Extremism Is Self-Undermining, but Populism Is Not

Matteo Bertoli[†]

June 2026

Abstract

Does holding office strengthen or weaken the long-term appeal of right-wing extremism and populism? Although populism and extreme-right ideology are often conflated, this paper separates the two traits using supervised text classification of candidate platforms to construct continuous candidate-level scores. It then exploits close mayoral elections in Italian municipalities, extending close-election politician-characteristic regression discontinuity to continuous treatments. Electing an extreme-right mayor produces clear electoral backlash. At the next municipal election voters shift their support toward more moderate candidates. Electing a populist mayor generates no comparable punishment. The most plausible explanation for these differences lies in behavior in office. Extreme-right mayors govern as ideologues, pursuing welfare retrenchment even at electoral cost, while populist mayors appear to act as office-seekers, with welfare expansion concentrated where re-election incentives remain.

Populist and extreme-right movements have reshaped political competition across established democracies (Inglehart and Norris, 2016; Rooduijn, 2015; Noury and Roland, 2020). From Trump and Brexit to the sustained advances of parties such as the Rassemblement National in France, the AfD in Germany, and Vox in Spain, they have won offices across levels of government in a number of countries and regions and increasingly shape policy agendas, whether by governing directly or by pulling mainstream competition in their direction. Whether these movements can win is no longer uncertain (Kitschelt and McGann, 1995; Mudde, 2007; Ivarsflaten, 2008; Kriesi et al., 2008; Arzheimer, 2009; Rydgren, 2007; Golder,

[†]Department of Political Science, University of Rochester. Email: mbertol2@ur.rochester.edu. I am deeply grateful to my advisor, Alexander Lee, for his guidance throughout this project, and to Anastassios Kalandrakis, Anderson Frey, Bonnie Meguid, and Curt Signorino, for their generous feedback. I thank the many Italian municipal administrations that shared their electoral platforms, without which the measurement effort behind this paper would not have been possible. This research was supported by the Democracy Center at the University of Rochester. All remaining errors are my own.

2016; Norris and Inglehart, 2019; Muis and Immerzeel, 2017; de Vries and Hobolt, 2020; Meguid, 2005; Abou-Chadi, 2016). The open question is what happens once they have won: do these movements undermine themselves, or does incumbency deepen their appeal?

Existing work usually bundles populism and extreme-right ideology together under the "anti-establishment" label. However, the two are conceptually distinct political tendencies with different implications for governing and voter behavior. Populism, whether understood as a thin ideology (Mudde, 2004), a discursive frame (Hawkins, 2009), or a mobilization strategy (Weyland, 2001), revolves around a moralized opposition between a virtuous people and a corrupt elite. It is compatible with programmatic agendas across the political spectrum and orthogonal to left–right positioning. Extreme-right ideology, by contrast, is defined by its substantive content: nativism, authoritarianism, and a hierarchical conception of social order (Mudde, 2007). The two traits are often thought to co-occur—Mudde (2007)'s "populist radical right" is the canonical composite—but they are logically independent.

Whether incumbency helps or hurts extreme-right and populist officeholders is unclear from the existing literature (Van Spanje, 2011; Riera and Pastor, 2022; Bergh and Kärnä, 2024; Albertazzi and McDonnell, 2015; Gromadzki, Sałach and Brzeziński, 2024). Like other incumbents, extreme-right and populist officeholders can benefit from name recognition, office resources, and local networks, especially where parties are strong, term limits are absent, and governing outputs are visible and attributable to voters (De Benedetto, 2020; Freier, 2015; Klašnja and Titiunik, 2017; Duch and Stevenson, 2008). Unlike other incumbents, populist and extreme-right officeholders also risk losing the outsider appeal that brought them to power and the programmatic ambiguity that broadened their electoral appeal (Grzymala-Busse and Nalepa, 2019).

Three limitations of existing empirical approaches have made it difficult to estimate what the effects of incumbency are for populist and extreme-right actors. The first is a problem of causal identification: much existing evidence is correlational, relying on cross-country comparisons where the selection of populist and extreme-right actors into office is confounded with systematic differences in institutions, party systems, coalition opportunities, and economic conditions (Van Spanje, 2011; Akkerman and De Lange, 2012; Riera and Pastor, 2022; Bergh and Kärnä, 2024). Second, existing measures are binary and conflate the two traits this paper seeks to separate. Studies that identify populist and extreme-right politicians through party labels, expert lists, or party-family classifications usually bundle populism and extreme-right ideology into a single anti-establishment category and impose uniformity across candidates; they cannot, for instance, distinguish a candidate who deploys populist rhetoric from one

who merely belongs to a party classified as populist (Mudde, 2004; Weyland, 2001; Rooduijn et al., 2019; Guriev and Papaioannou, 2022). Third, even when politician characteristics are measured continuously, the available designs have been developed mainly for ideology, not populism. While a large and sophisticated literature has developed ways to measure ideology from legislative behavior and donations (Poole and Rosenthal, 1985; Poole, 2005; McCarty, Poole and Rosenthal, 2006; Bonica, 2014), no such approach is possible for populism: roll-call votes and donor coalitions encode ideological position, not the moralized people-versus-elite rhetoric that defines populism.

To address these obstacles, the paper assembles a new dataset of municipal candidate platforms from Italy and analyzes them using a supervised text-as-data approach. Italy offers an unusually favorable setting, with populist and extreme-right actors both numerous across its municipalities. Because no unified repository of municipal platforms exists, the dataset was constructed through extensive direct outreach to municipalities across Italy to obtain full-text candidate programs. These texts were then linked to official candidate records, election returns, and administrative measures of municipal policy outputs from the 2019 and 2024 local elections. This dataset responds directly to the measurement problem: rather than inferring populism or extreme-right ideology from party labels, the paper measures both traits from the candidates' own campaign texts using supervised text classification. Focusing on a single country holds constant the broader institutional environment, reducing the cross-context heterogeneity that confounds much existing evidence. This dataset addresses the limitations identified above: it allows measuring populism and ideology at the candidate level, estimating the effects of incumbency through a close-election design, and avoiding arbitrary treatment thresholds.

First, to address the identification problem, the paper exploits quasi-random variation in who governs across razor-thin elections for Italy's directly elected mayors. The strategy builds on the gradient pairwise close-race RD (Bertoli and Hazlett, 2025), in which elections are signed by which candidate scores higher within each race, so that every close contest contributes to identification regardless of where both candidates fall in the absolute score distribution. The treatment is the 2019 winner's populism or ideology score, depending on the design, taken from their campaign platforms, with outcomes measured over the subsequent mandate.

Second, to address the conflation problem directly, the paper measures populism and extreme-right ideology as separate, continuously valued attributes of individual candidates rather than assuming them from partisan affiliation. Both dimensions are operationalized by manually

producing 6,643 paragraph-level labels following detailed codebooks, then fine-tuning four XLM-RoBERTa models on those annotations to scale the full corpus of candidate platforms. Using this transformer-based scaling, the analysis constructs continuous candidate-level measures that separately capture (i) ideological position and (ii) populist discourse, even where party labels are unavailable or misleading. In the full sample, the overall populism measure correlates only weakly with ideological extremism ($r = 0.20$); among the top quartile of candidates on this measure of populism, left/right composition is essentially a coin toss (50.4 / 49.6). Populism and ideological extremism are thus two distinct concepts rather than two labels for a single underlying construct. Because populism is itself multi-dimensional, the analysis also separates it into two component measures: anti-elite rhetoric and people-centrism (Jagers and Walgrave, 2007). These measures make it possible to assess whether the results depend on the overall populism score or instead reflect one of its underlying dimensions.

Third, the paper introduces a continuous-treatment extension to close-election politician-characteristic regression discontinuity (PCRD) designs. The standard approach in PCRD (Hall, 2015) collapses the intensity of the treatment — how extreme or populist the winning candidate is — into a binary variable. To sharpen the signal, it drops pairs of candidates whose scores on the trait are too close. The continuous-treatment extension employed here uses the elected mayor's score directly as the treatment, so the estimate gives the effect per unit of ideology or populism without dropping any pair. This adapts the fuzzy-RD framework of Hahn, Todd and Van der Klaauw (2001), Dong, Lee and Gou (2023), and Ebenstein et al. (2017) to PCRD.

The findings show that the effects of electing an extreme-right mayor and a populist mayor diverge. In razor-thin elections, electing more ideologically extreme-right mayors is followed by electoral backlash, measured as a leftward shift in the vote-weighted ideology of candidates in the subsequent municipal election. By contrast, electing more populist mayors produces no comparable backlash. At the municipal level, populist support is essentially unchanged, neither punished nor broadly reinforced, while in subsequent national and European elections the vote share of the main populist party rises. The pattern is therefore asymmetric: extreme-right incumbency is punished, whereas populist incumbency is not, and is associated with gains for the populist party at the national level. The paper also considers a single composite measure that combines populism and extreme-right ideology, reproducing the canonical "populist radical right" widely used in the literature (Mudde, 2007). Treating the two traits jointly adds nothing: the opposing effects cancel and the design recovers nothing, casting doubt on whether the populist radical right is a coherent single category at the candidate level.

To explain why these effects diverge, the analysis points to a mechanism rooted in the office-seeker versus policy-seeker distinction. Extreme-right incumbents govern as policy-seekers: their agenda centers on identity, sovereignty, and security, and they have no ideological commitment to local redistribution. None of these commitments translates naturally into the kind of visible welfare expansion that generates credit with municipal voters. Ideologically motivated incumbents thus pursue their ideological commitments even at electoral cost; in particular, they cut municipal welfare spending. Populist incumbents govern as office-seekers. A constitutive claim to govern visibly for "the people" and against elites makes welfare spending an attractive policy option. Shorter political horizons and a greater capacity to externalize blame onto elites or opponents also make populists more inclined to exploit fiscal tools that shift costs forward while delivering immediate visible benefits (Funke, Schularick and Trebesch, 2023). Higher spending should therefore be observed only when re-election incentives are binding.

The empirical analysis based on municipal expenditure data supports this account. Barely elected extreme-right mayors reduce welfare provision. Barely elected populist mayors show a positive but statistically imprecise effect on average. Populist mayors tend to expand welfare when facing re-election, especially as the election nears, but not when term-limited, though this contrast is imprecisely estimated. Extreme-right mayors, by contrast, cut welfare spending more under term limits, when lame-duck mayors implement their agenda free of electoral constraints.

Taken together, these findings show that electoral accountability operates differently for populist and extreme-right politicians. Holding office is self-undermining for extreme-right ideologues but not for populist office-seekers, because the two types govern differently and generate different policy outputs. This asymmetry helps explain why previous work that conflates populism and extreme-right ideology has produced mixed results: voters may punish extreme-right governing records while sustaining populist ones (Van Spanje, 2011; Riera and Pastor, 2022; Bergh and Kärnä, 2024).

This paper contributes to research on incumbency, retrospective accountability, and the electoral consequences of anti-establishment government. A large literature shows that voters use incumbents' governing records to update their support for parties and politicians (Key, 1966; Fiorina, 1981; Ferejohn, 1986; Fearon, 1999; Duch and Stevenson, 2008; Healy and Malhotra, 2013), while work on incumbency shows that office can generate electoral advantages through visibility, resources, and local networks (Gelman and King, 1990; Lee, 2008; Freier, 2015; De Benedetto, 2020), though it can also become a liability where parties are weak or term

limits bind (Klašnja and Titiunik, 2017). Recent studies ask whether these dynamics also apply to populist and extreme-right actors, but the evidence remains mixed (Van Spanje, 2011; Albertazzi and McDonnell, 2015; Riera and Pastor, 2022; Bergh and Kärnä, 2024; Gromadzki, Sałach and Brzeziński, 2024), partly because this work studies party families or broad anti-establishment categories that combine populism with extreme-right ideology. This paper shows that the electoral consequences of incumbency depend not simply on whether anti-establishment actors govern, but on which component of anti-establishment politics structures their behavior in office. By separating populist rhetoric from extreme-right ideology, it explains why incumbency can be self-undermining for some anti-establishment actors but not for others. Because these dynamics cannot be cleanly separated or identified at the national level, the analysis turns to the setting where they can: close municipal elections in Italy.

The remainder of the paper develops the theoretical framework (Section 1), introduces the Italian case (Section 2), describes the data and measures (Sections 3–4), outlines the research design (Section 5), and presents evidence on policy and electoral responses (Sections 6–7).

1 Theoretical Framework

1.1 Populism and Extreme-Right Ideology

Populism and extreme-right ideology are analytically distinct phenomena, even though they frequently co-occur in contemporary politics. Populism is commonly understood as a thin ideology that frames politics as a moral conflict between a virtuous people and a corrupt elite and demands that policy reflect the "will" of the former (Mudde, 2004; Mudde and Kaltwasser, 2017). Other proposed definitions treat populism as a political strategy (Weyland, 2001), a discursive logic (Laclau, 2005), or a political style (Moffitt, 2016). Independently of the specific definition, populism can be associated with both left- and right-wing ideologies (Kaltwasser et al., 2017; Stanley, 2008). The same actor may deploy populist rhetoric more or less intensively depending on the political context, the identity of the opponent, and the audience being addressed (see also Jones and Menon, 2024, on the distinction between opportunist and true-believer populists). Following Rooduijn and Pauwels (2011), the thin ideology of populism is commonly understood to consist of two components. Anti-elite rhetoric (AER) denotes the discursive construction of a corrupt, self-serving elite as the antagonist of the people. This elite might be accused of corruption, self-enrichment, clientelism, or capture by powerful interests. People-centrism (PC) denotes the discursive construction of "the people" as a virtuous community whose will should guide political action. This might translate into demands for participatory budgeting, binding referendums, or formal accountability

mechanisms that bypass representative institutions. Although the two are conceptually linked — people-centrism is at the core of any populist appeal (Jagers and Walgrave, 2007) — they vary independently in practice. A politician may claim to represent the "will" of the people without accusing elites of wrongdoing, or attack elite corruption without invoking the people as the legitimate source of her authority (Ernst, Engesser and Esser, 2017). The analysis examines both the overall populism measure and its two sub-dimensions separately.

Extreme-right ideology, in contrast to the thin ideology of populism just described, is a substantively thick doctrine centered on nativism, authoritarian conceptions of order, and the protection of national community boundaries. This is often accompanied by welfare chauvinism and a desert-based view of redistribution that prioritizes native, law-abiding, and contribution-making insiders over other groups (Mudde, 2007; Ennser-Jedenastik, 2016; Schumacher and Van Kersbergen, 2016). The distinction between the two traits is observable among real-world political parties. Figure 1 plots European parties using the Chapel Hill Expert Survey (CHES), one of the most widely used sources of cross-national data on party positions, to measure left–right ideology and populism (Bakker et al., 2019; Rovny et al., 2025), with party families classified by PopuList v2.0 (Rooduijn et al., 2019). Populism does not increase monotonically along the left–right axis. It is instead weakly U-shaped, rising toward both ideological extremes ($R^2 = 0.19$). Parties with similar ideological positions display substantial variation in their populism scores, and there are very populist parties on both the left and the right. This empirical pattern reflects the conceptual independence of the two dimensions. Figure 1 is purely illustrative; the analysis uses no expert-survey data, constructing all ideology and populism measures from the candidate platforms themselves (Section 4).

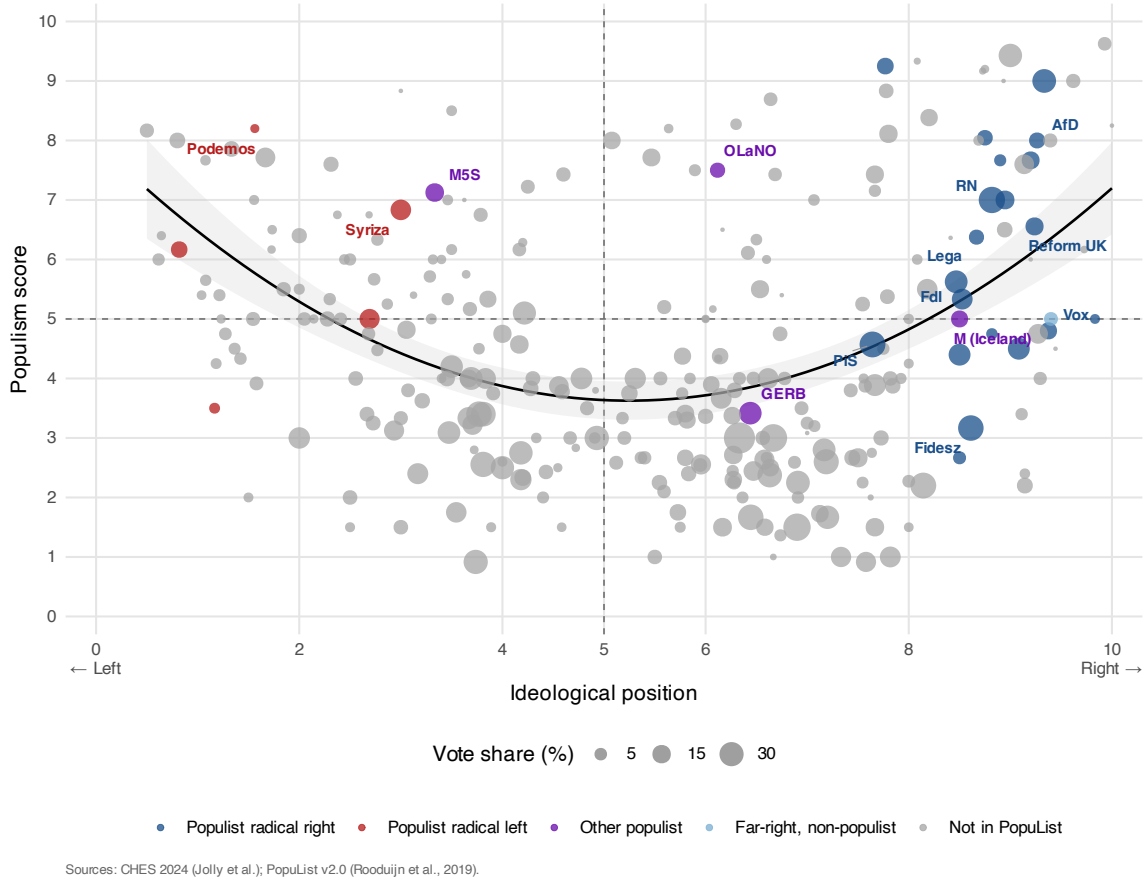

Figure 1: Populism and ideology across European parties. Each point represents a European party in the CHES 2024 expert survey, positioned by left–right ideological placement (x-axis, `lrgen`) and anti-elite populism score (y-axis, `people_vs_elite`), both on a 0–10 scale. Point size is proportional to vote share. Colours indicate party family according to PopuList v2.0 (Rooduijn et al., 2019): populist radical right (dark blue), populist radical left (red), other populist (purple), far-right non-populist (light blue), and parties not covered by PopuList (grey). The curve is a quadratic fit with 95% confidence interval ($R^2 = 0.19$, $N = 279$). *Sources*: CHES 2024 (Rovny et al., 2025); PopuList v2.0 (Rooduijn et al., 2019).

1.2 The Bundled Type

The distinction developed above matters because much of the literature is interested not only in populism and extreme-right ideology separately, but also in their combination as a joint political type. The canonical example is the populist radical right (Mudde, 2007), a category that combines a people-versus-elite appeal with a substantive ideological agenda centered on nativism, authoritarianism, and national community boundaries. This composite is useful for describing many contemporary parties, especially at the party-family level. It is less clear, however, that it travels cleanly to the candidate level, where politicians may combine the two traits in different degrees. This paper therefore treats the populist-radical-right type as an empirical possibility rather than as a maintained assumption.

If populism and extreme-right ideology were simply two manifestations of the same underlying political type, measuring them separately would add little. Candidates who score high on one dimension would also tend to score high on the other, and their effects in office should point in the same direction. If, instead, the two traits capture different political logics, then bundling them together risks averaging across distinct forms of behavior.

The candidate-level measures developed in this paper make this distinction observable. In this setting, populism and extreme-right ideology are only weakly related. Populism is weakly negatively correlated with signed ideology in the full candidate pool ($r = -0.13$) and essentially uncorrelated among the leading candidates in close races ($r = 0.03$). Among the 25 percent most right-wing close-race candidates, only 31.2 percent also fall within the most populist quartile, only modestly above the 25 percent expected if the two traits were statistically independent. High levels of extreme-right ideology therefore do not generally coincide with high levels of populism at the candidate level.

The implication is theoretical as well as empirical. A bundled category cannot distinguish between politicians whose behavior is driven mainly by extreme-right ideological commitments and politicians whose behavior is driven mainly by populist incentives to demonstrate responsiveness to the people. The remainder of the theory therefore develops separate expectations for the two traits. Appendix J later returns to the bundled type directly by recombining populism and extreme-right ideology into a single index and showing that this composite obscures, rather than clarifies, the patterns recovered when the traits are measured separately.

1.3 Incentives and Constraints in Governance

This subsection develops a framework for explaining how populist and extreme-right politicians behave in office and how voters respond to their governing records. It maps populism and extreme-right ideology onto the classic distinction between office-seeking and policy-seeking politicians (Downs, 1957; Wittman, 1983; Calvert, 1985). Extreme-right politicians are expected to behave as policy-seekers, pursuing programmatic commitments even when doing so is electorally costly. Whether populists govern as office-seekers, adapting policy to sustain electoral support, or as policy-seekers sincerely committed to a populist worldview is treated as an open question, examined through governing behavior under term limits. These motivations generate different governing strategies and different responses to electoral constraints. The framework then draws on retrospective accountability (Key, 1966; Fiorina, 1976; Healy and Malhotra, 2009) to explain how citizens translate the benefits and costs

produced by each type into subsequent political choices. Together, these elements yield contrasting predictions about governing behavior, the effect of term limits, and electoral feedback.

A long tradition in political science distinguishes politicians by what motivates them to seek and hold office. Office-seekers value the extrinsic rewards of winning and retaining power and therefore adjust their behavior to electoral demand. Policy-seekers, by contrast, value office primarily as a means of implementing substantive commitments and may pursue those commitments even when doing so carries electoral costs. In the classic formulation, office-seekers formulate policy in order to win office, whereas policy-seekers win office in order to make policy (Downs, 1957; Wittman, 1983; Calvert, 1985).

Ideologically extreme politicians are generally understood in the literature as policy-seekers, and their willingness to pursue policies that may be electorally costly follows from this intrinsic motivation. Hall (2019) offers a clear illustration of this understanding: extremists are taken to be so intrinsically motivated that they are willing to sacrifice their own wellbeing, bearing the personal costs of candidacy, in order to implement their preferred policies. On this basis, Hall develops a theory of selection in which rising costs of running for office push moderates out of the candidate pool first, since the policy stakes of winning are too small to justify the sacrifice, leaving a pool that is increasingly ideologically extreme. Politicians selected on policy motivation can then be expected, once in office, to pursue their ideological objectives even when electoral incentives point in another direction.

As for populists, the literature is more complex. A first trait that is usually stressed is their comparatively short time horizon. Populist politicians tend to prioritize policies that produce immediate and highly visible benefits, even when their fiscal or institutional costs are deferred into the future or shifted onto other actors (Dornbusch and Edwards, 1990; Funke, Schularick and Trebesch, 2023). In this respect, they differ from traditional mainstream politicians, who can build electoral support gradually through a reputation for competence, fiscal prudence, administrative efficiency, and reliable service provision (Ferejohn, 1986; Brender and Drazen, 2008). Because such a reputation develops over time and can support a longer political career, mainstream politicians generally have stronger incentives to preserve long-term governing performance. Populists, by contrast, rely more heavily on repeated and visible demonstrations that they are responding directly to the demands of ordinary citizens against an unresponsive or corrupt elite. Their governing strategy is therefore more strongly oriented toward short-term political returns (Strom, 1990).

As to whether populists are office-seekers or policy-seekers, the literature is less clear.

Under the ideational account, populism may reflect a sincere commitment to a worldview that divides society between a virtuous people and a corrupt elite and demands that politics express the popular will (Mudde, 2004; Hawkins and Rovira Kaltwasser, 2017; Mudde and Kaltwasser, 2017). Populist politicians may therefore be intrinsically motivated to seek office in order to act on that worldview. The strategic account, by contrast, treats populism as a tool for winning and exercising power, portraying populist politicians as opportunistic office-seekers who adapt their policy choices to electoral circumstances (Weyland, 2001). These accounts may generate similar behavior while re-election incentives remain active, but they diverge once those incentives are removed. An office-seeker should pursue electorally useful behavior only when re-election is possible, and not pursue it once that incentive is removed, whereas a policy-seeker should continue, or even intensify, the pursuit of sincerely held commitments.

These differences in governing behavior should, in turn, generate different electoral reactions through retrospective accountability: voters observe the benefits and costs produced during an incumbent's term and update their subsequent political choices accordingly (Key, 1966; Fiorina, 1976; Healy and Malhotra, 2009). What matters is not whether incumbents faithfully implement specific campaign promises, but the governing record voters experience once they are in office. If populist politicians deliver visible benefits, voters should sustain or increase their support for populist candidates. If extreme-right politicians pursue programmatic policies that impose visible costs on broad groups of citizens, voters should shift away from extreme-right alternatives.

This argument does not require campaign platforms to provide a complete or binding account of the policies incumbents will later pursue. Whether politicians fulfill specific campaign pledges is itself contested (Stokes, 2001; Thomson et al., 2017). Here, platforms matter because they reveal political type rather than because they enumerate every subsequent policy choice. Extreme-right candidates may campaign primarily on identity, sovereignty, or public order without explicitly announcing welfare retrenchment, while populist candidates may invoke popular responsiveness without specifying distributive expansion. Once in office, however, the motivations associated with each type shape how incumbents respond to governing incentives and which policy instruments they choose.

Retrospective accountability may extend beyond the individual incumbent. Voters can use a governing record not only to evaluate a particular officeholder, but also to update their beliefs about the broader political type that incumbent represents (Stiers, 2019). Observing visible distributive expansion under a populist mayor may therefore affect support for other

populist candidates or parties, just as experiencing retrenchment under an extreme-right mayor may reduce support for extreme-right alternatives more broadly. The electoral consequences of governing may thus appear both in the incumbent's own political arena and in subsequent elections involving other candidates who share the same political traits.

Taken together, these arguments connect political traits to motivations, governing behavior, and electoral feedback. Extreme-right politicians are expected to behave as policy-seekers, pursuing ideological commitments even when electoral incentives counsel restraint. Populist politicians are expected to favor short-term and visible benefits, while the term-limit analysis (Section 7.2) examines whether this behavior reflects sincere ideological commitment or opportunistic office-seeking. Retrospective accountability then converts the benefits and costs produced by each governing style into subsequent electoral support. The following subsections develop these expectations separately for populist and extreme-right incumbents and derive the corresponding predictions for governing behavior, term limits, and voter response.

1.4 Extreme-Right Politicians in Office

Extreme-right politicians enter office as intrinsically motivated policy-seekers. Extreme-right ideology is a substantively thick doctrine that specifies a concrete programmatic agenda. Politicians who hold this ideology seek office in order to implement it (Hall, 2019), even at the cost of electoral viability (Hall, 2015). For extreme-right politicians, high levels of spending are ideologically undesirable. They pursue policies consistent with their ideological commitment to fiscal consolidation, administrative order, and restrictions on redistribution. This programmatic fiscal retrenchment hurts many voters, who may then punish the extreme-right governing style electorally; the punishment attaches to the experienced retrenchment, not to ideology per se, so an ideological agenda whose implementation delivers visible benefits would by the same mechanism escape punishment.

Term limits carry a different implication for extreme-right incumbents. Because fiscal retrenchment is ideologically motivated rather than electorally strategic, re-election concerns can only restrain it; removing them should *intensify* rather than eliminate the cuts. Term-limited extreme-right mayors, freed from electoral pressure, should implement their preferred policies more aggressively, a lame-duck effect. Where the term-limit comparison speaks to the motivational accounts for populists, for the extreme right it sharpens a single prediction: retrenchment should be strongest among term-limited mayors. Hypothesis 1 summarizes the predictions, and Figure 2 illustrates the causal chain.

Hypothesis 1 (Extreme right).

H1a. *(Electoral response — main prediction)* Extreme-right incumbency generates electoral backlash; electoral support for extreme-right candidates and parties declines in subsequent elections.

H1b. *(Governing behavior)* Extreme-right politicians in office implement their programmatic agenda at the expense of distributive demand.

H1c. *(Term limits)* The retrenchment intensifies under term limits.

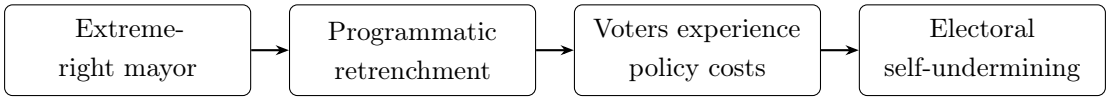

Figure 2: Extreme-right self-undermining mechanism.

1.5 Populist Politicians in Office

Because populism does not prescribe a fixed position on redistribution or public spending, populist incumbents have considerable flexibility in the policies they use to demonstrate responsiveness to ordinary citizens. Municipal officeholders can expand welfare provision in ways that generate visible and readily attributable benefits for residents, while allowing part of the fiscal burden to fall on future budgets or other levels of government (Rodden, 2002, 2006; Rogoff, 1987; Shi and Svensson, 2006). Populist incumbents should therefore be likely to pursue distributive expansion.[1]

Visible distributive benefits should, in turn, sustain or increase support for populist politicians and parties. Unlike policies that impose immediate costs on broad groups of voters, welfare expansion provides a governing record that citizens may reward. Populist incumbency should therefore not generate electoral backlash and may instead reinforce support for populist alternatives.

The term-limit comparison distinguishes between the competing motivational accounts introduced above. If distributive expansion is driven by office-seeking incentives, it should be strongest where re-election remains possible. If it instead reflects a sincere populist commitment, it should persist when electoral constraints are removed.

[1]The electoral mechanism requires that these benefits be experienced. A populist incumbent who lacks the administrative or fiscal capacity to deliver visible expansion should receive no corresponding electoral reward. Heterogeneity in governing capacity should therefore attenuate the estimated relationship rather than produce it. The fiscal results in Section 7.2 assess whether distributive expansion materializes where re-election incentives remain active.

Hypothesis 2 summarizes these expectations, and Figure 3 illustrates the corresponding causal chain.

Hypothesis 2 (Populism).

H2a. *(Electoral response — main prediction)* Populist incumbency does not generate electoral backlash; electoral support for populist candidates and parties is sustained or rises in subsequent elections.

H2b. *(Governing behavior)* Populist politicians in office pursue visible distributive expansion.

H2c. *(Term limits)* If populist distributive expansion is driven by office-seeking incentives, it should be strongest where re-election remains possible; if it reflects a sincere populist commitment, it should persist under term limits.

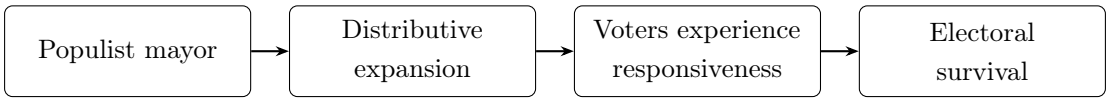

Figure 3: Populist electoral-survival mechanism.

2 Institutional and Political Setting

The costs and benefits of governing for extreme-right and populist actors are a general phenomenon, but they are extraordinarily hard to identify at the national or cross-national level, where institutions, party systems, coalition dynamics, and economic conditions vary together and confound any comparison. Subnational variation within a single country removes those confounds by holding the institutional environment fixed. Beyond removing confounds, Italian municipalities offer two things national data cannot: enough independent variation in populism and extreme-right ideology to separate two traits that party labels bundle together, and enough close races to identify the effect of electing one type of incumbent over another.

The rest of this section describes that setting: the parties that populate the populist and extreme-right space, and the electoral rules that make close mayoral races a usable source of variation.

2.1 The Italian Political Landscape

Italian politics realigned deeply after the early 1990s. The collapse of the traditional mass parties produced a fragmented, candidate-centered system in which new political forces rose and declined rapidly. By the time of the 2019 municipal elections — the anchor cohort

for this study — the party landscape was organized around several poles. The centre-left was anchored by the Partito Democratico (PD), the social-democratic party that governed through much of the 2010s. It was flanked by the smaller Alleanza Verdi e Sinistra (AVS), a left-wing coalition of environmentalist and socialist currents. On the centre-right, Forza Italia (FI), founded by Silvio Berlusconi, occupied a liberal-conservative position and typically governed in coalition with its right-wing allies, the Lega and Fratelli d'Italia (FdI).

The populist and extreme-right space, most relevant to this study, was shaped by three forces. The Movimento 5 Stelle (M5S), founded in 2009 by Beppe Grillo as an anti-corruption, anti-establishment movement, had become the largest single party in the 2018 general election. M5S deliberately refused left–right positioning. Its appeal rested on a people-versus-elite frame, direct democracy, and a rejection of professional politics, amounting to a post-ideological populism with no stable programmatic content. The Lega, originally a northern regionalist movement centered on fiscal federalism and anti-Southern sentiment, transformed under Matteo Salvini into a national extreme-right populist party built on anti-immigration, Euroskepticism, and law-and-order themes. It peaked at 34% in the 2019 European elections. Fratelli d'Italia, heir to the post-fascist Movimento Sociale Italiano, combined nationalist, sovereigntist, and socially conservative commitments with a strong anti-establishment appeal, and would go on to win the 2022 general election. Beyond these parliamentary forces, two openly neo-fascist parties, CasaPound and Forza Nuova, fielded candidates in municipal elections, providing an extreme ideological anchor for validating the text-based measures introduced below.

2.2 Electoral Rules and Municipal Powers

Italian municipal government centers on a directly elected mayor and council, and an executive committee (giunta) the mayor appoints. Elections are held every five years and are staggered across municipalities, typically within a spring electoral window. Each mayoral candidate is tied to one or more electoral lists competing for council seats. These lists may be local branches of national parties, purely *civic lists* (local parties or political entities with no national affiliation), or coalitions combining both. Regardless of list type, each mayoral candidate must submit a single administrative platform (programma amministrativo). This requirement produces one authoritative campaign text per candidate, making municipal elections unusually well suited for measuring candidate-level ideology and populist rhetoric directly from platforms.

The prevalence of civic lists is a defining feature of Italian municipal politics and creates a major limitation for party-label research designs. In the present data, nearly 78%

of mayoral candidates lack a national party label. This is consistent with prior evidence: roughly three-quarters of mayors in small Italian municipalities are affiliated with civic lists rather than national parties (Gamalerio, 2020). Designs that require national party affiliation (Bellodi, Morelli and Vannoni, 2024; Bordignon and Colussi, 2020) therefore lose a large share of the municipal candidate universe and retain a potentially non-representative subset of candidates, a selection concern examined in Section 4. This is especially consequential for close-election designs, where only races with a one-sided party-label contrast can identify an effect, as detailed in Section 4.[2]

Electoral rules vary by municipality size. In municipalities with fewer than 15,000 inhabitants, the candidate receiving a plurality of votes is elected mayor. In larger municipalities, an absolute majority is required: if no candidate receives more than 50% in the first round, the top two candidates advance to a runoff held two weeks later.

Beyond elections, Italian municipalities hold substantial governing powers, with authority over primary education, local policing, waste management, roads and infrastructure, and social services. They also have meaningful fiscal autonomy and manage procurement for goods and services. Mayors concentrate executive authority within this system, appointing and dismissing the giunta and exercising substantial discretion over hiring, dismissal, and task assignment for top bureaucratic positions. This combination of candidate-centered elections, mandatory platforms, fiscal responsibility, and mayoral authority makes Italian municipalities well suited for studying how populist and extreme-right candidates govern once in office.

3 Data

3.1 Data sources

The core dataset is a candidate-level corpus of municipal electoral platforms, one platform for each mayoral candidate. Each platform is linked to official municipal election returns, candidate records, and municipality-level public finance data. The election returns identify the full candidate set, winners and runners-up, vote totals, vote shares, turnout, and the winner–runner-up margin. The public finance data measure municipal revenues and expenditures during the subsequent governing term; they supply the fiscal outcomes, welfare spending above all, on which the framework's governing-behavior predictions are tested. Together,

[2]The absence of national party labels does not make these races ideologically illegible to voters: platforms are public by law, campaigns in small municipalities rely on direct contact, and the positions that define the extreme right locally are recognizable without a party logo. The electoral mechanism also does not require voters to read platforms, since they experience the governing record and update accordingly.

these sources connect campaign rhetoric, close-election outcomes, fiscal choices in office, and downstream electoral responses.

