mean dep. var. (2021)	0.0920	0.1662			
SD dep. var. (2021)	0.0755	0.0889			
Panel C: Excess					
Treatment	-0.0712**	0.0396			
	(0.0324)	(0.0304)			
p-value	0.0361	0.2024			
Wild-bootstrap p-value	0.0429	0.2225			
Mean dep. var. (2021)	0.3416	0.4680			
SD dep. var. (2021)	0.1522	0.1567			
Observations	277	299	301	301	301
N LGA clusters	30	30	30	30	30

Notes: The table reports regression estimates of the effect of Liberal Party absence from estimating Equation (1). The table uses the same specification as Table 2 but with standard errors and the wild cluster bootstrap clustered at the LGA level. **Panel A** provides the estimates using the actual first preference vote shares as the outcome variables. Column (1) uses the actual first preference vote share for Labor as the outcome variable. Column (2) uses the actual first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. Column (3) uses the share of invalid votes as the outcome variable. Column (4) uses the share of voters that turned out as the outcome variable. Column (5) uses the actual first preference vote share for Liberals as the outcome variable. **Panel B** uses the same specification as Panel A but with the expected transferred vote shares as the outcome variables (only available for Labor and Other). The estimate reported is the absolute value of the estimate from Equation (1). **Panel C** uses the same specification as Panel A but with the implied vote shares as the outcome variables (only available for Labor and Other). All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the LGA level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 7: Decomposition of Effect on Vote Shares (Polling Places)

	(1) Labor Share	(2) Other Share	(3) Invalid Share	(4) Turnout	(5) Liberal Share
Panel A: Actual change					
Treatment	-0.00198	0.0933***	0.0126***	-0.0142	-0.121***
	(0.0109)	(0.0134)	(0.00193)	(0.0110)	(0.0104)
p-value	0.8563	< 0.0001	< 0.0001	0.2001	< 0.0001
Wild-bootstrap p-value	0.8689	< 0.0001	< 0.0001	0.2422	< 0.0001
Mean dep. var. (2021)	0.1190	0.1497	0.0261	0.4040	0.1198
SD dep. var. (2021)	0.0849	0.1026	0.0192	0.1798	0.0926
Panel B: Expected transferred votes					
Treatment	0.0390***	0.0811***			
	(0.00963)	(0.00985)			
p-value	0.0001	< 0.0001			
Wild-bootstrap p-value	0.0002	< 0.0001			
Mean dep. var. (2021)	0.0402	0.0798			
SD dep. var. (2021)	0.0459	0.0683			
Panel C: Excess					
Treatment	-0.0410***	0.0123			
	(0.0152)	(0.0137)			
p-value	0.0086	0.3741			
Wild-bootstrap p-value	0.0123	0.3804			
Mean dep. var. (2021)	0.1592	0.2295			
SD dep. var. (2021)	0.1143	0.1431			
Observations	3952	4162	4179	4179	4179
N ward clusters	86	86	86	86	86

Notes: The table reports regression estimates of the effect of Liberal Party absence from estimating Equation (1)at the polling place level, with polling place fixed effects in place of ward fixed effects. **Panel A** provides the estimates using the actual first preference vote shares as the outcome variables. Column (1) uses the actual first preference vote share for Labor as the outcome variable. Column (2) uses the actual first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. Column (3) uses the share of invalid votes as the outcome variable. Column (4) uses the share of voters that turned out as the outcome variable. Column (5) uses the actual first preference vote share for Liberals as the outcome variable. **Panel B** uses the same specification as Panel A but with the expected transferred vote shares as the outcome variables (only available for Labor and Other). The estimate reported is the absolute value of the estimate from Equation (1). **Panel C** uses the same specification as Panel A but with the implied vote shares as the outcome variables (only available for Labor and Other). All outcomes are as a share of enrolled voters. Polling place counts exclude votes not cast at a polling place on election day, such as pre-poll, postal and declaration votes, so vote shares and turnout are lower than at the ward level. Observations are at the polling place-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 8: Robustness to Alternative Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Population-weighted			Benchmarks		Double machine learning		
	OLS	OLS, state control	DR-DiD	OLS	Plugin lasso	CV lasso	Random forest	Learner menu
Panel A: Labor share (excl. Liberal)								
Treatment	-0.0712***	-0.0706***	-0.0827**	-0.0770***	-0.0752***	-0.0574**	-0.0538**	-0.0659***
	(0.0228)	(0.0253)	(0.0317)	(0.0243)	(0.0244)	(0.0233)	(0.0234)	(0.0248)
p-value	0.0025	0.0066	0.0112	0.0021	0.0028	0.0158	0.0243	0.0095
Mean dep. var. (2021)	0.3416	0.3361	0.3370	0.3458	0.3458	0.3458	0.3458	0.3458
SD dep. var. (2021)	0.1522	0.1559	0.1520	0.1582	0.1582	0.1582	0.1582	0.1582
Observations	277	271	142	271	271	271	271	271
N ward clusters	85	83	71	83	83	83	83	83
Panel B: Other share (excl. Liberal)								
Treatment	0.0396*	0.0435*	0.0564*	0.0380*	0.0370	0.0343	0.0350	0.0388*
	(0.0224)	(0.0244)	(0.0303)	(0.0224)	(0.0224)	(0.0221)	(0.0224)	(0.0230)
p-value	0.0800	0.0773	0.0662	0.0937	0.1016	0.1249	0.1210	0.0954
Mean dep. var. (2021)	0.4680	0.4768	0.4764	0.4856	0.4856	0.4856	0.4856	0.4856
SD dep. var. (2021)	0.1567	0.1579	0.1584	0.1644	0.1644	0.1644	0.1644	0.1644
Observations	299	293	160	293	293	293	293	293
N ward clusters	88	86	80	86	86	86	86	86
Panel C: Invalid share								
Treatment	0.0230***	0.0232***	0.0250***	0.0247***	0.0243***	0.0198***	0.0205***	0.0208***
	(0.00297)	(0.00304)	(0.00288)	(0.00324)	(0.00323)	(0.00290)	(0.00275)	(0.00266)
p-value	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001
Mean dep. var. (2021)	0.0414	0.0392	0.0390	0.0373	0.0373	0.0373	0.0373	0.0373
SD dep. var. (2021)	0.0140	0.0113	0.0111	0.0112	0.0112	0.0112	0.0112	0.0112
Observations	301	295	162	295	295	295	295	295
N ward clusters	88	86	81	86	86	86	86	86
Panel D: Turnout								
Treatment	-0.0107*	-0.0106*	-0.00601	-0.00499	-0.00446	-0.00501	-0.00568*	-0.00666**
	(0.00617)	(0.00607)	(0.00496)	(0.00346)	(0.00349)	(0.00315)	(0.00316)	(0.00289)
p-value	0.0873	0.0849	0.2295	0.1536	0.2054	0.1153	0.0764	0.0236
Mean dep. var. (2021)	0.8360	0.8363	0.8360	0.8375	0.8375	0.8375	0.8375	0.8375
SD dep. var. (2021)	0.0488	0.0504	0.0505	0.0484	0.0484	0.0484	0.0484	0.0484
Observations	301	295	162	295	295	295	295	295
N ward clusters	88	86	81	86	86	86	86	86
Population weights	Yes	Yes	Yes					
State Coalition vote share		Yes	Yes		Yes	Yes	Yes	Yes
Census controls					Yes	Yes	Yes	Yes
Political controls					Yes	Yes	Yes	Yes
Squares & interactions					Yes	Yes		Yes

Notes: The table reports regression estimates of the effect of Liberal Party absence on implied vote shares, invalid votes and turnout under alternative specifications. **Panel A** provides the estimates using the implied first preference vote share for Labor as the outcome variable. Column (1) uses the same specification as Table 2. Column (2) uses the same specification as Column (1) but with the state Coalition vote share as a control, as in Equation (2). Column (3) uses the doubly robust DiD estimator (Sant'Anna and Zhao, 2020) on the 2021 and 2024 elections only, controlling for the change in the state Coalition vote share between the 2019 and 2023 state elections. Columns (4)–(8) limit the sample to ward-years in which all baseline characteristics are available. Column (4) uses the same specification as Column (1) but on this sample. Columns (5)–(8) present double machine-learning estimates of a partially linear model (Chernozhukov et al., 2018), with the ward and election fixed effects partialled out first. Column (5) uses a plug-in lasso. Column (6) uses a cross-validated lasso. Column (7) uses a random forest. Column (8) uses a stacked combination of OLS, cross-validated lasso, cross-validated ridge, random forest and gradient boosting. The controls included in each column are indicated in the final rows. **Panel B** uses the same specification as Panel A but with the implied first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. **Panel C** uses the same specification as Panel A but with the share of invalid votes as the outcome variable. **Panel D** uses the same specification as Panel A but with the share of voters that turned out as the outcome variable. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. In Columns (1)–(3), observations are weighted by population, including for the pre-treatment means and standard deviations. Standard errors clustered at the ward level are in parentheses. Standard p-values are reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 9: Robustness to Alternative Representativeness Assumptions

	(1) Remove Liberals + scale	(2) Remove Liberals + No scale	(3) Restrict sample + assume no change	(4) Extra restricted sample
Panel A: Labor				
Treatment	-0.0712***	-0.00215	-0.0165**	-0.0134**
	(0.0228)	(0.0151)	(0.00633)	(0.00663)
p-value	0.0025	0.8872	0.0109	0.0469
Wild-bootstrap p-value	0.0038	0.8825	0.0121	0.0523
Mean dep. var. (2021)	0.3416	0.2687	0.0963	0.0829
SD dep. var. (2021)	0.1522	0.1153	0.0410	0.0361
Observations	277	277	277	277
N ward clusters	85	85	85	85
Panel B: Other				
Treatment	0.0396*	0.145***	0.0321*	0.0331**
	(0.0224)	(0.0216)	(0.0167)	(0.0149)
p-value	0.0800	< 0.0001	0.0580	0.0283
Wild-bootstrap p-value	0.0831	< 0.0001	0.0662	0.0361
Mean dep. var. (2021)	0.4680	0.3403	0.1627	0.1361
SD dep. var. (2021)	0.1567	0.1335	0.0687	0.0614
Observations	299	299	299	297
N ward clusters	88	88	88	88

Notes: The table reports regression estimates of the effect of Liberal Party absence on implied vote shares under alternative constructions of the implied vote shares, from estimating Equation (1). **Panel A** provides the estimates using the implied first preference vote share for Labor as the outcome variable. Column (1) uses the same outcome as Panel C of Table 2: the implied vote share if Liberals are removed from all rankings, with transfers from Liberal ballots scaled so that exhausted ballots transfer in the same ratios as non-exhausted ballots. Column (2) uses the same outcome as Column (1) but without adjusting for ballot exhaustion, so the transfers are not scaled and exhausted Liberal ballots transfer to no one. Column (3) restricts the count to ballots ranking at least two distinct groups and, for treated wards in 2024, adds to each group the 2021 share of enrolled voters whose ballots ranked only the Liberals and that group. Column (4) counts only ballots that rank at least two distinct non-Liberal groups, which avoids this problem. **Panel B** uses the same specification as Panel A but with the implied first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 10: Exposed Control Units

	(1) Labor Share	(2) Other Share	(3) Invalid Share	(4) Turnout	(5) Liberal Share
Panel A: Actual change					
Treatment	0.00672	0.189***	0.0236***	-0.0127*	-0.228***
	(0.0136)	(0.0222)	(0.00301)	(0.00709)	(0.0161)
Exposed controls	-0.0371	0.0262	0.00309	-0.0106	-0.00171
	(0.0248)	(0.0282)	(0.00217)	(0.00728)	(0.0171)
p-value	0.6235	< 0.0001	< 0.0001	0.0774	< 0.0001
Wild-bootstrap p-value	0.6282	< 0.0001	< 0.0001	0.0821	< 0.0001
Exposed controls p-value	0.1394	0.3559	0.1592	0.1476	0.9204
Exposed controls wild-bootstrap p-value	0.2277	0.4452	0.2257	0.2072	0.9188
Mean dep. var. (2021)	0.2497	0.3018	0.0414	0.8360	0.2540
SD dep. var. (2021)	0.1094	0.1255	0.0140	0.0488	0.1055
Panel B: Expected transferred votes					
Treatment	0.0918***	0.138***			
	(0.0188)	(0.0152)			
Exposed controls	0.0337	-0.0337			
	(0.0275)	(0.0209)			
p-value	< 0.0001	< 0.0001			
Wild-bootstrap p-value	< 0.0001	< 0.0001			
Exposed controls p-value	0.2225	0.1102			
Exposed controls wild-bootstrap p-value	0.3022	0.2051			
Mean dep. var. (2021)	0.0920	0.1662			
SD dep. var. (2021)	0.0755	0.0889			
Panel C: Excess					
Treatment	-0.0851***	0.0509**			
	(0.0240)	(0.0239)			
Exposed controls	-0.0708	0.0599			
	(0.0466)	(0.0434)			
p-value	0.0006	0.0362			
Wild-bootstrap p-value	0.0009	0.0409			
Exposed controls p-value	0.1322	0.1709			
Exposed controls wild-bootstrap p-value	0.2225	0.2588			
Mean dep. var. (2021)	0.3416	0.4680			
SD dep. var. (2021)	0.1522	0.1567			
Observations	277	299	301	301	301
N ward clusters	85	88	88	88	88

Notes: The table reports regression estimates of the effect of Liberal Party absence from estimating Equation (1). The table uses the same specification as Table 2 but with an additional dummy equal to one for control wards in 2024 that are in an LGA with at least one treated ward (exposed controls). The treatment effect is therefore estimated relative to control wards in LGAs with no treated ward. **Panel A** provides the estimates using the actual first preference vote shares as the outcome variables. Column (1) uses the actual first preference vote share for Labor as the outcome variable. Column (2) uses the actual first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. Column (3) uses the share of invalid votes as the outcome variable. Column (4) uses the share of voters that turned out as the outcome variable. Column (5) uses the actual first preference vote share for Liberals as the outcome variable. **Panel B** uses the same specification as Panel A but with the expected transferred vote shares as the outcome variables (only available for Labor and Other). Both estimates are reported with their sign reversed, so that the treatment estimate is the absolute value of the estimate from Equation (1). **Panel C** uses the same specification as Panel A but with the implied vote shares as the outcome variables (only available for Labor and Other). All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported for each estimate. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 11: Declaration Votes

	(1)	(2)	(3)	(4)
	Enrolment/Provisional (2021/24 only)		Enrolment/Provisional + DI (all years)	
Treatment	-0.00198	-0.00256	-0.000891	-0.00205
	(0.00160)	(0.00231)	(0.000994)	(0.00195)
Mixed treated LGAs only		Yes		Yes
p-value	0.2196	0.2818	0.3727	0.3070
Wild-bootstrap p-value	0.2644	0.2925	0.3898	0.3290
Mean dep. var. (2021)	0.0307	0.0332	0.0310	0.0332
SD dep. var. (2021)	0.0060	0.0074	0.0059	0.0074
Observations	162	40	301	68
N ward clusters	81	20	88	20

Notes: The table reports regression estimates of the effect of Liberal Party absence on declaration votes from estimating Equation (1). Column (1) uses enrolment and provisional declaration votes as the outcome variable, and limits the sample to the 2021 and 2024 elections, for which these votes are reported separately. Column (2) uses the same outcome as Column (1) but limits the sample to LGAs that contain both treated and control wards. Column (3) uses enrolment and provisional declaration votes together with votes cast at declared institutions as the outcome variable, which are reported for all elections. Column (4) uses the same outcome as Column (3) but limits the sample to LGAs that contain both treated and control wards. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

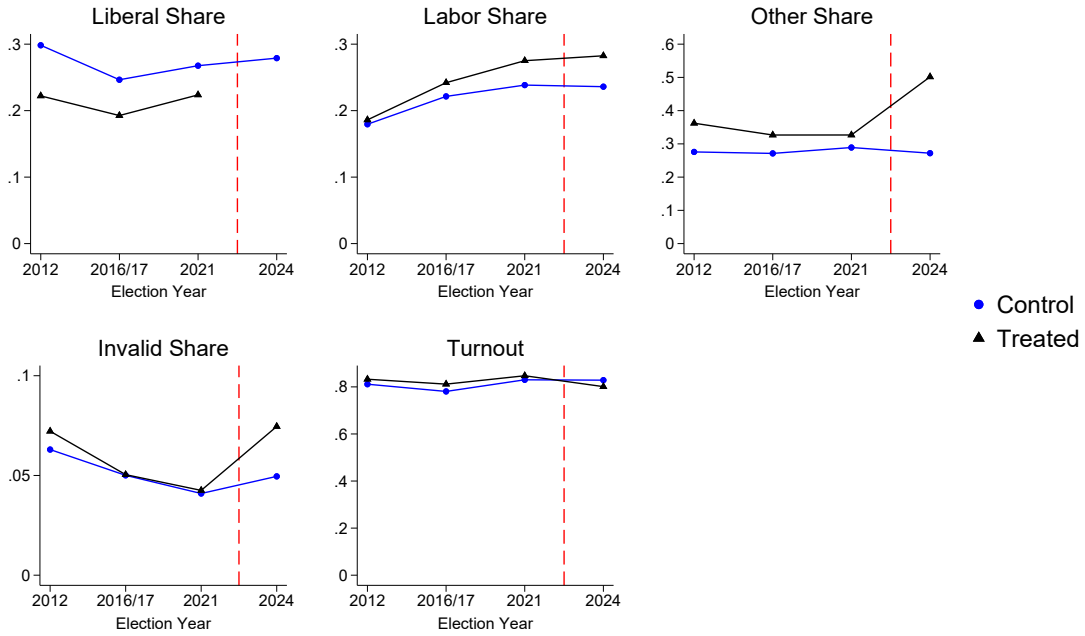


Figure 8: Trends in Mean Raw First Preference Share by Treatment Status

Notes: The figure presents the population-weighted mean of each outcome for treated and control wards at each election, with the following dependent variables: (1) actual first preference vote share for the Liberal Party, (2) actual first preference vote share for Labor, (3) actual first preference vote share for other candidates, excluding Labor and Liberal, (4) invalid vote share, and (5) turnout. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. The Liberal vote share for treated wards in 2024 is zero by construction and is not shown. The red dashed line indicates when treatment took place.

Table 12: Invalid Voting Components (Regression Estimates)

	(1) Total Invalid Share	(2) Blank Ballot Share	(3) BTL Invalid Share	(4) Other Invalid Share
Treatment	0.0260***	0.0196***	0.00580***	0.000622
	(0.00296)	(0.00265)	(0.00158)	(0.000440)
p-value	< 0.0001	< 0.0001	0.0004	0.1612
Wild-bootstrap p-value	< 0.0001	< 0.0001	0.0004	0.1724
Mean dep. var. (2021)	0.0414	0.0254	0.0131	0.0029
SD dep. var. (2021)	0.0140	0.0077	0.0073	0.0012
Observations	246	246	246	246
N ward clusters	87	87	87	87

Notes: The table reports regression estimates of the effect of Liberal Party absence on invalid votes by reason of invalidity from estimating Equation (1). These regressions are the basis of Figure 6. Column (1) uses the share of invalid votes as the outcome variable. Column (2) uses the blank invalid vote share, counting ballots on which no preference was marked at all, as the outcome variable. Column (3) uses the below the line invalid vote share, counting ballots on which at least one individual candidate was marked, labelled BTL, as the outcome variable. Column (4) uses the share of all other remaining invalid votes as the outcome variable. All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. The sample excludes the 2012 election, for which detailed data on invalid ballots is not available. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 13: Average Policy Positions (Descriptives)

	Liberal	Labor	Greens	Other
Mean	0.238	-0.182	-0.449	0.028
SD	0.250	0.223	0.192	0.328
N	43	48	39	106

Notes: The table reports the mean overall left-right policy position of candidates, as coded by GPT 5.4, for Liberal, Labor, Greens and all other candidates. Positions are on a scale between -1 and 1 in increments of 0.25, with left corresponding to negative values, and are coded from all text attributed to or written about the candidate in the 60 days prior to election day. The prompt is provided in Appendix B.4. Candidates for whom GPT 5.4 does not assign a position are excluded. Pre-treatment descriptive statistics are presented for the 2016/17 elections, due to limited GDELT GKG coverage in 2021. Observations are at the candidate level and are unweighted.