3.1.1 Municipal electoral platforms

Italy does not provide a centralized archive of municipal electoral platforms. Instead, these documents are managed locally by municipal administrations. Under the Consolidated Text on Local Authorities (TUEL),[3] each mayoral candidate must submit an administrative platform as part of the formal candidacy package.[4] The TUEL further requires that the platform be publicly disclosed by posting it on the municipal noticeboard (albo pretorio), the official public bulletin for municipal acts.[5] In addition, citizens may request copies of these documents directly from municipal offices and, when platforms are not readily accessible, may rely on Italy's freedom-of-information right of civic access (accesso civico generalizzato).[6]

In practice, municipalities do not systematically preserve electoral platforms online. Documents are often removed following website redesigns, routine content pruning, or changes in administrative software. The platform corpus was assembled through direct outreach to municipal administrations over approximately six months. Requests were sent primarily via certified email (PEC, posta elettronica certificata) and addressed to electoral offices, citizen relations offices (URP), the mayor and vice mayor, and the segretario comunale (chief administrative officer). When municipalities required a formal procedure, an accesso civico generalizzato request was submitted using the municipality's standard form, with identification provided in accordance with local requirements. Platforms were collected from official municipal records whenever available. When municipal offices did not respond or reported that the documents were unavailable, the corresponding platform remained missing from the corpus. Appendix A reports platform recovery rates and documents the geographic and size gradient in recovery.

Italian municipal elections are staggered across years and held within a single annual spring election window.[7] The analysis starts from the 2019 election round, the largest annual cohort, when approximately 3,700 municipalities went to the polls out of roughly 7,900 nationwide. To measure within-municipality change over one full term (the statutory five-year mayoral mandate),[8] the sample is restricted to municipalities that held mayoral elections in both

[3] Testo Unico degli Enti Locali (TUEL), D.Lgs. 267/2000.
[4] Art. 71(2) for municipalities with up to 15,000 residents; Art. 73(2) for larger municipalities.
[5] The posting requirement appears in TUEL Arts. 71(2) and 73(2).
[6] *Accesso civico generalizzato*, Art. 5(2) of D.Lgs. 33/2013, as amended by D.Lgs. 97/2016.
[7] Law 182/1991, Art. 1 establishes the ordinary election period between April 15 and June 15.
[8] TUEL, D.Lgs. 267/2000, Art. 51(1).

2019 and 2024, yielding a universe of 3,592 municipalities. Within this universe, 6,336 candidate platforms were recovered across the 2019 and 2024 election rounds, covering 1,412 municipalities; Appendix A reports recovery rates by macro-area and population size.

This restriction implicitly excludes municipalities where the 2019 mandate was cut short by early termination of the mayoral term—a small subset of 192 municipalities, or 5.1% of the 2019 total in available regions.[9] The pair-based RD design further requires at least two candidates in the 2019 election; uncontested 2019 races (274 municipalities in the universe) are therefore excluded. The sample covers all regions for which official election results are publicly available online; Valle d'Aosta, Trentino-Alto Adige, and Sicily are excluded because their electoral data are not centrally available. Sardinia is included in the universe but its 2019 mayoral elections, held on a special calendar separate from the main spring window, covered only a small number of municipalities. Platform recovery within the cohort is non-random: recovery rates vary across regions and municipality size classes. To address this, the main electoral results are re-estimated with inverse probability weighting based on region and municipality size (Appendix A); the headline findings are robust to this re-weighting.

3.1.2 Election results and fiscal outcomes

Election results (Eligendo). Official election returns come from *Eligendo*, the Ministry of the Interior's election-results portal.[10] Eligendo provides standardized municipality-level tabulations for administrative elections, including the full candidate set, first-round and (when applicable) runoff results, vote totals and vote shares, and turnout. These data are used to identify winners and runners-up, construct the winner–runner-up margin, and link each platform to its corresponding candidate and election.

Municipal public finance (SIOPE). The fiscal outcomes used to measure governing behavior in office are drawn from SIOPE (*Sistema Informativo sulle Operazioni degli Enti Pubblici*), a transaction-based system that records cash payments and receipts of public administrations using a harmonized coding scheme.[11] It reports municipal revenues (*entrate*) and expenditures (*spese*) by year and economic category.[12] The SIOPE panel is matched to

[9] Under Italian law, a mayoral term ends before its statutory five-year duration in several circumstances: voluntary resignation of the mayor (TUEL Art. 53); a motion of no confidence (*mozione di sfiducia*) approved by the absolute majority of assigned councilors, which simultaneously dissolves the council (TUEL Art. 52); simultaneous resignations of more than half of assigned councilors (TUEL Art. 141(1)(b)); failure to approve the budget within statutory deadlines (TUEL Art. 141(1)(c)); or permanent impediment, removal, or death of the mayor (TUEL Art. 141(1)(a)). In all cases, the council is dissolved by presidential decree on proposal of the Minister of the Interior, and new elections are called for the following spring window.

[10] https://elezionistorico.interno.gov.it/eligendo

[11] https://www.rgs.mef.gov.it/VERSIONE-I/e_government/amministrazioni_pubbliche/siope/

[12] https://www.siope.it/

the election and platform data, restricting the sample to municipalities that elected a mayor in 2019 and extracting fiscal outcomes over the 2019–2024 mandate.

4 Measuring Ideology and Populism

The introduction identified two measurement limitations of party-label approaches: conflation of populism with extreme-right ideology, and the absence of continuous candidate-level measures of populism. To address both, this section builds continuous, candidate-level measures of ideology and populism directly from the platforms themselves. Two further issues specific to close-election designs warrant emphasis before turning to the measurement approach.

First, party labels impose a binding sample restriction under close-election designs. The PCRD estimand requires a meaningful contrast between winner and runner-up; under a binary party-label classification, this is only possible when exactly one of the two top candidates carries the relevant label. Municipalities where both leading candidates are affiliated with populist or extremist parties, or where neither is, must be dropped. For example, Bracco et al. (2018) employ a regression discontinuity design that restricts the sample to mayoral races involving a Lega Nord candidate, discarding municipalities with no Lega Nord candidate. The candidate-level continuous scores developed in this section do not share this limitation: a within-pair contrast exists in every election.

Second, party-label restrictions introduce a selection concern. Extremist or populist politicians may strategically run as independents or local-party candidates where overt ideological branding is electorally costly, making observability endogenous to local conditions. Studies of Italian municipal politics that restrict the analysis to candidates affiliated with national parties (Bellodi, Morelli and Vannoni, 2024) therefore work on a selected subset. These candidates are plausibly more disciplined, better resourced, and more tightly embedded in party organizations than their unlabelled counterparts. Such party-label designs therefore work best in contexts where party labels are universal among candidates. In the Italian case, where about three-quarters of candidates run as independents or affiliated with a local party (civic lists), this requirement discards a large share of the candidate universe (Gamalerio, 2020).

4.1 Why Supervised Classification?

Text-based measurement of ideology and populism can be approached through unsupervised scaling, dictionary counts, or supervised classification. This paper adopts supervised classification. The chosen classifier is a multilingual transformer (XLM-RoBERTa; Conneau

et al., 2020; Liu et al., 2019), fine-tuned to assign each paragraph to the relevant category. It captures contextual meaning beyond word counts, distinguishing, for instance, whether a reference to "the people" is part of a populist appeal or a routine administrative formulation. Outputs are retained as continuous probability scores rather than hard class assignments, preserving measurement uncertainty. In cross-lingual political text classification benchmarks, XLM-RoBERTa outperforms both multilingual BERT and DeBERTa (Timoneda and Vallejo Vera, 2025; Sebők et al., 2025); fine-tuned classifiers of this type also outperform LLM prompting once a sufficient labeled training set is available (Bosley et al., 2023; Wang, Qu and Ye, 2024).

Supervised classification is preferred because the alternatives have known limitations in this setting. Unsupervised methods such as Wordfish and Wordscore (Slapin and Proksch, 2008; Laver, Benoit and Garry, 2003) recover corpus-specific latent dimensions that need not correspond to theoretically defined constructs (Rodriguez and Spirling, 2022; Rheault and Cochrane, 2020). In heterogeneous municipal platforms, the principal axes of variation may reflect administrative style rather than ideology or populism. Dictionary-based approaches (Pauwels, 2011; Young and Soroka, 2012) prove brittle because identical lexical items serve administrative, generic, or antagonistic functions depending on context. A dictionary that treats words such as "security" as ideological markers would conflate road safety and the structural safety of school buildings with policing or surveillance framed around immigrants and undeserving outsiders. Similarly, a dictionary that counts references to "citizens," "participation," or "referendums" as evidence of people-centrism would be unable to distinguish citizens framed as service recipients from citizens framed as decision-makers with authority to overrule the council. The problem is therefore not that the relevant words are rare or noisy, but that the same words change meaning with context, which word counts alone cannot capture.

Four task-specific classifiers are estimated — one each for ideology, populism, AER, and PC. Training paragraphs are drawn from (i) municipal platforms from the 2015–2017 election cycles and (ii) national-level platforms from the Manifesto Project. Model selection and early stopping use a development set constructed from the target 2019 and 2024 municipal corpus, disjoint from training data. Paragraph-level predictions are aggregated to candidate-level scores. Full codebooks, annotation protocol, training dynamics, per-class performance, and a reliability check in which the labeled samples are independently re-coded are reported in Appendices B–B.5.

4.2 Measuring Ideology

> *"In allocating scarce public resources, priority must be given to Italian citizens and long-standing residents who have contributed to the community through work and respect for the rules. Social housing and welfare cannot become incentives for uncontrolled immigration or tools that penalize those who have always lived here."*
>
> *Source: Municipal electoral platform (local civic list), Italy.*

This paragraph is a paradigmatic example of the extreme-right category used in the ideology codebook. The policy domain is welfare, but the ideological signal does not come from welfare policy as such. It comes from the exclusionary justification attached to it: scarce public resources are said to belong first to Italian citizens and long-standing residents, while immigration is framed as an illegitimate drain on social housing and welfare. The classifier is trained to recognize this kind of substantive ideological framing, rather than simply counting references to welfare, legality, or immigration.

The classifier assigns each paragraph to one of five ordered categories: *extreme left*, *left*, *centre*, *right*, and *extreme right*. *Extreme left* texts articulate maximalist or systemic left positions: explicit opposition to privatization, strong labor-protection claims, demands for public ownership or direct provision of services. *Left* texts advance progressive but non-maximalist agendas—inclusive education, anti-discrimination policies, recognition of social diversity that are framed in terms of equity and pluralism rather than systemic rupture. *Centre* is the default category: managerial, technocratic, or pragmatic language that emphasizes service provision, infrastructure, and efficiency without an explicit left or right orientation. *Right* texts emphasize institutional order, legality, and authority as well as cooperation with law enforcement. *Extreme right* texts are defined by exclusionary nationalism: welfare chauvinism and anti-immigration rhetoric that explicitly prioritize citizens over outsiders in access to public resources. The analysis covers the full left–right spectrum, though estimated electoral effects emerge primarily on the right. The detailed codebook is reported in Appendix S.1, and illustrative examples in Appendix B.4.

4.3 Measuring Populism

> *Just as any abandoned good becomes a breeding ground for the worst infestations of insects and rodents, causing its progressive rotting and decay, so too has local politics over the last twenty years—hell-bent on gutting the public coffers to honor their backroom deals, without the slightest concern for the issues that matter to ordinary citizens—brought us to a state of total and utter degradation.*
>
> *Source: Municipal electoral platform (local civic list), Italy.*

This paragraph is a paradigmatic example of hard populism. Its populist content does not come from a specific policy proposal, nor from its position on the left–right spectrum. It comes from the moralized antagonism it constructs. Local politicians are portrayed as a corrupt and self-serving elite that has gutted public resources through backroom deals, while ordinary citizens appear as the neglected victims of that corruption. The passage therefore captures the core textual signal the classifier is trained to detect: a people-versus-elite frame in which political conflict is presented not as ordinary disagreement over policy, but as a struggle between a betrayed public and an unresponsive or corrupt governing class.

Populism is notoriously hard to measure. Even its basic definition is contested, spanning thin-ideology (Mudde, 2004; Mudde and Kaltwasser, 2017), strategic (Weyland, 2001), discursive (Laclau, 2005), and stylistic (Moffitt, 2016) accounts. Because it is defined by framing rather than policy content, the same words can be populist in one context and innocuous in another. To avoid resting the results on a single operationalization, the paper measures populism in three ways. The main measure is a holistic populism index, described below. In addition, two measures capture populism's two constitutive components separately, anti-elite rhetoric (AER) and people-centrism (PC); these serve as a diagnostic check on whether the populism results are robust to how the concept is operationalized. AER captures the intensity of accusations against a corrupt, self-serving elite. PC captures the degree to which ordinary citizens are framed as a virtuous community whose will should guide political decisions.

The index is coded at the paragraph level on a three-level scale, designed to capture populist rhetoric as comprehensively as possible, from explicit antagonism to milder appeals that a binary classification would miss. *Hard* populism is defined by explicit antagonistic claims that either attribute intentional wrongdoing to political elites (corruption, clientelism, capture by powerful interests) or assert exclusive representation of the popular will. *Soft* populism captures people-centric orientations or less harsh anti-elite rhetoric that stops short of attributing intentional wrongdoing. It includes appeals to moral renewal, localism, or proximity to citizens, and portrayals of elites as distant or unresponsive rather than corrupt. *Non-populist* is the default category: it covers technical, managerial, or descriptive policy language, as well as any other political text that lacks a people-versus-elite framing. A nativist or exclusionary policy demand, for instance, is ideologically extreme but not populist, because populism is defined by the moral opposition between people and elite, not by policy content. The index is a holistic classification: annotators judge whether a paragraph's overall framing conveys a moral people-versus-elite opposition.

The two sub-dimensions are coded analogously; their full codebooks appear in Appen-

dices S.18 and S.25, and illustrative examples for all three measures in Appendix B.4.

4.4 Validating the Measures

Having defined the ideology, populism, AER, and PC measures, two checks indicate they are well measured. First, the classifiers perform well out of sample: all four reach macro-F1 between 0.75 and 0.82 on a held-out municipal development set, with accuracy up to 0.86, at or above benchmarks reported for supervised political-text classifiers, including models built on the same XLM-RoBERTa architecture (Sebők et al., 2025; Timoneda and Vallejo Vera, 2025) (full per-class results in Appendix B). Second, the resulting candidate-level scores recover familiar party-level patterns. This exercise is not used for estimation. It is a face-validity check: among candidates whose list names can be matched to one identifiable national party, the scores should place parties in roughly expected positions while also preserving the distinction between ideology and populism. Figure 4 and Table 1 confirm both expectations. The recovered party centroids reproduce the parties' established left–right ordering. From left to right, candidates on communist- or socialist-named lists (CS) and AVS occupy the furthest-left positions, followed by PD and M5S just left of center, then FI, FdI, and Lega on the center-right, with CasaPound and Forza Nuova furthest right. This ordering matches the parties' relative placement in the Chapel Hill Expert Survey (Bakker et al., 2019; Rovny et al., 2025), an external benchmark the scores were not trained on. The populism axis is largely orthogonal to ideology, confirming the conceptual independence of the two dimensions: M5S ranks highest on populism among the major national parties at a near-centrist ideological position, consistent with its identity as a post-ideological populist movement. Lega, by contrast, scores only modestly. Its municipal platforms take a more administrative tone, focused on local services, infrastructure, and territorial management rather than the anti-elite antagonism of its national campaigns.[13] This is not entirely surprising: the Lega's posture is known to shift between antagonistic opposition nationally and pragmatic government subnationally (Pirro, 2024).

[13] A random sample of 50 Lega municipal platforms was manually reviewed to rule out classifier error. Many Lega municipal candidates are pragmatic administrators whose platforms focus on local services and territorial management, with little of the people-versus-elite framing of the party's national campaigns. Party labels would classify all of them identically; text-based scores capture that carrying a Lega endorsement does not make a municipal candidate populist.

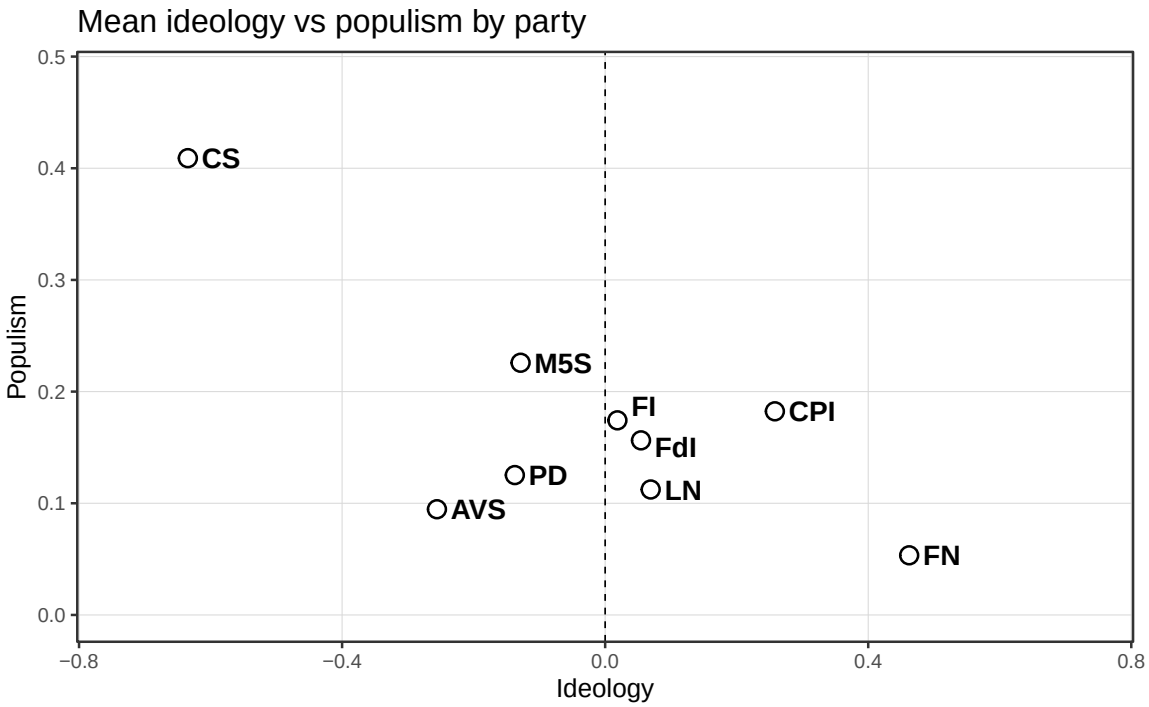

Figure 4: Party centroids in ideology–populism space. Each point is a party centroid (mean scores across affiliated candidates in the sample). Ideology (x-axis) and populism (y-axis) are continuous scores produced by the XLM-RoBERTa classifier. The dashed vertical line marks the ideological center. CS = candidates affiliated with lists whose name contains "communist" or "socialist"; AVS = Alleanza Verdi e Sinistra; PD = Partito Democratico; M5S = Movimento 5 Stelle; FI = Forza Italia; FdI = Fratelli d'Italia; LN = Lega; CPI = CasaPound; FN = Forza Nuova.

Table 1: Mean scores by party. Mean classifier scores for candidates whose electoral list name matches exactly one identifiable national party, aggregated across the 2019 and 2024 electoral cycles, ordered left to right by mean ideology. Standard deviations in parentheses. AER = anti-elite rhetoric; PC = people-centrism.

Party	*N*	Ideology	Populism	AER	PC
CS	62	-0.63	0.41	0.40	0.10
		(0.37)	(0.23)	(0.10)	(0.09)
AVS	6	-0.26	0.09	0.28	0.05
		(0.12)	(0.07)	(0.02)	(0.03)
PD	190	-0.14	0.13	0.27	0.09
		(0.10)	(0.09)	(0.03)	(0.05)
M5S	194	-0.13	0.23	0.28	0.20
		(0.15)	(0.12)	(0.04)	(0.15)
FI	20	0.02	0.17	0.28	0.11
		(0.16)	(0.17)	(0.07)	(0.13)
FdI	27	0.05	0.16	0.28	0.09
		(0.09)	(0.12)	(0.04)	(0.08)
LN	202	0.07	0.11	0.26	0.07
		(0.25)	(0.10)	(0.04)	(0.05)
CPI	17	0.26	0.18	0.32	0.07
		(0.41)	(0.11)	(0.05)	(0.05)
FN	8	0.46	0.05	0.25	0.06
		(0.48)	(0.04)	(0.03)	(0.06)

The party centroids above are based on candidates with national-party labels. One might worry that candidates without such a label add little. Their platforms might contain no clear ideological positions and little populist rhetoric. The data say otherwise. Candidates running as independents or under local parties (civic lists) show populism comparable to nationally affiliated candidates (mean populism 0.16 vs 0.15). Their ideological extremism is substantial, if somewhat lower (mean $|ideology|$ 0.11 vs 0.15). They span the full range of populism and ideology in the sample.

5 Close-Election Research Design

This paper estimates the effect of electing mayors with higher levels of extreme-right ideology or populist rhetoric using close-election politician-characteristic regression discontinuity (PCRD) designs. The unit of analysis is the municipality. The design compares close races in which the higher-scoring candidate barely wins with races in which that candidate barely loses. Near the electoral threshold, which of the two leading candidates wins is treated as

quasi-random, following the logic of close-election designs (Hall, 2015; Marshall, 2024; Bertoli and Hazlett, 2025).

5.1 Traditional Gradient PCRD

The gradient PCRD signs the winner–runner-up vote margin by the within-race score comparison between the two leading candidates. Positive values indicate that the higher-scoring candidate wins; negative values indicate that the higher-scoring candidate loses. The discontinuity at zero then estimates the effect of barely electing the higher-scoring candidate rather than the lower-scoring candidate.

The limitation is that this estimator is binary. As Hall (2015) notes, the treatment indicator "stands in for a variety of treatment intensities or doses": a race in which a far-right candidate narrowly defeats a centrist and one in which two near-centrist candidates face each other enter as the same threshold-crossing contrast. Existing applications address this by imposing score-gap calipers (Hall, 2015; Meisels, 2025). A caliper is a minimum within-pair score-gap threshold: races in which the two leading candidates have nearly identical scores are dropped, so that the close-election contrast captures a substantive switch in the elected mayor's ideology or populism rather than a nominal sign change between candidates whose positions are essentially the same. The main gradient specification uses a small caliper, dropping only the bottom 5% of within-pair score gaps, just enough to exclude contrasts too small to be informative. Sensitivity to stricter cutoffs is reported in Appendix D.5.

For ideology, the design is implemented on two separate samples. The right-side sample (S^R) consists of races with at least one right-leaning candidate among the top two, with the margin signed so that positive values indicate victory by the more right-wing candidate. The left-side sample (S^L) is constructed symmetrically for races with at least one left-leaning candidate among the top two. Both samples are reported. For populism, which has no left–right direction, the design compares the more populist candidate with the less populist candidate in each race.

5.2 Continuous-Treatment PCRD

The continuous-treatment PCRD keeps the same close-election instrument but replaces the binary treatment with the elected mayor's continuous ideology or populism score. The estimand is a local dose-response: how outcomes change as the elected mayor's score increases, induced by barely winning a close race. This adapts continuous-treatment fuzzy RD logic to the politician-characteristic setting, building on Hahn, Todd and Van der Klaauw

(2001); Dong, Lee and Gou (2023) and on related applications to geographic discontinuities (Ebenstein et al., 2017).

Treatment intensity enters the estimator directly. Races with larger within-pair score differences generate larger first-stage shifts in the elected mayor's score; races with nearly identical candidates generate little movement in the continuous treatment. Because the treatment is the elected mayor's score itself, rather than a binary indicator for whether the higher-scoring candidate wins, no score-gap caliper is required. The continuous-treatment estimator thus preserves information that the binary gradient contrast collapses, while the gradient estimator remains useful as the reduced-form close-election effect.

All estimates use local-linear RD specifications with triangular kernels. The main specifications include region fixed effects and log municipal population as pre-treatment covariates. The appendix reports the formal estimands, implementation details, bandwidth choices, diagnostics, density and balance tests, and sensitivity to caliper restrictions. The two estimators should be read together: gradient PCRD gives the reduced-form effect of barely electing the higher-scoring candidate, while continuous-treatment PCRD gives the corresponding dose-response per unit of mayoral ideology or populism.

The credibility of the close-election design rests on two conditions. The candidates on each side of the cutoff should be comparable on pre-determined traits, and the bundled nature of the treatment should be read as part of the estimand rather than a violation of it (Bertoli and Hazlett, 2025). A related concern, *compensating differentials* (Marshall, 2024), is that close-race candidates may systematically differ on other dimensions. Following Bertoli and Hazlett (2025), this does not threaten the internal validity of the local PCRD estimate, though it limits extrapolation beyond the close-race regime. Appendix E reports balance tests on eight pre-determined traits and a description of the bundle each side of the cutoff selects. The candidates are balanced under the right-wing and people-centrism designs. The anti-elite-rhetoric design shows some imbalance and is therefore treated as descriptive.

6 Electoral Results of Extremism and Populism

This section tests whether voters respond differently to extreme-right and populist incumbency. The theoretical framework predicts an asymmetric pattern (Hypotheses H1a and H2a). Extreme-right mayors should be electorally self-undermining: by pursuing programmatic retrenchment, they should reduce subsequent support for the extreme right. Populist mayors should not generate the same backlash: by pursuing visible distributive expansion, they

should sustain, and possibly reinforce, subsequent support for populist candidates and parties.

Electoral Responses to Incumbency

The municipal results show a clear asymmetry. Barely electing a more right-wing mayor in 2019 moves the next local election away from the right. Barely electing a more populist mayor, by contrast, does not produce a comparable decline in populist support. The local electorate therefore appears to punish extreme-right incumbency, but not populist incumbency.

The main outcome is the political center of gravity of the 2024 mayoral race. For each municipality, candidate ideology and populism scores are weighted by the vote share each candidate receives. This vote-weighted measure captures both parts of the electoral response: whether voters shift their support across candidates, and whether different kinds of candidates enter the race. A marginal fringe candidate barely affects the measure, while a candidate who attracts substantial support moves it more strongly.

Table 2 reports the ideology results. In municipalities where one of the two leading candidates in 2019 was right-leaning, barely electing the more right-wing candidate shifts the 2024 race's ideological center of gravity leftward. The gradient PCRD estimate is -0.058 ($p = 0.002$), and the CT-PCRD estimate is -0.271 ($p = 0.013$). The first stage is strong (0.204, $p < 0.001$), meaning that close victories by the more right-wing candidate produce a large change in the ideology of the elected mayor. In the symmetric left-side sample, the estimates are close to zero and statistically insignificant. The effect is therefore not a generic backlash against ideological extremism. It is concentrated on the right.

The size of the first stage helps interpret the result substantively. Just below the cutoff, where the more right-wing candidate barely lost, the elected mayor's ideology score is -0.062, close to the ideological center. Just above the cutoff, where the more right-wing candidate barely won, the elected mayor's score rises to $+0.144$. This is well to the right of the Lega party mean ($+0.07$) and roughly halfway toward CasaPound ($+0.26$; Table 1). The local close-race contrast therefore does not simply compare two moderate candidates. It moves municipalities from near-centrist mayors to mayors with substantively extreme-right profiles.

Table 2: Effect of Barely Electing a More Extreme Mayor on the Vote-Weighted Ideology in 2024.

Outcome: vote-weighted ideology in 2024

	Right sample ($\mathcal{S}^R$)		Left sample ($\mathcal{S}^L$)	
	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	-0.058^{***}	-0.271^{**}	0.026	0.165
SE	(0.019)	(0.109)	(0.019)	(0.130)
p-value	0.002	0.013	0.151	0.204
First stage	—	0.204^{***}	—	0.159^{***}
SE		(0.032)		(0.023)
p-value		0.000		0.000
N (bandwidth)	278	368	411	359
N (total)	773	814	779	820
Bandwidth	0.152	0.188	0.229	0.182

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. N (bandwidth): observations within the MSE-optimal bandwidth on each side of the cutoff. N (total): full directional sample used for estimation. Gradient PCRD columns apply a first-stage intensity filter that drops races with |diff| below the 5th percentile, following Hall (2015); CT-PCRD columns use the full sample, since the Wald-ratio structure downweights low-dose races automatically. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The same pattern does not appear as cleanly for populism. Barely electing a more populist mayor in 2019 does not produce a stable municipal-level shift across the paper's three populism measures. The overall populism score and anti-elite rhetoric remain essentially unchanged in the 2024 race, while people-centrism increases. Because the positive result is confined to one subdimension, the municipal evidence should be read cautiously: it does not show a broad reinforcement of populist support, but it also provides no evidence of populist backlash. A backlash hidden by noise would be unlikely to leave this pattern across all three measures (full estimates in Appendix F, Table 29).

The distributional evidence reinforces this interpretation. The vote-weighted average could, in principle, hide changes in different parts of the candidate distribution. To test this, the treatment is kept fixed, while the outcome is changed from the average score of the 2024 race to the vote share received by candidates above different ideology or populism cutoffs. This shows whether incumbency affects support only for moderately high-scoring candidates or also for the most extreme and most populist candidates. The results show that barely electing a more right-wing mayor reduces the vote share of high-ideology candidates across

the distribution, including the most extreme-right candidates. By contrast, barely electing a more populist mayor leaves the vote share of high-populism candidates essentially unchanged. The municipal backlash against extreme-right incumbency is therefore not an artifact of averaging, and the populism null is not driven by offsetting movements across different parts of the race (Appendix H, Figure 40).

Several additional checks support the main result. The ideology finding is similar when the outcome is measured as a simple unweighted average across candidates, indicating that the effect reflects not only how voters allocate support among candidates but also which candidates enter the race (Appendix G). Standard close-election diagnostics following Bertoli and Hazlett (2025), including density continuity, covariate balance, predicted treatment, and predicted potential outcomes, are reported in Appendix D and apply to all results in this section. Because right-wing bare winners are less often prior incumbents, the headline regressions are also re-estimated with both-side incumbency indicators; the coefficients are essentially unchanged (Appendix E). A party-label benchmark, matching the conventional approach in much of the literature, points in the same direction for the extreme-right dimension. Using Lega/FdI affiliation, the estimates are negative for the combined Lega/FdI vote in 2022 and 2024 and for 2024 municipal candidate branding, although the latter two are imprecise (Appendix K). No analogous benchmark is available for populism because M5S is nearly absent as a formal municipal label in the close-race sample.

The appendix also extends the analysis beyond the municipal arena by testing whether local incumbency spills over into national and European elections. Barely electing a more extreme-right mayor reduces Lega vote shares in the 2024 European election. For populism, the main populism measure predicts higher Movimento 5 Stelle vote shares in both the 2022 parliamentary election and the 2024 European election.[14] The people-centrism measure shows the same pattern, while anti-elite rhetoric is weaker but directionally consistent. These spillover results reinforce the central asymmetry: extreme-right incumbency reduces support for the extreme right, while populist incumbency does not generate backlash and is associated with higher support for the main populist party in national-level contests (Appendix I).

A final diagnostic shows what is lost when populism and extreme-right ideology are collapsed into a single measure. The two traits are recombined into a single equal-weight populist–extreme-right index, mimicking the populist radical right composite of the literature (Mudde, 2007), and the close-election design is run on this index alone, over all close races. The design

[14]The 2022 estimate attenuates once pre-treatment M5S support is controlled, whereas the 2024 estimate is robust to that control (Appendix D, Table 24). People-centrism, the populism design balanced on pre-treatment covariates, shows its clearest M5S effect in 2022.

recovers nothing: the vote-weighted electoral effect is insignificant under both estimators (-0.140, $p = 0.421$ gradient; -0.064, $p = 0.570$ CT-PCRD; Appendix J).

7 Mechanisms: Governing as Office-Seekers or Ideologues

The electoral results show an asymmetry: extreme-right incumbency is followed by electoral backlash, while populist incumbency is not. This section asks whether the two types also govern differently. The framework predicts that extreme-right mayors pursue programmatic retrenchment even when it is electorally costly, and that populist mayors pursue visible distributive expansion (Hypotheses H1b and H2b). Whether the expansion reflects office-seeking or sincere policy commitment is the question the term-limit comparison is designed to illuminate.

Two diagnostics follow from this argument. First, the policy record should move in the direction predicted by the theory: extreme-right incumbency should reduce targeted welfare provision, while populist incumbency should not produce comparable retrenchment. Second, term limits offer a descriptive contrast between office-seeking and policy-seeking behavior. If office-motivated, populist expansion should appear mainly when re-election is possible; extreme-right retrenchment, being ideologically motivated, should intensify when mayors are freed from electoral pressure.

7.1 Fiscal Policy as Governing Behavior

If the electoral effects operate through retrospective accountability, the two types of mayors should leave different governing records. Municipal fiscal policy provides the most systematic way to observe that record.[15] The analysis uses SIOPE expenditure data during the 2019 mayoral term, covering 2022 and 2023 (full years) and 2024 in two cumulative windows — through March (24M3) and June (24M6), annualized — that fall under the outgoing administration.[16] Each outcome is measured in budget-share and per-capita form, with the 2018 value serving as both a placebo and a pre-trend covariate.

The revenue side provides little evidence of a mechanism. Barely electing a more right-wing or more populist mayor does not systematically change the tax burden on residents.[17] The

[15] Municipalities of Sardinia do not have SIOPE data in the sample; the two other smallest regions, Molise and Basilicata, are additionally omitted for covariate stability (a combined 12 close-race municipalities). This convention is applied throughout the rest of the paper.

[16] Fiscal years 2020 and 2021 are excluded because COVID-related emergency spending and transfers distort municipal fiscal patterns; 2018 serves as the pre-treatment baseline.

[17] The outcome is the total of own-source municipal taxes paid directly by residents (SIOPE classification 1.01.01): the property tax (IMU), the municipal income-tax surcharge (*addizionale comunale IRPEF*),

estimates are small and unstable across post-treatment periods, and the pre-treatment placebo does not suggest a meaningful pre-existing difference.[18] The main behavioral margin is therefore expenditure, not taxation (Appendix L).

The expenditure analysis focuses on welfare provision. This is the policy area most directly connected to the theory. Welfare spending produces visible benefits, is politically salient to recipients, and can be adjusted during a mayoral term. The main outcome is targeted welfare spending: municipal expenditure directed to residents selected by need, disability, income, or other eligibility criteria. Two nested versions are used: a Core measure (W_c) and a broader Full measure (W_c^{++}). Universalist services, such as childcare, school meals, and school transport (U_c), are analyzed separately as a contrast category, since they are open to all residents (Korpi and Palme, 1998; Ferrera, 1996; Esping-Andersen, 1990).[19]

Figure 5 shows the central pattern. Barely electing a more right-wing mayor is followed by lower targeted welfare spending. The estimates are consistently negative across both definitions of targeted welfare and across both budget-share and per-capita measures. This pattern fits the policy-seeking account: extreme-right mayors do not merely lose support after governing; they also govern in a way that plausibly generates backlash, by reducing targeted welfare provision.[20]

The contrast with universalist services shows that the retrenchment is peculiar to targeted spending. Extreme-right incumbency produces no comparable cuts in services open to all residents; the reductions fall on benefits reserved for selected recipients. This points to distributive priorities rather than a general contraction of local expenditure (Appendix N).

---

the waste charge (TARI/TASI), the tourist tax (*imposta di soggiorno*), the levy on occupation of public spaces (TOSAP/COSAP), construction-amnesty payments (*condoni edilizi*), and minor own-source levies (*altre imposte*). Excluded are user fees for municipal services (kindergarten, school meals, transit, library), which are service prices rather than extractive levies; fines (*multe e sanzioni*), which are punitive penalties triggered by specific behaviors; intergovernmental transfers from the State, regions, and the European Union; state-allocated tax shares and perequative transfers (*compartecipazioni*, *Fondo Solidarietà Comunale*); and capital revenue (asset sales and capital grants).

[18] Under CT-PCRD, the people-centrism design produces a marginally significant 24M6 EUR-per-capita estimate; the gradient PCRD equivalent is null.

[19] W_c includes items explicitly classified as social assistance in the SIOPE descriptor; W_c^{++} adds items whose targeting is established by named eligibility laws (L. 104/1992 on disability, D.Lgs. 468/1997 on socially-useful work, L. 62/2000 on need-based study grants, and L. 328/2000-authorized assessment for residential placements).

[20] The retrenchment pattern is also visible under the Lega/FdI label treatment: post-treatment welfare estimates are negative across both definitions of targeted welfare, while the 2018 placebo shows no comparable pattern (Appendix K).

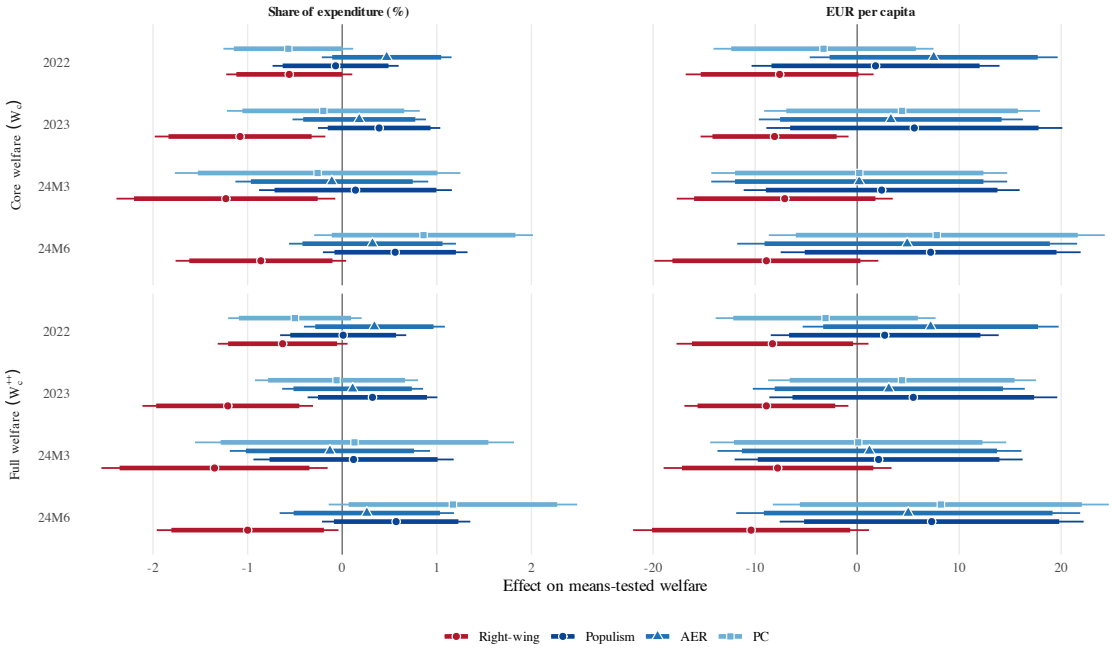

Figure 5: Effect of treatment on means-tested welfare expenditure (Gradient PCRD). Points are bias-corrected gradient PCRD coefficients; thick segments are 90% confidence intervals, thin segments 95%. Red denotes the right-wing ($\mathcal{S}^R$) design; blues denote the populism designs (PopT, AER, PC). Outcomes are means-tested welfare in budget-share (%) and EUR-per-capita form, for the Core (W_c) and Full (W_c^{++}) measures. The 2018 placebo is omitted; full numerical estimates are in Appendix Table 42.

The full-sample populism pattern is less decisive. The estimates generally point more toward expansion than retrenchment, especially in per-capita terms, but they are not stable enough to support a strong average effect. If the expansion is electorally motivated, it should be strongest where re-election remains possible (Hypothesis 2c); a full sample that mixes term-limited and re-election-eligible mayors would then dilute it. The term-limit split therefore provides the sharper test.

The single index that bundles populism and extreme-right ideology is equally uninformative: just as it recovers no electoral feedback (Section 6), it yields no detectable fiscal effect. The mixed index captures nothing on either margin; the effects appear only when the two traits are measured separately (Appendix J).

The next subsection presents the term-limit comparison.

7.2 Office-Seekers and Ideologues Under Term Limits

This comparison groups municipalities by whether the elected mayor was term-limited; because the difference between the two subgroup estimates cannot be estimated precisely, it

is read as descriptive heterogeneity rather than as a causal interaction. For populist mayors, this is where the two accounts would diverge: if welfare expansion is electorally motivated, it should be strongest among mayors who can seek another term; if it reflects a sincere populist commitment, it should persist among mayors who cannot. For extreme-right mayors, the policy-seeking account yields a single prediction. Their programmatic agenda is not designed to maximize short-term electoral support; re-election incentives can only discipline how far they go, since mayors who expect to face voters again have reason to soften, delay, or limit policies that impose visible costs on residents. Term limits weaken that constraint. Once re-election is no longer available, extreme-right mayors should be freer to pursue their ideological preferences in office. If those preferences favor welfare retrenchment, the cuts should be strongest among term-limited mayors.

The term-limit split therefore speaks to both types descriptively. For populists it bears on the motivational accounts: expansion concentrated among mayors who can run again would point to office-seeking, while expansion that survives term limits would point to sincere commitment. For the extreme right it bears on the lame-duck pattern, in which retrenchment is strongest where electoral constraints no longer bind.

A mayor is classified as term-limited if the 2019 term is the last permitted under Law n. 35/2022, the regime in force for the bulk of the observation window.[21]

[21]The legal regime governing term limits changed during the mayors' term. The 2019 election was held under the Delrio Law (Legge 56/2014), which set the three-term exemption below 3,000 inhabitants on the 2011 legal population. In April 2022, Law n. 35/2022 raised this threshold to 5,000 inhabitants (2021 legal population), reprieving 66 mayors in the 3,001–5,000 bracket. Because this reform occurred after the 2019 close election but was external to which candidate barely won, the resulting variation in term-limit exposure is treated as a mid-term institutional change in the incentive environment rather than as an outcome of the 2019 treatment. D.L. 7/2024, in force from 30 January 2024 and definitively converted by L. 38/2024 on 25 March 2024, further raised the exemption to 15,000 inhabitants and removed mandate limits entirely for municipalities at or below 5,000. The main classification reflects L. 35/2022, the operative regime through all of 2022 and 2023 (the bulk of the observation window); mayors in the 5,001–15,000 bracket who were second-term remain coded as term-limited under this classification despite the early-2024 reprieve, which affects only the 24M3 and 24M6 windows. The 291 mayors classified as term-limited under L. 35/2022 were also term-limited under Delrio at the 2019 election. Appendix P reports the same decomposition under the Delrio classification, which assigns the 66 mayors in the 3,001–5,000 bracket to the term-limited group despite their April-2022 reprieve and consequent perception of being not term-limited for the bulk of the observed period; this contamination attenuates the term-limited group's coefficient, as predicted.

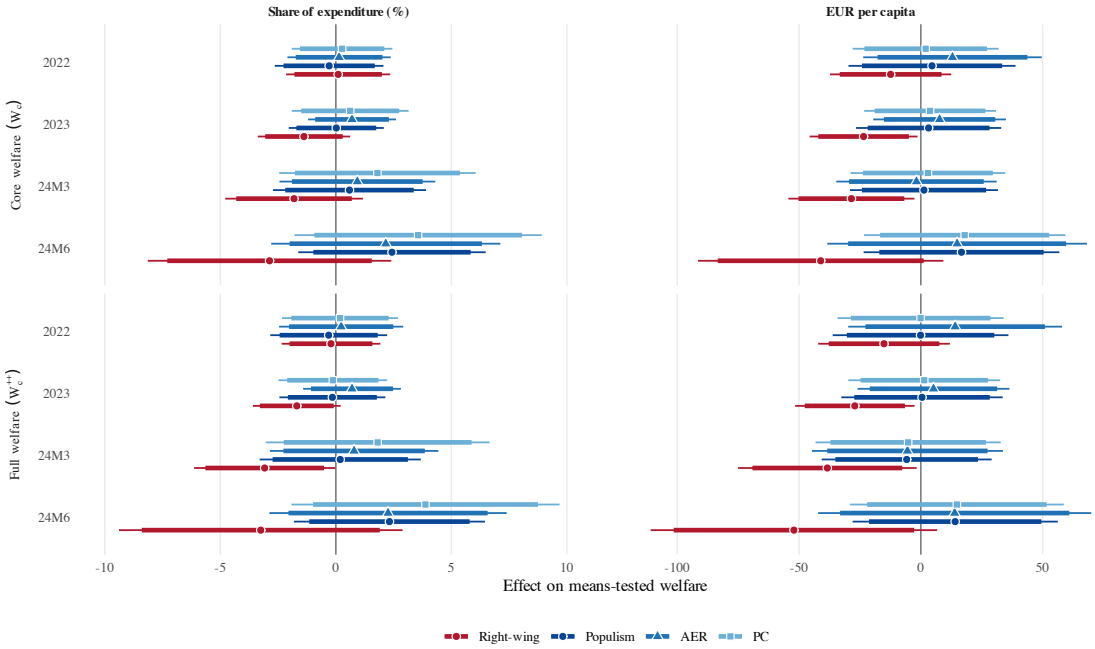

Figure 6: Effect of treatment on means-tested welfare expenditure — term-limited mayors (Gradient PCRD). Points are bias-corrected gradient PCRD coefficients; thick segments are 90% confidence intervals, thin segments 95%. Red denotes the right-wing ($\mathcal{S}^R$) design; blues denote the populism designs (PopT, AER, PC). The 2018 placebo is omitted; full numerical estimates are in Appendix Table 43.

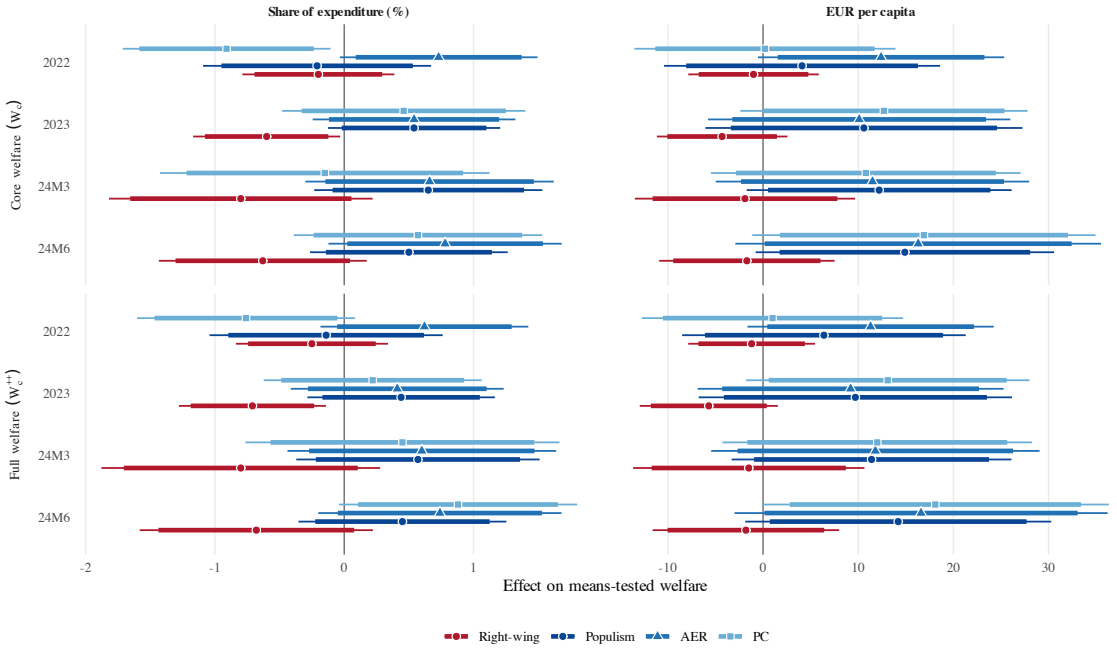

Figure 7: Effect of treatment on means-tested welfare expenditure — not-term-limited mayors (Gradient PCRD). Points are bias-corrected gradient PCRD coefficients; thick segments are 90% confidence intervals, thin segments 95%. Red denotes the right-wing ($\mathcal{S}^R$) design; blues denote the populism designs (PopT, AER, PC). The horizontal scale is set to the range of the not-term-limited estimates and differs from Figure 6. The 2018 placebo is omitted; full numerical estimates are in Appendix Table 44.

Figures 6 and 7 show this split. Among term-limited mayors, the right-wing estimates become more negative. Welfare retrenchment is therefore strongest when extreme-right mayors are least constrained by future electoral concerns. The cuts in the term-limited subsample are far larger than among mayors who can still seek re-election — several times the magnitude — the lame-duck amplification the theory predicts.[22] This is the pattern expected from policy-seeking incumbents: when the electoral cost is removed, programmatic preferences become more visible in policy.

The populism estimates behave differently. Among term-limited mayors, they remain statistically indistinguishable from zero. Among mayors who can still seek re-election, per-capita welfare spending rises as the election approaches, reaching marginal significance in the pre-election windows (Core means-tested welfare per capita: +12.2 EUR at 24M3, $p \approx 0.09$; +14.9 EUR at 24M6, $p \approx 0.06$), whereas term-limited populists show no comparable pre-

[22]The retrenchment is larger among term-limited mayors in every post-treatment period, consistent with the lame-duck pattern; the term-limited and not-term-limited estimates are imprecise on these small subsamples, so the comparison is read descriptively. Pooling the post-treatment periods into a single mean outcome, the difference-in-coefficients test (each cell's full-sample bandwidth held fixed across the two disjoint subsets, estimated separately within each) is significant in per-capita terms (Core −40 EUR, $p = 0.024$; Full −48 EUR, $p = 0.020$) but not in budget-share terms.

election increase. The two groups differ in the direction the office-seeking account predicts — expansion concentrated where and when re-election is at stake — though the difference between them is itself imprecisely estimated and should be read as suggestive rather than decisive.[23] With that caveat, the pattern is more consistent with the strategic account of populism (Weyland, 2001), in which populists expand welfare when doing so can help them remain in office, than with a sincere commitment to the populist worldview (Mudde, 2004; Hawkins and Rovira Kaltwasser, 2017).

Taken together, the term-limit patterns are suggestive of a distinction between political strategy and ideological commitment. Populist welfare expansion appears conditional on electoral incentives. Extreme-right welfare retrenchment intensifies when those incentives are removed. The same fiscal instrument is therefore most consistent with two different motivations: office-seeking among populist mayors and policy-seeking among extreme-right mayors.

For extreme-right mayors, the analysis has established two results separately: they generate electoral backlash, and they reduce welfare provision, especially when they no longer face re-election. The remaining question is whether these two outcomes are connected. The descriptive mediation analysis in Appendix Q suggests that they are. Where the cuts are deeper, the subsequent backlash is larger.

8 Conclusion

This paper asks whether holding office reinforces or undermines populist and extreme-right movements once they win it. Building separate, continuous measures of populism and ideology from candidates' own platform texts, and exploiting quasi-random variation in who governs across razor-thin elections, the analysis finds that the two dimensions carry opposite electoral consequences.

This divergence stays invisible unless the traits are measured apart. Lumping populism and ideological extremism together does not produce a single "anti-establishment incumbency" effect; it blends two divergent responses into an average that reflects neither. This is not merely conceptual. When the text-based measures of right-wing extremism and populism are combined into a single index and run through the same design, their opposite effects on electoral support and on fiscal behavior cancel, and the estimates shrink toward zero. Bundled

[23] The not-term-limited and term-limited estimates are imprecise on these small subsamples; the office-seeking reading describes the direction and timing of the not-term-limited expansion rather than a precisely estimated group difference.

measures may produce mixed and inconclusive findings for the same reason. Separating the dimensions is not a measurement refinement at the margin; it changes whether incumbency looks self-reinforcing or self-undermining.

Drawing this distinction required measuring something the existing toolkit could not. Continuous candidate scores had existed only for ideology, built from roll-call votes or campaign contributions. For populism, the fallback was party labels, which conflate populist rhetoric with ideological position and cannot score the many candidates who run without a usable label. Measuring both traits directly from platform texts yields continuous, candidate-level scores for each, recovers people-versus-elite appeals from the platform itself, and extends to independents and local lists that party-based measures cannot reach. A second advance lies on the design side: the close-election threshold instruments for the elected mayor's continuous score, so that races electing higher-scoring mayors contribute proportionally more to identification, without imposing an arbitrary caliper.

In the Italian case, populist and extreme-right mayors govern in sharply different ways, and voters respond to them in opposite directions. After an extreme-right candidate barely wins, the local electorate shifts its support toward less extreme alternatives. No comparable punishment follows the election of a populist mayor; if anything, the national results point in the opposite direction, with support for the Movimento 5 Stelle rising modestly in subsequent elections. The administrative record helps explain this asymmetry. Extreme-right incumbents retrench welfare provision, and the cuts become larger under term limits, when re-election constraints are removed. Populist incumbents, by contrast, expand welfare mainly when they remain eligible for re-election, with the expansion becoming most visible late in the term as the next election approaches; among term-limited mayors, the estimates remain close to zero. The term-limit split therefore speaks to the competing motivational accounts. The extreme-right pattern is consistent with policy-seeking incumbents who pursue their programmatic preferences more freely once electoral constraints are lifted. The populist pattern is instead more consistent with an office-seeking interpretation: policy responds to the opportunity to win and retain power rather than to a sincerely held commitment to a people-versus-elite worldview (Weyland, 2001).

The electoral results also help distinguish the governing-record account from alternative interpretations of incumbency. A normalization account holds that electoral success reduces the stigma attached to previously marginalized political actors, so representation should increase their subsequent support regardless of what they do in office (Valentim, 2021). This interpretation is less directly applicable to populism in the Italian setting, where populist

actors were already electorally prominent and broadly integrated into mainstream political competition. The absence of backlash against populist mayors is nevertheless compatible with a further legitimizing effect of incumbency. The extreme-right results, however, run in the opposite direction. After an extreme-right candidate takes office, voters shift away from extreme-right alternatives rather than becoming more willing to support them. Moreover, the mediation analysis in Appendix Q links part of this backlash to the governing record itself: municipalities experiencing deeper welfare retrenchment also exhibit larger subsequent shifts away from the extreme right. Because stigma reduction operates through the fact of election rather than through what incumbents do in office, this performance-dependent pattern is difficult to reconcile with a pure normalization account.

Several limitations bound these conclusions. The analysis covers a single country and a single electoral cohort, and the estimates may not transfer directly to other settings. The sample is not random, since southern and smaller municipalities were less likely to supply platforms, though inverse-probability weighting by region and size leaves the headline results unchanged. The text-based measures carry classification error, and the fiscal records capture cash-flow behavior rather than every channel through which mayors affect voter welfare. The estimates are also local to marginal winners and are silent on incumbents who win by wide margins. Because marginal winners face the most binding re-election incentives, the close-race sample is the setting in which populist office-seeking, and the welfare expansion it predicts, should be most visible.

While the clean identification recovered here is local in the sense described above, the mechanism it uncovers is not: it turns on the balance between populism and extreme-right ideology within an incumbent, and nothing about that is specific to Italy. Politicians often combine traits, blending populist and ideological appeals in varying degrees, and the argument is not that candidates fall neatly into one category, but that the dominant trait shapes how they govern. Where the extreme-right ideological component dominates, incumbents tend toward welfare retrenchment even at electoral cost; the Trump administration's 2025 Medicaid cuts, enacted despite broad public opposition, are consistent with this pattern.[24] Where the populist component dominates, governing turns instead toward visible, attributable benefits; Poland's Law and Justice, with its Family 500+ child transfer, fits this second pattern (Gromadzki, Sałach and Brzeziński, 2024). This is why bundling the two traits under

[24]The 2025 One Big Beautiful Bill Act reduced federal Medicaid spending by an estimated $911 billion over a decade, with roughly ten million more people projected to lose coverage, and cut the SNAP food-stamp program; KFF polling found the law broadly unpopular, with 63% holding an unfavorable view. See KFF, "Public Views on Recent Tax and Budget Legislation," July 2025, https://www.kff.org/medicaid/kff-health-tracking-poll-public-views-on-recent-tax-and-budget-legislation/.

a single "anti-establishment" or "populist radical right" label risks obscuring why incumbency becomes self-reinforcing in some cases and self-undermining in others.

Future research should build trait-level scores in other settings to test whether conflicting findings in the literature reflect genuine cross-national heterogeneity or the bundling of two opposing effects into one estimate. The distinction is not only methodological: the two traits imply different things about whether voters can hold incumbents to account.

The implication for democratic accountability is mixed. Voters appear to punish extreme-right ideological rigidity, but the same mechanism does not seem to discipline the populist use of visible, re-election-timed spending. Whether accountability can constrain a governing model that wins support through benefits voters can see, timed to the electoral calendar, is the harder question this paper leaves open.

References

Abou-Chadi, Tarik. 2016. "Niche Party Success and Mainstream Party Policy Shifts—How Green and Radical Right Parties Differ in Their Impact." *British Journal of Political Science* 46(2):417–436.

Akkerman, Tjitske and Sarah L De Lange. 2012. "Radical right parties in office: Incumbency records and the electoral cost of governing1." *Government and Opposition* 47(4):574–596.

Albertazzi, Daniele and Duncan McDonnell. 2015. *Populists in power*. Routledge.

Arzheimer, Kai. 2009. "Contextual Factors and the Extreme Right Vote in Western Europe, 1980–2002." *American Journal of Political Science* 53(2):259–275.

Bakker, Ryan, Liesbet Hooghe, Seth Jolly, Gary Marks, Jonathan Polk, Jan Rovny, Marco Steenbergen and Milada Anna Vachudova. 2019. "Chapel Hill expert survey." *Version 1.2* .

Baron, Reuben M. and David A. Kenny. 1986. "The Moderator-Mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic, and Statistical Considerations." *Journal of Personality and Social Psychology* 51(6):1173–1182.

Bellodi, Luca, Massimo Morelli and Matia Vannoni. 2024. "A costly commitment: Populism, economic performance, and the quality of bureaucracy." *American Journal of Political Science* 68(1):193–209.

Bergh, Andreas and Anders Kärnä. 2024. Populist parties' popularity post power. Technical Report 1491 IFN Working Paper.

Bertoli, Andrew and Chad Hazlett. 2025. "Seeing like a district: Understanding what close-election designs for leader characteristics can and cannot tell us." *Political Analysis* pp. 1–19.

Bonica, Adam. 2014. "Mapping the ideological marketplace." *American Journal of Political Science* 58(2):367–386.

Bordignon, Massimo and Tommaso Colussi. 2020. "Dancing with the Populist: New Parties, Electoral Rules and Italian Municipal Elections.".

Bosley, Mitchell, Musashi Jacobs-Harukawa, Hauke Licht and Alexander Hoyle. 2023. Do we still need BERT in the age of GPT? Comparing the benefits of domain-adaptation and in-context-learning approaches to using LLMs for Political Science Research. In *2023 Annual Meeting of the Midwest Political Science Association (MPSA)*.

Bracco, Emanuele, Maria De Paola, Colin P Green and Vincenzo Scoppa. 2018. "The effect of far right parties on the location choice of immigrants: Evidence from Lega Nord Mayors." *Journal of Public Economics* 166:12–26.

Brader, Ted, Nicholas A. Valentino and Elizabeth Suhay. 2008. "What Triggers Public Opposition to Immigration? Anxiety, Group Cues, and Immigration Threat." *American Journal of Political Science* 52(4):959–978.

Brender, Adi and Allan Drazen. 2008. "How Do Budget Deficits and Economic Growth Affect Reelection Prospects? Evidence from a Large Panel of Countries." *American Economic Review* 98(5):2203–2220.

Calvert, Randall L. 1985. "Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence." *American Journal of Political Science* 29(1):69–95.

Cattaneo, Matias D., Michael Jansson and Xinwei Ma. 2020. "Simple Local Polynomial Density Estimators." *Journal of the American Statistical Association* 115(531):1449–1455.

Conneau, Alexis, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th annual meeting of the association for computational linguistics*. pp. 8440–8451.

De Benedetto, Marco Alberto. 2020. "Personal or Partisan Incumbency Advantage? Evidence from an Electoral Reform at the Local Level in Italy." *The B.E. Journal of Economic Analysis & Policy* 20(1):20180125.

de Vries, Catherine E. and Sara B. Hobolt. 2020. *Political Entrepreneurs: The Rise of Challenger Parties in Europe*. Princeton, NJ: Princeton University Press.

Dong, Yingying, Ying-Ying Lee and Michael Gou. 2023. "Regression discontinuity designs with a continuous treatment." *Journal of the American Statistical Association* 118(541):208–221.

Dornbusch, Rudiger and Sebastian Edwards. 1990. "Macroeconomic populism." *Journal of Development Economics* 32(2):247–277.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper and Row.

Duch, Raymond M and Randolph T Stevenson. 2008. *The economic vote: How political and economic institutions condition election results*. Cambridge University Press.

Ebenstein, Avraham, Maoyong Fan, Michael Greenstone, Guojun He and Maigeng Zhou. 2017. "New evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River Policy." *Proceedings of the National Academy of Sciences* 114(39):10384–10389.

Ennser-Jedenastik, Laurenz. 2016. "A welfare state for whom? A group-based account of the Austrian Freedom Party's social policy profile." *Swiss Political Science Review* 22(3):409–427.

Ernst, Nicole, Sven Engesser and Frank Esser. 2017. "Bipolar populism? The use of anti-elitism and people-centrism by Swiss parties on social media." *Swiss political science review* 23(3).

Esping-Andersen, Gøsta. 1990. *The Three Worlds of Welfare Capitalism*. Princeton, NJ: Princeton University Press.

Fearon, James D. 1999. Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance. In *Democracy, Accountability, and Representation*, ed. Adam Przeworski, Susan C. Stokes and Bernard Manin. Cambridge: Cambridge University Press pp. 55–97.

Ferejohn, John. 1986. "Incumbent performance and electoral control." *Public Choice* 50(1):5–25.

Ferrera, Maurizio. 1996. "The 'Southern Model' of Welfare in Social Europe." *Journal of European Social Policy* 6(1):17–37.

Fiorina, Morris P. 1976. "Retrospective voting in American national elections." *NSF Award* 76(7602083):2083.

Fiorina, Morris P. 1981. *Retrospective Voting in American National Elections*. New Haven: Yale University Press.

Freier, Ronny. 2015. "The Mayor's Advantage: Causal Evidence on Incumbency Effects in German Mayoral Elections." *European Journal of Political Economy* 40:16–30.

Funke, Manuel, Moritz Schularick and Christoph Trebesch. 2023. "Populist leaders and the economy." *American Economic Review* 113(12):3249–3288.

Gamalerio, Matteo. 2020. "Do national political parties matter? Evidence from Italian municipalities." *European Journal of Political Economy* 63:101862.

Gelman, Andrew and Gary King. 1990. "Estimating Incumbency Advantage Without Bias." *American Journal of Political Science* 34(4):1142–1164.

Golder, Matt. 2016. "Far Right Parties in Europe." *Annual Review of Political Science* 19:477–497.

Gromadzki, Jan, Katarzyna Sałach and Michał Brzeziński. 2024. "When populists deliver on their promises: The electoral effects of a large cash transfer programme in Poland." *Economica* 91(361):320–345.

Grzymala-Busse, Anna and Monika Nalepa. 2019. "Clear winners: Programmatic ambiguity and the electoral punishment of populists." *Unpublished manuscript. Retrieved from https://www. monikanalepa. com/uploads/6/6/3/1/66318923/5feb2019. pdf* .

Guriev, Sergei and Elias Papaioannou. 2022. "The political economy of populism." *Journal of Economic literature* 60(3):753–832.

Hahn, Jinyong, Petra Todd and Wilbert Van der Klaauw. 2001. "Identification and estimation of treatment effects with a regression-discontinuity design." *Econometrica* 69(1):201–209.

Hall, Andrew B. 2015. "What happens when extremists win primaries?" *American Political Science Review* 109(1):18–42.

Hall, Andrew B. 2019. *Who Wants to Run? How the Devaluing of Political Office Drives Polarization*. Chicago: University of Chicago Press.

Hawkins, Kirk A. 2009. "Is Chávez Populist? Measuring Populist Discourse in Comparative Perspective." *Comparative Political Studies* 42(8):1040–1067.

Hawkins, Kirk A. and Cristóbal Rovira Kaltwasser. 2017. "The Ideational Approach to Populism." *Latin American Research Review* 52(4):513–528.

Healy, Andrew and Neil Malhotra. 2009. "Myopic voters and natural disaster policy." *American Political Science Review* 103(3):387–406.

Healy, Andrew and Neil Malhotra. 2013. "Retrospective Voting Reconsidered." *Annual Review of Political Science* 16:285–306.

Imai, Kosuke, Luke Keele and Dustin Tingley. 2010. "A General Approach to Causal Mediation Analysis." *Psychological Methods* 15(4):309–334.

Imai, Kosuke and Teppei Yamamoto. 2013. "Identification and Sensitivity Analysis for Multiple Causal Mechanisms: Revisiting Evidence from Framing Experiments." *Political Analysis* 21(2):141–171.

Inglehart, Ronald F and Pippa Norris. 2016. "Trump, Brexit, and the rise of populism: Economic have-nots and cultural backlash.".

Ivarsflaten, Elisabeth. 2008. "What unites right-wing populists in Western Europe? Re-examining grievance mobilization models in seven successful cases." *Comparative Political Studies* 41(1):3–23.

Jagers, Jan and Stefaan Walgrave. 2007. "Populism as political communication style: An empirical study of political parties' discourse in Belgium." *European journal of political research* 46(3):319–345.

Jones, Pauline and Anil Menon. 2024. Varieties of Populists: Paths to Power and Implications for Regime Stability. In *Still the Age of Populism? Re-examining Theories and Concepts*, ed. Michael Bernhard, Amie Kreppel and Carlos de la Torre. London: Routledge.

Kaltwasser, Cristóbal Rovira, Paul A Taggart, Paulina Ochoa Espejo and Pierre Ostiguy. 2017. *The Oxford handbook of populism*. Oxford University Press.

Key, Valdimer Orlando. 1966. *The responsible electorate: Rationality in presidential voting 1936–1960*. Harvard University Press.

Kitschelt, Herbert and Anthony J. McGann. 1995. *The Radical Right in Western Europe: A Comparative Analysis*. Ann Arbor: University of Michigan Press.

Klašnja, Marko and Rocío Titiunik. 2017. "The Incumbency Curse: Weak Parties, Term Limits, and Unfulfilled Accountability." *American Political Science Review* 111(1):129–148.

Korpi, Walter and Joakim Palme. 1998. "The Paradox of Redistribution and Strategies of Equality: Welfare State Institutions, Inequality, and Poverty in the Western Countries." *American Sociological Review* 63(5):661–687.

Kriesi, Hanspeter, Edgar Grande, Romain Lachat, Martin Dolezal, Simon Bornschier and Timotheos Frey. 2008. *West European Politics in the Age of Globalization.* Cambridge: Cambridge University Press.

Laclau, Ernesto. 2005. *On populist reason.* Verso.

Laver, Michael, Kenneth Benoit and John Garry. 2003. "Extracting policy positions from political texts using words as data." *American political science review* 97(2):311–331.

Lee, David S. 2008. "Randomized experiments from non-random selection in U.S. House elections." *Journal of Econometrics* 142(2):675–697.

Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer and Veselin Stoyanov. 2019. "Roberta: A robustly optimized bert pretraining approach." *arXiv preprint arXiv:1907.11692* .

Marshall, John. 2024. "Can close election regression discontinuity designs identify effects of winning politician characteristics?" *American Journal of Political Science* 68(2):494–510.

McCarty, Nolan, Keith T. Poole and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches.* Walras-Pareto Lectures Cambridge, MA: MIT Press.

McCrary, Justin. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics* 142(2):698–714.

Meguid, Bonnie M. 2005. "Competition Between Unequals: The Role of Mainstream Party Strategy in Niche Party Success." *American Political Science Review* 99(3):347–359.

Meisels, Mellissa. 2025. "Everything in moderation? The effect of extremist nominations on individual and corporate PAC fundraising." *Political Science Research and Methods* pp. 1–17.

Moffitt, Benjamin. 2016. *The global rise of populism: Performance, political style, and representation.* Stanford University Press.

Mudde, Cas. 2004. "The populist zeitgeist." *Government and opposition* 39(4):541–563.

Mudde, Cas. 2007. *Populist radical right parties in Europe.* Vol. 22 Cambridge university press Cambridge.

Mudde, Cas and Cristóbal Rovira Kaltwasser. 2017. *Populism: A very short introduction.* oxford university Press.

Muis, Jasper and Tim Immerzeel. 2017. "Causes and consequences of the rise of populist radical right parties and movements in Europe." *Current Sociology* 65(6):909–930.

Norris, Pippa and Ronald Inglehart. 2019. *Cultural backlash: Trump, Brexit, and authoritarian populism*. Cambridge University Press.

Noury, Abdul and Gerard Roland. 2020. "Identity politics and populism in Europe." *Annual Review of Political Science* 23(1):421–439.

Pauwels, Teun. 2011. "Measuring populism: A quantitative text analysis of party literature in Belgium." *Journal of Elections, Public Opinion and Parties* 21(1):97–119.

Pirro, Andrea L. P. 2024. "Performing (during) the Coronavirus Crisis: The Italian Populist Radical Right between National Opposition and Subnational Government." *Government and Opposition* 59(4):1092–1108.

Poole, Keith T. 2005. *Spatial Models of Parliamentary Voting*. Analytical Methods for Social Research Cambridge: Cambridge University Press.

Poole, Keith T. and Howard Rosenthal. 1985. "A Spatial Model for Legislative Roll Call Analysis." *American Journal of Political Science* 29(2):357–384.

Rheault, Ludovic and Christopher Cochrane. 2020. "Word embeddings for the analysis of ideological placement in parliamentary corpora." *Political analysis* 28(1):112–133.

Riera, Pedro and Marco Pastor. 2022. "Cordons sanitaires or tainted coalitions? The electoral consequences of populist participation in government." *Party Politics* 28(5):889–902.

Rodden, Jonathan. 2002. "The dilemma of fiscal federalism: Grants and fiscal performance around the world." *Am. J. Pol. Sci.* 46:670.

Rodden, Jonathan. 2006. *Hamilton's paradox: the promise and peril of fiscal federalism*. Vol. 2 Cambridge University Press Cambridge.

Rodriguez, Pedro L and Arthur Spirling. 2022. "Word embeddings: What works, what doesn't, and how to tell the difference for applied research." *The Journal of Politics* 84(1):101–115.

Rogoff, Kenneth S. 1987. "Equilibrium political budget cycles.".

Rooduijn, Matthijs. 2015. "The rise of the populist radical right in Western Europe." *European view* 14(1):3–11.

Rooduijn, Matthijs, Stijn Van Kessel, Caterina Froio, Andrea Pirro, Sarah De Lange, Daphne Halikiopoulou, Paul Lewis, Cas Mudde and Paul Taggart. 2019. "The PopuList: An overview of populist, far right, far left and Eurosceptic parties in Europe.".

Rooduijn, Matthijs and Teun Pauwels. 2011. "Measuring populism: Comparing two methods of content analysis." *West European Politics* 34(6):1272–1283.

Rovny, Jan, Jonathan Polk, Ryan Bakker, Liesbet Hooghe, Seth Jolly, Gary Marks, Marco Steenbergen and Milada Anna Vachudova. 2025. "The 2024 Chapel Hill Expert Survey on political party positioning in Europe: Twenty-five years of party positional data." *Electoral Studies* 97:102981.

Rydgren, Jens. 2007. "The Sociology of the Radical Right." *Annual Review of Sociology* 33:241–262.

Schumacher, Gijs and Kees Van Kersbergen. 2016. "Do mainstream parties adapt to the welfare chauvinism of populist parties?" *Party Politics* 22(3):300–312.

Sebők, Miklós, Ákos Máté, Orsolya Ring, Viktor Kovács and Richárd Lehoczki. 2025. "Leveraging Open Large Language Models for Multilingual Policy Topic Classification: The Babel Machine Approach." *Social Science Computer Review* .