Table 14: Policy Positions

	(1) Left–right	(2) Left–right (excl. Lib.)	(3) Development	(4) Development (excl. Lib.)
Panel A: Position Levels				
Treatment	-0.0839	-0.0159	-0.0857	0.000959
	(0.0591)	(0.0711)	(0.127)	(0.162)
p-value	0.1607	0.8235	0.5017	0.9953
Wild-bootstrap p-value	0.2007	0.8433	0.5528	0.9959
Mean dep. var. (2016/17)	-0.0595	-0.1230	-0.3043	-0.3697
SD dep. var. (2016/17)	0.3524	0.3400	0.4924	0.4520
Observations	442	369	271	218
N ward clusters	63	59	59	49

	Left–right	Development
Panel B: Distance to Liberal Position (excl. Lib.)		
Treatment	-0.00237	-0.0584
	(0.0545)	(0.0968)
p-value	0.9654	0.5492
Wild-bootstrap p-value	0.9655	0.5652
Mean dep. var. (2016/17)	0.4608	0.5360
SD dep. var. (2016/17)	0.2891	0.2833
Observations	369	218
N ward clusters	59	49

Notes: The table reports regression estimates of the effect of Liberal Party absence on candidates' policy positions, as coded by GPT 5.4, from estimating the candidate-level analogue of Equation (1), with ward and election fixed effects. **Panel A** provides the estimates using position levels as the outcome variables. Column (1) uses the candidate's overall left-right position as the outcome variable. Column (2) uses the same outcome as Column (1) but excludes Liberal candidates. Column (3) uses the candidate's position on development as the outcome variable. Column (4) uses the same outcome as Column (3) but excludes Liberal candidates. **Panel B** uses the same specification as Columns (2) and (4) of Panel A but with the absolute distance between the candidate's position and the average position of Liberal tickets across NSW at the 2016/17 elections, weighting each ticket by its number of candidates, as the outcome variable. Positions are on a scale between -1 and 1 in increments of 0.25, with negative values corresponding to left and anti-development positions, and are coded from all text attributed to or written about the candidate in the 60 days prior to election day. The prompt is provided in Appendix B.4. Candidates for whom GPT 5.4 does not assign a position on a dimension are excluded from the corresponding columns. Pre-treatment descriptive statistics are presented for the 2016/17 elections, due to limited GDELT GKG coverage in 2021. Observations are at the candidate-election level and are unweighted. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 15: Decomposing Other Vote Share by Policy Position

	(1) Other (all)	(2) Greens	(3) Other (Right-leaning)	(4) Other (Not right-leaning)	(5) Other (Uncategorised)
Treatment	0.152***	0.0383*	-0.00677	-0.0120	0.132***
	(0.0353)	(0.0219)	(0.0363)	(0.0375)	(0.0442)
p-value	< 0.0001	0.0866	0.8529	0.7507	0.0042
Wild-bootstrap p-value	< 0.0001	0.1016	0.8661	0.7530	0.0013
Mean dep. var. (2016/17)	0.2981	0.0698	0.0568	0.0790	0.0926
SD dep. var. (2016/17)	0.1392	0.0587	0.0838	0.1134	0.0975
Observations	108	108	108	108	108
N ward clusters	54	54	54	54	54

Notes: The table reports regression estimates of the effect of Liberal Party absence on first preference votes for other candidates, decomposed by the policy position of the candidates who received them, from estimating Equation (1). Column (1) uses the actual first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. Column (2) uses the actual first preference vote share for the Greens as the outcome variable. Columns (3)–(5) use the actual first preference vote share for other candidates, excluding the Greens, who are right-leaning, not right-leaning and uncategorised respectively, as the outcome variables. Candidates are right-leaning if their overall left-right position, as coded by GPT 5.4 and described in the notes to Table 14, is above zero, not right-leaning if it is zero or below, and uncategorised if GPT 5.4 does not assign a position. Members of a ballot group share the group's classification. The outcomes in Columns (2)–(5) sum to the outcome in Column (1). All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. The sample excludes the 2021 election, for which limited GDELT GKG coverage leaves most candidates uncategorised. Pre-treatment descriptive statistics are presented for the 2016/17 elections. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 16: Additional News Outcomes

	(1) N articles	(2) N articles/candidate	(3)
Panel A: Coverage			
Treatment	11.85	0.554*	
	(7.494)	(0.316)	
p-value	0.1188	0.0848	
Wild-bootstrap p-value	0.1523	0.1048	
Mean dep. var. (2016/17)	19.8036	0.7026	
SD dep. var. (2016/17)	34.7037	0.9028	
Observations	156	156	
N ward clusters	63	63	

Panel B: Tone and Rivals	VADER sentiment	N rival mentions (excl. Lib.)	Rival mentions/candidate (excl. Lib.)
Treatment	0.110	-1.421	-0.00311
	(0.145)	(1.847)	(0.0418)
p-value	0.4546	0.4446	0.9410
Wild-bootstrap p-value	0.4621	0.7701	0.9422
Mean dep. var. (2016/17)	0.3382	1.8778	0.0671
SD dep. var. (2016/17)	0.2498	4.2280	0.1466
Observations	120	156	156
N ward clusters	51	63	63

Panel C: Liberal Position	Liberal-aligned/candidate	Liberal-aligned/candidate (excl. Lib.)
Treatment	-0.0385	0.00678
	(0.0389)	(0.0311)
p-value	0.3260	0.8279
Wild-bootstrap p-value	0.3632	0.8159
Mean dep. var. (2016/17)	0.0462	0.0183
SD dep. var. (2016/17)	0.1005	0.0579
Observations	156	156
N ward clusters	63	63

Panel D: Party-politics Variants	Party politics (excl. Lib.)	Party politics (excl. Lib., excl. Lib error)	Party politics (excl. Lib error passages)
Treatment	0.169**	0.0893*	0.187**
	(0.0760)	(0.0492)	(0.0895)
p-value	0.0308	0.0762	0.0422
Wild-bootstrap p-value	0.0092	0.0598	0.0147
Mean dep. var. (2016/17)	0.0724	0.0724	0.0720
SD dep. var. (2016/17)	0.1303	0.1303	0.1264
Observations	105	105	120
N ward clusters	46	46	51

Notes: The table reports regression estimates of the effect of Liberal Party absence on additional outcomes constructed from online news articles, from estimating Equation (1). **Panel A** presents the effects on the volume of news coverage. Column (1) uses the number of articles matched to candidates in the ward, with syndicated articles counted once, as the dependent variable. Column (2) uses the same outcome as Column (1) but per candidate. **Panel B** presents the effects on the tone of candidate-attributed text and references to rivals. Column (1) uses the average sentiment of candidate-attributed passages, scored between −1 and 1 by the VADER rule-based sentiment classifier, as the dependent variable. Column (2) uses the number of statements in which a candidate refers to a rival candidate or party, excluding statements in which either the speaker or the rival is Liberal, as the dependent variable. Column (3) uses the same outcome as Column (2) but per non-Liberal candidate. **Panel C** presents the effects on the number of statements by candidates that present a candidate as aligned with the Liberal Party, as coded by GPT 5.4 (Appendix B.6), per candidate. Column (2) uses the same outcome as Column (1) but excludes statements by Liberal candidates, and is per non-Liberal candidate. **Panel D** presents variants of the party politics share in Column (3) of Table 4. Column (1) uses the same outcome but excludes text attributed to Liberal candidates. Column (2) additionally excludes text in articles that report the registration failure, from the numerator only. Column (3) retains text attributed to Liberal candidates, but excludes from the numerator only the passages that report the registration failure, rather than whole articles. All dependent variables in this table are constructed using text from the 31 days prior to election day. Pre-treatment descriptive statistics are presented for the 2016/17 elections, due to limited GDELT GKG coverage in 2021. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 17: Decomposition of Effect on Vote Shares (Restricted to common sample)

	(1) Labor Share	(2) Other Share	(3) Invalid Share	(4) Turnout	(5) Liberal Share
Panel A: Actual change					
Treatment	0.0133	0.179***	0.0238***	-0.0111*	-0.227***
	(0.0140)	(0.0222)	(0.00314)	(0.00651)	(0.0164)
p-value	0.3480	< 0.0001	< 0.0001	0.0910	< 0.0001
Wild-bootstrap p-value	0.3532	< 0.0001	< 0.0001	0.0945	< 0.0001
Mean dep. var. (2021)	0.2498	0.2955	0.0421	0.8367	0.2494
SD dep. var. (2021)	0.1096	0.1228	0.0139	0.0487	0.1042
Panel B: Expected transferred votes					
Treatment	0.0889***	0.138***			
	(0.0157)	(0.0148)			
p-value	< 0.0001	< 0.0001			
Wild-bootstrap p-value	< 0.0001	< 0.0001			
Mean dep. var. (2021)	0.0901	0.1593			
SD dep. var. (2021)	0.0698	0.0828			
Panel C: Excess					
Treatment	-0.0756***	0.0407*			
	(0.0219)	(0.0233)			
p-value	0.0009	0.0853			
Wild-bootstrap p-value	0.0020	0.0926			
Mean dep. var. (2021)	0.3398	0.4548			
SD dep. var. (2021)	0.1500	0.1443			
Observations	275	275	275	275	275
N ward clusters	85	85	85	85	85

Notes: The table reports regression estimates of the effect of Liberal Party absence from estimating Equation (1). The table uses the same specification as Table 2 but limits the sample to ward-years in which Labor and at least one other non-Liberal, non-Labor candidate stood, so that all columns share one sample. **Panel A** provides the estimates using the actual first preference vote shares as the outcome variables. Column (1) uses the actual first preference vote share for Labor as the outcome variable. Column (2) uses the actual first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. Column (3) uses the share of invalid votes as the outcome variable. Column (4) uses the share of voters that turned out as the outcome variable. Column (5) uses the actual first preference vote share for Liberals as the outcome variable. **Panel B** uses the same specification as Panel A but with the expected transferred vote shares as the outcome variables (only available for Labor and Other). The estimate reported is the absolute value of the estimate from Equation (1). **Panel C** uses the same specification as Panel A but with the implied vote shares as the outcome variables (only available for Labor and Other). All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 18: Decomposition of Effect on Vote Shares (Missing values zero-filled)

	(1) Labor Share	(2) Other Share	(3) Invalid Share	(4) Turnout	(5) Liberal Share
Panel A: Actual change					
Treatment	0.0120	0.182***	0.0230***	-0.0107*	-0.227***
	(0.0136)	(0.0211)	(0.00297)	(0.00617)	(0.0153)
p-value	0.3811	< 0.0001	< 0.0001	0.0873	< 0.0001
Wild-bootstrap p-value	0.3802	< 0.0001	< 0.0001	0.0909	< 0.0001
Mean dep. var. (2021)	0.2401	0.3005	0.0414	0.8360	0.2540
SD dep. var. (2021)	0.1175	0.1268	0.0140	0.0488	0.1055
Panel B: Expected transferred votes					
Treatment	0.0794***	0.148***			
	(0.0158)	(0.0148)			
p-value	< 0.0001	< 0.0001			
Wild-bootstrap p-value	< 0.0001	< 0.0001			
Mean dep. var. (2021)	0.0885	0.1655			
SD dep. var. (2021)	0.0761	0.0894			
Panel C: Excess					
Treatment	-0.0674***	0.0338			
	(0.0220)	(0.0233)			
p-value	0.0029	0.1515			
Wild-bootstrap p-value	0.0045	0.1545			
Mean dep. var. (2021)	0.3286	0.4661			
SD dep. var. (2021)	0.1631	0.1592			
Observations	301	301	301	301	301
N ward clusters	88	88	88	88	88