Shi, Min and Jakob Svensson. 2006. "Political budget cycles: Do they differ across countries and why?" *Journal of public economics* 90(8-9):1367–1389.

Slapin, Jonathan B and Sven-Oliver Proksch. 2008. "A scaling model for estimating time-series party positions from texts." *American Journal of Political Science* 52(3):705–722.

Slothuus, Rune. 2008. "More than Weighting Cognitive Importance: A Dual-Process Model of Issue Framing Effects." *Political Psychology* 29(1):1–28.

Stanley, Ben. 2008. "The thin ideology of populism." *Journal of political ideologies* 13(1):95–110.

Stiers, Dieter. 2019. "Beyond the Distinction Between Incumbency and Opposition: Retrospective Voting at the Level of Political Parties." *Party Politics* 25(6):805–816.

Stokes, Susan C. 2001. *Mandates and Democracy: Neoliberalism by Surprise in Latin America.* Cambridge: Cambridge University Press.

Strom, Kaare. 1990. "A behavioral theory of competitive political parties." *American journal of political science* pp. 565–598.

Thomson, Robert, Terry Royed, Elin Naurin, Joaquín Artés, Rory Costello, Laurenz Ennser-Jedenastik, Mark Ferguson, Petia Kostadinova, Catherine Moury, François Pétry and Katrin Praprotnik. 2017. "The Fulfillment of Parties' Election Pledges: A Comparative Study on the Impact of Power Sharing." *American Journal of Political Science* 61(3):527–542.

Timoneda, Joan C. and Sebastián Vallejo Vera. 2025. "BERT, RoBERTa or DeBERTa? Comparing Performance Across Transformer Models in Political Science Text." *Journal of Politics* 87(1):347–364.

Valentim, Vicente. 2021. "Parliamentary representation and the normalization of radical right support." *Comparative political studies* 54(14):2475–2511.

Van Spanje, Joost. 2011. "Keeping the rascals in: Anti-political-establishment parties and their cost of governing in established democracies." *European Journal of Political Research* 50(5):609–635.

Wang, Yu, Wen Qu and Xin Ye. 2024. "Selecting between BERT and GPT for text classification in political science research." *arXiv preprint arXiv:2411.05050* .

Weyland, Kurt. 2001. "Clarifying a contested concept: Populism in the study of Latin American politics." *Comparative politics* pp. 1–22.

Wittman, Donald. 1983. "Candidate Motivation: A Synthesis of Alternative Theories." *American Political Science Review* 77(1):142–157.

Young, Lori and Stuart Soroka. 2012. "Affective news: The automated coding of sentiment in political texts." *Political communication* 29(2):205–231.

Appendices

A Platform Recovery and Sample Construction

Electoral platforms were collected through direct outreach to municipal administrations over approximately six months. The universe of municipalities comprises all municipalities that held mayoral elections in both 2019 and 2024, restricted to regions whose electoral data were available online. Table 3 reports recovery rates overall and broken down by macro-area and population size class.

Table 3: Platform Recovery: Universe and Recovered Municipalities

	Universe		Recovered		
	N	Share (%)	*N*	Share (%)	Recovery rate (%)
Panel A: By macro-area					
North	2,503	69.7	1,104	78.2	44.1
Centre	530	14.8	215	15.2	40.6
South	559	15.6	93	6.6	16.6
Total	3,592	100.0	1,412	100.0	39.3
Panel B: By population size					
$<$ 1,000	983	27.4	206	14.6	21.0
1,000–4,999	1,719	47.9	618	43.8	36.0
5,000–14,999	683	19.0	442	31.3	64.7
$\geq$ 15,000	207	5.8	146	10.3	70.5

Notes: Universe comprises all municipalities that held mayoral elections in both May–June 2019 and 2024, restricted to regions whose electoral data were available online. Sardinia is included in the South macro-area. A municipality is classified as recovered if at least one 2019 candidate platform was successfully obtained. Population size classes are based on 2021 census data.

Recovery rates are higher in the North and Centre than in the South, and increasing in municipality size. The gradient by size is expected: larger municipalities have more administrative capacity, are more likely to maintain accessible digital records, and were more likely to respond to outreach requests. The geographic gradient likely reflects similar differences in administrative infrastructure.

Inverse Probability Weighting

Because platform recovery is correlated with municipality size and geography, the analysis sample may not be representative of the full universe. To assess sensitivity to non-random recovery, all principal specifications are re-estimated with inverse probability weights (IPW), constructed from a logit model of recovery status on region dummies and log population estimated on the full universe of 3,592 municipalities. Each recovered municipality is weighted by the inverse of its predicted recovery probability, trimmed at the 5th percentile.

Tables 4 and 5 report the results for gradient PCRD and CT-PCRD, respectively. The two headline findings are robust. The $\mathcal{S}^R$ ideology backlash is stable under IPW (gradient -0.049, $p = 0.021$; CT-PCRD -0.295, $p = 0.016$). The PC $\rightarrow$ M5S spillover strengthens under

reweighting in both estimators and both elections — the gradient Camera estimate increases from 0.014 ($p = 0.022$) to 0.029 ($p < 0.001$), and the CT-PCRD European 2024 estimate increases from 0.092 ($p = 0.122$) to 0.132 ($p = 0.026$) — indicating that the unweighted estimates are, if anything, conservative. The one partial exception is the Lega European 2024 CT-PCRD result, which weakens under IPW ($p = 0.208$, versus $p = 0.067$ unweighted), though the gradient estimate remains marginally significant ($p = 0.087$). Non-significant designs remain non-significant throughout.

Table 4: IPW Robustness: Gradient PCRD

	Unweighted		IPW-weighted	
Design	Coeff.	N_h	Coeff.	N_h
Panel A: Municipal 2024 outcomes (vote-weighted)				
Right-wing $\to Y_{2024}^{\text{ideo}}$	-0.058^{***}	278	-0.049^{**}	208
	[$p = 0.002$]		[$p = 0.021$]	
Panel B: Camera 2022 outcomes (national party vote shares)				
Right-wing $\to$ Lega	-0.000	270	-0.002	169
	[$p = 0.970$]		[$p = 0.881$]	
Right-wing $\to$ FdI	-0.003	259	0.005	235
	[$p = 0.831$]		[$p = 0.755$]	
Populism $\to$ M5S	0.014^{*}	257	0.011	531
	[$p = 0.071$]		[$p = 0.170$]	
AER $\to$ M5S	0.009	381	0.009	494
	[$p = 0.198$]		[$p = 0.295$]	
PC $\to$ M5S	0.014^{**}	388	0.029^{***}	320
	[$p = 0.022$]		[$p < 0.001$]	
Panel C: European 2024 outcomes (national party vote shares)				
Right-wing $\to$ Lega	-0.020^{**}	216	-0.017^{*}	233
	[$p = 0.029$]		[$p = 0.087$]	
Right-wing $\to$ FdI	0.017	327	0.025	216
	[$p = 0.201$]		[$p = 0.158$]	
Populism $\to$ M5S	0.013^{**}	236	0.014^{**}	389
	[$p = 0.034$]		[$p = 0.027$]	
AER $\to$ M5S	0.006	452	0.010	482
	[$p = 0.199$]		[$p = 0.163$]	
PC $\to$ M5S	0.008^{*}	338	0.022^{***}	241
	[$p = 0.088$]		[$p = 0.001$]	

Notes: Each row reports the bias-corrected robust gradient PCRD (sharp RD) estimate with p-values in brackets. IPW weights are constructed from a logit model of platform recovery on region dummies and log population, estimated on the full universe of 3,592 municipalities. Recovered municipalities are weighted by the inverse of their predicted recovery probability (trimmed at the 5th percentile). All specifications include region fixed effects and log population as covariates. N_h denotes effective observations within the MSE-optimal bandwidth. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: IPW Robustness: CT-PCRD (Fuzzy RD)

	Unweighted		IPW-weighted	
Design	Coeff.	N_h	Coeff.	N_h
Panel A: Municipal 2024 outcomes (vote-weighted)				
Right-wing $\to Y_{2024}^{\text{ideo}}$	-0.271^{**}	368	-0.295^{**}	365
	$[p = 0.013]$		$[p = 0.016]$	
Panel B: Camera 2022 outcomes (national party vote shares)				
Right-wing $\to$ Lega	-0.006	387	0.048	387
	$[p = 0.883]$		$[p = 0.387]$	
Right-wing $\to$ FdI	-0.032	367	0.020	394
	$[p = 0.595]$		$[p = 0.783]$	
Populism $\to$ M5S	0.148^{**}	407	0.117	557
	$[p = 0.025]$		$[p = 0.144]$	
AER $\to$ M5S	0.154	479	0.197	510
	$[p = 0.370]$		$[p = 0.417]$	
PC $\to$ M5S	0.197^{**}	420	0.237^{***}	458
	$[p = 0.012]$		$[p = 0.002]$	
Panel C: European 2024 outcomes (national party vote shares)				
Right-wing $\to$ Lega	-0.085^{*}	378	-0.063	385
	$[p = 0.067]$		$[p = 0.208]$	
Right-wing $\to$ FdI	0.057	386	0.131^{*}	394
	$[p = 0.354]$		$[p = 0.068]$	
Populism $\to$ M5S	0.135^{***}	375	0.144^{**}	406
	$[p = 0.005]$		$[p = 0.021]$	
AER $\to$ M5S	0.144	529	0.271	463
	$[p = 0.259]$		$[p = 0.215]$	
PC $\to$ M5S	0.092	412	0.132^{**}	413
	$[p = 0.122]$		$[p = 0.026]$	

Notes: Each row reports the bias-corrected robust CT-PCRD (fuzzy RD, 2SLS) estimate with p-values in brackets. IPW weights are constructed from a logit model of platform recovery on region dummies and log population, estimated on the full universe of 3,592 municipalities. Recovered municipalities are weighted by the inverse of their predicted recovery probability (trimmed at the 5th percentile). All specifications include region fixed effects and log population as covariates. N_h denotes effective observations within the MSE-optimal bandwidth. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B Measurement and Classification

B.1 Model Training and Validation

XLM-RoBERTa (XLM-R) is fine-tuned to assign codebook-based labels at the paragraph level. The unit of analysis is a paragraph extracted from political platforms, and XLM-R is trained as a supervised classifier using human-coded labels as ground truth. Four separate task-specific classifiers are estimated (one per labeling task). To construct the labeled training corpus, candidate paragraphs are drawn from (i) a sample of local/municipal platforms from the 2015–2017 election cycles and (ii) national-level platforms from the Manifesto Project. Because the target categories—especially the most extreme classes—are rare in unconstrained text, a purely random labeling strategy would require an impractically large number of annotations to obtain sufficient positive cases.

To make manual coding feasible while retaining both positive and negative examples, a large language model (LLM) is used only as a screening device to prioritize paragraphs for annotation under severe class imbalance. In each annotation round, the LLM assigns paragraphs a coarse, task-specific screening flag (plausibly relevant vs. plausibly irrelevant to the target codebook features), and paragraphs are sampled from both strata to ensure coverage of both likely positives and likely negatives. This screening step is used strictly to allocate annotation effort: all sampled paragraphs are subsequently labeled by humans under the pre-specified codebook, and only human labels enter training, model selection, or evaluation. Because four task-specific classifiers are estimated, the screening prompt can be updated across rounds to better target whichever labels are currently under-represented for a given task; these prompt changes affect only which paragraphs are sent to annotation.

Model selection and early stopping rely on a task-specific municipal development set constructed from the 2019 and 2024 municipal platform corpus (the target domain). The development set is disjoint from the training data and is built directly from the target municipal corpus (not from earlier local platforms or national texts). Because label prevalence in the target domain is extremely low, the development set is constructed using the same iterative, class-targeted screening-and-annotation procedure used to expand the training set. As a result, the development set is intentionally class-enriched relative to the raw target corpus, while still including paragraphs sampled from both screening strata to retain negative examples. Targeted municipal annotation continues until the pre-specified development-set size (and minimum per-class support) is reached; once the development set is filled, its membership is frozen and reused unchanged for all subsequent fine-tuning runs, model comparisons, and early-stopping decisions.

Finally, OCR artifacts are intentionally retained during preprocessing. Because the model is ultimately applied to OCR-derived municipal platforms, aggressive de-noising could create a train–test mismatch by discarding malformed but informative tokens that systematically appear at inference. Training therefore proceeds on paragraphs spanning the full distribution of OCR quality observed in the raw data to promote robustness to OCR-induced noise. For model selection and early stopping, however, evaluation is restricted to paragraphs exceeding a pre-specified minimum text-quality threshold—the same threshold later applied when scoring the full corpus—so that tuning reflects performance on the subset of text that enters the downstream analysis.

B.2 Classifier Performance

Performance is evaluated on the municipal development set described above. For each classifier, training dynamics are reported across epochs (train loss, validation loss, macro-F1, and accuracy on the development set), followed by the confusion matrix and per-class precision, recall, and F1 at the best epoch (selected by macro-F1 on the development set).

Ideology

The best epoch is epoch 6, achieving macro-F1 of 0.767 and accuracy of 0.780. Performance is strong across all five classes; the Left and Extreme Right categories show somewhat lower F1 scores, reflecting greater ambiguity at ideological boundaries.

Epoch	Train Loss	Val. Loss	Macro-F1	Accuracy
1	1.636	1.709	0.095	0.172
2	1.043	0.794	0.704	0.725
3	0.655	0.736	0.724	0.744
4	0.445	0.757	0.739	0.753
5	0.336	0.767	0.755	0.768
6	**0.243**	**0.809**	**0.767**	**0.780**
7	0.161	0.816	0.747	0.772
8	0.143	0.891	0.731	0.753
9	0.102	0.914	0.750	0.764

Table 6: Ideology classifier: training dynamics. Train loss, validation loss, macro-F1, and accuracy on the development set across epochs. The bold row indicates the selected best epoch.

	Pred. EL	Pred. L	Pred. C	Pred. R	Pred. ER
True EL	99	18	5	1	2
True L	9	99	14	4	6
True C	5	21	167	11	5
True R	1	17	23	213	10
True ER	1	5	4	18	61

Table 7: Ideology classifier: confusion matrix (best epoch). Rows are true classes; columns are predicted classes. EL = Extreme Left; L = Left; C = Centre; R = Right; ER = Extreme Right.

Class	Precision	Recall	F1	Support
EL (Extreme Left)	0.861	0.792	0.825	125
L (Left)	0.619	0.750	0.678	132
C (Centre)	0.784	0.799	0.791	209
R (Right)	0.862	0.807	0.834	264
ER (Extreme Right)	0.726	0.685	0.705	89
Macro average	0.770	0.767	0.767	819

Table 8: Ideology classifier: per-class performance (best epoch). Precision, recall, and F1 by class on the development set.

Populism

The best epoch is epoch 6, achieving macro-F1 of 0.790 and accuracy of 0.824. The None class is classified with highest F1 (0.883), while the Soft class has the lowest F1, reflecting its intermediate position between the two extremes.

Epoch	Train Loss	Val. Loss	Macro-F1	Accuracy
1	1.034	1.217	0.275	0.288
2	0.601	0.724	0.726	0.752
3	0.388	0.547	0.768	0.799
4	0.212	0.772	0.745	0.750
5	0.149	0.676	0.772	0.795
6	**0.098**	**0.611**	**0.790**	**0.824**
7	0.062	0.658	0.776	0.817
8	0.058	0.692	0.775	0.815
9	0.053	0.730	0.765	0.812
10	0.052	0.698	0.786	0.826

Table 9: Populism classifier: training dynamics. Train loss, validation loss, macro-F1, and accuracy on the development set across epochs. The bold row indicates the selected best epoch.

	Pred. None	Pred. Soft	Pred. Hard
True None	290	21	8
True Soft	34	86	14
True Hard	14	6	79

Table 10: Populism classifier: confusion matrix (best epoch). Rows are true classes; columns are predicted classes. None = no populist rhetoric; Soft = people-centric appeals or weak anti-elite framing; Hard = people-vs-elite conflict or explicit accusations.

Class	Precision	Recall	F1	Support
None	0.858	0.909	0.883	319
Soft	0.761	0.642	0.696	134
Hard	0.782	0.798	0.790	99
Macro average	0.800	0.783	0.790	552

Table 11: Populism classifier: per-class performance (best epoch). Precision, recall, and F1 by class on the development set.

Anti-Elite Rhetoric (AER)

The best epoch is epoch 4, achieving macro-F1 of 0.815 and accuracy of 0.857. The None class is classified with highest F1 (0.924), while the Soft and Hard anti-elite classes show lower but comparable F1 (0.740 and 0.780, respectively), reflecting the difficulty of distinguishing degrees of antagonism toward political elites.

Epoch	Train Loss	Val. Loss	Macro-F1	Accuracy
1	1.090	1.174	0.126	0.234
2	0.977	0.890	0.627	0.674
3	0.667	0.578	0.770	0.832
4	**0.478**	**0.455**	**0.815**	**0.857**
5	0.228	0.584	0.800	0.848
6	0.121	0.546	0.814	0.853

Table 12: AER classifier: training dynamics. Train loss, validation loss, macro-F1, and accuracy on the development set across epochs. The bold row indicates the selected best epoch.

	Pred. None	Pred. Soft	Pred. Hard
True None	305	25	7
True Soft	14	97	18
True Hard	4	11	71

Table 13: AER classifier: confusion matrix (best epoch). Rows are true classes; columns are predicted classes. None = no anti-elite rhetoric; Soft = distancing or neglect framing; Hard = explicit accusations of elite wrongdoing.

Class	Precision	Recall	F1	Support
None	0.944	0.905	0.924	337
Soft	0.729	0.752	0.740	129
Hard	0.740	0.826	0.780	86
Macro average	0.804	0.828	0.815	552

Table 14: AER classifier: per-class performance (best epoch). Precision, recall, and F1 by class on the development set.

People-Centrism (PC)

The best epoch is epoch 5, achieving macro-F1 of 0.749 and accuracy of 0.748. The Hard people-centrism class is classified with highest F1 (0.795), while the Soft class shows lower F1 (0.676), reflecting the difficulty of distinguishing weak participatory framing from both purely technocratic language and strong citizen-empowerment rhetoric.

Epoch	Train Loss	Val. Loss	Macro-F1	Accuracy
1	0.816	0.766	0.615	0.653
2	0.433	0.614	0.736	0.735
3	0.247	0.776	0.745	0.745
4	0.205	1.099	0.670	0.682
5	**0.149**	**0.797**	**0.749**	**0.748**
6	0.109	0.935	0.707	0.704
7	0.105	0.754	0.742	0.742
8	0.086	1.055	0.720	0.720
9	0.054	0.925	0.727	0.726
10	0.104	0.755	0.681	0.679

Table 15: People-Centrism classifier: training dynamics. Train loss, validation loss, macro-F1, and accuracy on the development set across epochs. The bold row indicates the selected best epoch.

	Pred. None	Pred. Soft	Pred. Hard
True None	98	7	2
True Soft	38	72	5
True Hard	10	19	70

Table 16: People-Centrism classifier: confusion matrix (best epoch). Rows are true classes; columns are predicted classes. None = no people-centrism; Soft = rhetorical appeals or non-binding consultation; Hard = binding participatory mechanisms.

Class	Precision	Recall	F1	Support
None	0.671	0.916	0.775	107
Soft	0.735	0.626	0.676	115
Hard	0.909	0.707	0.795	99
Macro average	0.772	0.750	0.749	321

Table 17: People-Centrism classifier: per-class performance (best epoch). Precision, recall, and F1 by class on the development set.

B.3 Aggregating Paragraph Predictions to Candidate Scores

The measurement models operate at the paragraph level. Consider a given electoral platform (document) d with N_d retained paragraphs after applying the minimum text-quality threshold described above. For each paragraph $i \in \{1, \ldots, N_d\}$, the classifier outputs a probability distribution over K ordered classes, denoted p_{ik} for $k = 1, \ldots, K$.

From probabilities to paragraph scores. Each paragraph-level probability distribution is converted into a single continuous score by taking the expected value of an ordinal coding of the classes:

$$s_i \;=\; \sum_{k=1}^{K} w_k \, p_{ik}, \tag{1}$$

where w_k assigns numeric values to the ordered classes. Evenly spaced codings are used: $w \in \{-2, -1, 0, 1, 2\}$ for the five-class ideology measure and $w \in \{0, 1, 2\}$ for the three-class measures (Populism, AER, and PC), so that intermediate values reflect model uncertainty and genuinely mixed paragraph content.

From paragraphs to a candidate score. The candidate-level score is the average paragraph score across all paragraphs in the platform:

$$S_d \;=\; \frac{1}{N_d} \sum_{i=1}^{N_d} s_i. \tag{2}$$

Equation (2) has a natural interpretation: S_d is the expected intensity of the trait for a paragraph drawn at random from the platform. Since each candidate submits a single platform, platform-level and candidate-level scores coincide. This probability-based aggregation avoids discretization noise from collapsing each paragraph to a hard label before aggregation and preserves within-document variation in model confidence.[25]

B.4 Illustrative Examples: Ideology, Populism, Anti-Elite Rhetoric, and People-Centrism

Ideology. Ideology classifies each paragraph into one of five ordered categories (extreme left, left, centre, right, extreme right); see Section 4 for the full definitions. Table 18 reports illustrative excerpts drawn from the corpus.

[25] In implementation, (2) is algebraically equivalent to first averaging class probabilities across paragraphs and then applying the same class weights, but the paragraph-level formulation clarifies the object being measured.

Category	Excerpt from Electoral Platform
Extreme Left Maximalist left positions: opposition to privatization, strong labor protections, public ownership, redistribution framed as social justice.	*"The municipality must actively defend workers' rights, oppose precarious employment, and prevent the outsourcing of essential public services. Privatization weakens labor protections, lowers wages, and undermines social justice. Public administration must guarantee stable employment and dignified working conditions for all those who deliver services to the community."*
Left Progressive but non-maximalist: inclusive education, anti-discrimination, social diversity framed as equity and pluralism.	*"We aim to promote educational projects designed to foster an understanding of the complexity of our society, which is expressed through differences in origin, gender, religion, or sexual orientation—differences that should be recognized as an added value rather than as a problem to be solved."*
Centre Managerial, technocratic, pragmatic language: service provision, infrastructure, efficiency, modernization without ideological orientation.	*"Redevelopment and efficiency improvements of public lighting across the entire municipal territory, through the replacement of existing poles with new LED lighting systems. This will allow not only for significant energy savings, but also for modernization and a noticeable improvement in the aesthetic quality of public lighting."*
Right Institutional order, legality, authority: law-and-order frames stressing enforcement and cooperation with law enforcement.	*"A vigilant, responsible, and rigorous administration can curb illegal presences through strengthened controls, cooperation with law-enforcement agencies, and respect for rules as the foundation of civil coexistence."*
Extreme Right Exclusionary nationalism, hierarchical deservingness: welfare chauvinism prioritizing citizens over outsiders.	*"In allocating scarce public resources, priority must be given to Italian citizens and long-standing residents who have contributed to the community through work and respect for the rules. Social housing and welfare cannot become incentives for uncontrolled immigration or tools that penalize those who have always lived here."*

Table 18: Ideology Classification: Excerpts from Electoral Platforms. Excerpts are drawn from actual municipal electoral platforms (translated from Italian).

Populism (Composite). The composite populism index is coded on a three-level scale (Hard, Soft, Non-populist); see Section 4 for the full definitions. Table 19 provides illustrative

examples for each category.

Category	Excerpt from Electoral Platform
Hard Explicit accusations of corruption, clientelism, capture, betrayal, or "privileges for the few"; claims to exclusive representation of popular will.	*"Behind this apparent calm lies an opaque society and economy, where clans, friends, friends of friends, and last-minute hangers-on compete for the favors of the powers that be. In this jungle of laws, by-laws, and regulations, it is easy for a simple rule to become a labyrinth for some and a highway for others. When, for too long, a right turns into a privilege, the healthy forces of a community lose momentum. In a democratic society, alternation in power is a foundational principle; when it freezes for decades, it becomes a troubling anomaly, suggesting a system blocked by reciprocal clientelistic logics."*
Soft People-centric appeals or weak elite-distance framing without hard accusations; moral renewal, bottom-up localism, or participatory ideals.	*"Over the past five years we have endured decisions imposed by the municipal administration without any real consultation or involvement of the population. This approach distances citizens from institutions and deepens distrust toward the authority that should be closest to them. We firmly believe in participation in public and political life, because the town belongs to those who live in it, not to those who administer it."*
None Technical, managerial, or descriptive policy language; may criticize inefficiency without framing politics as a people-vs-elite conflict.	*"Central to this reorganization of urban mobility is the relocation of through traffic outside the city. At present, vehicular traffic passing through the area overwhelms the historic center, generating serious problems in terms of congestion, road safety, and air pollution. The objective, to be achieved in a relatively short timeframe, is the construction of an alternative road connection to the Adriatic state highway, whose coastal route between [city] and [city] is wholly inadequate to handle traffic volumes, especially during the summer period."*

Table 19: Populism Classification: Excerpts from Electoral Platforms. Excerpts are drawn from actual municipal electoral platforms (translated from Italian); municipality names are redacted.

Anti-Elite Rhetoric. Anti-Elite Rhetoric captures the intensity of antagonistic claims directed at political elites or institutions, independently of any claim about citizen participation or decision authority. Table 20 provides illustrative examples for each category.

Category	Excerpt from Electoral Platform
Hard Explicit accusations attributing intentional wrongdoing to elites: corruption, self-enrichment, clientelism, betrayal, or capture by powerful interests.	*"The city of [name redacted] suffers the effects of desertification, decay, and annihilation caused by short-sighted and ideologically driven political visions, too often subordinated to the wishes of powerful interests and of the 'usual few,' who for nearly thirty years have seized control of public affairs in order to extract private advantages."*
Soft Distance, neglect, or detachment without hard accusations; representatives portrayed as removed from everyday life or inattentive to local needs.	*"Politics has become distant from the real problems of ordinary people, and political representatives, even at the local level, no longer live among the people nor have a genuine perception of what the population actually needs in everyday life in the city of [name redacted]."*
None Technical or managerial language; problems attributed to resource constraints or administrative complexity, not to political actors.	*"It will be essential to establish an effective working group on mobility, which should play a priority role within the strategic development plan for the territory. While awaiting more sustainable transport solutions, it will be important to improve and streamline existing transport services, for example by simplifying ticket purchases and making timetables more accessible."*

Table 20: Anti-Elite Rhetoric Classification: Excerpts from Electoral Platforms. Excerpts are drawn from actual municipal electoral platforms (translated from Italian); municipality names are redacted.

People-Centrism. People-Centrism captures whether citizens are framed as political agents who should directly influence or constrain municipal decision-making. Table 21 provides illustrative examples for each category.

Category	Excerpt from Electoral Platform
Hard Binding mechanisms transferring decision authority to citizens: participatory budgeting, binding referenda, recall procedures, or institutionalized citizen oversight of budgets and spending.	*"The people may delegate a limited portion of their sovereignty to their representatives, but they must always retain the freedom to modify the rules of that delegation. Through binding citizen-initiated referenda, political representation and parties will continue to exist, but they will be subordinated to the decisions of the citizens of the city of [name redacted]."*
Soft Consultative or non-binding involvement: participatory forums, promises of inclusion, or advisory bodies preserving representative discretion.	*"New participatory bodies will be created in the city's outlying areas, composed of individuals designated by local residents. These bodies will perform consultative and propositional functions vis-à-vis the municipal administration, serving as an additional channel through which inhabitants of decentralized areas can express their views and needs to local authorities in the city of [name redacted]."*
None Citizens framed as service recipients or information users; social participation without governance connection.	*"The city has a low share of urban gardens, while it records higher values with respect to equipped green spaces and school gardens. The proportion of outdoor sports areas is more limited. Over recent years, several issues have emerged concerning the maintenance conditions of existing green spaces and the distribution of different types of green areas across the urban territory."*

Table 21: People-Centrism Classification: Excerpts from Electoral Platforms. Excerpts are drawn from actual municipal electoral platforms (translated from Italian); municipality names are redacted.

B.5 Codebook Reliability via Independent Re-Coding

The reliability of the hand-coded training labels is assessed by an independent second coding of the full set of labeled paragraphs. Because the dimensions require following long, judgment-heavy codebooks on noisy OCR text, a large language model (GPT-5.1) was used as the second coder: it re-coded every labeled paragraph following the same codebooks reproduced in this appendix, blind to the original human annotations, judging each paragraph from its text alone. This check asks whether the codebooks are explicit enough that an independent coder, applying them mechanically, reproduces the human labels.

The re-coding covers every labeled paragraph: 2,196 for ideology, 1,301 for people-centrism, and 1,496 each for populism and anti-elite rhetoric. Agreement is summarized by Cohen's κ (Table 22). Because all four scales are ordinal (ideology runs from extreme left to extreme right, and the populism, AER, and PC scales run from none through soft to hard), the headline statistic is the quadratic-weighted κ, which credits near-misses (e.g., *left* vs. *extreme left*) more than far ones (e.g., *extreme left* vs. *extreme right*). Unweighted κ and raw percent

agreement are reported alongside.

Table 22: Inter-rater agreement: LLM second coder vs. human annotations.

Dimension	N	% agreement	κ	Weighted κ
Ideology (5-class)	2,196	79.8	0.741	0.898
People-centrism (PC)	1,301	78.6	0.674	0.792
Populism (ternary)	1,496	81.6	0.692	0.775
Anti-elite rhetoric (AER)	1,496	85.3	0.745	0.818

Note: Agreement between the human annotations and an independent LLM re-coding of the same paragraphs, following the codebooks in this appendix and blind to the human labels. κ is Cohen's kappa; the weighted column uses quadratic weights, appropriate for the ordinal scales. N is the number of human-labeled paragraphs for each dimension.

Agreement is substantial across all four dimensions. The quadratic-weighted κ ranges from 0.775 for populism to 0.898 for ideology (people-centrism 0.792, anti-elite rhetoric 0.818); because the scales are ordinal, disagreement falls overwhelmingly on adjacent categories, so the weighted κ exceeds the unweighted κ for every dimension.

C Formal Design and Estimation Details

This appendix provides the formal specification for the gradient pairwise close-race regression discontinuity design and its continuous-treatment extension. The main text gives the intuition; here the running variable, estimands, estimating equations, directional implementation for ideology, and identifying assumptions are defined.

C.1 Setup

Let municipalities be indexed by i. In each 2019 mayoral election, let $c_i^{(1)}$ and $c_i^{(2)}$ denote the winner and runner-up, with vote shares

$$\mathrm{VS}_{c_i^{(1)}} > \mathrm{VS}_{c_i^{(2)}}.$$

Each candidate has a measured score T_c, where T_c denotes the trait analyzed in a given specification: ideology, populism, anti-elite rhetoric (AER), or people-centrism (PC). The winner–runner-up vote margin is

$$\Delta_i = \mathrm{VS}_{c_i^{(1)}} - \mathrm{VS}_{c_i^{(2)}} > 0,$$

and the within-pair score gap is

$$G_i = \left|T_{c_i^{(1)}} - T_{c_i^{(2)}}\right|.$$

The signed running variable is

$$M_i = \Delta_i \cdot \operatorname{sign}\left(T_{c_i^{(1)}} - T_{c_i^{(2)}}\right).$$

Thus $M_i > 0$ if and only if the higher-scoring candidate wins, $M_i < 0$ if and only if the lower-scoring candidate wins. The vote margin Δ_i is observed directly from official election returns; the sign depends on which candidate has the higher trait score.

The threshold indicator is

$$Z_i = \mathbb{1}\{M_i \geq 0\},$$

so $Z_i = 1$ when the higher-scoring candidate wins and $Z_i = 0$ otherwise. The score of the elected mayor is

$$E_i = T_{c_i^{(1)}}.$$

The outcome of interest is denoted Y_i (a 2024 electoral outcome, a national-election vote share, or a municipal fiscal outcome, depending on the specification). Pre-treatment covariates are denoted X_i.

C.2 Gradient PCRD

The gradient PCRD (Bertoli and Hazlett, 2025) estimates the effect of barely electing the higher-scoring rather than the lower-scoring candidate. Let $Y_i(1)$ denote the potential outcome if the higher-scoring candidate wins and $Y_i(0)$ the potential outcome if the lower-scoring candidate wins. The estimand is

$$\tau^G = \mathbb{E}[Y_i(1) - Y_i(0) \mid M_i = 0]\,.$$

Under the close-election continuity assumption, this is identified by the discontinuity in the conditional expectation of Y_i at the cutoff:

$$\tau^G = \lim_{m \downarrow 0} \mathbb{E}[Y_i \mid M_i = m] - \lim_{m \uparrow 0} \mathbb{E}[Y_i \mid M_i = m].$$

Estimation uses local linear RD with separate slopes on each side of the cutoff. Within bandwidth h, the estimating equation is

$$Y_i \;=\; \alpha \;+\; \tau^G Z_i \;+\; \beta_- \, M_i \, \mathbb{1}\{M_i < 0\} \;+\; \beta_+ \, M_i \, \mathbb{1}\{M_i \geq 0\} \;+\; X_i'\gamma \;+\; \varepsilon_i, \qquad |M_i| \leq h,$$

where β_- is the slope to the left of the cutoff and β_+ is the slope to the right. The coefficient τ^G is the discontinuity at $M_i = 0$: the reduced-form effect of crossing from a lower-scoring to a higher-scoring winner. Observations are weighted by a triangular kernel,

$$K(M_i/h) \;=\; (1 - |M_i/h|)\, \mathbb{1}\{|M_i| \leq h\}.$$

Because the binary contrast Z_i does not scale with the magnitude of the within-pair score gap G_i, the gradient design treats all threshold crossings identically.

5% caliper and measurement-error robustness. The main specification imposes a minimal score-gap restriction

$$G_i > \kappa,$$

where κ is set to the 5th percentile of the empirical distribution of G_i within the directional sample, following the close-election extremism literature (Hall, 2015; Meisels, 2025). This drops pairs whose within-pair contrast is negligible. The caliper also serves as a robustness check against measurement error in the running variable. Because M_i is signed using *estimated* scores rather than observed ground truth, classification noise can misassign the within-pair sign for some municipalities, placing them on the wrong side of the cutoff. The vote margin Δ_i is observed precisely; only the sign — which candidate scores higher — depends on the classifier. Sign misclassification *attenuates* the RD estimate by blurring the discontinuity, biasing results toward zero rather than generating spurious effects. Progressively restricting the sample to pairs with larger G_i removes the municipalities most vulnerable to sign misassignment.

C.3 Continuous-Treatment PCRD

The continuous-treatment PCRD uses the same threshold but replaces the binary treatment with the elected mayor's continuous score E_i. The threshold indicator Z_i instruments for E_i: crossing the cutoff changes whether the higher- or lower-scoring candidate becomes mayor, which generates a discontinuity in the elected mayor's score.

Continuous-treatment fuzzy RD is rooted in Hahn, Todd and Van der Klaauw (2001) and Dong, Lee and Gou (2023), and has been applied to geographic discontinuities by

Ebenstein et al. (2017). The present paper is the first, to the best of the author's knowledge, to transpose this framework to the PCRD setting, where the electoral threshold instruments for the continuous score of the elected mayor rather than for a binary winner indicator.

The reduced-form discontinuity is

$$\tau^Y = \lim_{m \downarrow 0} \mathbb{E}[Y_i \mid M_i = m] - \lim_{m \uparrow 0} \mathbb{E}[Y_i \mid M_i = m],$$

and the first-stage discontinuity is

$$\tau^E = \lim_{m \downarrow 0} \mathbb{E}[E_i \mid M_i = m] - \lim_{m \uparrow 0} \mathbb{E}[E_i \mid M_i = m].$$

The CT-PCRD estimand is the local Wald ratio

$$\tau^{CT} = \frac{\tau^Y}{\tau^E},$$

interpreted as the local dose-response effect of a one-unit increase in the elected mayor's score, induced by barely winning a close election.

Estimation uses the same local linear RD structure as the gradient PCRD, with separate slopes on each side of the cutoff. The first-stage equation is

$$E_i \;=\; \pi_0 \;+\; \pi_1 Z_i \;+\; \pi_- \, M_i \, \mathbb{1}\{M_i < 0\} \;+\; \pi_+ \, M_i \, \mathbb{1}\{M_i \geq 0\} \;+\; X_i'\delta \;+\; \nu_i,$$

where π_1 is the discontinuity in the elected mayor's score at the cutoff, π_- is the left-side slope, and π_+ is the right-side slope. The reduced-form equation for Y_i has the same structure:

$$Y_i \;=\; \alpha \;+\; \rho_1 Z_i \;+\; \rho_- \, M_i \, \mathbb{1}\{M_i < 0\} \;+\; \rho_+ \, M_i \, \mathbb{1}\{M_i \geq 0\} \;+\; X_i'\gamma \;+\; u_i.$$

The CT-PCRD estimator is the Wald ratio

$$\widehat{\tau}^{CT} = \frac{\widehat{\rho}_1}{\widehat{\pi}_1},$$

implemented as local linear two-stage least squares with Z_i as the excluded instrument for E_i and the same bandwidth h, triangular kernel weights, and covariates as the gradient PCRD.

The relationship between the two estimators is

$$\tau^Y \;=\; \tau^{CT} \cdot \tau^E.$$

The gradient PCRD estimates the reduced-form threshold effect; the CT-PCRD rescales that reduced form by the first-stage jump in the elected mayor's score. CT-PCRD coefficients are therefore generally larger in magnitude, because they are expressed per unit of the candidate score rather than per threshold crossing.

C.4 Directional Implementation for Ideology

Ideology is a signed measure: right-leaning candidates have positive scores and left-leaning candidates have negative scores. A single global implementation would mix substantively different contrasts (a race between two left-wing candidates would not contribute meaningfully to the estimated effect of barely electing a more right-wing mayor). Ideology effects are therefore estimated separately on right- and left-directional subsamples.

For the right-wing design, define

$$\mathcal{S}^R = \big\{\, i : I_{c_i^{(1)}} > 0 \text{ or } I_{c_i^{(2)}} > 0 \,\big\},$$

where I_c is the ideology score of candidate c. Within $\mathcal{S}^R$, the score $T_c = I_c$ and the running variable is

$$M_i^R = \Delta_i \cdot \operatorname{sign}\big(I_{c_i^{(1)}} - I_{c_i^{(2)}}\big),$$

so $M_i^R > 0$ whenever the more right-leaning candidate wins. For the left-wing design, define

$$\mathcal{S}^L = \big\{\, i : I_{c_i^{(1)}} < 0 \text{ or } I_{c_i^{(2)}} < 0 \,\big\}.$$

Within $\mathcal{S}^L$, the score is negated, $T_c = -I_c$, so

$$M_i^L = \Delta_i \cdot \operatorname{sign}\big(I_{c_i^{(2)}} - I_{c_i^{(1)}}\big),$$

and $M_i^L > 0$ whenever the more left-leaning candidate wins. Both the gradient and continuous-treatment specifications are estimated separately in $\mathcal{S}^R$ and $\mathcal{S}^L$, with the corresponding elected-mayor treatments $E_i^R = I_{c_i^{(1)}}$ and $E_i^L = -I_{c_i^{(1)}}$.

C.5 Implementation for Populism, AER, and People-Centrism

The populism measures (composite populism, anti-elite rhetoric, and people-centrism) are non-signed: higher values always indicate more of the relevant trait. Estimation therefore uses the full close-races sample for each measure, with T_c instantiated to the corresponding score, M_i defined as above, and $E_i = T_{c_i^{(1)}}$.

C.6 Identification Assumptions

The identifying assumption is the standard continuity condition for close-election RD, adapted to the pairwise candidate-characteristic setting.

- **Continuity:** potential outcomes evolve smoothly through the cutoff. Formally, $\mathbb{E}[Y_i(d) \mid M_i = m]$ is continuous in m at $m = 0$ for all values of the elected mayor's score d.
- **No manipulation:** candidates cannot precisely control the final signed running variable M_i near the cutoff. Manipulation would require sorting exactly around the threshold in favor of the higher- or lower-scoring candidate, not simply winning by a larger or smaller margin.
- **Exclusion (CT-PCRD):** barely crossing the threshold affects outcomes only through the score of the elected mayor E_i.
- **Relevance (CT-PCRD):** the threshold must generate a nonzero first-stage discontinuity in the elected mayor's score, $\tau^E \neq 0$ (testable via the first stage).

These conditions are assessed empirically through density tests at the cutoff and covariate-balance checks.

C.7 Covariates, Bandwidths, and Inference

The main specifications include region fixed effects and log municipal population as pre-treatment covariates,

$$X_i = [\,\text{Region FE}_i,\ \log(\text{Population}_{i,2018})\,].$$

Covariates are not required for identification but improve precision and adjust for residual imbalance in finite samples. All estimates use local linear specifications with triangular kernel weights; bandwidths are selected using the MSE-optimal procedure implemented in `rdrobust`. Reported standard errors and p-values are robust bias-corrected RD quantities. For CT-PCRD, the first stage, reduced form, and local IV estimate use the same running-variable orientation, bandwidth logic, kernel, and covariate set.

D RD Validity and Robustness Checks

D.1 Density Continuity Tests

Density continuity tests check whether the density of the running variable is continuous at the cutoff: a jump would indicate sorting — elections where the ideologically more extreme (or more populist) candidate barely won being systematically more or less frequent than those where the more moderate (or less populist) candidate barely won — which would threaten the identifying assumption of local exchangeability around the cutoff. Both tests are reported on the 5%-trimmed subsample used in the main gradient PCRD specification. Following the recommendation of Bertoli and Hazlett (2025) for gradient PCRD designs, the same tests are also reported on the full directional sample of close races prior to any caliper restriction, to maximize statistical power; the conclusion is identical across both samples, with all four designs passing density continuity at conventional levels.

Ideology design ($\mathcal{S}^R$)

On the 5%-trimmed subsample ($N = 778$), the `rddensity` test of Cattaneo, Jansson and Ma (2020) applied to the ideology running variable M_i in $\mathcal{S}^R$ yields $T = -1.438$ ($p = 0.151$); the McCrary (2008) test yields $\theta = -0.161$ ($z = -0.988$, $p = 0.323$). Following Bertoli and Hazlett (2025), the same tests on the full directional sample ($N = 819$) yield similar conclusions (`rddensity`: $T = -1.514$, $p = 0.130$; McCrary: $\theta = -0.221$, $z = -1.327$, $p = 0.185$). No evidence of sorting at the ideology cutoff in either sample (Figure 8).

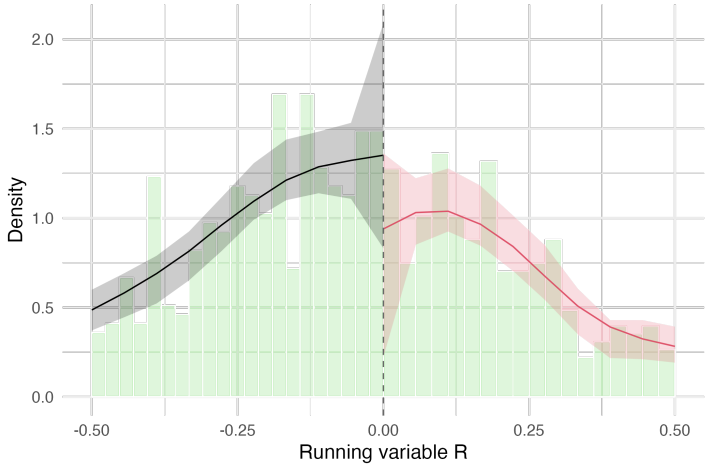

Figure 8: Density continuity test (Cattaneo, Jansson and Ma, 2020) for the running variable M_i in the right sample ($\mathcal{S}^R$). The histogram shows the empirical distribution of M_i on the full sample; the curves are local polynomial density estimates on each side of the cutoff with 95% confidence bands. Trimmed sample ($N = 778$): `rddensity` $T = -1.438$, $p = 0.151$; McCrary $\theta = -0.161$, $p = 0.323$. Full sample ($N = 819$): `rddensity` $T = -1.514$, $p = 0.130$; McCrary $\theta = -0.221$, $p = 0.185$.

Populism design

On the 5%-trimmed subsample ($N = 1{,}028$), the `rddensity` test (Cattaneo, Jansson and Ma, 2020) applied to the populism running variable yields $T = 0.770$ ($p = 0.441$); the McCrary (2008) test yields $\theta = 0.109$ ($z = 0.754$, $p = 0.451$). Following Bertoli and Hazlett (2025), the same tests on the full directional sample ($N = 1{,}083$) yield similar conclusions (`rddensity`: $T = 0.804$, $p = 0.421$; McCrary: $\theta = 0.086$, $z = 0.621$, $p = 0.535$). No evidence of sorting at the populism cutoff in either sample (Figure 9).

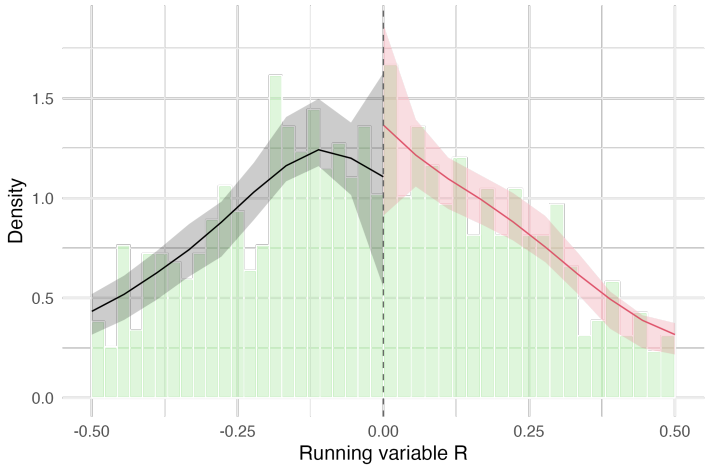

Figure 9: Density continuity test (Cattaneo, Jansson and Ma, 2020) for the populism ternary running variable M_i^{pop}. Histogram on the full sample. Trimmed sample ($N = 1{,}028$): `rddensity` $T = 0.770$, $p = 0.441$; McCrary $\theta = 0.109$, $p = 0.451$. Full sample ($N = 1{,}083$): `rddensity` $T = 0.804$, $p = 0.421$; McCrary $\theta = 0.086$, $p = 0.535$.

Anti-Elite Rhetoric (AER) design

On the 5%-trimmed subsample ($N = 1{,}028$), the `rddensity` test (Cattaneo, Jansson and Ma, 2020) applied to the AER running variable yields $T = -0.590$ ($p = 0.555$); the McCrary (2008) test yields $\theta = -0.086$ ($z = -0.603$, $p = 0.547$). Following Bertoli and Hazlett (2025), the same tests on the full directional sample ($N = 1{,}083$) yield similar conclusions (`rddensity`: $T = -0.473$, $p = 0.637$; McCrary: $\theta = -0.041$, $z = -0.280$, $p = 0.779$). No evidence of sorting at the AER cutoff in either sample (Figure 10).

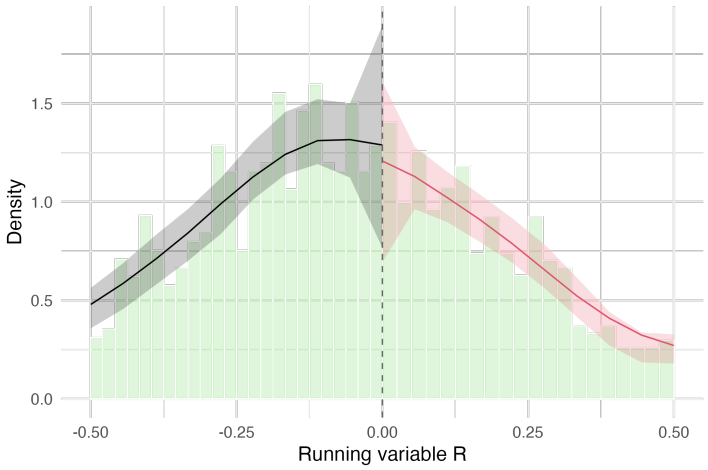

Figure 10: Density continuity test (Cattaneo, Jansson and Ma, 2020) for the anti-elite rhetoric (AER) running variable. Histogram on the full sample. Trimmed sample ($N = 1{,}028$): `rddensity` $T = -0.590$, $p = 0.555$; McCrary $\theta = -0.086$, $p = 0.547$. Full sample ($N = 1{,}083$): `rddensity` $T = -0.473$, $p = 0.637$; McCrary $\theta = -0.041$, $p = 0.779$.

People-Centrism (PC) design

On the 5%-trimmed subsample ($N = 1{,}028$), the `rddensity` test (Cattaneo, Jansson and Ma, 2020) applied to the PC running variable yields $T = 0.259$ ($p = 0.795$); the McCrary (2008) test yields $\theta = -0.080$ ($z = -0.528$, $p = 0.597$). Following Bertoli and Hazlett (2025), the same tests on the full directional sample ($N = 1{,}083$) yield similar conclusions (`rddensity`: $T = 0.349$, $p = 0.727$; McCrary: $\theta = -0.068$, $z = -0.475$, $p = 0.635$). No evidence of sorting at the PC cutoff in either sample (Figure 11).

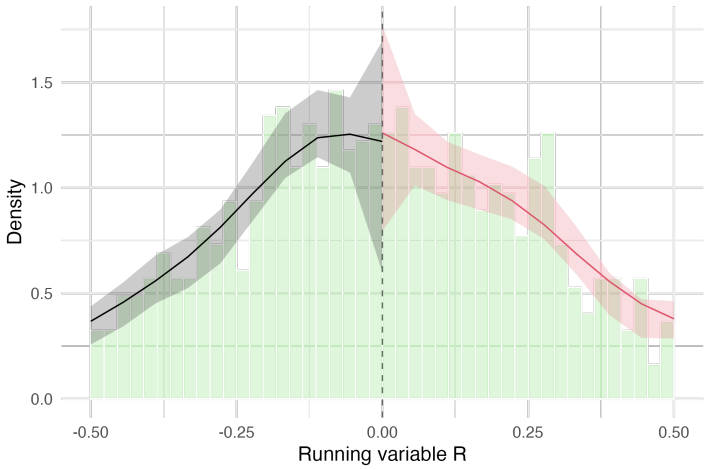

Figure 11: Density continuity test (Cattaneo, Jansson and Ma, 2020) for the people-centrism (PC) running variable. Histogram on the full sample. Trimmed sample ($N = 1{,}028$): `rddensity` $T = 0.259$, $p = 0.795$; McCrary $\theta = -0.080$, $p = 0.597$. Full sample ($N = 1{,}083$): `rddensity` $T = 0.349$, $p = 0.727$; McCrary $\theta = -0.068$, $p = 0.635$.

D.2 Covariate Balance

Each district-level covariate (log population, urbanization, 2014 incumbent on the 2019 ballot, and region fixed effects) and each pre-treatment electoral outcome (Lega and M5S vote shares from the 2018 Camera and 2014 European elections) is used as the outcome in an `rdrobust` call with the respective design's running variable, following the univariate balance test of Bertoli and Hazlett (2025) (their Figure 4).

Each balance figure shows the 5%-trimmed sample (top panel) and the full directional sample (bottom panel).

Ideology design ($\mathcal{S}^R$)

Toscana is the only district-level cell flagged at 5% ($\hat{\tau} = -0.101$, $p = 0.036$ trim; $p = 0.037$ full); Marche and urbanization are marginal at 10% (Figure 12). All four lagged electoral outcomes pass cleanly in both samples (smallest $p = 0.352$): the $\mathcal{S}^R$ cutoff predicts neither past Lega nor past M5S vote shares. With 22 covariates tested, one 5% rejection is within chance expectations.

Table 23 adds ISTAT urbanization as a covariate on the full close-races universe (no region drops; region fixed effects absorb level differences, and the lagged outcomes pass balance cleanly). The vote-weighted ideology result is essentially unchanged (gradient $p = 0.003$; CT-PCRD $p = 0.012$); the Lega European 2024 estimates strengthen under both estimators (gradient $p = 0.007$; CT-PCRD $p = 0.048$).

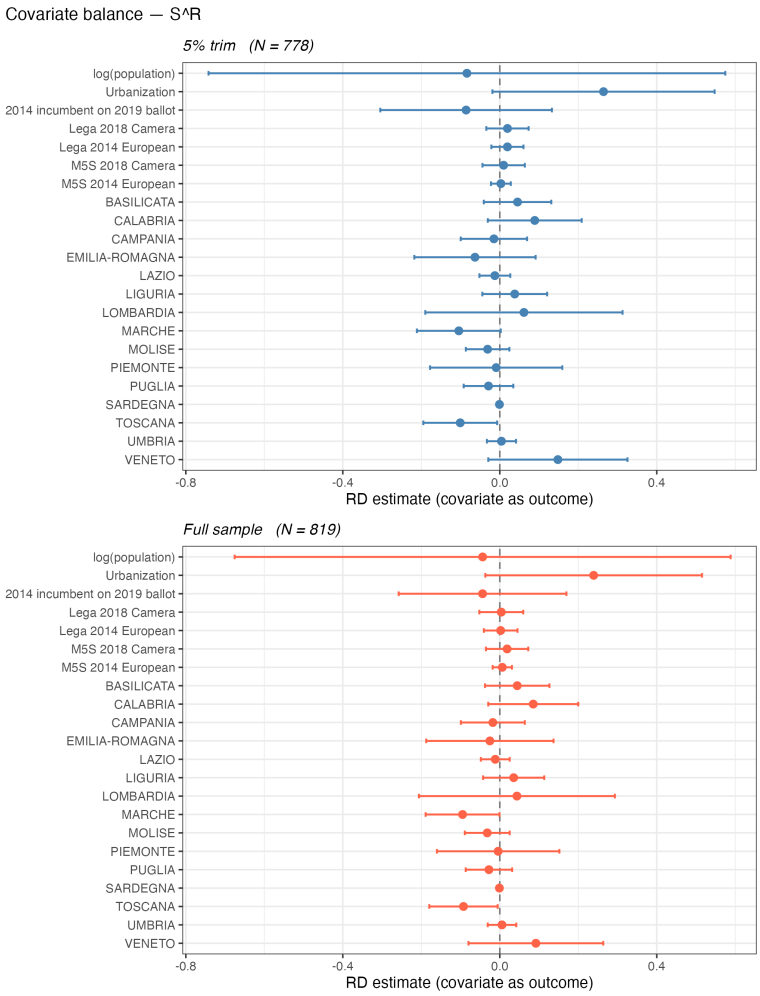

Figure 12: Covariate balance for the ideology design ($\mathcal{S}^R$). Top panel: 5%-trimmed sample ($N = 778$). Bottom panel: full directional sample ($N = 819$). Each point is the `rdrobust` estimate using that covariate as the outcome, with M_i as the running variable; bars are 95% robust confidence intervals.

Table 23: Robustness: Controlling for Urbanization — Ideology Design.

	VW Ideology 2024		*Lega European 2024*	
	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	-0.055^{***}	-0.274^{**}	-0.024^{**}	-0.105^{**}
SE	(0.019)	(0.108)	(0.010)	(0.049)
p-value	0.005	0.011	0.013	0.033
First stage	—	0.202^{***}	—	0.206^{***}
SE		(0.031)		(0.032)
p-value		0.000		0.000
N (bandwidth)	269	385	227	349
N (total)	773	814	778	819
Bandwidth	0.143	0.197	0.121	0.181

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. Same specification as the main $\mathcal{S}^R$ result with ISTAT urbanization grade added as an additional covariate (urbanization is marginally flagged at the 10% level in the covariate balance test). No regions are dropped: region fixed effects absorb level differences, and the lagged Lega and M5S vote shares pass balance cleanly across both samples. Gradient PCRD is estimated on the 5%-trimmed sample, consistent with the main specification; CT-PCRD uses the full sample. Covariates: region fixed effects, log municipal population, and urbanization grade. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Populism design

Piemonte is flagged at 5% trim only ($\hat{\tau} = 0.140$, $p = 0.027$; $p = 0.204$ full); Toscana and Umbria are marginal at 10%. Lagged outcomes split between Lega (both samples balanced, $p \geq 0.235$) and M5S, where the 2014 European share fails in both samples ($\hat{\tau} = 0.032$, $p = 0.006$ trim, $p = 0.016$ full) and the 2018 Camera share is borderline in trim ($\hat{\tau} = 0.041$, $p = 0.048$; $p = 0.116$ full) (Figure 13). Municipalities barely won by populism-ternary candidates had $\approx$ 3–4 percentage points higher pre-treatment M5S support, direction- and lag-consistent.

Table 24 adds the format-matched lagged M5S vote share as a covariate on the full close-races universe: the 2018 Camera M5S share for the M5S Camera 2022 outcome and the 2014 European M5S share for the M5S European 2024 outcome. The Camera 2022 effect collapses once the pre-treatment M5S sorting is absorbed (gradient $p = 0.337$; CT-PCRD $p = 0.337$), confirming that much of the baseline estimate reflected pre-existing M5S imbalance. The European 2024 effect attenuates from the baseline but the CT-PCRD remains significant ($\hat{\tau} = 0.080$, $p = 0.020$) and the gradient PCRD remains marginal ($p = 0.063$). The popT M5S results are therefore reported as suggestive corroboration of the People-Centrism design,

which provides the clean identification for M5S vote-share outcomes.

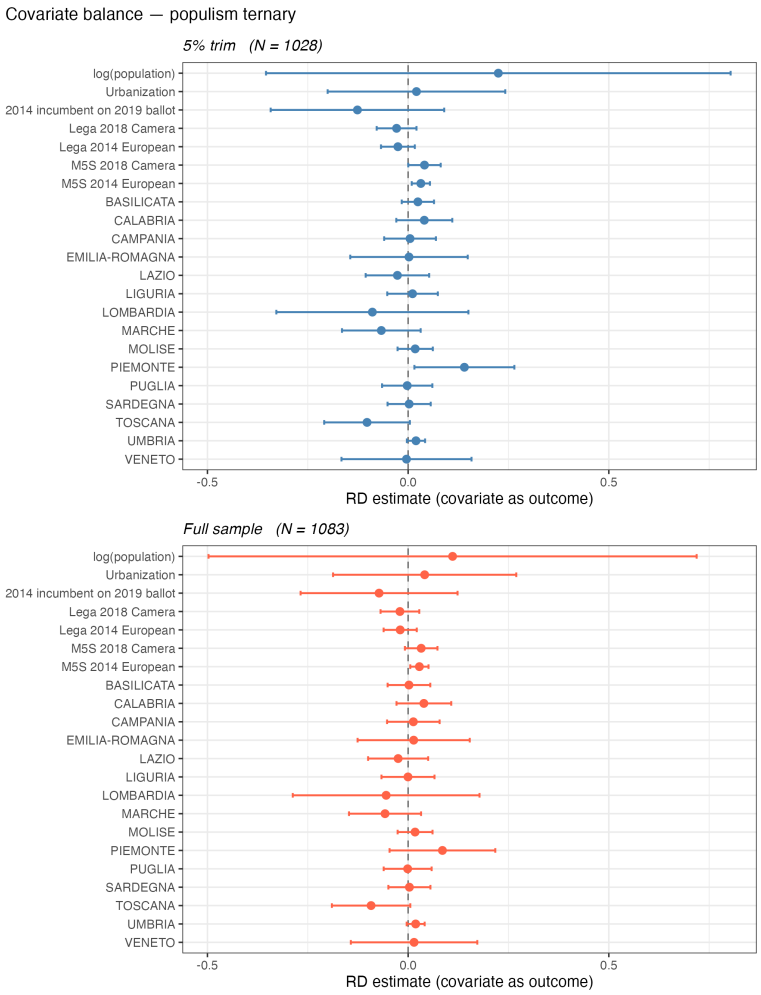

Figure 13: Covariate balance for the populism design. Top panel: 5%-trimmed sample ($N = 1{,}028$). Bottom panel: full directional sample ($N = 1{,}083$). Each point is the `rdrobust` estimate using that covariate as the outcome with M_i^{pop} as the running variable; bars are 95% robust confidence intervals.

Table 24: Robustness: Controlling for Lagged M5S Vote Share — Populism Design.

	M5S Camera 2022		*M5S European 2024*	
	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	0.004	0.039	0.007*	0.080**
SE	(0.005)	(0.040)	(0.004)	(0.034)
p-value	0.337	0.337	0.063	0.020
First stage	—	0.112***	—	0.109***
SE		(0.029)		(0.026)
p-value		0.000		0.000
N (bandwidth)	162	330	386	391
N (total)	1,016	1,070	1,011	1,065
Bandwidth	0.065	0.127	0.159	0.151

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. Same specification as the main populism result with the format-matched pre-treatment M5S vote share added as an additional covariate: the 2018 Camera M5S share for the M5S Camera 2022 outcome, and the 2014 European M5S share for the M5S European 2024 outcome (each lagged share is flagged in the covariate balance test for the populism-ternary design). No regions are dropped. The Camera 2022 effect collapses once the pre-treatment M5S sorting is absorbed; the European 2024 effect attenuates but the CT-PCRD remains significant at the 5% level and the Gradient PCRD remains marginal at the 10% level. Covariates: region fixed effects, log municipal population, and the format-matched lagged M5S vote share. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Anti-Elite Rhetoric design

No district-level covariate is significant at 5%; Puglia is marginal at 10%. Three of four lagged electoral outcomes fail in both samples: Lega 2018 Camera ($\hat{\tau} = -0.089$, $p = 0.002$ trim, $p = 0.001$ full), Lega 2014 European ($p = 0.044$ / $p = 0.029$), and M5S 2018 Camera ($\hat{\tau} = 0.076$, $p = 0.004$ / $p = 0.002$); only M5S 2014 European passes (Figure 14). AER-treated municipalities had systematically lower pre-treatment Lega support and higher pre-treatment M5S support, indicating pre-treatment sorting at the AER cutoff. For this reason AER is reported throughout the paper for descriptive comparison rather than as an independent identification.

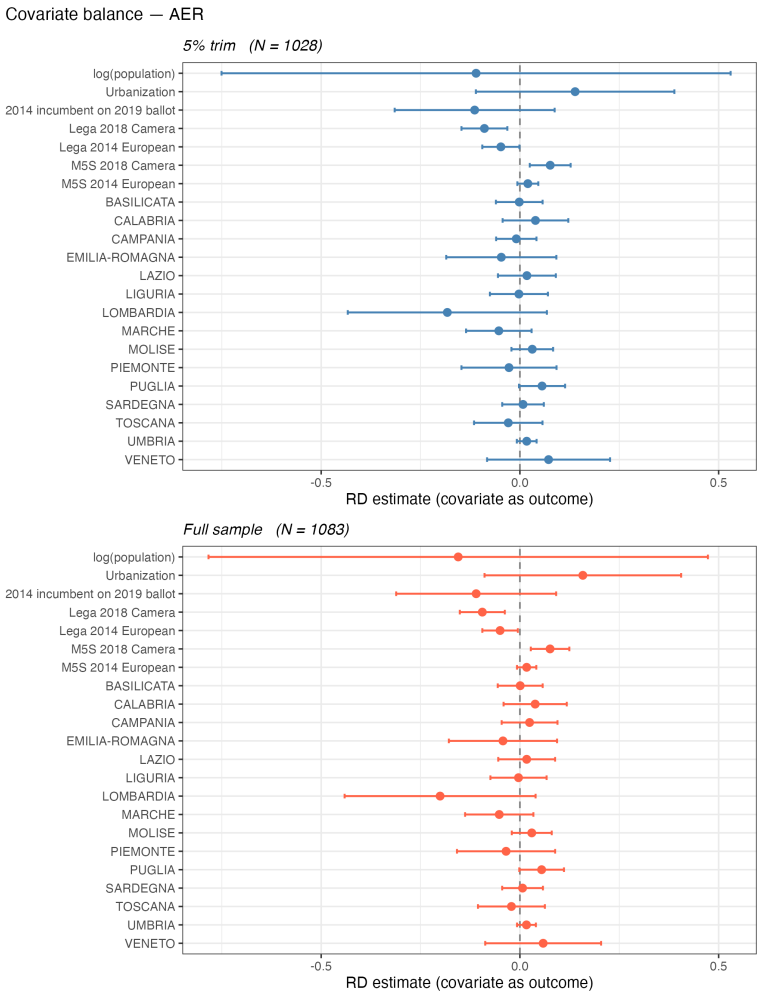

Figure 14: Covariate balance for the anti-elite rhetoric design. Top panel: 5%-trimmed sample ($N = 1{,}028$). Bottom panel: full directional sample ($N = 1{,}083$). Each point is the `rdrobust` estimate using that covariate as the outcome with the AER running variable; bars are 95% robust confidence intervals.

People-Centrism design

No covariate is significant at 5% trim; Toscana is marginal at 10% ($p = 0.097$) and shifts to 5% in the full sample ($p = 0.046$). All four lagged electoral outcomes pass cleanly in both samples (smallest $p = 0.137$) (Figure 15): the PC cutoff predicts neither past Lega nor past M5S vote shares, identifying M5S vote-share treatment effects without the pre-treatment sorting that flags the popT design.

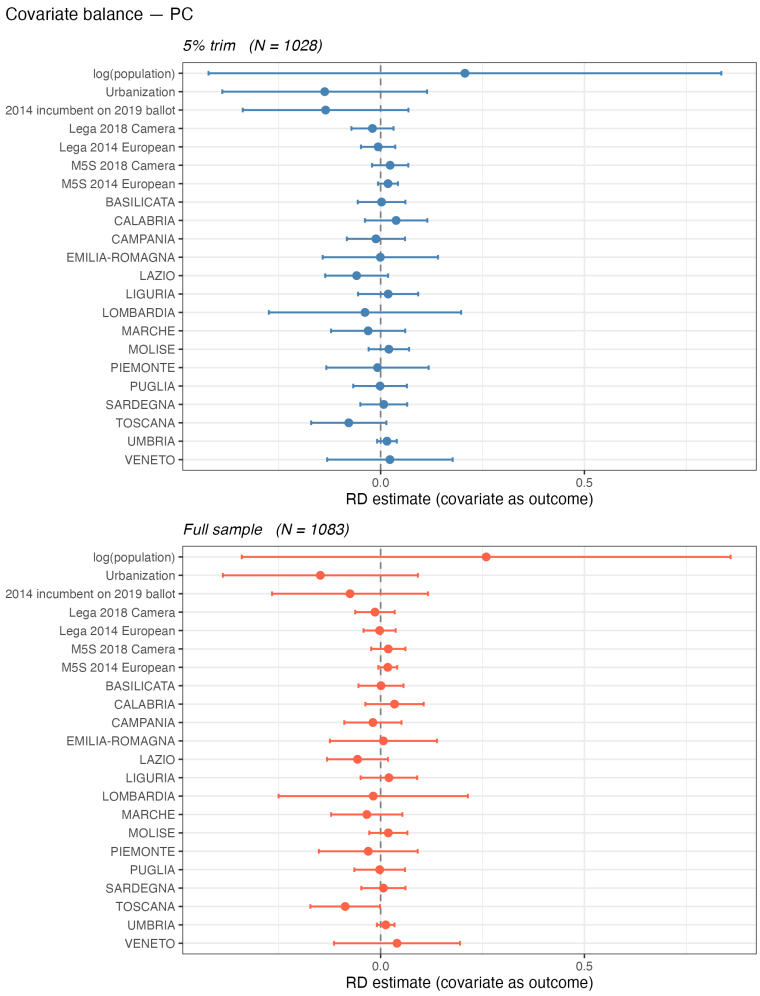

Figure 15: Covariate balance for the people-centrism design. Top panel: 5%-trimmed sample ($N = 1{,}028$). Bottom panel: full directional sample ($N = 1{,}083$). Each point is the `rdrobust` estimate using that covariate as the outcome with the PC running variable; bars are 95% robust confidence intervals.

D.3 Predicted Treatment

Bertoli and Hazlett (2025) recommend testing whether a prediction of the binary treatment constructed from pre-election characteristics alone exhibits a discontinuity at the cutoff. For each design, $\hat{D} \equiv \Pr(D_i = 1 \mid X_i)$ is constructed by fitting a logistic regression of the treatment indicator on X_i (region fixed effects, log municipal population, ISTAT urbanization grade, 2014 incumbent on 2019 ballot, and the 2018 Camera intrinsic-strength share — Lega + FdI + FI for the ideology design, M5S for the populism designs — popT, AER, and PC) using observations away from the cutoff ($|M_i| > 0.10$); the fitted model then generates predicted probabilities for the full sample. Each figure shows the 5%-trimmed sample (top panel) and the full directional sample (bottom panel).

Ideology design ($\mathcal{S}^R$)

$\hat{D}$ shows no discontinuity at $M_i = 0$ (Figure 16): trim $\hat{\tau} = 0.036$ ($p = 0.358$, $N = 768$, pseudo-$R^2 = 0.060$); full $\hat{\tau} = 0.029$ ($p = 0.438$, $N = 809$).

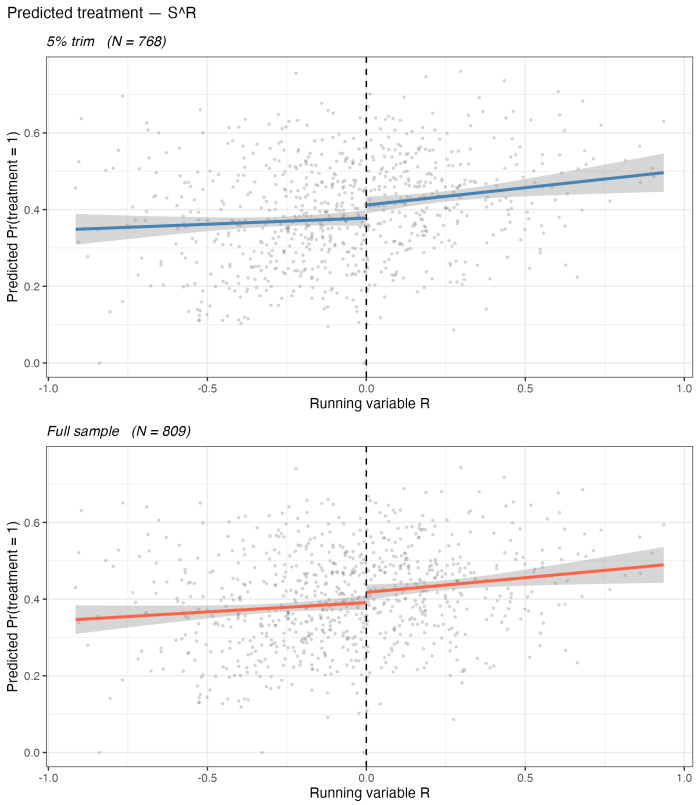

Figure 16: Predicted treatment $\hat{D}$ against M_i ($\mathcal{S}^R$). Top panel: 5%-trimmed sample. Bottom panel: full directional sample. Lines are local linear fits with 95% confidence bands.

Populism design

$\hat{D}^{pop}$ shows no discontinuity (Figure 17): trim $\hat{\tau} = -0.004$ ($p = 0.837$, $N = 1{,}013$, pseudo-$R^2 = 0.023$); full $\hat{\tau} = 0.003$ ($p = 0.916$, $N = 1{,}067$).

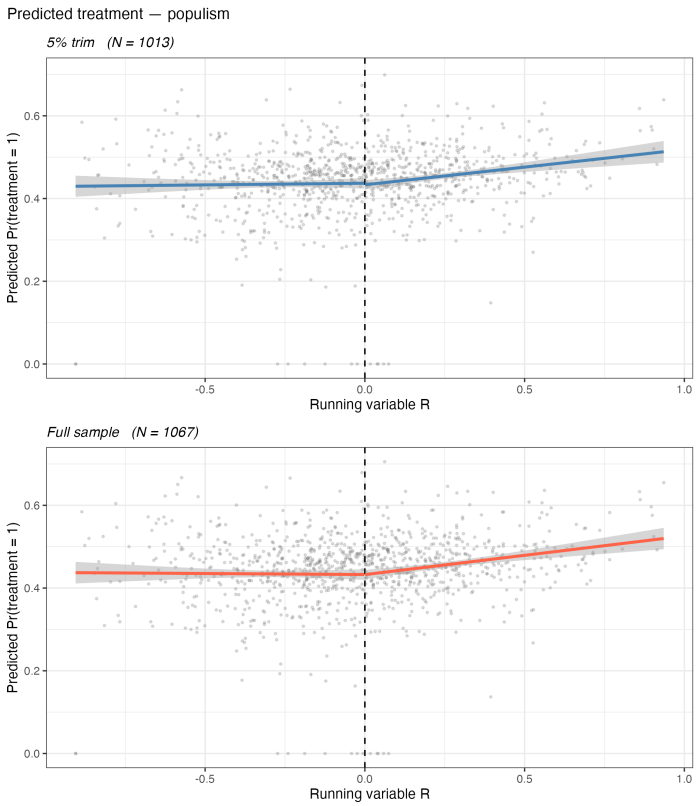

Figure 17: Predicted treatment $\hat{D}^{pop}$ against M_i^{pop}. Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

Anti-Elite Rhetoric design

$\hat{D}^{AER}$ shows no discontinuity (Figure 18): trim $\hat{\tau} = 0.032$ ($p = 0.320$, $N = 1{,}013$, pseudo-$R^2 = 0.042$); full $\hat{\tau} = 0.027$ ($p = 0.291$, $N = 1{,}067$).