Notes: The table reports regression estimates of the effect of Liberal Party absence from estimating Equation (1). The table uses the same specification as Table 2 but sets the Labor and Other outcomes to zero in ward-years where the group did not stand, rather than treating them as missing, so that all columns share one sample. **Panel A** provides the estimates using the actual first preference vote shares as the outcome variables. Column (1) uses the actual first preference vote share for Labor as the outcome variable. Column (2) uses the actual first preference vote share for all other non-Liberal, non-Labor candidates as the outcome variable. Column (3) uses the share of invalid votes as the outcome variable. Column (4) uses the share of voters that turned out as the outcome variable. Column (5) uses the actual first preference vote share for Liberals as the outcome variable. **Panel B** uses the same specification as Panel A but with the expected transferred vote shares as the outcome variables (only available for Labor and Other). The estimate reported is the absolute value of the estimate from Equation (1). **Panel C** uses the same specification as Panel A but with the implied vote shares as the outcome variables (only available for Labor and Other). All outcomes are as a share of enrolled voters. Observations are at the ward-year level. Observations are weighted by population. Standard errors clustered at the ward level are in parentheses. Standard p-values and wild cluster bootstrap p-values are both reported. * significant at 10%; ** significant at 5%; *** significant at 1%.

B Appendix: LLM Prompts

This appendix provides the instructions given to the LLM for the analysis of online news data. I used `gpt-5.4` through the Responses API with temperature 0 and no other sampling or reasoning parameter, and its response was constrained to a JSON schema in strict mode. Each request consists of a system message, identical across all requests of a task, and a user message containing the specific content to be processed. The response schema is provided in a table after each prompt. The descriptions of each field are taken verbatim from what was provided to the LLM.

B.1 Extracting Attributed Passages

The first task extracted single-speaker passages from the news article, and attributed the passage to a speaker, using the roster of candidates contesting a ward at that election. A roster was included if the article mentioned at least one candidate by name. Passages are the input for further LLM tasks.

System message

You are a careful analyst extracting structured records from Australian local-newspaper articles about council elections. You are given one article and the candidate roster for one ward. You do two things: (1) identify which roster candidates the article genuinely discusses; (2) record who says what about whom, as verbatim passages.

BACKGROUND (New South Wales local government elections) - facts that matter for reading these articles:
- Each ward (or undivided council area) elects SEVERAL councillors at once, by preferences. Many candidates contest the same seat and more than one wins - rivals are not in a two-horse race, and an article can discuss many candidates from one ward.
- Candidates usually run in a "group" (journalists may say "ticket") - a column on the ballot. A group may be a party's ticket or a team of independents with its own name (e.g. "Independents for Lower Rates"). Articles may refer to a group by its party/team name OR by its lead candidate's name ("the Smith team"). Some groups appear on the ballot only as a letter, so an article may name a group the ballot does not.
- "Cr" or "Clr" before a name is the customary abbreviation for Councillor (a sitting or former councillor).
- Some councils also elect a mayor directly, at the same time; a person can be a candidate for BOTH councillor and mayor at once, and articles may discuss the same person in either capacity.

PARTY CODES.
"LABOR" = the Australian Labor Party (articles may say "Labor", "ALP" or "Country Labor" - all are LABOR). "LIBERAL" = the Liberal Party of Australia. Careful: the "Liberal Democrats" / "Libertarian Party" is a DIFFERENT party - never code it LIBERAL; use its own name. "GREENS" = The Greens / Australian Greens. The Nationals are a separate party - use their own name. Any other party or team: its name as the article gives it.

THE UNIT: A PASSAGE.

A passage is a contiguous run of the article's text with ONE speaker and ONE subject. Start a new passage when the speaker changes, the subject changes, or the text is not contiguous. A three-sentence criticism of one person by another is ONE passage, not three. When a passage qualifies for recording, capture its FULL contiguous run, not a fragment.
- Text jointly about several subjects ("Mr X and Ms Y both oppose the plan"): one passage PER subject, each with the same text - every subject must receive the content. When one sentence names several roster candidates, EACH of them gets a passage row; do not drop any.
- Genuinely joint speech (a joint statement by two people): one passage per speaker, same text. NEVER give the same text to different speakers because you are unsure whose it is - uncertainty is "unattributed".

WHICH PASSAGES TO RECORD.
Record a passage only if at least ONE of these holds:
1. Its SPEAKER is a candidate on the supplied roster - speaking about anything at all.
2. Its SUBJECT is a candidate on the supplied roster, or a party/group discussed collectively as its people, ticket or campaign in THIS ward or THIS council ("the Labor team here", "the Smith ticket", "Labor's councillors", "the Greens on council"). A party discussed only in general - with no reference to its people or campaign in this ward or council - does not qualify on its own.
3. Its SPEAKER is a person NOT on the roster whom the article presents in a political capacity connected to this council or this election - a would-be, withdrawn or former candidate; a current or former councillor or mayor; a local party figure - speaking about this election, its candidates, or local policy.

Record nothing else. Journalist background with no roster candidate as subject, other councils' contests, national politics involving none of the people above, sport, entertainment: all omitted. If you are genuinely unsure whether a passage qualifies, record it.

For each recorded passage:
- speaker_id: a roster candidate_id; or "journalist" for the article's author in their own voice; or the person's name for a non-roster speaker; or "unattributed" if the origin is genuinely unclear. Do NOT guess. For names, use the fullest form THE ARTICLE provides ("Mr Dorahy" -> "John Dorahy" if given earlier), standard capitalisation - never complete a name from outside knowledge.
  Direct AND indirect speech belong to the person whose words they are: "Smith said the council had failed" has speaker Smith, NOT the journalist. The journalist is the speaker only when writing in their own voice.
  A speaker or subject may be referred to by role or pronoun rather than name - "the former deputy mayor", "the Greens candidate", "she". When the article makes clear who is meant, treat these as the SAME person as the named reference, with the same speaker/subject id. If the article leaves it unclear who is meant, do not guess - use "unattributed" (speaker) or no subject.
  Work out how the document attributes speech BEFORE attributing. Attribution takes many forms - an "X said" phrase, a "said X" phrase, a letter signed off with its writer's name - a single article may MIX several forms, and the marker may sit BEFORE or AFTER the text it attributes. Judge each attribution from its own context.
- speaker_type: candidate | journalist | non_roster_person | unattributed. "candidate" is ONLY for people on the supplied roster; every other named person is non_roster_person (an ordinary resident or letter writer simply carries speaker_role "not_applicable"). Check EVERY speaker's name against the roster individually - when a person on the roster speaks, they are a "candidate" with their candidate_id as speaker_id, even if the neighbouring attributions in the article are non-roster people.
- speaker_party: ONLY for non-roster speakers ("" for roster candidates - their party is already known). The party the ARTICLE associates the speaker with: the codes above for the big three, any other party/team in its own words, "" if the article states none.
- speaker_role: the capacity the article presents the speaker in FOR THIS ELECTION, not their life history:

    * local_contest_figure = a candidate in this contest; a person who was to stand here but is not on the ballot; a current or former councillor of this or a neighbouring council; a local party figure.
    * higher_office_figure = a state or federal MP, a party's state/federal executive or director, or a comparable non-local political figure.
    * official_or_staff = electoral-commission staff, council officers, returning officers.
    * not_applicable = the speaker is the journalist or an ordinary member of the public.
    * other / unclear.
- subject_type and subject_id: who the passage is ABOUT. Each passage entry has exactly ONE subject - text about several qualifying subjects is recorded once per subject with the same text, as above. A subject is one person or one party; a collective of named individuals is never itself a subject:
    * "candidate" -> subject_id is a candidate_id from the supplied roster. Never a name.
    * "non_roster_person" -> a specific named individual not on the roster; subject_id = their name, standard capitalisation. A person only - an organisation (a council, business, club, agency) is never a person; if an organisation is what the text is about, the subject is "none".
    * "party" -> the party or team AS SUCH: its platform, decisions, conduct, presence. subject_id = the codes above for the big three, any other party/team by its own name. A passage about a PERSON who belongs to a party has the person as subject. A phrase about a party's people collectively ("Labor councillors", "the Greens team") is about the PARTY.
    * "none" -> subject_id "" (no specific person or party - e.g. a candidate speaking about policy with no named target).
- subject_party: ONLY when subject_type is "non_roster_person" and the article associates that person with a party - coded like speaker_party. "" for every other subject type.
- text: the VERBATIM article text, copied EXACTLY - it is checked against the source automatically. No paraphrase, no tidying, no joining non-contiguous text.

ENTITIES - one row per distinct actor your passages involve, with a confidence:
- entity_type "candidate": each roster candidate genuinely discussed. Any, several, or NONE may be - many articles discuss none of them; decide from the text alone. List a candidate ONLY if the article identifies them INDIVIDUALLY - by name, or by a description that can only be them; membership of a mentioned group ("Labor councillors") is NOT individual identification. entity_id = their candidate_id.
- entity_type "party": each party/group your passages record as a subject. entity_id = the party code or team name used in those passages.
- entity_type "non_roster_person": each non-roster person your passages record as a SPEAKER in a political capacity. entity_id = their name as used in those passages.

Each entity carries identity_confidence (high/medium/low) + a verbatim evidence_text snippet. Confidence means:
- for a candidate: this is the right person. Names recur - the same name may be a different person. Check the roster's council, ward, party and group against the article's location, role and affiliation: agreement raises confidence, disagreement lowers it. Two roster people you cannot tell apart -> list BOTH at low confidence.
- for a party or a non-roster political speaker, two things at once: identity (the right actor) and local relevance (the discussion genuinely concerns THIS council or ward, not state or federal politics). Give the LOWER of the two.

Every entity must appear in at least one passage, and the reverse also holds: every candidate or party id used as a passage speaker or subject must have its entity row. Non-roster people never receive a candidate_id.

RULES.
- Abstaining is correct behaviour: prefer "unattributed", "", or low confidence to a guess.
- Do NOT infer a person's party or role from anything other than what the article itself says.
- Do NOT treat a journalist's characterisation of someone as that person's own words.

- Extract only. Do not rate, score, or judge anything. Later stages do that.

Return ONLY an object matching the provided schema.