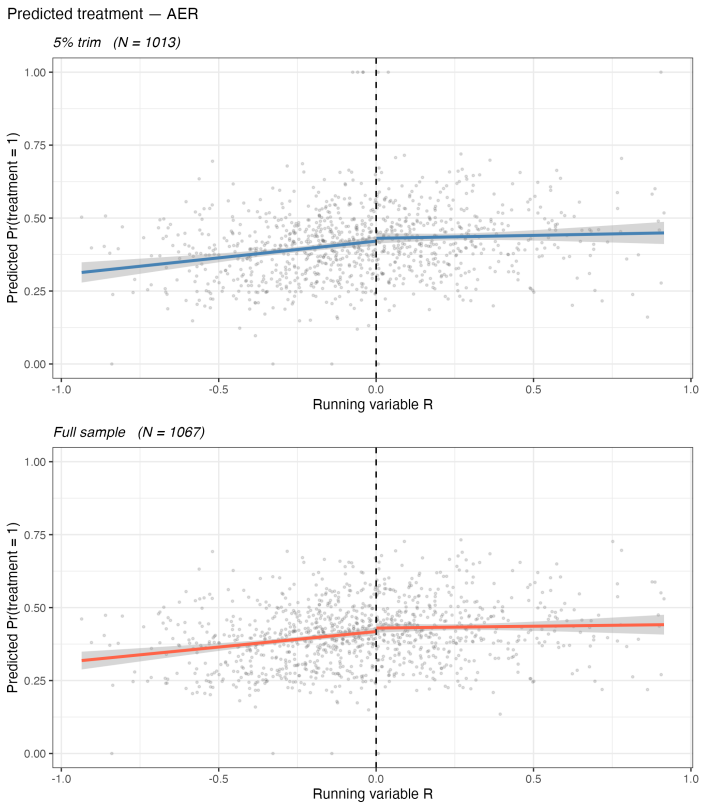

Figure 18: Predicted treatment $\hat{D}^{AER}$ against M_i^{AER}. Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

People-Centrism design

$\hat{D}^{PC}$ shows no discontinuity (Figure 19): trim $\hat{\tau} = -0.021$ ($p = 0.161$, $N = 1{,}012$, pseudo-$R^2 = 0.010$); full $\hat{\tau} = -0.018$ ($p = 0.166$, $N = 1{,}067$).

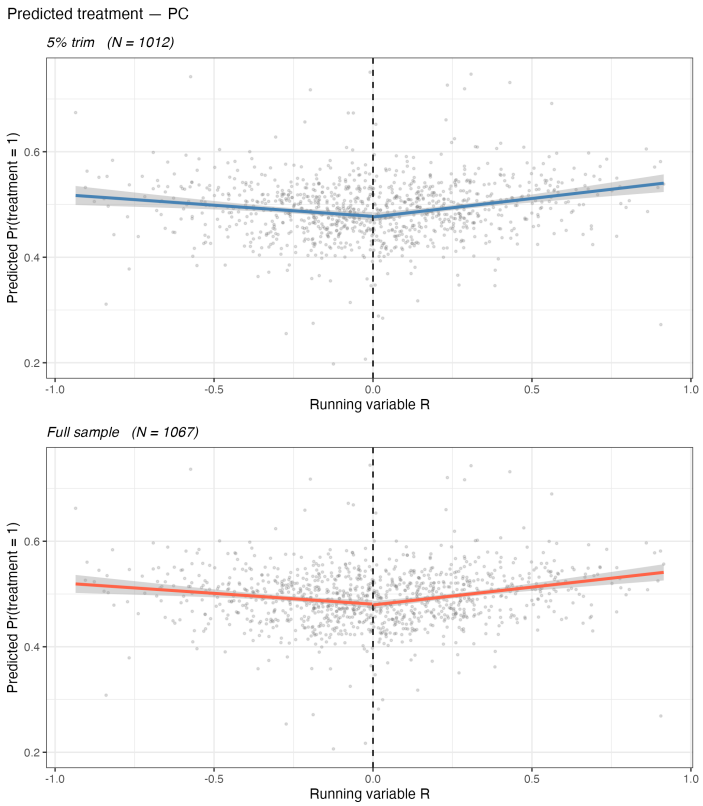

Figure 19: Predicted treatment $\hat{D}^{PC}$ against M_i^{PC}. Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

D.4 Predicted Potential Outcomes

The predicted potential outcomes test is more directly linked to bias than the predicted treatment test (Bertoli and Hazlett, 2025). The idea is to ask: conditional on pre-election municipality characteristics, do the two sides of the cutoff look as if they would have produced the same outcome regardless of who won?

$\hat{Y}_i(0)$ is approximated by regressing Y on the baseline covariate vector (region fixed effects, log population, urbanization, 2014 incumbent on 2019 ballot) augmented by a pre-treatment intrinsic-strength control (2018 Camera Lega + FdI + FI share for ideology designs; 2018 Camera M5S share for populism designs), using observations on the control side of the cutoff with $|M_i| \leq 0.10$. The fitted model generates predicted values $\hat{Y}_i(0)$ for the full sample and an RD discontinuity is tested. $\hat{Y}_i(1)$ is constructed analogously using treated-side near-cutoff observations. Each figure shows the 5%-trimmed sample (top panel) and the full directional sample (bottom panel); a discontinuity in either panel — especially with both $\hat{Y}(0)$ and $\hat{Y}(1)$ pointing in the same direction — would indicate pre-treatment imbalance.

Ideology design ($\mathcal{S}^R$)

Main outcome (vote-weighted ideology). $\hat{Y}(0)$: trim $\hat{\tau} = 0.009$ ($p = 0.505$, $N_{\text{train}} = 97$); full $\hat{\tau} = 0.002$ ($p = 0.914$). $\hat{Y}(1)$: trim $\hat{\tau} = 0.000$ ($p = 0.980$, $N_{\text{train}} = 84$); full $\hat{\tau} = -0.004$ ($p = 0.664$). All four pass (Figures 20–21).

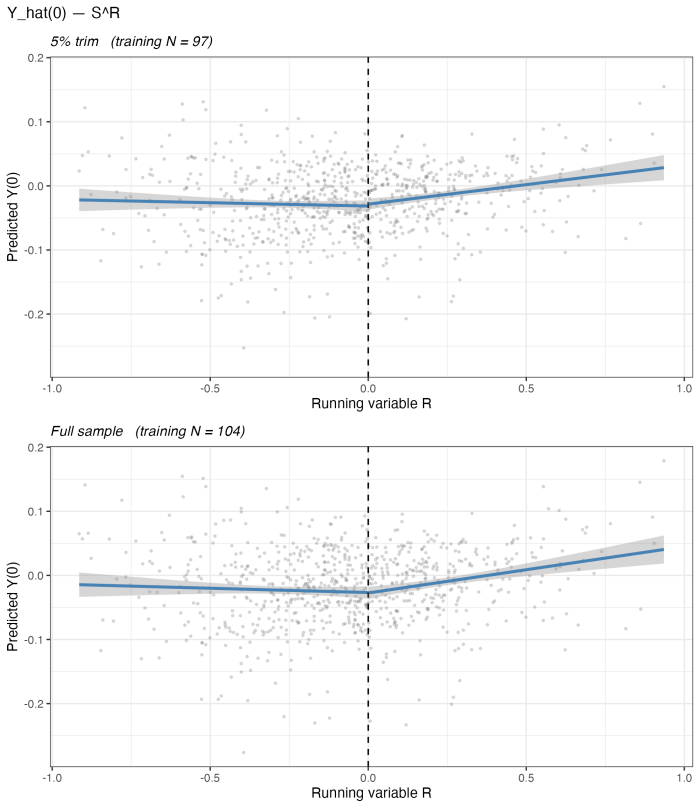

Figure 20: Predicted potential outcome $\hat{Y}(0)$ against M_i ($\mathcal{S}^R$, vote-weighted ideology 2024). Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

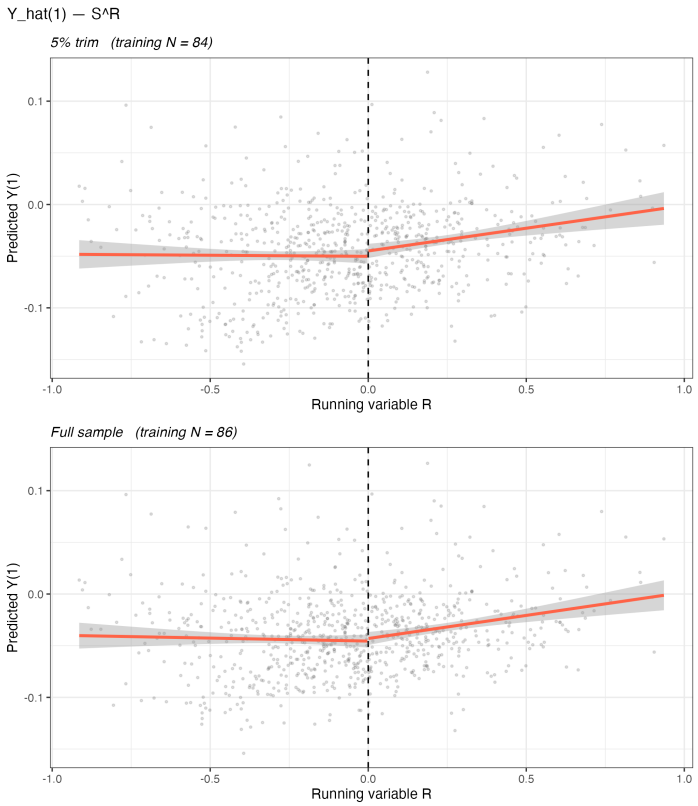

Figure 21: Predicted potential outcome $\hat{Y}(1)$ against M_i ($\mathcal{S}^R$, vote-weighted ideology 2024). Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

Lega European 2024 outcome. $\hat{Y}(0)$: trim $\hat{\tau} = 0.016$ ($p = 0.254$); full $\hat{\tau} = 0.008$ ($p = 0.545$). $\hat{Y}(1)$: trim $\hat{\tau} = 0.016$ ($p = 0.149$); full $\hat{\tau} = 0.008$ ($p = 0.457$). All four pass (Figures 22–23).

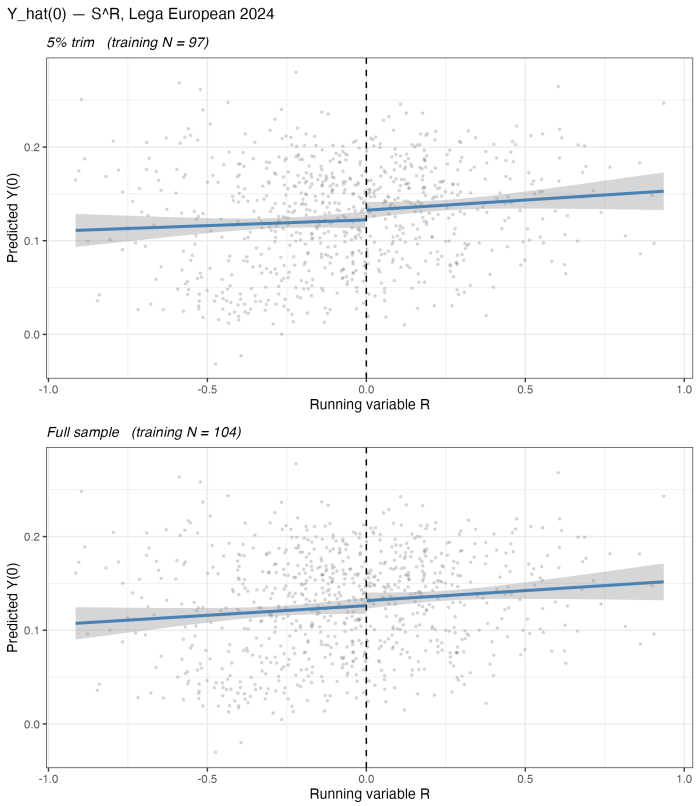

Figure 22: Predicted potential outcome $\hat{Y}(0)$ against M_i ($\mathcal{S}^R$, Lega EU 2024 outcome). Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

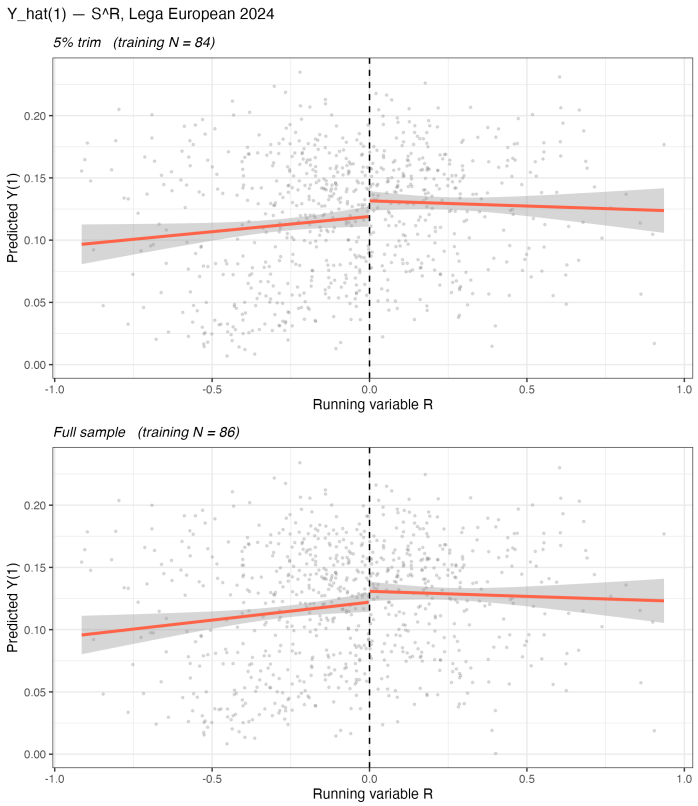

Figure 23: Predicted potential outcome $\hat{Y}(1)$ against M_i ($\mathcal{S}^R$, Lega EU 2024 outcome). Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

Populism composite design — M5S vote share

Camera 2022. $\hat{Y}^{pop}(0)$: trim $\hat{\tau} = 0.023$ ($p = 0.098$); full $\hat{\tau} = 0.016$ ($p = 0.266$). $\hat{Y}^{pop}(1)$: trim $\hat{\tau} = 0.023$ ($p = 0.151$); full $\hat{\tau} = 0.017$ ($p = 0.305$). **European 2024.** $\hat{Y}^{pop}(0)$: trim $\hat{\tau} = 0.016$ ($p = 0.104$); full $\hat{\tau} = 0.013$ ($p = 0.189$). $\hat{Y}^{pop}(1)$: trim $\hat{\tau} = 0.016$ ($p = 0.089$); full $\hat{\tau} = 0.013$ ($p = 0.179$).

The four trimmed-sample p-values cluster between 0.09 and 0.15 with both $\hat{Y}^{pop}(0)$ and $\hat{Y}^{pop}(1)$ pointing positive — a stable but borderline imbalance. The full-sample placebos pass (smallest $p = 0.179$). The popT M5S finding is reported here as suggestive corroboration of the clean PC-design M5S result.

People-Centrism design — M5S vote share

Camera 2022. $\hat{Y}^{PC}(0)$: trim $\hat{\tau} = 0.009$ ($p = 0.558$); full $\hat{\tau} = 0.006$ ($p = 0.709$). $\hat{Y}^{PC}(1)$: trim $\hat{\tau} = 0.010$ ($p = 0.575$); full $\hat{\tau} = 0.006$ ($p = 0.715$). All four pass (Figures 24–25).

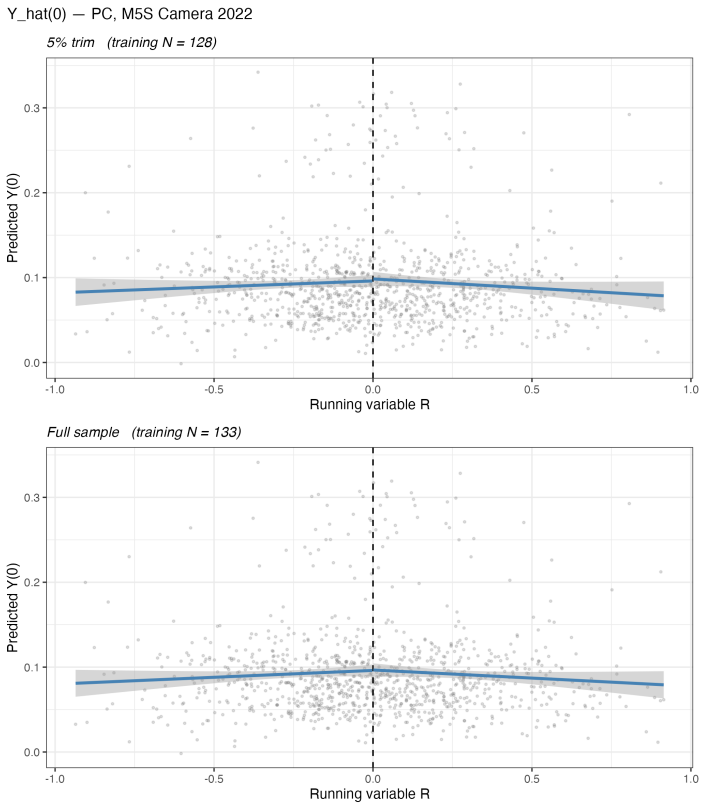

Figure 24: Predicted potential outcome $\hat{Y}^{PC}(0)$ against M_i^{PC} (M5S Camera 2022 outcome). Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

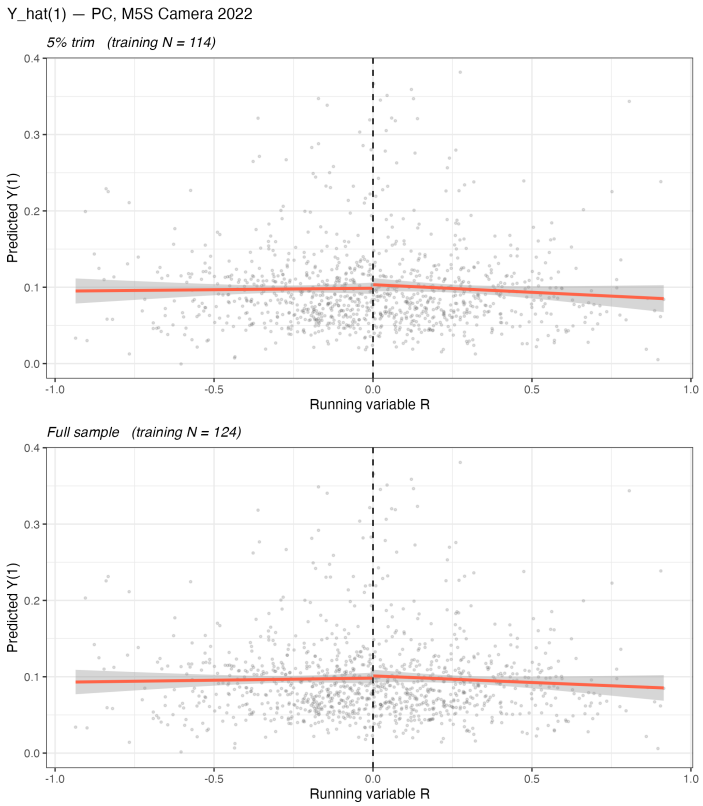

Figure 25: Predicted potential outcome $\hat{Y}^{PC}(1)$ against M_i^{PC} (M5S Camera 2022 outcome). Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

European 2024. $\hat{Y}^{PC}(0)$: trim $\hat{\tau} = 0.006$ ($p = 0.559$); full $\hat{\tau} = 0.004$ ($p = 0.711$). $\hat{Y}^{PC}(1)$: trim $\hat{\tau} = 0.008$ ($p = 0.419$); full $\hat{\tau} = 0.006$ ($p = 0.536$). All four pass (Figures 26–27).

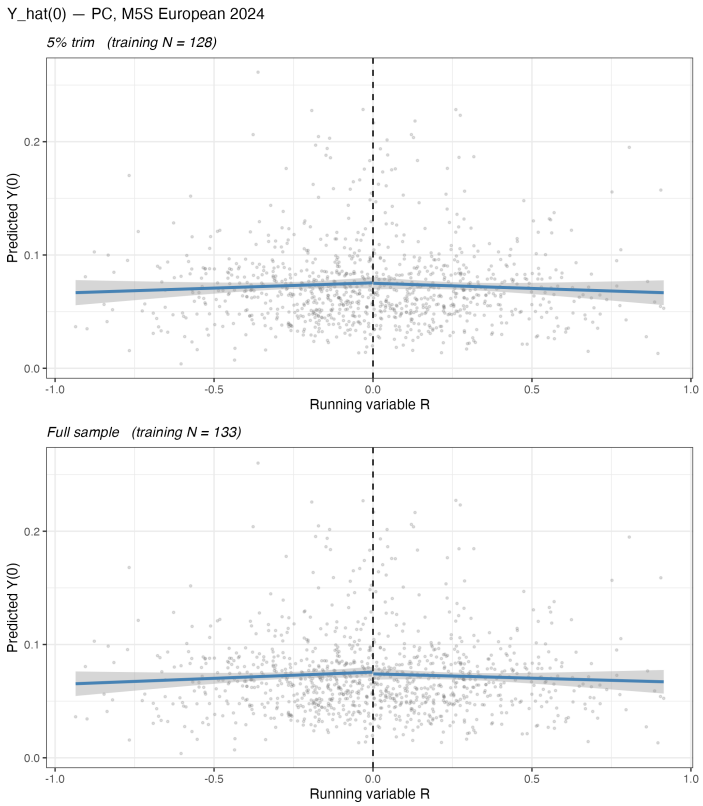

Figure 26: Predicted potential outcome $\hat{Y}^{PC}(0)$ against M_i^{PC} (M5S European 2024 outcome). Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

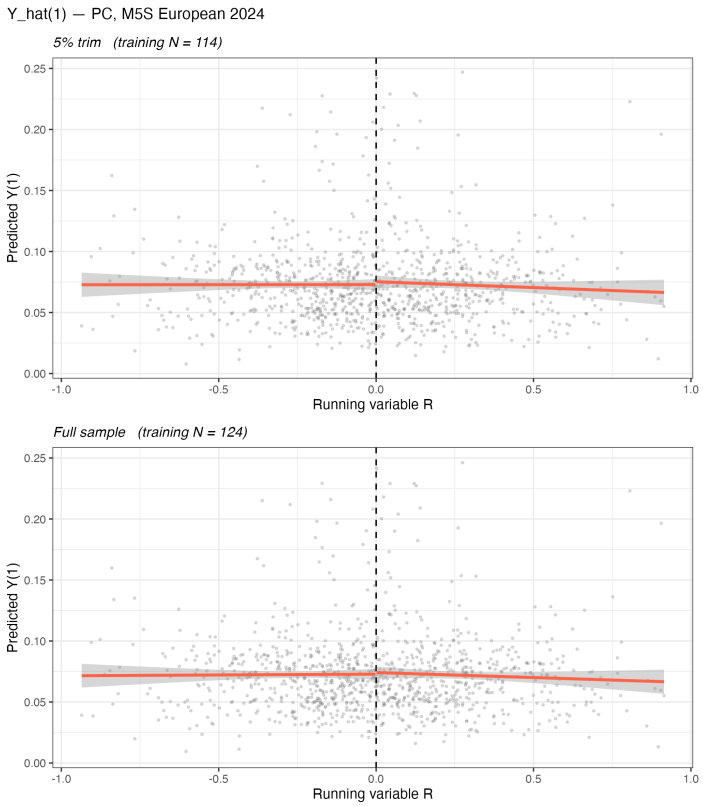

Figure 27: Predicted potential outcome $\hat{Y}^{PC}(1)$ against M_i^{PC} (M5S European 2024 outcome). Top panel: 5%-trimmed sample. Bottom panel: full directional sample.

D.5 Caliper Sensitivity

The main gradient PCRD specification applies a 5th-percentile caliper on the within-pair characteristic gap, whereas the CT-PCRD specification uses the full directional sample. Figures 28–33 report how the gradient PCRD estimate and its 95% robust confidence interval change as the minimum within-pair characteristic gap is progressively tightened from no restriction (0th percentile, full sample) to the 90th percentile of the respective difference distributions. Covariates are region fixed effects and log municipal population throughout.

Ideology design ($\mathcal{S}^R$) — vote-weighted ideology 2024

The estimate is negative across the entire caliper range, with the largest magnitudes at the restrictive p_{70} and p_{80} ($\hat{\tau} = -0.099$ at both). It is significant at the 5% level at the lower percentiles p_0, p_{10}, p_{20}, and p_{30} (where N is largest) and again at p_{70} ($p = 0.020$, $N = 244$) and p_{80} ($p = 0.028$, $N = 163$); the intermediate p_{40} is marginal ($p = 0.086$), while p_{50}, p_{60}, and the most restrictive p_{85} ($N = 122$) are not significant as the effective sample shrinks.

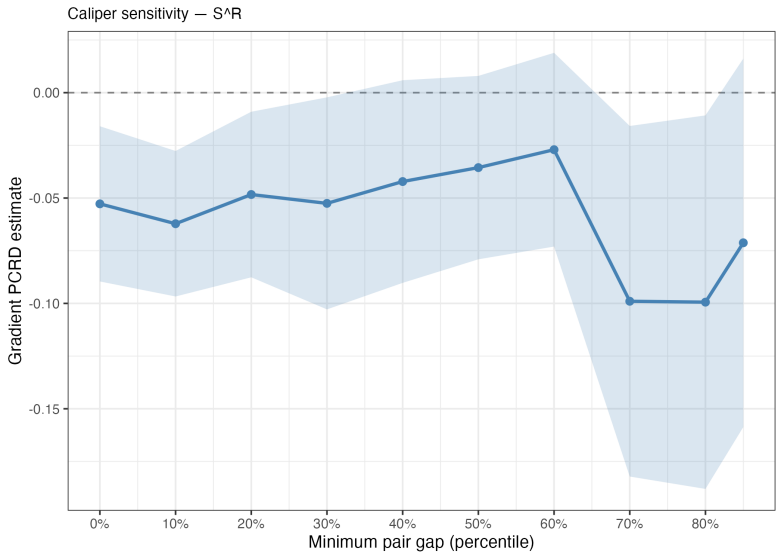

Figure 28: Caliper sensitivity for the gradient PCRD in $\mathcal{S}^R$. Each point is the robust bias-corrected RD estimate when the sample is restricted to municipality pairs whose ideology gap exceeds the indicated percentile of the ideology-difference distribution (0% = full sample); shaded band is the 95% robust confidence interval.

Ideology design ($\mathcal{S}^R$) — Lega European 2024

The estimate is negative across the entire caliper range, significant at the 5% level at p_0, p_{20}, p_{30}, p_{80}, and p_{85} (marginally at p_{10} and p_{40}); the middle range (p_{50}–p_{70}) is attenuated and not statistically significant. The largest magnitude appears at p_{80} ($\hat{\tau} = -0.093$, $p < 0.001$, $N = 163$), where the most ideologically polarized pairs are isolated (Figure 29).

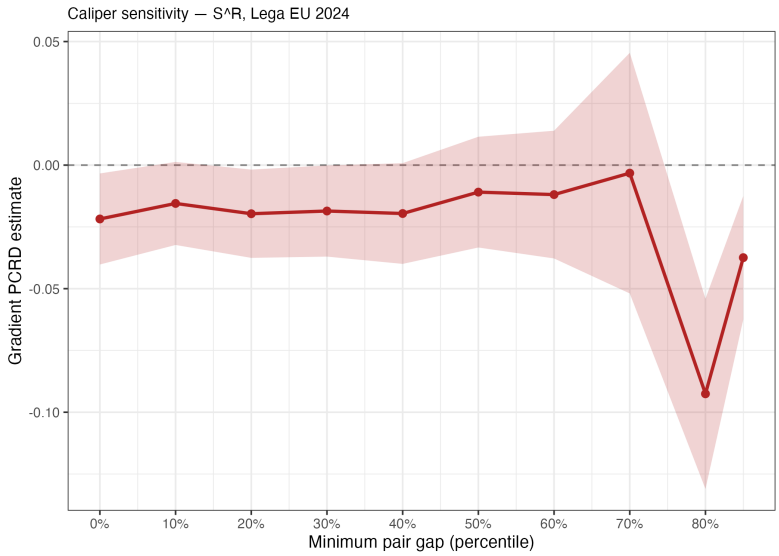

Figure 29: Caliper sensitivity for the gradient PCRD in $\mathcal{S}^R$, outcome: Lega vote share 2024 European election. Each point is the robust bias-corrected RD estimate when the sample is restricted to municipality pairs whose ideology gap exceeds the indicated percentile of the ideology-difference distribution (0% = full sample); shaded band is the 95% robust confidence interval.

Populism composite design — M5S Camera 2022 and European 2024

For the M5S Camera 2022 outcome, the estimate is positive across the entire caliper range, significant at the 5% level at p_0 through p_{40} and at p_{60}, marginal at p_{50}, and non-significant at p_{70}–p_{85} as the effective sample shrinks (Figure 30). For the M5S European 2024 outcome the pattern is similar: positive throughout, significant at p_0–p_{30}, p_{60}, and p_{85}, marginal at p_{50} and p_{70}, with a single non-significant cell at p_{40} (Figure 31).

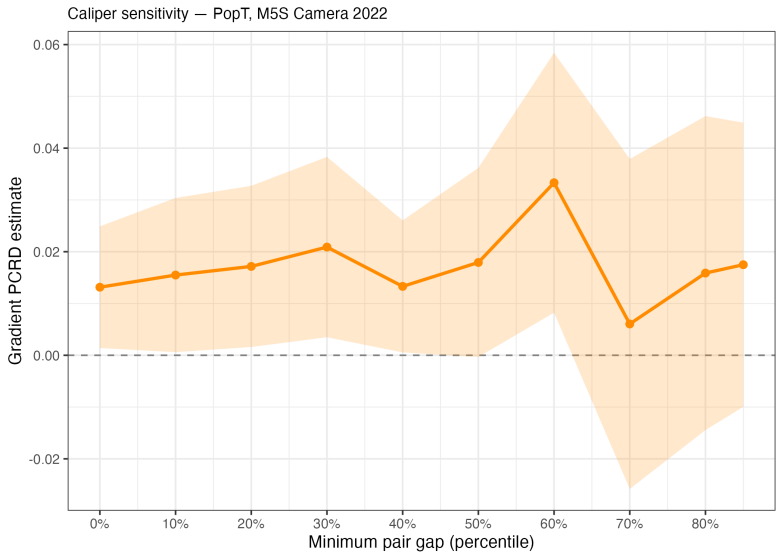

Figure 30: Caliper sensitivity for the gradient PCRD in the populism composite design, outcome: M5S Camera 2022 vote share.

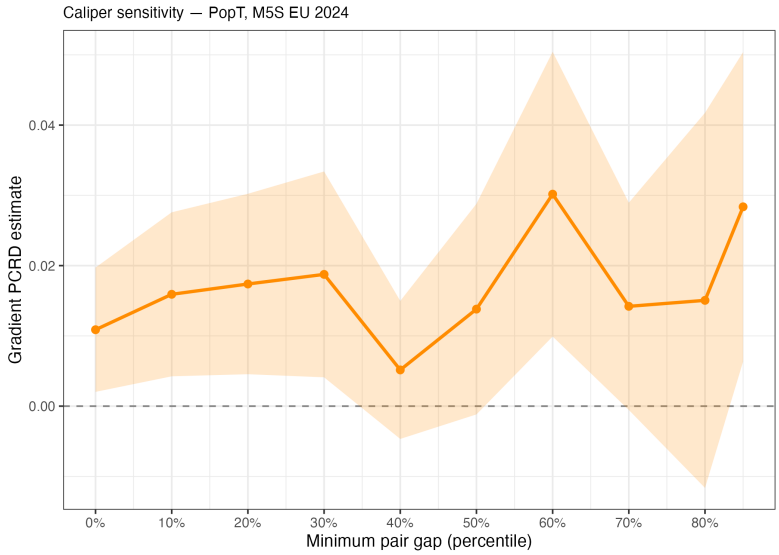

Figure 31: Caliper sensitivity for the gradient PCRD in the populism composite design, outcome: M5S European 2024 vote share.

People-Centrism design — M5S Camera 2022

The estimate is positive and significant at the 5% level at most caliper percentiles (p_0, p_{10}, p_{20}, p_{50}, p_{60}, p_{70}, p_{80}, p_{85}), weakening only at p_{30}–p_{40} (Figure 32).

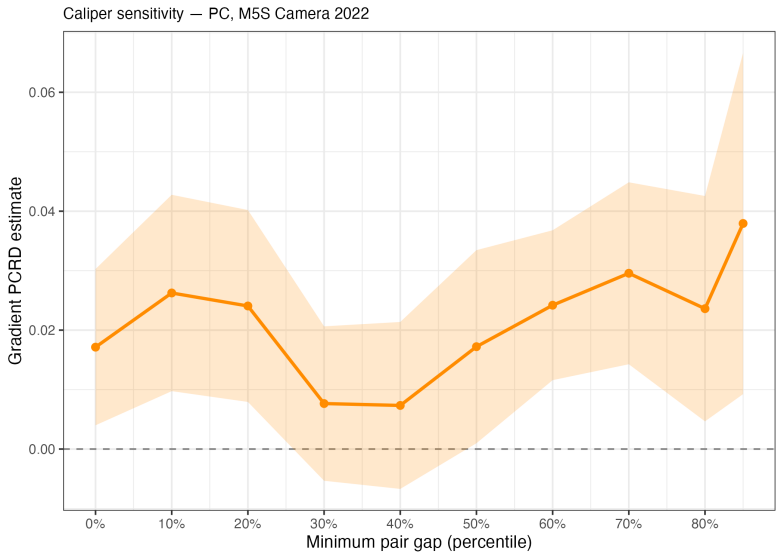

Figure 32: Caliper sensitivity for the gradient PCRD in the PC design, outcome: M5S vote share 2022 parliamentary election. Each point is the robust bias-corrected RD estimate when the sample is restricted to municipality pairs whose PC gap exceeds the indicated percentile of the PC-difference distribution (0% = full sample); shaded band is the 95% robust confidence interval.

People-Centrism design — M5S European 2024

The estimate is positive and significant at the 5% level at intermediate caliper percentiles (p_{10}, p_{20}, p_{50}, p_{60}, p_{70}, p_{80}), but is not significant at the unrestricted p_0 caliper ($p = 0.107$) or at p_{30}–p_{40}, mirroring the fragility of this estimate documented elsewhere in this appendix (Figure 33).

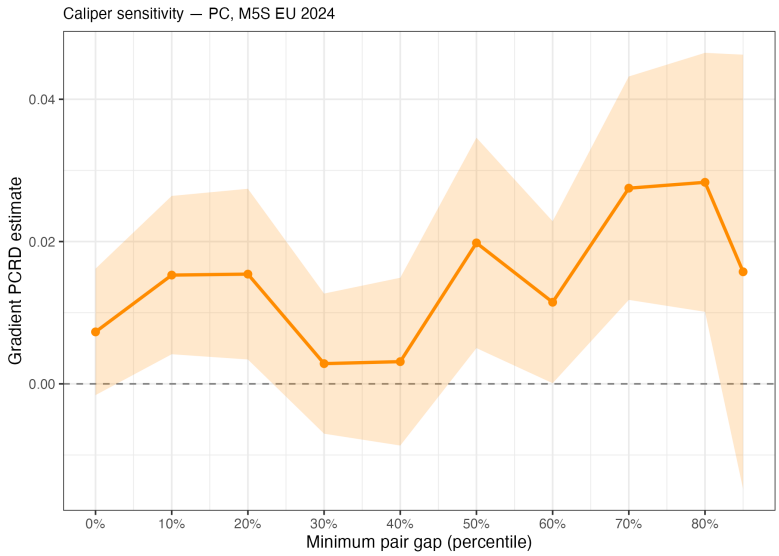

Figure 33: Caliper sensitivity for the gradient PCRD in the PC design, outcome: M5S vote share 2024 European election. Each point is the robust bias-corrected RD estimate when the sample is restricted to municipality pairs whose PC gap exceeds the indicated percentile of the PC-difference distribution (0% = full sample); shaded band is the 95% robust confidence interval.

D.6 Bandwidth Sensitivity

All designs are re-estimated at 0.50×, 0.75×, 1.00×, 1.25×, 1.50×, and 2.00× the MSE-optimal bandwidth, scaling the bias-correction bandwidth b by the same factor. Covariates are region fixed effects and log municipal population throughout.

Ideology design ($\mathcal{S}^R$) — vote-weighted ideology 2024 ($h = 0.164$)

The estimate is negative and significant at the 5% level from ×0.75 through ×2.00; at ×0.50 the estimate is marginally significant ($p = 0.069$) (Figure 34).

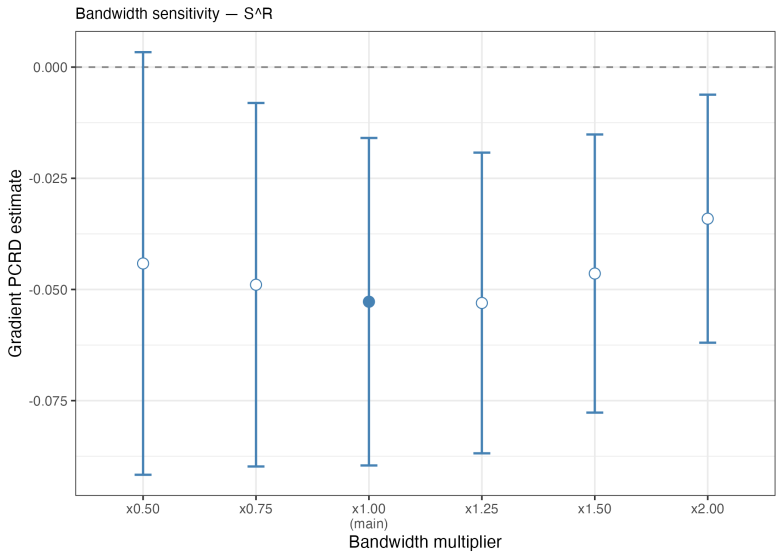

Figure 34: Bandwidth sensitivity for the gradient PCRD in $\mathcal{S}^R$. Each point is the robust bias-corrected RD estimate at the indicated bandwidth multiplier (relative to the MSE-optimal $h = 0.164$); bars are 95% robust confidence intervals. Filled circle marks the main specification ($\times 1.00$). Covariates: region fixed effects and log population.

Ideology design ($\mathcal{S}^R$) — Lega European 2024 ($h = 0.130$)

The estimate is negative and significant at the 5% level from $\times 0.50$ through $\times 1.25$; it weakens at wider bandwidths ($p = 0.067$ at $\times 1.50$, $p = 0.113$ at $\times 2.00$), consistent with a local effect that is diluted as observations farther from the cutoff are included (Figure 35).

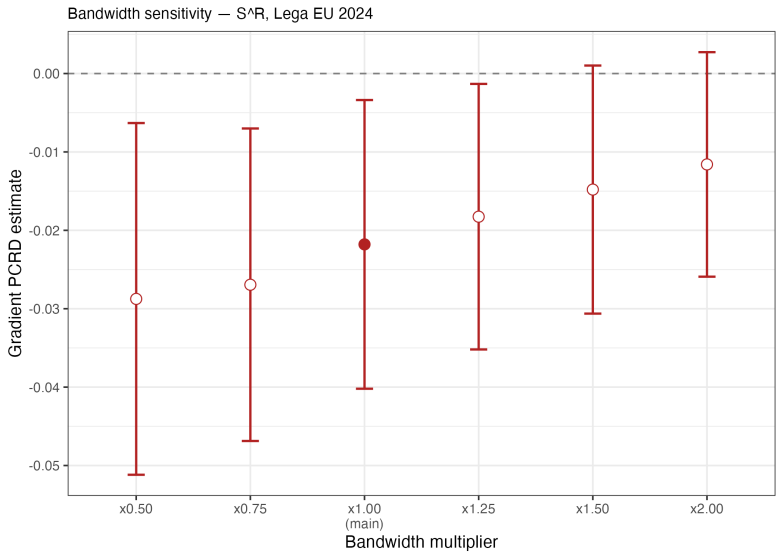

Figure 35: Bandwidth sensitivity for the gradient PCRD in $\mathcal{S}^R$, outcome: Lega vote share 2024 European election. Each point is the robust bias-corrected RD estimate at the indicated bandwidth multiplier (relative to the MSE-optimal $h = 0.130$); bars are 95% robust confidence intervals. Filled circle marks the main specification ($\times 1.00$). Covariates: region fixed effects and log population.

Populism composite design — M5S Camera 2022 and European 2024

For the M5S Camera 2022 outcome ($h = 0.201$), the estimate is positive at all bandwidths, significant at $\times 0.75$ and $\times 1.00$ (marginal at $\times 0.50$ and $\times 1.25$); it loses significance at $\times 1.50$ and $\times 2.00$ ($p = 0.132$ and $p = 0.143$), consistent with a local effect diluted as wider bandwidths admit observations farther from the cutoff (Figure 36). For the M5S European 2024 outcome ($h = 0.184$) the pattern is similar: significant at $\times 0.50$ through $\times 1.00$ (all $p \leq 0.016$), marginal at $\times 1.25$, and not significant at $\times 1.50$ or $\times 2.00$ (Figure 37).

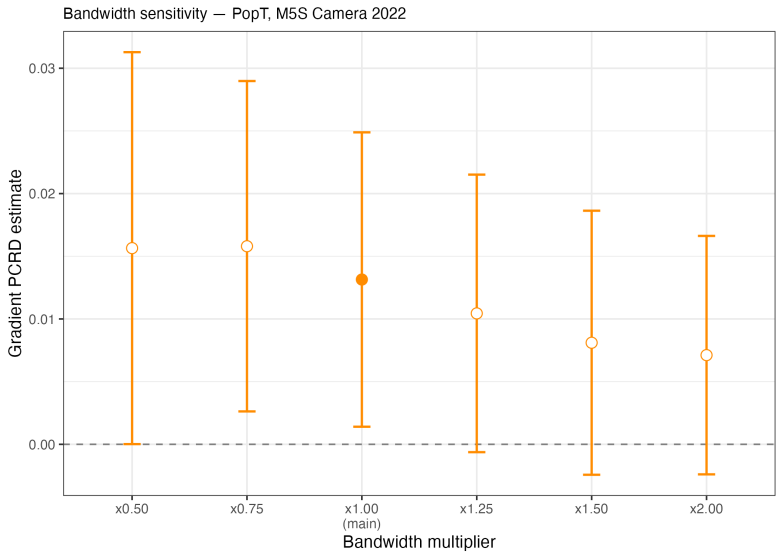

Figure 36: Bandwidth sensitivity for the gradient PCRD in the populism composite design, outcome: M5S Camera 2022 vote share. Covariates: region fixed effects and log population.

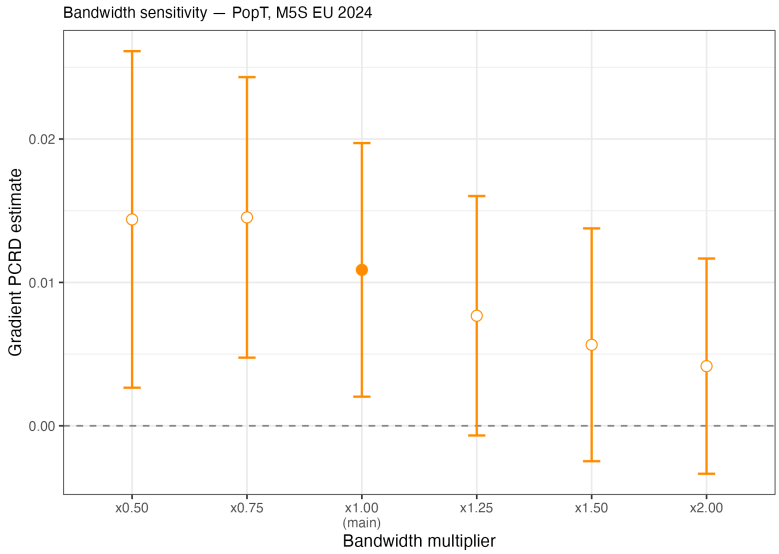

Figure 37: Bandwidth sensitivity for the gradient PCRD in the populism composite design, outcome: M5S European 2024 vote share. Covariates: region fixed effects and log population.

People-Centrism design — M5S Camera 2022 ($h = 0.129$)

The estimate is positive and significant at the 5% level from $\times 0.50$ through $\times 1.50$; at the widest bandwidth $\times 2.00$ the estimate is positive but imprecise ($p = 0.127$), consistent with a local effect diluted as observations farther from the cutoff are included (Figure 38).

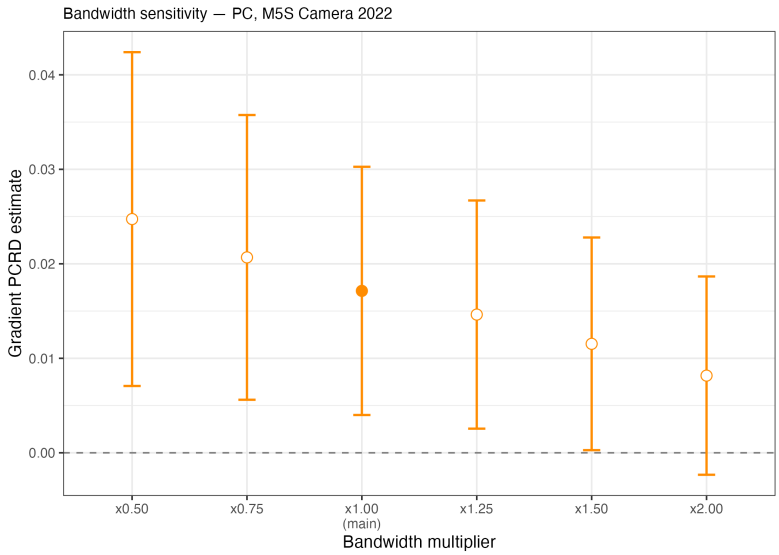

Figure 38: Bandwidth sensitivity for the gradient PCRD in the PC design, outcome: M5S vote share 2022 parliamentary election. Each point is the robust bias-corrected RD estimate at the indicated bandwidth multiplier (relative to the MSE-optimal $h = 0.129$); bars are 95% robust confidence intervals. Filled circle marks the main specification ($\times 1.00$). Covariates: region fixed effects and log population.

People-Centrism design — M5S European 2024 ($h = 0.131$)

The estimate is positive but fragile to bandwidth choice: it is marginally significant at the narrow bandwidths $\times 0.50$ ($p = 0.053$) and $\times 0.75$ ($p = 0.050$), but is not significant at the MSE-optimal bandwidth ($p = 0.107$ at $\times 1.00$) or at wider bandwidths ($p = 0.270$ at $\times 1.25$, $p = 0.485$ at $\times 1.50$, $p = 0.608$ at $\times 2.00$); the result mirrors the fragility of this estimate documented elsewhere in this appendix (Figure 39).

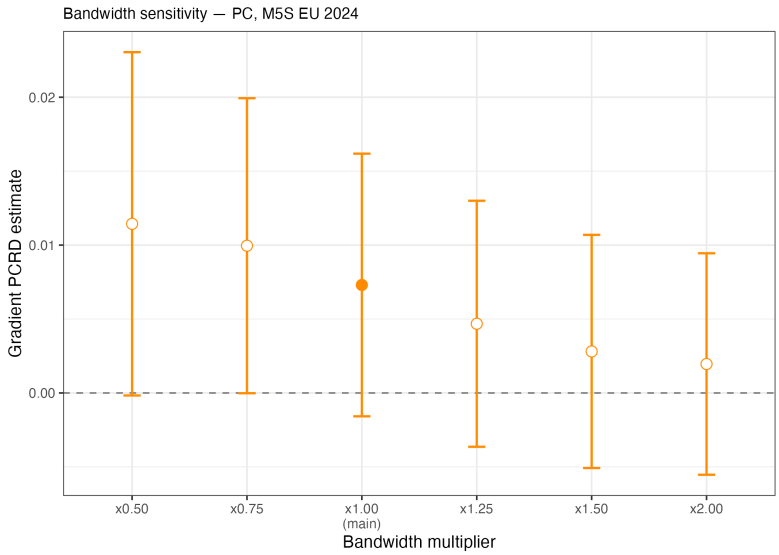

Figure 39: Bandwidth sensitivity for the gradient PCRD in the PC design, outcome: M5S vote share 2024 European election. Each point is the robust bias-corrected RD estimate at the indicated bandwidth multiplier (relative to the MSE-optimal $h = 0.131$); bars are 95% robust confidence intervals. Filled circle marks the main specification ($\times 1.00$). Covariates: region fixed effects and log population.

E Candidate Characteristics and Incumbency Controls

The electoral and fiscal results rest on the close-election design. Two checks in this section address its validity and the interpretation of the treatment label. A related concern, *compensating differentials* (Marshall, 2024), is that close-race candidates may systematically differ on other dimensions. Following Bertoli and Hazlett (2025), this does not threaten the internal validity of the local PCRD estimate, though it limits extrapolation beyond the close-race regime.

The first concern is identification. Right-wing candidates who barely win may systematically differ on pre-determined traits (age, education, prior experience) from right-wing candidates who barely lose, and the same for populist candidates. If so, the close-race contrast would conflate the type effect with these candidate-level differences. Bertoli and Hazlett (2025) note that the close-election design implies candidate-level balance on such traits as a falsification test. Let Type A denote the more right-wing (or more populist) of the two leading candidates in a race, and Type B the other. Table 25 reports balance tests separately for the Type-A and Type-B candidate of each pair, on eight pre-determined traits. No significant imbalance is found in the right-wing ($\mathcal{S}^R$) or people-centrism (PC) designs. The populism (PopT) design shows one isolated significant trait (Type-A age), within the range expected by chance under multiple testing across 16 trait-by-side comparisons. The anti-elite-rhetoric (AER) design shows three significant imbalances. It is therefore reported

as a descriptive comparison rather than as an independent identification.

The second concern is the bundle. Bertoli and Hazlett (2025) make explicit that PCRDs identify the effect of districts electing a bundled candidate type, not the pure attribute effect, and that bundling, correctly characterized, is not a violation of identification but part of the estimand itself (see also Marshall, 2024). The bundle-comparison tests reported below describe what bundle each side of the cutoff selects. Under $\mathcal{S}^R$, right-wing bare winners are significantly less likely to be the prior incumbent, defined as the holder of the 2014–2019 mayoral mandate. Their coalitions are not broader than those of narrow losers but tilt more toward national parties and away from local civic groups. This rules out the alternative that they win marginally by assembling unusually broad or heterogeneous coalitions. The national-party tilt is unsurprising: being further to the right is more strongly associated with national-party affiliation. The substantive concern is incumbency: voters could be reacting to the loss of an incumbent rather than to the election of a more right-wing candidate. Because incumbency is the bundled trait most likely to confound the right-wing effect, the headline electoral and fiscal regressions are augmented with both-side incumbency indicators (Inc_A, Inc_B). These indicators denote whether the Type-A and Type-B candidates, respectively, were the prior incumbent of the municipality. The headline coefficients are essentially unchanged across both electoral estimators and all sixteen fiscal cells (Tables 27 and 28). The loss-of-incumbency channel is therefore not the source of the headline result.

"Bare-winning" and "bare-losing" refer to candidates whose pair landed just to one or the other side of the cutoff $M_i = 0$ — so close that local randomization is supposed to hold. If, on a pre-determined trait (age, prior office, education, ...), bare-winning Type-A candidates differ systematically from bare-losing Type-A candidates, then the bare-win indicator is correlated with that trait at the cutoff: the discontinuity in 2024 outcomes could in principle reflect the trait rather than the treatment, and the local-randomization assumption that licenses the RD is in tension with the data. A clean candidate balance test (no discontinuity in any pre-determined trait) is therefore a direct check on identification, parallel to the district-level balance test of Section D.2 but at the candidate level.

Eight candidate traits are considered: age, female, an indicator for being born in the municipality of candidacy, a binary indicator for having held the immediate 2014–2019 mayoral mandate of this municipality (*prior incumbent*), an education-level ordinal index taking seven equally spaced values (0 no schooling, 1 primary, 2 middle school, 3 high school, 4 bachelor, 5 master, 6 specialization / PhD), a managerial / professional occupation indicator, the share of pure civic groups among the political *entities* (parties and civic groups) backing the candidate, and the total count of those entities. In Italian municipal elections, voters

choose a candidate together with one or more named ballot lists. Each list is decomposed into its constituent entities — a list whose name combines several parties (for instance a joint Lega–Forza Italia–Fratelli d'Italia list) counts as one entity per named party. Each entity is then classified as either a national party (Lega, Forza Italia, Fratelli d'Italia, Partito Democratico, M5S, etc.) or a pure civic group (no national-party affiliation). In comuni below the 15,000-inhabitant threshold a candidate is typically backed by a single list, while above the threshold candidates are supported by coalitions of several lists; the entity-level decomposition makes the measure comparable across both regimes. All eight come from the 2019 mayoral candidate registry merged on surname, name, date of birth and *comune*. Both tests use `rdrobust` with the design's running variable M_i (linear local polynomial, triangular kernel) on the 5%-trimmed sample.

Candidate balance (identification check)

For each design and each trait, two regressions are run — the trait of the Type-A candidate and the trait of the Type-B candidate are each used as the outcome with M_i as the running variable. A non-zero discontinuity at the cutoff means bare-winning and bare-losing candidates of the same type differ systematically on that trait, which would threaten identification.

Ideology design ($\mathcal{S}^R$): no trait is significant at the 5% or 10% level on either side. With sixteen tests, this is the strongest possible balance result.

Populism ternary design: one Type-A trait flags — age ($\hat{\tau} = +7.55$, $p = 0.007$). All other traits, on both sides, are clean.

Anti-Elite Rhetoric design: three traits flag at 5% (Type-A born-locally $p = 0.008$, Type-A civic share $p = 0.050$, Type-B managerial $p = 0.015$). This is the most heavily flagged design, consistent with the district-level imbalance of Section D.2: AER is reported as descriptive comparison rather than as an independent identification.

People-Centrism design: no trait is significant at 5%; three traits are marginal at 10% (Type-A born-locally $p = 0.052$, Type-A civic share $p = 0.054$, Type-B civic share $p = 0.077$). Together with the clean district-level balance, the PC design carries no candidate-level red flag.

Table 25: Candidate balance.

	Design			
Trait	$\mathcal{S}^R$	Pop. ternary	AER	PC
Panel A: Type-A trait, 5%-trimmed sample				
age	+1.483	+7.553***	+2.304	+2.961
female	-0.068	+0.075	-0.077	+0.104
born locally	+0.117	+0.125	+0.238***	+0.198*
prior incumbent	-0.098	-0.056	-0.069	-0.102
education level (0-6)	-0.408	-0.105	-0.417*	-0.288
prof managerial	-0.065	-0.181	+0.004	-0.146
civic share	+0.039	+0.047	+0.156*	+0.125*
n entities	+0.328	+0.192	+0.447	-0.138
Panel B: Type-A trait, full sample				
age	+0.622	+5.357**	+2.242	+1.571
female	-0.017	+0.088	-0.080	+0.075
born locally	+0.109	+0.084	+0.267***	+0.163*
prior incumbent	-0.096	-0.051	-0.078	-0.097
education level (0-6)	-0.318	-0.140	-0.321	-0.279
prof managerial	-0.076	-0.167	+0.072	-0.129
civic share	+0.056	+0.054	+0.163**	+0.057
n entities	+0.359	+0.142	+0.447	+0.031
Panel C: Type-B trait, 5%-trimmed sample				
age	-1.684	+1.264	-0.506	+0.716
female	-0.066	+0.047	-0.109	+0.042
born locally	-0.103	-0.126	-0.063	-0.077
prior incumbent	-0.033	-0.038	-0.013	-0.068
education level (0-6)	-0.263	+0.163	+0.053	+0.029
prof managerial	+0.098	+0.182	+0.264**	+0.165
civic share	-0.006	-0.084	-0.003	-0.115*
n entities	-0.340	+0.129	+0.064	-0.369
Panel D: Type-B trait, full sample				
age	-1.183	+1.830	+0.466	+0.209
female	-0.057	+0.063	-0.094	+0.061
born locally	-0.112	-0.155	-0.046	-0.084
prior incumbent	-0.008	-0.022	-0.011	-0.022
education level (0-6)	-0.253	+0.155	+0.051	+0.069
prof managerial	+0.082	+0.186	+0.296**	+0.204*
civic share	-0.008	-0.082	+0.000	-0.092
n entities	-0.304	-0.211	+0.041	-0.314

Note: Each cell is the `rdrobust` estimate of the discontinuity in the trait at the cutoff (linear local polynomial, triangular kernel). Panels A–B: trait of the Type-A (more right-wing / populist) candidate. Panels C–D: trait of the Type-B candidate. The 5%-trimmed sample drops pairs whose absolute ideological distance falls in the lowest 5% (the main specification); the full sample uses all close races. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Bundle comparison (descriptive)

Each candidate trait is regressed on M_i with `rdrobust` using the *winner*'s value as the outcome. Because the winner is Type A when $M_i > 0$ and Type B when $M_i < 0$, the discontinuity at $M_i = 0$ gives the difference in trait between Type-A bare winners and Type-B bare winners (Table 26).

Ideology design ($\mathcal{S}^R$): two traits differ substantially at the cutoff. Bare right-wing winners are about 29 percentage points less likely to be the prior incumbent of the municipality ($\hat{\tau} = -0.292$, $p = 0.001$) and come from coalitions whose civic-list share is 21 percentage points lower ($p = 0.003$). Female representation is marginally lower ($p = 0.067$).

Populism ternary design: no trait differs significantly at the cutoff; the prior-incumbent indicator is the closest ($\hat{\tau} = -0.106$, $p = 0.215$).

Anti-Elite Rhetoric design: female representation is 23 percentage points lower among AER bare winners ($p = 0.006$) and the prior-incumbent indicator is marginally lower ($\hat{\tau} = -0.137$, $p = 0.093$); other traits are clean.

People-Centrism design: no trait differs at the cutoff. PC and non-PC bare winners have indistinguishable trait bundles.

Table 26: Bundle comparison.

	Design			
Trait	$\mathcal{S}^R$	Pop. ternary	AER	PC
Panel A: 5%-trimmed sample				
age	-1.561	-0.246	+2.196	+0.251
female	-0.180*	+0.027	-0.233***	+0.154*
born locally	+0.035	-0.031	+0.138	+0.016
prior incumbent	-0.292***	-0.106	-0.137*	-0.050
education level (0-6)	-0.320	-0.073	-0.326	-0.049
prof managerial	+0.032	-0.022	-0.000	+0.047
civic share	-0.211***	-0.077	+0.043	+0.038
n entities	+0.513	+0.150	+0.385	-0.503
Panel B: full sample				
age	-1.357	+0.387	+2.519	-0.736
female	-0.144	+0.047	-0.227***	+0.138
born locally	-0.010	-0.054	+0.134	-0.004
prior incumbent	-0.167**	-0.113	-0.162**	-0.056
education level (0-6)	-0.345	-0.017	-0.276	-0.078
prof managerial	+0.007	-0.013	+0.003	+0.044
civic share	-0.215***	-0.054	+0.032	-0.009
n entities	+0.532	-0.206	+0.396	-0.367

Note: Each cell is the `rdrobust` estimate of the discontinuity in the winner's trait at the cutoff (linear local polynomial, triangular kernel). A positive $\hat{\tau}$ means that bare-winning Type-A (more right-wing / populist) candidates score higher than bare-winning Type-B candidates on the trait. The 5%-trimmed sample drops pairs whose absolute ideological distance falls in the lowest 5% (the main specification); the full sample uses all close races. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Robustness: ideology effect with incumbency controls

The bundle pattern in the $\mathcal{S}^R$ design — right-wing bare winners are about 29 percentage points less likely to be the prior incumbent of the municipality — raises a substantive question about the headline result. If voters punish "losing the incumbent" rather than "electing a right-wing candidate," then the negative effect of $\mathcal{S}^R$ on subsequent electoral and fiscal outcomes could be partially or wholly attributable to a new-mayor channel rather than to ideology.

To address this, the headline regressions are augmented with incumbency indicators on both

sides of the candidate pair:

$$\text{Incumbent}_{A,i} = \mathbf{1}\{\text{Type-A candidate was the prior elected mayor of municipality } i\},$$
$$\text{Incumbent}_{B,i} = \mathbf{1}\{\text{Type-B candidate was the prior elected mayor of municipality } i\}.$$

Both variables are pre-treatment candidate-pair characteristics, fixed before the 2019 election. They are also balanced at the cutoff under all four designs (preceding subsection). The augmented specification is

$$Y_i = \alpha + \tau D_i + f(M_i) + \delta_A \,\text{Incumbent}_{A,i} + \delta_B \,\text{Incumbent}_{B,i} + X_i'\beta + \varepsilon_i,$$

where $D_i = 1$ if Type A wins, M_i is the signed margin, $f(M_i)$ is the RD adjustment, X_i is the headline covariate set (region fixed effects and log municipal population for the electoral specifications; the 2018 baseline of the outcome for the fiscal specifications), and $\text{Incumbent}_{A,i}$, $\text{Incumbent}_{B,i}$ control for the pre-election incumbency structure of the candidate pair.

Electoral outcomes. For both vote-weighted ideology 2024 and Lega European 2024, the augmented specification leaves the headline coefficient essentially unchanged (Table 27). The vote-weighted ideology effect moves by less than 2% on either estimator ($-0.058 \to -0.059$, $p = 0.003$ on the Gradient; $-0.271 \to -0.272$, $p = 0.010$ on the CT-PCRD). The Lega European effect is unchanged on the Gradient estimator ($-0.020 \to -0.020$, $p = 0.036$) and grows in magnitude on the CT-PCRD ($-0.085 \to -0.127$, $p = 0.024$). Adding both-side incumbency indicators therefore does not absorb the right-wing electoral effect.

Table 27: Robustness: $\mathcal{S}^R$ headline with both-side incumbency controls.

	VW Ideology 2024		Lega European 2024	
	Gradient	CT-PCRD	Gradient	CT-PCRD
Headline	-0.058***	-0.271**	-0.020**	-0.085*
	(0.002)	(0.013)	(0.028)	(0.067)
+ Inc_A + Inc_B	-0.059***	-0.272**	-0.020**	-0.127**
	(0.003)	(0.010)	(0.036)	(0.024)

Note: Robust bias-corrected RD estimates; robust p-values in parentheses below the point estimate. The Headline row uses the main-text $\mathcal{S}^R$ specification (region fixed effects, log municipal population). The augmented row adds the two pre-treatment indicators $\text{Inc}_{A,i}$ and $\text{Inc}_{B,i}$ for whether the Type-A and Type-B candidate, respectively, was the prior incumbent of municipality i (holder of the 2014–2019 mayoral mandate). Gradient PCRD is estimated on the 5%-trimmed sample; CT-PCRD uses the full sample. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Fiscal outcomes. For the means-tested welfare cuts under the $\mathcal{S}^R$ design, the augmented specification preserves — and in several cells modestly strengthens — the headline coefficient

(Table 28). Across sixteen outcome–year combinations (W_c and W_c^{++}, share and per-capita, four years), most cells move by less than 5%; the largest movements are in the late-mandate windows (W_c per-capita 24M3 grows by ∼11%, W_c^{++} share 24M3 by ∼8%), all in the direction of larger retrenchment. The cell-by-cell significance pattern is preserved with two minor reinforcements (W_c per-capita 24M6 gains $p < 0.10$ and W_c^{++} share 2023 sharpens to $p < 0.01$).

Table 28: Robustness: $\mathcal{S}^R$ fiscal headline with both-side incumbency controls.

	2022	2023	24M3	24M6
Wc — Narrow means-tested welfare				
share (×100) (headline)	-0.56^{*}	-1.08^{**}	-1.23^{**}	-0.86^{*}
+ Inc_A + Inc_B	-0.58^{*}	-1.09^{**}	-1.35^{**}	-0.89^{*}
per capita (headline)	−7.56	-8.10^{**}	−7.06	−8.88
+ Inc_A + Inc_B	−7.48	-7.99^{**}	−7.87	-9.42^{*}
Wc++ — Broad means-tested welfare				
share (×100) (headline)	-0.63^{*}	-1.21^{***}	-1.35^{**}	-1.00^{**}
+ Inc_A + Inc_B	-0.66^{*}	-1.21^{***}	-1.46^{**}	-1.01^{**}
per capita (headline)	-8.30^{*}	-8.91^{**}	−7.77	-10.38^{*}
+ Inc_A + Inc_B	-8.14^{*}	-8.85^{**}	−8.48	-10.66^{*}

Note: Each cell is the `rdrobust` estimate of the $\mathcal{S}^R$ effect on the corresponding welfare outcome. Within each layer (Wc and Wc++), the rows alternate between the headline specification (region fixed effects, log municipal population, 2018 baseline of the outcome) and the same specification augmented with the two pre-treatment incumbency indicators $\text{Inc}_{A,i}$ and $\text{Inc}_{B,i}$. Share outcomes are scaled by 100 (percentage-point units). The right-wing share-form 2023 cells inherit the EUR-per-capita bandwidth, matching the convention used in Table 42. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Verdict. The bundle of traits that comes with a bare right-wing victory does include a less-incumbent candidate, but controlling directly for both-side incumbency in the headline specification leaves the right-wing coefficient essentially unchanged across both electoral estimators and all sixteen fiscal cells. The loss-of-incumbency channel is therefore not the source of the headline result; voters are responding to the ideological extremism of the elected mayor, not to the displacement of the prior incumbent.

F Vote-Weighted Municipal Electoral Results: Full Estimates

This appendix reports the vote-weighted populism specification referenced in Section 6. The evidence is inconclusive: the estimates are not corroborated across measures or estimators.[26]

[26]The gradient PCRD on the composite populism index is marginally negative (−0.029, $p = 0.078$), but the CT-PCRD on the same index is insignificant (−0.199, $p = 0.290$, with a strong first stage of 0.105, $p < 0.001$);

The informative populism signal appears instead in the national-election spillover results (Table 35).

Table 29: Effect of Barely Electing a More Populist Mayor on the Vote-Weighted Populism in 2024.

Outcome: vote-weighted populism (composite) in 2024		
	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	-0.029^{*}	-0.199
SE	(0.017)	(0.188)
p-value	0.078	0.290
First stage	—	0.105^{***}
SE		(0.025)
p-value		0.000
N (bandwidth)	453	430
N (total)	1020	1074
Bandwidth	0.183	0.168

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. N (bandwidth): observations within the MSE-optimal bandwidth on each side of the cutoff. N (total): full sample. Gradient PCRD applies a first-stage intensity filter that drops races with |diff| below the 5th percentile, following Hall (2015); CT-PCRD uses the full sample, since the Wald-ratio structure downweights low-dose races automatically. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G Alternative Municipal Outcome: Simple Average Scores

Tables 30–31 replicate the main ideology and populism specifications using the simple (unweighted) average of candidate-level scores as the outcome, rather than the vote-weighted average. For ideology, the self-undermining pattern is confirmed: the right-sample gradient PCRD estimate is -0.077 ($p = 0.002$) and the CT-PCRD estimate is -0.354 ($p = 0.005$), both significant and directionally consistent with the main specification. The left sample remains insignificant. For populism, the gradient PCRD is marginally significant (-0.032, $p = 0.041$) but the CT-PCRD remains imprecise (-0.238, $p = 0.192$), consistent with the fragile marginal-negative pattern documented in the main specification.

the anti-elite-rhetoric sub-dimension is null under both estimators; and the people-centrism sub-dimension is marginally *positive* under the gradient PCRD ($+0.025$, $p = 0.060$) but insignificant under CT-PCRD ($+0.186$, $p = 0.222$; AER and PC tables in Appendix D).

Table 30: Robustness: Simple Average Ideology in 2024.

Outcome: simple average ideology in 2024

	Right sample ($\mathcal{S}^R$)		Left sample ($\mathcal{S}^L$)	
	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	-0.077^{***}	-0.354^{***}	0.045^{*}	0.239
SE	(0.025)	(0.128)	(0.024)	(0.152)
p-value	0.002	0.005	0.067	0.116
First stage	—	0.203^{***}	—	0.157^{***}
SE		(0.031)		(0.023)
p-value		0.000		0.000
N (bandwidth)	277	385	340	372
N (total)	773	814	779	820
Bandwidth	0.149	0.197	0.182	0.188

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. Outcome: simple (unweighted) average of the candidate-level ideology score across all 2024 candidates. Gradient PCRD applies a first-stage intensity filter that drops races with |diff| below the 5th percentile; CT-PCRD uses the full sample. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 31: Robustness: Simple Average Populism in 2024.

Outcome: simple average populism (composite) in 2024

	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	-0.032^{**}	-0.238
SE	(0.016)	(0.182)
p-value	0.041	0.192
First stage	—	0.105^{***}
SE		(0.025)
p-value		0.000
N (bandwidth)	409	410
N (total)	1020	1074
Bandwidth	0.167	0.161

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. Outcome: simple (unweighted) average of the candidate-level populism (composite) score. Gradient PCRD applies a first-stage intensity filter that drops races with |diff| below the 5th percentile; CT-PCRD uses the full sample. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Tables 32 and 33 decompose the vote-weighted populism outcome into its two sub-dimensions. Anti-elite rhetoric (AER) shows no significant effect under either estimator. People-centrism (PC) is marginally significant under the gradient PCRD (0.025, $p = 0.060$) but not under the CT-PCRD (0.186, $p = 0.222$), broadly consistent with the national-election results reported in the main text.

Table 32: Alternative Outcome: Vote-Weighted Anti-Elite Rhetoric (AER).

Outcome: vote-weighted AER in 2024

	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	−0.004	0.008
SE	(0.006)	(0.143)
p-value	0.526	0.953
First stage	—	0.042***
SE		(0.007)
p-value		0.000
N (bandwidth)	537	496
N (total)	1020	1074
Bandwidth	0.229	0.190

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. Outcome: vote-weighted anti-elite rhetoric (AER) in 2024. Both the running variable and the CT-PCRD treatment are AER-specific. Gradient PCRD applies a first-stage intensity filter that drops races with |diff| below the 5th percentile; CT-PCRD uses the full sample. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 33: Alternative Outcome: Vote-Weighted People-Centrism (PC).

Outcome: vote-weighted PC in 2024		
	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	0.025*	0.186
SE	(0.013)	(0.152)
p-value	0.060	0.222
First stage	—	0.073***
SE		(0.015)
p-value		0.000
N (bandwidth)	505	374
N (total)	1020	1074
Bandwidth	0.211	0.142

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. Outcome: vote-weighted people-centrism (PC) in 2024. Both the running variable and the CT-PCRD treatment are PC-specific. Gradient PCRD applies a first-stage intensity filter that drops races with |diff| below the 5th percentile; CT-PCRD uses the full sample. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

H Electoral Support Across the Distribution

The threshold sweep fixes a score cutoff at a percentile of the 2019 distribution and measures the share of 2024 votes won by candidates above it. The p_{70} threshold, for instance, is the score below which 70% of 2019 candidates fall. Sweeping the cutoff across percentiles traces where in the distribution any effect operates.