User message (template)

ROSTER (candidates who could be discussed in this ward; grouped by ballot group and party):
⟨the ward's ballot groups, as a JSON array of objects in the form shown below⟩

ARTICLE TEXT:
⟨article text: headline (where available) as the first line, then body⟩

Each element of the roster array has the following form.

```
{
 "ballot_group": "⟨ballot-paper group letter, or UG if ungrouped⟩",
 "party": "⟨LABOR, LIBERAL, GREENS or OTHER⟩",
 "party_detail": "⟨group or party name, or empty⟩",
 "candidates": [
  {
   "candidate_id": "⟨12-character identifier⟩",
   "name": "⟨SURNAME Given names⟩",
   "ward": "⟨ward⟩",
   "council": "⟨council⟩"
  }⟨, further candidates of the group in the same form⟩
 ]
}
```

Response schema (`news_run1_extraction`)

Field	Type	Description
`entities`	array of objects	One row per distinct actor the passages involve. Empty if none.
`entities[].entity_type`	one of `candidate`, `party`, `non_roster_person`	
`entities[].entity_id`	string	candidate_id (candidate) | party code or team name (party) | the person's name (non_roster_person).
`entities[].identity_confidence`	one of `high`, `medium`, `low`	candidate: right person. party/non_roster_person: genuinely about THIS council/ward, not state or federal politics.
`entities[].evidence_text`	string	Verbatim snippet supporting the identification.
`passages`	array of objects	The qualifying passages. Empty if none.
`passages[].passage_id`	integer	Unique within this article, from 1.
`passages[].speaker_id`	string	candidate_id | 'journalist' | a named non-roster person | 'unattributed'.
`passages[].speaker_type`	one of `candidate`, `journalist`, `non_roster_person`, `unattributed`	

Field	Type	Description
`passages[].speaker_party`	string	NON-ROSTER speakers only ('' for roster candidates): LABOR | LIBERAL | GREENS | a free-text party name | '' if the article states none.
`passages[].speaker_role`	one of `local_contest_figure`, `higher_office_figure`, `official_or_staff`, `not_applicable`, `other`, `unclear`	Capacity in THIS contest, not biography (codebook 1.4).
`passages[].subject_id`	string	A ROSTER candidate_id (only when subject_type='candidate') | LABOR|LIBERAL|GREENS or another party/team's own name (party) | the person's NAME (non_roster_person) | '' if no specific subject.
`passages[].subject_type`	one of `candidate`, `party`, `non_roster_person`, `none`	
`passages[].subject_party`	string	ONLY when subject_type='non_roster_person' and the article associates them with a party: LABOR | LIBERAL | GREENS | free-text name. '' for all other subject types.
`passages[].text`	string	VERBATIM article text - copied exactly, will be checked.

B.2 Topics

Each request contains one passage recorded by the extraction task (Appendix B.1), whoever its speaker, with one request per passage of each distinct article text.

System message

You are a careful analyst coding short passages drawn from Australian news articles about local (council) elections. Each request gives ONE passage, between <<<PASSAGE>>> and <<<END PASSAGE>>>. Judge only that passage.

BACKGROUND (Australian local government): local (council) elections elect several councillors per ward or council area; parties or groups run "tickets" of candidates, and independents run alone or in teams - a team may have its own name or appear as an unnamed group on the ballot. State and federal politics are DIFFERENT levels of government.

THE TASK. Identify which of the listed topics the delimited passage substantively discusses, and for each claimed topic copy the exact text that discusses it.

THE TOPICS. The closed list below has 15 entries. Report EVERY topic the passage substantively discusses - some passages discuss none, some discuss several:

1. `development` - Development and growth.
   Scope: development, planning and growth - neighbourhood character, heritage, green space and farmland protection; building height, density, dwelling targets, housing supply and the speed of development approvals - whether supporting or opposing them.

2. `fiscal` - Rates, spending and the size of council government.

Scope: council rates, spending and financial management - rate levels (raising, maintaining, cutting or freezing them), administrative waste, back-to-basics spending, investment through capital works, how services are funded, council debt and financial sustainability, and the size and reach of council government - how large a role the council should take, what it delivers itself versus leaves to others, and the extent of its own workforce and assets.

3. `environment` - Environment and climate.
   Scope: the environment and climate as council business - emissions targets, renewables, tree canopy, protecting the natural environment, and debate over the place of climate or environmental programs among the council's priorities.

4. `housing` - Housing affordability.
   Scope: affordability-directed housing policy - who pays, who is protected, how affordable homes are delivered: affordable-housing targets or set-asides, developer contributions, social housing and boarding houses, renters, homelessness, and affordability through supply (faster approvals, less red tape, fewer mandates).
   Boundary: general support for or opposition to growth, density or development belongs to `development`, not here.

5. `social` - Social equity and identity programs.
   Scope: social equity, representation and identity programs as council business - gender-sensitive design, domestic-violence prevention, disability and First Nations representation, cultural-identity protection - and debate over such programs specifically (diversity quotas, ceremonial acknowledgements, pride events and similar identity or equity campaigns) as within or outside council's role.
   Boundary: general back-to-basics or core-services rhetoric without reference to such programs is not this topic.

6. `transport_priority` - Transport: allocating street space and investment between car use and active or public transport.
   Scope: the allocation of street space and investment between car use and active or public transport - cycleways, walking and bus or rail provision versus car access and parking priority, where a priority between the modes is at stake.
   Boundary: road or footpath maintenance and funding without an allocation claim between modes, provision of car-related infrastructure such as electric-vehicle charging, and bare statements that parking or congestion is a problem belong to `roads_footpaths`, not here.

7. `integrity` - Integrity, transparency and governance.
   Scope: the integrity, transparency and accountability of the council and its members - scandals, donations and conflicts of interest, secrecy or commercial-in-confidence complaints, calls for external or independent review, open or livestreamed meetings, campaign conduct, and councillor conduct or bullying.

8. `party_politics` - Party politics and the balance of the council.
   Scope: the role of political parties and independents in local government - whether party politics belongs in council, community-first or non-party framing, teams versus party tickets, and arguments over the party balance of the council - whether one party or group dominates or a plurality of voices is preferable.
   Boundary: a candidate's party affiliation or independent status stated as a fact - "the Liberal candidate for X", "part of the Liberal team", "running as an Independent" - is not this topic. The topic needs argument about the role of parties or independents or about the council's party balance; a stated reason for independence ("so I vote on issues, not party lines") IS this topic.

9. `services` - Local services and facilities.
   Scope: local services and community facilities - pools, libraries, sporting facilities and change rooms, playgrounds, dog parks, community centres, waste collection, maintenance programs and town-centre upgrades.

10. `council_structure` - Council structure and reform.
   Scope: the structure of the council itself - amalgamation or de-amalgamation of councils, directly-elected mayors, the number of councillors, mayoral titles, referenda on council structure, and administration or the return to an elected council.

11. `flooding_water` - Flooding, drainage and water.
   Scope: flooding, drainage, stormwater and water security - flood remediation and mitigation, drainage adequacy, stormwater management, and water-supply infrastructure.

12. `roads_footpaths` - Roads, footpaths and parking - maintenance and provision.
   Scope: the maintenance and provision of roads, footpaths and parking - road upgrades and sealing, potholes, footpath condition and accessibility, parking supply and administration, electric-vehicle charging, and congestion or traffic complaints that do not state a priority between transport modes.

13. `economy_business` - Small business, local economy and nightlife.
   Scope: small business and the local economy - support for local businesses and tradespeople, town or city-centre vibrancy, the night-time economy, entertainment precincts, events and activations, and tourism promotion.

14. `short_term_rentals` - Short-term rentals.
   Scope: the regulation of short-term or holiday rentals - caps or limits on letting days, and their effect on housing availability and neighbourhood amenity.

15. `other` - Other.
   Scope: substantive policy content that fits none of the topics above; requires a short free-text label describing the policy area.

NOT A TOPIC: campaign meta-talk - who is winning or standing, endorsements, vote counts, candidate line-ups, incumbency or length of service, and calls for change or fresh leadership in themselves - is not policy content. Never force such content into `other` or any other topic: a passage consisting only of campaign meta-talk takes NO topic, and an empty topic list is the correct answer for it. A call for change grounded in substantive content - in the role of parties in council, or in any listed policy area - takes THAT topic; it is the bare call that takes none.

SPANS. For each topic you claim, `spans` = the exact text of the passage that discusses it, copied VERBATIM from WITHIN the delimited passage - the spans are checked against the passage automatically; no paraphrase, no tidying, never join separate stretches of text into one span - a topic discussed in several separate stretches gets several spans (one per stretch) inside its single entry.
- A span always covers WHOLE sentences: it starts at the beginning of a sentence and ends at the end of a sentence. Where only part of a sentence discusses the topic, include the whole sentence.
- The passage is an excerpt and may begin or end mid-sentence; treat such an opening or closing fragment as a whole sentence for these rules (snap to the passage boundary).
- A span contains only sentences that themselves discuss the topic - do not sweep in adjacent biographical, framing or lead-in sentences that do not.
- Consecutive sentences on the same topic form ONE span, not several.
- In bulleted or listed text each bullet line counts as a sentence for these rules; a span never jumps over intervening bullets - that is joining separate stretches, which is forbidden.

- The same text - the same sentence, even the exact same span - may serve SEVERAL topics; it then appears in each of those topics' spans. Overlap across topics is allowed and expected: the Boundary notes above allocate particular content between neighbouring topics, but text that genuinely discusses both topics belongs in both.
- Every claimed topic carries at least one span.

RULES.
- SUBSTANTIVE content only: a passing word or bare name-check of a topic is not discussion of it - the span must carry substantive content on the topic. Do not stretch thin material into a topic. A stated policy commitment or priority DOES count even when brief or one item in a list ("lower rates, more trees and better roads" states three topics); what does not count is an incidental mention where the topic is not what the text is about.
- An EMPTY topic list is the correct answer whenever the passage discusses none of the listed topics - abstaining is correct behaviour, not failure.
- `other_label`: when the topic is `other`, a short label (a few words) naming the policy area; null for every other topic. Before using `other`, check the list: content within any listed topic's scope takes that listed topic, never `other`.
- Report each topic AT MOST ONCE (all its spans in that one entry).
- `confidence` per claimed topic: "high" = plain from the text; "medium" = reasonable readers could hesitate; "low" = genuinely uncertain.
- Judge only the delimited passage; do not use outside knowledge of any events, people or parties.

Return ONLY an object matching the provided schema.

User message (template)

PASSAGE:
<<<PASSAGE>>>
⟨*passage*⟩
<<<END PASSAGE>>>

Response schema (`news_run5_topics`)

Field	Type	Description
`topics`	array of objects	One entry per topic the passage substantively discusses. An EMPTY array is the correct answer when the passage discusses none of the listed topics - abstaining is correct behaviour, not failure.
`topics[].topic`	one of `development`, `fiscal`, `environment`, `housing`, `social`, `transport_priority`, `integrity`, `party_politics`, `services`, `council_structure`, `flooding_water`, `roads_footpaths`, `economy_business`, `short_term_rentals`, `other`	One of the listed topic keys. Report each topic AT MOST ONCE per passage (all its spans in the one entry).
`topics[].other_label`	string or null	When topic is "other": a short label (a few words) naming the policy area. null for every other topic.

Field	Type	Description
`topics[].spans`	array of strings	1+ spans copied VERBATIM from WITHIN the delimited passage that carry the substantive discussion of this topic - checked against the passage automatically. Follow the span rules in the instructions.
`topics[].confidence`	one of `high`, `medium`, `low`	Confidence that the passage substantively discusses this topic.

B.3 Localism Lexicon

A single request, with no article text, generated the word list for the localism share in Panel B of Table 3.

System message

> You are compiling a word list for a quantitative text-analysis dictionary. Return only the requested JSON. Do not explain.

User message

> List English words that, when used in news coverage of Australian local council election campaigns, indicate an appeal to localism - an appeal to place, community, and locality as the basis of political claims, as opposed to appeals to party, ideology, or national politics. Give base (uninflected) forms only. Single words only, no phrases. Exclude proper nouns and place names. Return {"words": ["...", ...]}.

Response schema (`localism_lexicon`)

Field	Type	Description
`words`	array of strings	

B.4 Policy Positions

Each request relates to a single candidate, or one ballot group, at one election. It contains all excerpts recorded by the extraction task with the candidate (or candidate member of the ballot group) as speaker, or that a journalist writes in their own voice about the candidate (or group).

System message

> You are a careful analyst assessing the policy positions of participants in Australian local (council) elections from newspaper text. Each request gives numbered excerpts of news coverage concerning ONE entity - a candidate, or a group (ticket) of candidates running together. Judge only from these excerpts.
>
> BACKGROUND (Australian local government): these elections elect several councillors per ward or council area; parties or groups run "tickets" of candidates, and independents run alone or in teams - a team may have its own name or appear as an unnamed group on the ballot. State and federal politics are different levels of government.

WHAT THE EXCERPTS ARE. Each excerpt is either (a) the entity's own speech - quoted or reported - or (b) a journalist's description of the entity's platform or record, which counts as evidence about the entity. Together they are the entity's record in the coverage: weigh every excerpt as evidence about the entity.
- Excerpts are grouped by source article and appear in the order they occur within it. A marker line separates excerpts drawn from different articles, so consecutive excerpts with no marker between them come from the same article and may be read jointly as connected context.
- The entity's own statements of policy, priorities and criticism are the primary evidence.
- Where the entity criticises or praises OTHERS' policies or conduct, that reveals the entity's own priorities and positions - it is NEVER evidence about the other person's positions, and you are not assessing anyone but the named entity.
- Where the entity reports constituent demands approvingly or as its motivation ("residents want lower rates"), treat that as an implied policy intention of the entity - unless the entity distances itself from the demand.

PARTY LABELS ARE NOT POSITIONS. Party names, memberships and affiliations appearing in the excerpts are NOT evidence of any policy position - not for anyone, on any dimension. Judge ONLY from statements of policy and expressed priorities - what the entity says it will do, supports, opposes or prioritises - and never from outside knowledge of any party, place, person or event.

THE DIMENSIONS. Score the entity on each of the seven dimensions below. Every dimension runs from pole A to pole B; the poles are defined per dimension. Each definition lists that dimension's core issues; content on clearly connected issues counts too, but ONLY where it genuinely maps onto the two poles - never guess a dimension the excerpts do not speak to:

1. `lr_fiscal` - Rates, spending and the size of council government [feeds `left_right_overall`].
   Pole A (the LEFT pole): rates fund services and infrastructure - maintain or raise them rather than cut services; invest through capital works; an active, spending council; rate-cut promises are reckless.
   Pole B (the RIGHT pole): keep rates low or frozen; cut administrative waste; back-to-basics - ratepayer money for ratepayers' core needs; a lean council doing less.

2. `lr_environment` - Environment and climate [feeds `left_right_overall`].
   Pole A (the LEFT pole): climate action is core council business - emissions targets, renewables, tree canopy, protecting the environment ahead of unfettered growth.
   Pole B (the RIGHT pole): scepticism of council climate activism - core services placed ahead of climate programs; growth and the economy first.

3. `lr_housing` - Housing affordability [feeds `left_right_overall`].
   Pole A (the LEFT pole): mandated affordable-housing targets or set-asides; developers made to pay their fair share; protect social housing and boarding houses; champion renters.
   Pole B (the RIGHT pole): affordability through supply and a lighter touch - faster approvals, less red tape, fewer mandates on developers.
   Boundary: This dimension is about affordability-directed policy - who pays, who is protected, how affordable homes are delivered. General support for or opposition to growth, density or development belongs to `development`, not here; housing content that does not map onto these two poles is not_addressed.

4. `lr_social` - Social equity and identity programs [feeds `left_right_overall`].
   Pole A (the LEFT pole): equity and representation as council priorities - gender-sensitive design, domestic-violence prevention, disability and First Nations representation, cultural-identity protection.

Pole B (the RIGHT pole): rejection of such programs specifically - diversity quotas, ceremonial acknowledgements, pride events and similar identity or equity campaigns dismissed as outside council's proper role.
Boundary: General back-to-basics or core-services rhetoric without reference to such programs is NOT this pole.

5. `lr_transport` - Transport: allocating street space and investment between car use and active or public transport [feeds `left_right_overall`].
Pole A (the LEFT pole): shift street space and investment toward cycleways, walking and public transport - expanding bus, rail and active-transport provision; reduce car dependence. Naming public transport among other priorities, without a stated shift away from car provision, is not this pole.
Pole B (the RIGHT pole): protect or expand car access and parking as the priority - keeping or adding road space and parking for cars, resisting the reallocation of either to other modes. Building, upgrading or maintaining roads or footpaths, providing car infrastructure such as electric-vehicle charging, or calling parking or congestion a problem is not this pole without that resistance - that is service delivery.
Boundary: without a stated priority between modes, transport content is service delivery and this dimension is not_addressed.

6. `left_right_overall` - Overall left-right position.
Pole A: left.
Pole B: right.
How to score it: integrate everything in the excerpts that speaks to left versus right - the themed dimensions marked above and any other left-right-relevant content that fits none of them, provided that unthemed content is genuinely ideological: a stance that plainly signals left or right. Where you have scored no themed dimension and the unthemed content does not plainly signal an ideology, "not_addressed" (or a low-confidence score) is the right overall answer. It is a judgement in its own right, NOT a mechanical average of the themed scores. Its confidence should weigh breadth and strength together: a position inferred from a single theme usually warrants lower confidence than one corroborated across several themes - but a single stance that is itself emphatic, unambiguous and strongly ideological can still justify a confident overall placement.

7. `development` - Development and growth (not part of `left_right_overall`).
Pole A: anti-overdevelopment - protect neighbourhood character, heritage, green space and farmland; resist height, density and dwelling targets.
Pole B: pro-development - support growth, housing supply, density and faster approvals.
Weigh ALL development content in the excerpts together: support for some development and opposition to other development are both signals, and the bucket should reflect their overall balance - a strong signal one way can outweigh a weaker one the other way; genuinely even content is centre_mixed. Never code from one signal while ignoring counter-signals elsewhere in the excerpts. Administrative or process material about planning - tools, procedures, meeting mechanics - is not itself a position.

`development` must NOT move `left_right_overall` - in local government it is genuinely cross-cutting; score it separately on its own poles. Likewise, presenting as independent or community-first, or opposing party politics in council, is NOT left-right evidence by itself - independents span the whole spectrum. Never cite development content or independent/community-first framing as evidence for `left_right_overall`.

SCORING. For each dimension, score the entity in ONE category, returned in `bucket`. The categories run in order from pole A to pole B:

"strongly_pole_a" / "moderately_pole_a" / "somewhat_pole_a" / "slightly_pole_a" / "centre_mixed" / "slightly_pole_b" / "somewhat_pole_b" / "moderately_pole_b" / "strongly_pole_b"
(intensity falls from "strongly" to "slightly"; note "somewhat" is weaker than "moderately". "centre_mixed" sits exactly between the poles.)
- NOT ADDRESSED: where the excerpts do not speak to a dimension, the category is "not_addressed" and `evidence` is empty - abstaining is correct behaviour, not failure. Do not stretch thin material into a category. This includes directionally ambiguous material: content that touches a dimension's topic but does not genuinely favour either pole is "not_addressed", not "centre_mixed".
- ANTI-MIDPOINT: "centre_mixed" means the excerpts genuinely show a mixed or centrist position on that dimension. Never use it as a soft way of saying "not sure" - uncertainty belongs in confidence, and having no content on a dimension belongs in "not_addressed".
- NOT EVIDENCE, on any dimension:
  * support for or opposition to ONE specific local project - a vote, motion or campaign about a drainage scheme, a pool, a road upgrade, a single development - is not a rates or spending position, and not an ideological position in either direction, unless the excerpts tie it to an explicit wider stance on rates, spending or the council's role;
  * generic commitments nearly every candidate makes - delivering services, fixing roads and footpaths, upgrading parks and facilities, listening to the community - support no dimension. Support for local business is usually such a commitment too, except where it extends into cutting red tape, fees or regulation, which is genuine content on `lr_fiscal`;
  * content whose object the excerpts never identify ("the proposal", "this development", with no account of what it is) supports nothing: if you cannot tell what is being supported or opposed, do not score from it.
- CONFIDENCE per dimension - including a "not_addressed" judgement: "high" = plain from the excerpts; "medium" = reasonable readers could hesitate; "low" = genuinely uncertain.
- EVIDENCE per scored dimension: 1-3 spans copied VERBATIM from the excerpts and checked against them automatically - each span a contiguous run of text from a SINGLE excerpt; no paraphrase, no tidying, never join separate fragments into one span. Empty exactly when the dimension is "not_addressed".

Return ONLY an object matching the provided schema.

User message (template)

⟨entity line, in one of the forms below⟩

EXCERPT [1]:
⟨excerpt⟩

EXCERPT [2]:
⟨excerpt from the same article⟩

--- the excerpts below are from a different article ---
EXCERPT [3]:
⟨excerpt⟩

⟨further excerpts, numbered consecutively; the marker line precedes the first excerpt of each further article⟩

The entity description takes one of the following forms.

ENTITY under assessment: ⟨*SURNAME Given names*⟩ (surname first), a candidate in the council election.
ENTITY under assessment: a group (ticket) of candidates running together in the council election, presented in the coverage as "⟨*team or party name*⟩", consisting of: ⟨*member names, surname first*⟩.
ENTITY under assessment: a group (ticket) of candidates running together in the council election, consisting of: ⟨*member names, surname first*⟩.
ENTITY under assessment: the party or group "⟨*party or group name*⟩" contesting the council election.

Response schema (`news_run4_positions`)

Field	Type	Description
`dimensions`	object	
`dimensions.`⟨*key*⟩	object, one for each key: `lr_fiscal`, `lr_environment`, `lr_housing`, `lr_social`, `lr_transport`, `left_right_overall`, `development`	
`dimensions.`⟨*key*⟩`.bucket`	one of `strongly_pole_a`, `moderately_pole_a`, `somewhat_pole_a`, `slightly_pole_a`, `centre_mixed`, `slightly_pole_b`, `somewhat_pole_b`, `moderately_pole_b`, `strongly_pole_b`, `not_addressed`	The entity's category on this dimension. The nine scored values are listed in order from pole A to pole B (poles defined per dimension in the instructions); use "not_addressed" when the excerpts do not speak to the dimension.
`dimensions.`⟨*key*⟩`.confidence`	one of `high`, `medium`, `low`	Confidence in this dimension's category (or in the judgement that the dimension is not addressed).
`dimensions.`⟨*key*⟩`.evidence`	array of strings	1-3 contiguous spans copied VERBATIM from the numbered excerpts that best support the category - checked automatically. Empty exactly when the dimension is "not_addressed".

B.5 Registration-Failure Coverage

This task codes what information regarding the Liberal registration failure was provided in the article. It is run on all news articles from 2024 that either mention a candidate (as in the passages task), or which contain both the word “Liberal” or “Liberals” and, in any capitalisation, the string “council” (as in “councils” or “councillor”).

System message

You are a careful analyst recording article-level facts about texts drawn from Australian news websites. You are given ONE text (its headline, when available, is the first line). Many texts are local-newspaper articles about council elections; some are about something else entirely - record what THIS text actually contains, nothing more.

BACKGROUND (Australian local government) - facts that matter for reading these articles:
- Local (council) elections elect several councillors per ward or council area. Candidates must be registered/nominated by a deadline for their names to appear on the ballot; parties or groups run "tickets" of candidates.
- An independent electoral commission administers these elections (nominations, ballots, counting).
- Voting is compulsory for enrolled voters (non-voters face a fine); a voter can still lodge a blank or otherwise informal ballot.
- State and federal politics are DIFFERENT levels of government; an article can involve several levels at once.

PARTY CODES.
"LABOR" = the Australian Labor Party (articles may say "Labor", "ALP" or "Country Labor" - all are LABOR). "LIBERAL" = the Liberal Party of Australia. Careful: the "Liberal Democrats" / "Libertarian Party" is a DIFFERENT party - never code it LIBERAL; use its own name. "GREENS" = The Greens / Australian Greens. The Nationals are a separate party - use their own name. Any other party or team: its name as the article gives it.

EVIDENCE. Every positive coding carries `evidence`: text copied VERBATIM from the supplied article - it is checked against the source automatically; no paraphrase, no tidying. Each evidence value is ONE contiguous span. Judge every coding from the WHOLE article; the span is the single strongest piece of support, not the sole basis - prefer the span that best demonstrates the coding on its own, and if no single span carries it all, quote the strongest one. NEVER join separate passages together. If you cannot point to supporting text, the answer is the absent/false/empty value - abstaining is correct behaviour, not failure.

1. about_local_government + validity_note. True if the article contains ANY substantive local-government content (a council, councillors, a council election, council issues), even where that is not the main topic - e.g. a state-politics story drawing out local implications. False only when there is no substantive local-government content at all, or the text is junk / not a coherent article (then explain briefly in validity_note; else validity_note is ""). Careful: a party's "state council" or "federal council" is an internal party conference, NOT local government.

2. failure_reported (+ failure_affected_party, failure_evidence). Does the article report that candidates were absent from a contest, or a ballot affected, because of an error, irregularity or administrative failure in candidate registration or nomination? Any party, any contest, any cause - report what the article says. failure_affected_party = the party/group the article reports as affected - where one of the three party codes above applies, use it exactly; any other party/team in its own words; "" if none is stated. This is about candidates missing from the BALLOT through failures of process - a candidate merely losing, retiring, or choosing not to run is NOT a registration failure.
   failure_voter_consequence: when a failure is reported, does the article also state the consequence FOR VOTERS - that voters face fewer or no candidates to choose from, or an uncontested or one-sided contest? "stated" requires a voter-facing statement, with its own evidence span; a sentence that merely restates the absence ("the party is absent from ballots") is "not_stated". "not_applicable" when no failure is reported. Its evidence is "" unless the value is "stated".

3. Blame - the PRIMARY narrative (only when failure_reported; otherwise blamed_side="not_applicable", blamed_party="", blame_locus="not_applicable", blame_scope="not_applicable", and all blame evidence fields ""). Who does the article PRIMARILY place responsibility with - the main narrative of blame, judged from the whole article? A quoted claim counts toward the narrative; code the overall picture the article communicates, not any single voice.
- blamed_side: "affected_party" = the party/group whose candidates were affected is itself responsible; "electoral_authority" = the body administering the election; "other_institution" = another institution (a court, a government body); "another_party" = a DIFFERENT party is primarily blamed - then also set blamed_party (where one of the party codes above applies, use it exactly; else the party's own name) with blamed_party_evidence (its own span; both "" otherwise); "unclear" = blame is discussed but the primary target is unclear; "not_attributed" = the article reports the failure without attributing responsibility.
- blame_locus: is responsibility placed on the organisation and its processes ("organisational"), on named individuals ("named_individuals"), or "both"? "unclear" if attributed but not resolvable.
- blame_scope: does the article treat the failure as a one-off slip ("one_off") or as symptomatic of broader dysfunction of the blamed side ("broader_pattern")? "broader_pattern" requires blame_scope_evidence - a span where the article generalises beyond this failure ("" otherwise). "not_discussed" when blame is attributed but the article does not go beyond the failure itself. When blamed_side is "not_attributed", locus and scope are "not_applicable" and all blame evidence fields are "".

4. integrity_mentions. Does anyone in the article suggest that a failure was deliberate, engineered to disadvantage someone, or that the electoral process or the bodies administering it were partial or applied the rules unequally? One entry PER DISTINCT ACTOR raising such a concern (if any), each with claim_source (who raises it: "affected_party", "another_party", "electoral_authority_or_official", "journalist_or_article" = the article advancing it in its own voice (a concern quoted or reported from someone keeps that person's category), "other_or_unclear") and its own verbatim evidence span. Empty array = no such concern raised. Note: question 3 already records the article's MAIN narrative - here record concerns RAISED AT ALL, by whom, even in passing. NOT an integrity concern: saying voters have fewer choices or the contest is poorer (question 2's voter-consequence); criticising a party's competence or internal workings; criticising a rule while accepting it was applied impartially; lamenting that an error or absence is bad for democracy or for fair elections, without alleging bias, deliberate action, or unequal application of the rules.

5. protest_call (+ protest_evidence). Does the article contain a suggestion that voters should abstain, not vote AT ALL, vote informally, or leave the ballot blank? The suggested action must be one of those - refusing or spoiling the vote itself. "reported" = the article quotes or reports someone making such a call; "endorsed" = the article itself advances it; "absent" otherwise. NOT a protest call: urging a vote FOR anyone; urging a vote AGAINST a particular party ("do not vote for X" is ordinary campaigning); reporting turnout or informal-vote statistics; complaints without a suggestion that voters withhold or spoil their vote.

6. beneficiaries. Each party or campaign the article frames as gaining an ELECTORAL ADVANTAGE BECAUSE candidates or competitors are absent or the field is reduced (e.g. likely to win seats or a majority as a result) - one entry per party/campaign (where one of the party codes above applies, use it exactly; any other party/team in its own words) with its own verbatim evidence. Empty if none. Report the article's framing, not your own election analysis: a speaker campaigning for or against someone, or a group merely holding or joining an event, is NOT a framed beneficiary - the article must state the gain and tie it to the reduced field.

RULES.
- Evidence text is copied EXACTLY and will be checked. No span, no code.

- Do not use outside knowledge of any events; code only what THIS article contains.
- The article may be about something else entirely - then the fields simply take their absent/false values.

Return ONLY an object matching the provided schema.

User message (template)

ARTICLE TEXT:
⟨article text: headline (where available) as the first line, then body⟩

Response schema (`news_run2_article_context`)

Field	Type	Description
`about_local_government`	boolean	True if the article contains ANY substantive local-government content (a council, councillors, a council election, council issues), even where that is not the main topic. False only when there is no substantive local-government content at all, or the text is junk / not a coherent article.
`validity_note`	string	Brief note when the text is junk, truncated, or not a coherent article; else ''.
`failure_reported`	boolean	Does the article report that candidates were absent from a contest, or a ballot affected, because of an error, irregularity or administrative failure in candidate registration or nomination?
`failure_affected_party`	string	The party/group the article reports as affected: LABOR | LIBERAL | GREENS | another party/team's own name | '' if none stated or no failure reported.
`failure_evidence`	string	VERBATIM article text reporting the failure; '' iff failure_reported is false.
`failure_voter_consequence`	one of `stated`, `not_stated`, `not_applicable`	Does the article state the consequence FOR VOTERS of the failure - fewer or no candidates to choose from, an uncontested contest? A voter-facing statement, not a mere restatement of the absence. 'not_applicable' iff failure_reported is false.
`failure_voter_consequence_evidence`	string	VERBATIM voter-facing statement; '' unless failure_voter_consequence='stated'.
`blamed_side`	one of `affected_party`, `electoral_authority`, `other_institution`, `another_party`, `unclear`, `not_attributed`, `not_applicable`	Who the article PRIMARILY places responsibility with - the main narrative of blame, judged from the whole article. 'not_applicable' iff no failure reported; 'not_attributed' = failure reported, no one blamed.
`blamed_party`	string	ONLY when blamed_side='another_party': the blamed party (codes above / own name). Else ''.
`blamed_party_evidence`	string	VERBATIM text supporting blamed_party; '' iff blamed_party is ''.
`blame_locus`	one of `organisational`, `named_individuals`, `both`, `unclear`, `not_applicable`	Blame at the organisation/its processes vs named individuals. 'not_applicable' iff blamed_side is 'not_applicable' or 'not_attributed'.

Field	Type	Description
`blame_scope`	one of `one_off`, `broader_pattern`, `not_discussed`, `not_applicable`	One-off slip vs symptomatic of broader dysfunction of the blamed side. 'not_discussed' when blame is attributed but scope is not; 'not_applicable' iff blamed_side is 'not_applicable' or 'not_attributed'.
`blame_scope_evidence`	string	VERBATIM text supporting a broader_pattern reading; '' unless blame_scope='broader_pattern'.
`blame_evidence`	string	VERBATIM text supporting the primary blame coding; '' iff blamed_side is 'not_applicable' or 'not_attributed'.
`integrity_mentions`	array of objects	One entry per DISTINCT ACTOR raising a qualifying integrity concern (deliberate / partial / rules applied unequally, targeting the process or its institutions). Empty = no such concern raised.
`integrity_mentions[].claim_source`	one of `affected_party`, `another_party`, `electoral_authority_or_official`, `journalist_or_article`, `other_or_unclear`	Who raises the concern. journalist_or_article = the article's own voice; anyone quoted/reported keeps their own category.
`integrity_mentions[].evidence`	string	VERBATIM article text of/about this actor's claim - copied exactly, checked.
`protest_call`	one of `absent`, `reported`, `endorsed`	A suggestion that voters should abstain, not vote at all, vote informally, or leave the ballot blank. reported = the article quotes/reports someone making the call; endorsed = the article advances it in its own voice.
`protest_evidence`	string	VERBATIM text; '' iff protest_call='absent'.
`beneficiaries`	array of objects	Each party/campaign the article frames as gaining BECAUSE candidates or competitors are absent or the field is reduced. Empty if none; multiple allowed.
`beneficiaries[].party`	string	LABOR | LIBERAL | GREENS | any other party/team's own name.
`beneficiaries[].evidence`	string	VERBATIM article text supporting this entry - copied exactly, checked.

B.6 Party Alignment

This task takes as its input all passages about or by a person associated with the Liberal Party (as coded in the passages task), or which mention the Liberals. It combines these with the candidate or party that is the subject or speaker of the passage. The purpose is to measure how these candidates and groups are framed relative to the Liberals.

System message

You are a careful analyst coding short passages drawn from Australian news articles about local (council) elections. Each request gives ONE passage, between <<<PASSAGE>>> and <<<END PASSAGE>>>. Judge only that passage.

BACKGROUND (Australian local government): local (council) elections elect several councillors per ward or council area; parties or groups run "tickets" of candidates, and independents run alone or in teams - a team may have its own name or appear as an unnamed group on the ballot. State and federal politics are DIFFERENT levels of government.

PARTY CODES.
"LABOR" = the Australian Labor Party (articles may say "Labor", "ALP" or "Country Labor" - all are LABOR). "LIBERAL" = the Liberal Party of Australia. Careful: the "Liberal Democrats" / "Libertarian Party" is a DIFFERENT party - never code it LIBERAL; use its own name. "GREENS" = The Greens / Australian Greens. The Nationals are a separate party - use their own name. Any other party or team: its name as the article gives it.

THE TASK. The request names the FOCAL entity - a candidate, or a party/group contesting the council election. Report EVERY political party (or named political figure) that the delimited passage frames the focal entity as ALIGNED WITH or DIFFERENTIATED FROM - one entry per target, an empty list when it asserts neither. For each entry:
- "aligns": the passage presents the focal entity as politically similar to, aligned with, a replacement or substitute for, or the natural choice for supporters of, that target - OR states that the focal entity is or was a member of, part of, or associated with that party - OR shows the focal entity praising or endorsing that party's policies or positions (liking a party's policies frames the entity as like that party). A bare biographical association counts: "X, a former Example Party councillor, is running for Ward A" is "aligns" with the Example Party even with nothing more - including the background is itself a framing choice. A party label attached to the focal entity's own name ("Example Party councillor X", "the Example Party's Y") is a stated association and codes "aligns" with that party. This holds even where the sentence is a list or line-up of rival candidates: the same-list rule below stops you coding two listed candidates as aligned WITH EACH OTHER merely for being listed together - it never cancels a party label attached to the focal entity's own name. Jointly taking the same political action - for example voting the same way on a motion - is a minor but real form of political similarity: code it "aligns", with reduced confidence when it stands alone, higher when other content in the passage corroborates it. Likewise, two entities described as taking the SAME position are weakly aligned with each other, not differentiated.
- "differentiates": the passage sets the focal entity apart from that target - distancing, denial of association, explicit contrast ("X insists she is nothing like the Example Party", "X quit the Example Party in protest and rejects its policies") - or shows the focal entity criticising or rejecting that party's policies or positions ("X said the Example Party's policies are foolish" frames X as unlike that party, the mirror image of praise). A passage may differentiate the focal entity from SEVERAL targets at once ("X says Party A and Party B both ignore residents" = two entries).

Where one passage carries BOTH an association and a distancing toward the SAME target, code that target's NET framing - the direction the passage as a whole communicates - with evidence carrying that net reading and the tension reflected in a reduced confidence.

NOT alignment or differentiation: a party merely appearing in the passage without being connected to the focal entity; two names merely appearing in the same list or event (this is about the listed entities' relation to each other - a party label on the focal entity's own name still codes "aligns", per above); ELECTORAL COMPETITION - running against a party, taking seats or votes from it, eroding its position, challenging it, or pitching for change against it is a contest, NOT a statement that the focal entity is politically like or unlike it. And code ONLY what the passage asserts about the FOCAL entity: a statement about some other entity's conduct, decisions or affiliations is not a coding of the focal entity, however prominent it is in the passage.

TARGET. Each entry's target = who the focal entity is aligned with / differentiated from: where one of the party codes above applies, use it exactly; any other party or team, its name as the article gives it; and where the passage ties the focal entity to a named PERSON rather than a party (for example, presented as the successor or replacement of a named former councillor), that person's name exactly as the passage gives it. The target is NEVER the focal entity itself: where the passage states the focal entity's own affiliation, the target is the party or affiliation stated, not the focal entity's name.

EVIDENCE. Each entry carries evidence = ONE contiguous span copied VERBATIM from WITHIN the delimited passage that best carries THAT entry's framing - checked automatically; no paraphrase, never join separate fragments.

RULES.
- Judge only the delimited passage; context is for interpretation only.
- Do not use outside knowledge of any events, people, or parties; code only what this text contains.
- An empty list is the correct answer whenever the passage asserts no alignment or differentiation for the focal entity - abstaining is correct behaviour, not failure.
- confidence per entry: "high" = plain from the text; "medium" = reasonable readers could hesitate; "low" = genuinely uncertain.

Return ONLY an object matching the provided schema.

User message (template)

PASSAGE IN CONTEXT:
<<<PASSAGE>>>
⟨*passage*⟩
<<<END PASSAGE>>>

FOCAL entity: ⟨*entity name*⟩ (⟨*kind of entity*⟩).

Response schema (`news_run3_alignment`)

Field	Type	Description
`alignment_entries`	array of objects	One entry per party or named figure the delimited passage frames the FOCAL entity as aligned with or differentiated from. Empty array when the passage asserts neither about the focal entity.
`alignment_entries[].alignment`	one of `aligns`, `differentiates`	aligns = similar to / aligned with / a replacement or substitute for / the natural choice for supporters of / is or was a member of or associated with this target. differentiates = set apart / distanced / contrasted from this target. Mixed signals toward the SAME target: the passage's net framing.
`alignment_entries[].target`	string	Who: a party code (where one applies, exactly) | another party/team's own name | a named PERSON exactly as the passage gives them (person targets are resolved to parties in code after the run).
`alignment_entries[].evidence`	string	ONE contiguous span copied VERBATIM from WITHIN the delimited passage that best carries this entry's framing - checked automatically.
`alignment_entries[].confidence`	one of `high`, `medium`, `low`	Confidence in this entry.

C Appendix: Geographic Harmonisation

In this section, I explain the construction of the geographic crosswalk used to harmonise the geographic units to 2024 ward (or polling place) geography. I follow the spirit of Ferrara et al. (2024), using granular population data for population weights. Specifically, I use 2021 Census mesh block level population counts (Australian Bureau of Statistics, 2021). The general principle is that a count variable recorded before 2024 is split across the 2024 units it overlaps with, in proportion to the share of its population living in each. Where a mesh block must be split across geographies, its population is assigned in proportion to the relative areas of the segments.

C.1 Inputs

Ward boundaries are taken from the NSWEC's boundary files for each council, scraped from the NSWEC website. Where two ward polygons overlap, I remove the overlapping area from both. Polling place coordinates are taken from the NSWEC's polling place map for each election. Mesh block geographies are provided by Australian Bureau of Statistics (2021).

C.2 Polling Place Catchments

Voters are free to vote at any polling place within their ward. I construct polling place catchment areas on the basis that voters choose the closest polling place. I construct these catchment areas separately for each election, partitioning each ward into Voronoi cells.

Note that polling places and vote types that are not assigned a unique polling place geography (e.g., pre-polling and postal votes) are excluded from polling place level data. Additionally, enrolment counts are only available at ward level, and are disaggregated using the same principle as this crosswalk. Noise in this disaggregation means that, without top-coding, some vote share values would be above one.

C.3 Crosswalk

Let o index the units of an earlier election (wards, or catchments) and d the 2024 units of the same type. I intersect the earlier units, the 2024 units and the mesh blocks m, and compute their areas. Let N_m be the population of mesh block m, A_m its area, and A_{odm} the area of the part of m lying in both o and d. The population of the intersection of o and d, the crosswalk weight, and the harmonised value of a count X are

$$P_{od} = \sum_m N_m \frac{A_{odm}}{A_m}, \qquad w_{od} = \frac{P_{od}}{\sum_{d'} P_{od'}}, \qquad X_d = \sum_o w_{od} X_o. \tag{4}$$

In words, the population of the intersection of o and d is the sum of the mesh block populations contained in the intersection, with split mesh blocks having their population contributed in proportion to the split. The weight is then the proportion of the original geography's (o's) population that lies in that intersection. The counts for d are then the weighted sum of the original geography's counts.

I drop 2024 unit-years with population coverage below 95%. In nearly all such cases, coverage is near zero and the result of a negligible overlap of geographies, likely due to noise in the recording of the geometries. The primary assumption underlying this construction is that the count variables are spread in proportion to the (2021) population distribution — i.e., that voters are distributed like the population.