Figure 40 reports the gradient PCRD estimates at four evenly spaced thresholds (p_{50}, p_{60}, p_{70}, p_{80}) for ideology, populism, and the two component measures (AER and PC). For ideology, the self-undermining pattern extends across the distribution: estimates are consistently negative, significant at the 5% level at p_{50} ($p = 0.001$), p_{60} ($p = 0.010$), and p_{70} ($p = 0.032$), and marginal at p_{80} ($p = 0.069$). For populism on the composite index, point estimates oscillate around zero and do not reach significance at any threshold; people-centrism and anti-elite rhetoric are also insignificant at every threshold.

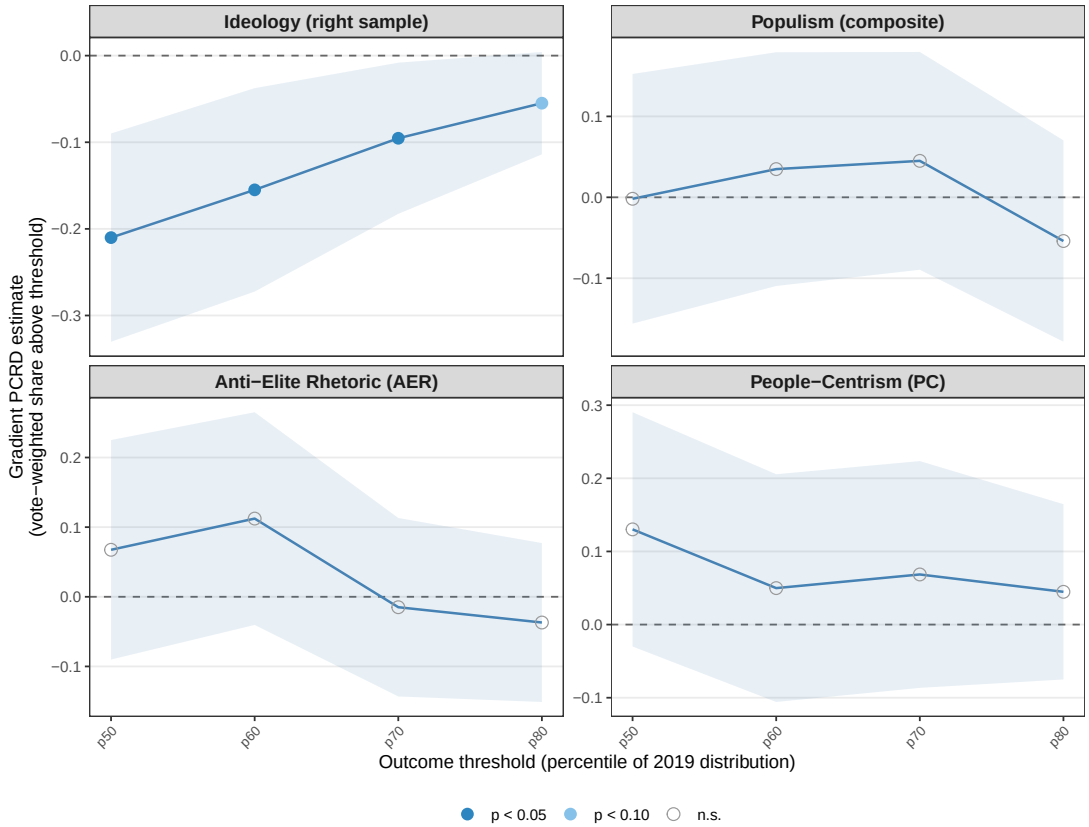

Figure 40: Threshold Sweep: Gradient PCRD. Each point is a gradient PCRD estimate of the effect of barely electing a more extreme or more populist mayor in 2019 on the share of votes going to 2024 candidates above the indicated percentile of the 2019 distribution. Top left: ideology (right sample); top right: populism (composite index); bottom left: anti-elite rhetoric (AER); bottom right: people-centrism (PC). Shaded bands are 95% robust confidence intervals. Filled dark circles indicate estimates significant at the 5% level; filled light circles at the 10% level; hollow circles are insignificant.

The CT-PCRD (2SLS) version shows a similar pattern. Figure 41 reports it for all four measures at the same four thresholds. The ideology sweep (top left) confirms the self-undermining pattern: estimates are significant at the 5% level at p_{50} (-0.942, $p = 0.008$) and p_{60} (-0.718, $p = 0.019$), and marginally significant at p_{70} (-0.378, $p = 0.075$). The populism composite sweep (top right) is positive at p_{50}–p_{70} but imprecise throughout. AER (bottom left) is marginally significant at p_{60} (3.256, $p = 0.067$) but null at the other thresholds. PC (bottom right) is null at all four thresholds.

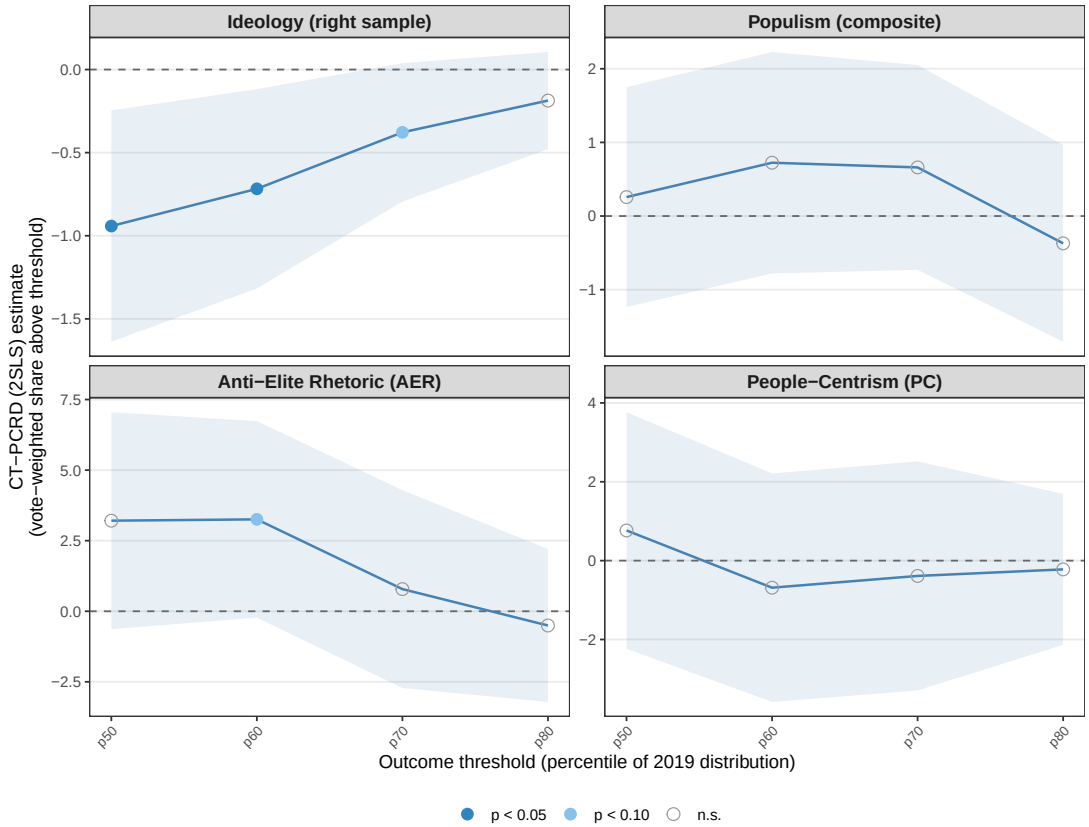

Figure 41: Threshold Sweep: CT-PCRD (2SLS). Each point is a CT-PCRD (fuzzy RD / 2SLS) estimate of the effect of a one-unit increase in the 2019 winner's score on the vote-weighted share of 2024 candidates above the indicated percentile of the 2019 distribution. Top left: ideology ($\mathcal{S}^R$); top right: populism (composite); bottom left: anti-elite rhetoric (AER); bottom right: people-centrism (PC). Shaded bands are 95% robust confidence intervals.

I National and European Elections

This appendix reports the national and European spillover results summarized in Section 6. Two national elections fell within the 2019–2024 mayoral term: the September 2022 election for the *Camera dei Deputati*, Italy's lower chamber, and the June 2024 European Parliament election, held concurrently with the 2024 municipal vote. These elections offer a cleaner test of spillover because voters choose among nationally defined party lists rather than locally assembled coalitions. The outcomes are pre-specified party vote shares corresponding to the political type analyzed in each design. For the ideology design, the relevant outcomes are Lega and Fratelli d'Italia (FdI), the two nationally competitive parties that occupy the extreme-right pole of Italian party competition in this period. More extreme organizations such as CasaPound and Forza Nuova do not provide comparable national-election outcomes because they do not consistently contest these elections as independent lists. For the populism

design, the outcome is Movimento 5 Stelle (M5S), the clearest nationally competitive populist party in the Italian system. Restricting the analysis to these pre-specified party outcomes avoids searching across all party vote shares.

Table 34 examines whether barely electing a more right-wing mayor affects right-wing parties' vote shares. No effects appear in the 2022 parliamentary election. In the 2024 European election, the Lega vote share falls significantly (gradient -0.020, $p = 0.029$; CT-PCRD -0.085, $p = 0.067$), consistent with the self-undermining pattern at the municipal level. Fratelli d'Italia is insignificant, with inconsistent sign across the two elections.

Table 34: Effect of Right-Wing Ideology Treatment ($\mathcal{S}^R$) on Right-Wing Party Vote Shares in National Elections.

	Gradient PCRD	CT-PCRD
Panel A: 2022 parliamentary election		
Lega	-0.000	-0.006
SE	(0.008)	(0.040)
p-value	0.970	0.884
FdI	-0.003	-0.032
SE	(0.014)	(0.060)
p-value	0.831	0.595
Panel B: 2024 European		
Lega	-0.020^{**}	-0.085^{*}
SE	(0.009)	(0.046)
p-value	0.029	0.067
FdI	0.017	0.057
SE	(0.013)	(0.062)
p-value	0.201	0.354

Note: Outcome: party vote share at the municipality level. Treatment and running variable from the $\mathcal{S}^R$ ideology design. Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. Gradient PCRD applies a first-stage intensity filter that drops races with |diff| below the 5th percentile, following Hall (2015); CT-PCRD uses the full sample. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 35 reports the effect of the populism treatment on M5S vote share. The sign is consistently positive across all three measures (populism, AER, and PC) in both elections. The overall populism measure is significant in both elections, though the 2022 parliamentary gradient is only marginal. People-centrism (PC) is significant in 2022 and marginal in 2024. Anti-elite rhetoric (AER) is directionally positive but imprecise. The CT-PCRD variant

shows a similar pattern.

Table 35: Effect of Populism Treatment on M5S Vote Share in National Elections.

	Gradient PCRD	CT-PCRD
Panel A: 2022 parliamentary election		
Populism	0.014*	0.148**
SE	(0.008)	(0.066)
p-value	0.071	0.025
AER	0.009	0.154
SE	(0.007)	(0.172)
p-value	0.198	0.370
PC	0.014**	0.197**
SE	(0.006)	(0.078)
p-value	0.022	0.012
Panel B: 2024 European		
Populism	0.013**	0.135***
SE	(0.006)	(0.048)
p-value	0.034	0.005
AER	0.006	0.144
SE	(0.005)	(0.128)
p-value	0.199	0.259
PC	0.008*	0.092
SE	(0.005)	(0.060)
p-value	0.088	0.122

Note: Outcome: Movimento 5 Stelle vote share at the municipality level. Each row uses its own running variable and treatment. Robust bias-corrected RD estimates; robust standard errors in parentheses; robust *p*-values below. Gradient PCRD applies a first-stage intensity filter that drops races with |diff| below the 5th percentile, following Hall (2015); CT-PCRD uses the full sample. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

J Bundled Populist–Extreme-Right Index

This appendix reports the full estimates for the bundled-index check described in Section 6, and states the construction of the composite index used throughout it.

Construction. For each candidate j, let I_j denote the signed ideology score and P_j the ternary populism score. Each trait is standardized over the pool of the two leading candidates across

the 2019 races,

$$\tilde{I}_j = \frac{I_j - \bar{I}}{s_I}, \qquad \tilde{P}_j = \frac{P_j - \bar{P}}{s_P},$$

where $(\bar{I}, s_I)$ and $(\bar{P}, s_P)$ are the mean and standard deviation of ideology and populism over that pool. The composite index is the equal-weight sum

$$C_j = \tilde{I}_j + \tilde{P}_j,$$

so that extreme-right ideology and populism each contribute one standard deviation per unit.

The index is the within-paper analogue of the populist radical right composite that party-family classifications implicitly assign (Mudde, 2007), and it behaves as one: among candidates in the top quartile of the index, 72.5 percent are above the candidate-pool median on both traits, with a mean ideology of +0.15 and mean populism of 0.26. The aggregation is candid about what it mixes: a high score identifies candidates high on both traits or high enough on either, and the sum cannot distinguish the conjunction from compensation between the two; 14.6 percent of the top quartile is left-leaning. The index defines a single ranking over all candidates, so the design is run on the full close-election sample: in each race the running variable is signed by the higher-scoring candidate on the index ($M_i = \mathit{margin}_i \cdot \text{sign}(C_{A,i} - C_{B,i})$ for the two leading candidates A and B), with no restriction by trait. The gradient PCRD compares races where the higher-C candidate barely wins with those where she barely loses, and the CT-PCRD uses the elected mayor's C as the continuous (fuzzy) treatment. The outcome in Table 36 is the vote-weighted composite of the 2024 race; the fiscal check applies the same full-sample treatment to welfare outcomes.

Table 36: Bundled Populist–Extreme-Right Index: Effect on the Vote-Weighted 2024 Outcome.

Outcome: vote-weighted populist–extreme-right index in 2024 (full close-election sample)

	Gradient PCRD (Sharp RD)	CT-PCRD (2SLS)
Estimate	−0.140	−0.064
SE	(0.174)	(0.112)
p-value	0.421	0.570
First stage	—	1.733***
SE		(0.267)
p-value		0.000
N (total)	1,074	1,074

Note: Robust bias-corrected RD estimates; robust standard errors in parentheses; robust p-values below. The composite index standardizes the candidate ideology and ternary populism scores (each to mean zero and unit standard deviation) and sums them with equal weight; the outcome is its vote-weighted average across the 2024 candidates, and the treatment is barely electing the higher-scoring candidate on the index, run on the full close-election sample. The separated extreme-right ideology design, run on its coherent right-side sample, produces a strongly significant leftward shift (−0.058, $p = 0.002$; Table 2); the bundled index registers nothing. Because the composite is built from standardized scores, its coefficients sit on a unit-variance outcome scale and are not directly comparable in magnitude to the separated estimate. Gradient PCRD applies the 5th-percentile |diff| filter following Hall (2015); CT-PCRD uses the full sample. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The bundled treatment is real: at the cutoff, barely electing the higher-index candidate raises the elected mayor's ideology by 0.141 ($p < 0.001$, 0.90 candidate-pool standard deviations) and populism by 0.085 ($p < 0.001$, 0.73 standard deviations), 69 and 81 percent of the dose the corresponding separated designs deliver on their own samples. The design therefore administers both traits at once. On the aggregate electoral outcome it recovers nothing (Table 36): the separated designs show that the ideological dose moves the subsequent race leftward while the populist dose moves it nowhere, and a bundled outcome blends the two responses into an estimate that registers neither.

The governing-behavior margin is no different (Table 37). The separated extreme-right design recovers a clear welfare retrenchment on its right-side sample, the Core welfare share falling by roughly one percentage point across the mid-mandate windows. Under the bundled treatment no retrenchment is detected: every post-treatment estimate is small and insignificant, with a null 2018 placebo. The index itself cannot say why; whether the populist component of the dose offsets the ideological one, or the contrast simply averages over races in which neither candidate is extreme-right, is indistinguishable within a single index. The separated designs

resolve it: the cuts follow the ideological dose, not the populist one.

Table 37: Bundled Index on the Fiscal Margin: Means-Tested Welfare, Gradient PCRD.

	Share (pp)					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Panel A: Core means-tested welfare (W_c)										
Extreme-right ($\mathcal{S}^R$)	−0.94	−0.56*	−1.08**	−1.23**	−0.86*	−11.38	−7.56	−8.10**	−7.06	−8.88
	(0.81)	(0.34)	(0.46)	(0.59)	(0.46)	(9.56)	(4.66)	(3.71)	(5.43)	(5.60)
Bundled (full sample)	−0.68	−0.22	−0.24	−0.42	−0.42	−25.19	−1.42	−1.29	−5.64	−4.26
	(0.77)	(0.32)	(0.31)	(0.45)	(0.37)	(21.41)	(6.00)	(6.77)	(5.98)	(7.04)
Panel B: Full means-tested welfare (W_c^{++})										
Extreme-right ($\mathcal{S}^R$)	−0.95	−0.63*	−1.21***	−1.35**	−1.00**	−11.62	−8.30*	−8.91**	−7.77	−10.38*
	(0.82)	(0.35)	(0.46)	(0.61)	(0.49)	(9.75)	(4.78)	(4.07)	(5.73)	(5.88)
Bundled (full sample)	−0.67	−0.35	−0.26	−0.49	−0.53	−25.09	−2.36	−2.03	−4.90	−4.39
	(0.77)	(0.35)	(0.30)	(0.47)	(0.40)	(21.39)	(5.91)	(6.72)	(6.26)	(7.20)

Note: Gradient PCRD estimates of the effect on means-tested welfare expenditure, in budget-share (percentage points) and EUR-per-capita form; robust bias-corrected standard errors in parentheses. The *extreme-right* row runs the ideology treatment on its coherent right-side ($\mathcal{S}^R$) sample; the *bundled* row runs the composite treatment on the full close-election sample, which is the only sample available to it (Construction). The extreme-right design recovers a clear retrenchment across the mid-mandate windows; the bundled index recovers nothing. Whether the populist component of the bundled treatment offsets the ideological one, or the contrast averages over races in which neither candidate is extreme-right, is indistinguishable within a single index. The 2018 placebo is null for both. The bundled per-capita placebo is large but imprecise; the share form is the primary specification. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Taken together, the bundled index delivers an impoverished picture: no aggregate electoral effect and no fiscal effect, despite administering a strong joint dose of both traits. This is not a defect of estimation but of aggregation. The bundled treatment moves both traits at once, so its estimate blends the ideological response with the populist one and can attribute the outcome to neither. The separated designs supply what the index cannot: the ideological dose drives the backlash and the retrenchment, while the populist dose drives neither. It is the reason the two traits are measured separately throughout the paper.

K Party-Label Benchmark

The paper's main results rely on text-based candidate scores. A natural benchmark is whether the extreme-right findings survive under the more traditional approach used in much of the literature: party labels. They do, but the exercise also illustrates the limitations that motivate the paper's measurement strategy.

In the Italian municipal context, party labels are usable only for the extreme-right side of the analysis. They impose a severe sample restriction. The text-based extreme-right design, $\mathcal{S}^R$, uses 814 close races, while the party-label version is limited to the 363 races in which exactly one of the two leading 2019 candidates ran under a Lega or Fratelli d'Italia label. The populist side cannot be replicated in this way. The Movimento 5 Stelle almost never reaches the final round of a municipal race, leaving no close-election contrast to exploit. This absence is itself an instance of the measurement problem discussed in Section 4: party labels make some theoretically relevant traits visible only when nationally organized parties happen to compete in the relevant races.

The extreme-right side nonetheless permits a useful benchmark. Lega and Fratelli d'Italia are the two main national extreme-right parties in this setting, so a binary party-label treatment is constructed, equal to one when exactly one of the two leading 2019 candidates ran under a Lega or FdI label. The two parties are pooled because FdI alone appears too rarely in close municipal races to estimate separately.[27] The design is otherwise identical to the main specification, and neither the treatment nor the outcomes use the classifier. The exercise recovers the same substantive pattern as the text-based analysis, though on a much smaller sample.

Table 38: Party-Label Benchmark: Effect of Barely Electing a Lega/FdI Mayor on Electoral Outcomes.

	2024 municipal Lega/FdI branding	Lega+FdI vote (Camera 2022)	Lega+FdI vote (European 2024)
Estimate	−0.128	−0.072***	−0.348
SE	(0.094)	(0.023)	(0.285)
p-value	0.170	0.001	0.221
N (bandwidth)	92	91	88
N (total)	363	363	363

Note: Sharp close-election RD estimates (robust bias-corrected); robust standard errors in parentheses, robust p-values below. Treatment: exactly one of the two leading 2019 candidates ran under a Lega or FdI label, identified in $N = 363$ close races (the full sample). The 2024 municipal outcome is the share of mayoral candidates running under a Lega/FdI label; the national outcomes are the combined Lega+FdI vote share. N (bandwidth) is the effective number of observations within the MSE-optimal bandwidth. Covariates: region fixed effects and log municipal population. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

[27] Adding the two minor extreme-right parties, CasaPound and Forza Nuova, to the label leaves the picture essentially unchanged. Because they almost never field municipal mayoral candidates, they add only three close races to the sample; the municipal branding outcome is similar, and the welfare retrenchment is materially unchanged.

Table 38 reports the electoral results. The first outcome is the share of 2024 mayoral candidates running under a Lega or FdI label. Barely electing a Lega/FdI mayor in 2019 lowers this outcome by 12.8 percentage points, although the estimate is imprecise ($p = 0.170$). The direction matches the text-based municipal backlash reported in the main results. This outcome should be interpreted cautiously, however, because a decline in Lega/FdI-labelled candidacies can mean either that voters shifted away from Lega/FdI candidates or that candidates who might otherwise have run chose not to carry the party brand. The outcome therefore conflates party presence with electoral support. The text-based measure avoids this ambiguity by scoring the content of each platform regardless of whether the candidate adopts a national label.

The national-election outcomes point in the same direction. Barely electing a Lega/FdI mayor reduces the combined Lega and FdI vote share in the 2022 parliamentary election by 7.2 percentage points ($p = 0.001$). The 2024 European Parliament estimate is also negative, though imprecise. Because these outcomes are measured in national and European elections, they are not affected by the composition of the subsequent municipal race. Their direction is consistent with the text-based spillover results reported in Appendix I.

The loss of precision in some specifications should not be interpreted only as a statistical inconvenience. Part of it is mechanical: the party-label design uses fewer than half as many close races as the text-based extreme-right design. But the label-based treatment may also be substantively noisier, since it pools candidates who share a national party tag but differ in the intensity and content of their platform positions. The benchmark cannot separate these two sources of imprecision. What it shows is more limited: even under this coarser and more restrictive measurement strategy, the estimated effects point in the same direction as the text-based results.

Table 39: Party-label benchmark: effect of barely electing a Lega/FdI mayor on means-tested welfare expenditure.

	Share %				EUR per capita			
	2018 plac.	2022	2023	24M6	2018 plac.	2022	2023	24M6
Core (W_c)	+0.12	−1.15	−1.28**	−1.67*	−0.8	−7.0	−7.4*	−3.5
	(1.48)	(1.04)	(0.60)	(0.98)	(15.9)	(5.6)	(4.5)	(7.3)
Full (W_c^{++})	+0.10	−2.74**	−1.55**	−0.54	−1.5	−6.3	−8.9*	−5.0
	(1.52)	(1.10)	(0.66)	(1.36)	(16.2)	(5.9)	(4.7)	(8.0)

Note: Sharp close-election RD estimates (robust bias-corrected). Treatment: exactly one of the two leading 2019 candidates ran under a Lega or FdI label, on the 363 close races with such a contrast. Outcome: means-tested welfare expenditure in budget-share form (% of total municipal expenditure, ×100) and EUR per capita (2018 population). The 2018 column is a pre-treatment placebo. Robust standard errors in parentheses; effective within-bandwidth N ranges from 55 to 195 across cells. Covariates: region fixed effects, log municipal population, and the 2018 baseline of the same outcome. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 39 reports the fiscal benchmark. Barely electing a Lega/FdI mayor reduces means-tested welfare spending, in line with the retrenchment pattern recovered by the main text-based strategy. The estimates are negative across both the Core (W_c) and Full (W_c^{++}) definitions and across both budget-share and per-capita forms, with several estimates reaching conventional significance levels. The 2018 placebo estimates are close to zero across all rows. Effective sample sizes are small and vary across cells, reflecting the restricted Lega/FdI close-race sample, but the fiscal results are substantively coherent with the paper's main mechanism: extreme-right incumbents reduce welfare provision even when the treatment is defined using conventional party labels.

The benchmark therefore supports two conclusions. First, where party labels are usable, they reproduce the paper's extreme-right pattern: barely electing a Lega/FdI mayor is followed by lower subsequent support for the extreme right and by welfare retrenchment consistent with the main mechanism. Second, the exercise confirms the limits of party-label measurement. Party labels apply only to the minority of candidates who run under a national tag, the municipal branding outcome conflates electoral support with candidate entry and party presentation, and the populist side cannot be studied this way at all. The party-label benchmark validates the text-based results where labels are available and confirms the necessity of text-based measurement where they are not.

L Revenue-Side Fiscal Results

Tables 40 and 41 report the gradient PCRD and CT-PCRD (fuzzy RD, 2SLS) estimates of each treatment's effect on the tax burden borne by residents (own-source municipal taxes,

SIOPE classification 1.01.01).

Table 40: Effect of treatment on tax burden on residents (Gradient PCRD).

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Right-wing	+1.84	−0.15	−0.81	+5.58	+2.89	−10.7	−17.4	+0.5	+7.3	+17.7
	(3.46)	(2.06)	(2.05)	(3.83)	(2.51)	(78.2)	(38.0)	(35.2)	(11.0)	(19.5)
Populism	+1.97	+2.64	−0.33	+2.89	+1.21	−25.0	−46.2	−93.0	+8.1	+14.4
	(2.81)	(1.84)	(1.99)	(2.32)	(1.72)	(93.1)	(45.5)	(66.3)	(7.8)	(14.9)
AER	+1.74	−0.05	+0.27	+1.57	+0.53	+23.1	−0.6	−6.7	+11.2	+7.4
	(3.03)	(1.98)	(2.06)	(2.49)	(1.79)	(97.3)	(36.1)	(53.9)	(9.4)	(16.1)
PC	+1.96	+2.25	+0.26	−1.31	+1.52	−14.9	+37.6	+3.5	+0.7	+22.1
	(2.64)	(2.09)	(2.21)	(1.78)	(1.69)	(90.2)	(47.2)	(60.7)	(7.4)	(15.2)

Note: Each cell reports the bias-corrected gradient PCRD estimate; robust standard errors in parentheses. Outcome: total of own-source municipal taxes paid directly by residents (SIOPE classification 1.01.01), in budget-share form (% of total municipal revenue) and EUR per capita (2018 population). Sample: $n = 1065$ close-race municipalities (excluding the three smallest regions for numerical stability). Covariates: region fixed effects, log municipal population, and the 2018 baseline of the same outcome. The 2018 placebo column uses the 2018 baseline as the outcome with no baseline control. The 24M3 and 24M6 windows use the cumulative monthly cash flows through March and June 2024, isolating the 2019 mayor's tenure. No estimate reaches significance at the 10% level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 41: Effect of treatment on tax burden on residents (CT-PCRD).

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Right-wing	+3.41	+3.09	+2.65	+18.25	+8.65	−334.9	−8.9	+115.3	+30.0	+65.3
	(17.73)	(10.29)	(9.99)	(16.09)	(11.26)	(565.7)	(172.2)	(161.9)	(50.6)	(91.7)
Populism	+28.95	+23.37	+0.64	+21.70	+8.21	−316.7	−315.4	−644.3	+92.8	+103.3
	(27.82)	(16.36)	(15.59)	(23.06)	(16.16)	(873.2)	(414.5)	(565.1)	(79.9)	(131.0)
AER	+26.68	+1.98	+10.73	+37.77	+6.93	+1269.4	−285.4	−894.2	+178.0	+93.2
	(76.12)	(45.88)	(47.86)	(63.18)	(43.97)	(2725.6)	(1018.9)	(1431.4)	(218.0)	(374.3)
PC	+39.19	+38.80	+8.98	+6.66	+31.66	−129.0	+679.8	+42.1	+43.9	+395.0*
	(38.98)	(28.51)	(30.62)	(39.57)	(26.67)	(1222.9)	(525.3)	(832.3)	(118.0)	(230.5)

Note: CT-PCRD (fuzzy RD, 2SLS) version of Table 40. The treatment instrument is the cutoff dummy and the endogenous treatment intensity is the winning candidate's design-specific score; CT-PCRD coefficients are therefore on the per-unit-treatment-score scale. Sample: $n = 1065$ close-race municipalities (excluding the three smallest regions for numerical stability). Covariates: region fixed effects, log municipal population, and the 2018 baseline of the same outcome. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

M Means-Tested Welfare Expenditure: Full Estimates

Tables 42–44 report the full gradient PCRD estimates underlying Figures 5–7: the full sample (Table 42), the term-limited subgroup (Table 43), and the not-term-limited subgroup (Table 44). Each cell reports the bias-corrected coefficient and robust standard error, including the 2018 placebo column omitted from the figures.

Table 42: Effect of treatment on means-tested welfare expenditure (Gradient PCRD).

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Panel A: Core welfare (W_c) — social-assistance anchor + L.328 case-management delivery										
Right-wing	−0.94	−0.56*	−1.08**	−1.23**	−0.86*	−11.4	−7.6	−8.1**	−7.1	−8.9
	(0.81)	(0.34)	(0.46)	(0.59)	(0.46)	(9.6)	(4.7)	(3.7)	(5.4)	(5.6)
Populism	−0.16	−0.07	+0.39	+0.14	+0.56	−19.7	+1.8	+5.6	+2.4	+7.2
	(0.79)	(0.34)	(0.33)	(0.52)	(0.39)	(25.1)	(6.2)	(7.4)	(6.9)	(7.5)
AER	−0.66	+0.47	+0.18	−0.11	+0.32	−31.1	+7.5	+3.3	+0.2	+4.9
	(0.70)	(0.35)	(0.36)	(0.52)	(0.45)	(22.5)	(6.2)	(6.6)	(7.4)	(8.5)
PC	−0.16	−0.57	−0.20	−0.26	+0.86	−8.8	−3.3	+4.4	+0.2	+7.8
	(0.75)	(0.35)	(0.52)	(0.77)	(0.59)	(25.5)	(5.5)	(6.9)	(7.4)	(8.4)
Panel B: Full welfare (W_c^{++}) — Core plus statutorily anchored items										
Right-wing	−0.95	−0.63*	−1.21***	−1.35**	−1.00**	−11.6	−8.3*	−8.9**	−7.8	−10.4*
	(0.82)	(0.35)	(0.46)	(0.61)	(0.49)	(9.7)	(4.8)	(4.1)	(5.7)	(5.9)
Populism	−0.11	+0.01	+0.32	+0.12	+0.57	−19.2	+2.7	+5.5	+2.1	+7.3
	(0.80)	(0.34)	(0.35)	(0.54)	(0.40)	(25.1)	(5.7)	(7.2)	(7.2)	(7.6)
AER	−0.62	+0.34	+0.11	−0.13	+0.26	−30.9	+7.2	+3.1	+1.2	+5.0
	(0.70)	(0.38)	(0.38)	(0.54)	(0.47)	(22.5)	(6.4)	(6.8)	(7.6)	(8.6)
PC	−0.19	−0.50	−0.06	+0.13	+1.17*	−9.7	−3.1	+4.4	+0.1	+8.2
	(0.77)	(0.36)	(0.44)	(0.86)	(0.67)	(25.3)	(5.5)	(6.7)	(7.4)	(8.4)

Note: Each cell reports the bias-corrected gradient PCRD estimate; robust standard errors in parentheses. Outcome: means-tested welfare expenditure, in budget-share form (% of total municipal expenditure) and EUR per capita (2018 population). Panel A (Core W_c) sums social-care residential and home-care contracts (.15.008, .15.009), L.328 case-management labor (.11.002), and *altri assegni e sussidi assistenziali* (.04.02.02.999). Panel B (Full W_c^{++}) adds disabled/elderly transport (.15.003), LSU (.12.002), study grants (.04.02.03.001), and psychiatric residential placements (.18.010). Sample: $n = 1065$ close-race municipalities (excluding the three smallest regions for numerical stability). Covariates: region fixed effects, log municipal population, and the 2018 baseline of the same outcome. The right-wing share-form 2023 cells (W_c and W_c^{++}) inherit the EUR-per-capita MSE-optimal bandwidth to eliminate a column-name sensitivity in rdrobust's bandwidth selector specific to those cells; all other cells use their own MSE-optimal bandwidth. The 24M3 and 24M6 windows use the cumulative monthly cash flows through March and June 2024, isolating the 2019 mayor's tenure. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 43: Effect of treatment on means-tested welfare expenditure — Term-limited mayors (Gradient PCRD).

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Panel A: Core welfare (W_c) — social-assistance anchor										
Right-wing	−4.78	+0.10	−1.38	−1.81	−2.87	−58.2	−12.4	−23.5**	−28.5**	−41.1
	(3.77)	(1.15)	(1.02)	(1.52)	(2.69)	(50.0)	(12.7)	(11.3)	(13.2)	(25.7)
Populism	−3.10	−0.29	+0.02	+0.59	+2.43	−27.4	+4.6	+3.2	+1.3	+16.7
	(3.04)	(1.20)	(1.05)	(1.69)	(2.07)	(32.1)	(17.5)	(15.2)	(15.5)	(20.5)
AER	−3.43	+0.14	+0.70	+0.93	+2.16	−41.2	+13.0	+7.7	−1.8	+14.9
	(2.46)	(1.14)	(0.97)	(1.72)	(2.53)	(30.4)	(18.7)	(13.9)	(16.8)	(27.2)
PC	−3.31	+0.27	+0.62	+1.80	+3.56	−29.6	+2.0	+3.8	+2.9	+18.0
	(3.25)	(1.11)	(1.29)	(2.17)	(2.73)	(38.0)	(15.3)	(13.8)	(16.2)	(21.1)
Panel B: Full welfare (W_c^{++}) — Core plus statutorily anchored items										
Right-wing	−4.82	−0.21	−1.69*	−3.08**	−3.25	−57.3	−15.1	−27.1**	−38.4**	−52.1*
	(3.80)	(1.09)	(0.97)	(1.56)	(3.13)	(50.8)	(13.8)	(12.5)	(18.7)	(30.0)
Populism	−2.95	−0.31	−0.15	+0.19	+2.32	−25.1	−0.1	+0.5	−5.8	+14.1
	(3.13)	(1.29)	(1.17)	(1.78)	(2.11)	(32.4)	(18.4)	(16.9)	(17.8)	(21.5)
AER	−3.57	+0.23	+0.70	+0.79	+2.26	−42.7	+14.1	+5.2	−5.5	+13.9
	(2.46)	(1.37)	(1.08)	(1.86)	(2.62)	(30.4)	(22.4)	(15.9)	(20.0)	(28.6)
PC	−3.37	+0.18	−0.13	+1.81	+3.88	−27.6	−0.1	+1.4	−5.2	+14.8
	(3.36)	(1.28)	(1.20)	(2.47)	(2.96)	(37.7)	(17.4)	(15.9)	(19.4)	(22.4)

Note: Each cell reports the bias-corrected gradient PCRD estimate; robust standard errors in parentheses. Outcome: means-tested welfare expenditure, share % (% of total municipal expenditure) and EUR per capita (2018 population). Sample: term-limited mayors under the L. 35/2022 classification (5,000-inhabitant threshold against the 2021 legal population), $n = 291$ (excluding the three smallest regions). Covariates: region FE, log municipal population, and 2018 baseline of the same outcome. Bandwidth held fixed at the MSE-optimal value computed on the Full sample. The right-wing share-form 2023 cells inherit the EUR-per-capita bandwidth to eliminate a column-name sensitivity in rdrobust specific to those cells. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 44: Effect of treatment on means-tested welfare expenditure — Not-term-limited mayors (Gradient PCRD).

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Panel A: Core welfare (W_c) — social-assistance anchor										
Right-wing	−0.13	−0.20	−0.60**	−0.80	−0.63	+4.2	−1.0	−4.3	−1.9	−1.7
	(0.79)	(0.30)	(0.29)	(0.52)	(0.41)	(8.6)	(3.5)	(3.5)	(5.9)	(4.7)
Populism	+0.49	−0.21	+0.54	+0.65	+0.50	−21.4	+4.1	+10.6	+12.2*	+14.9*
	(1.01)	(0.45)	(0.34)	(0.45)	(0.39)	(47.9)	(7.4)	(8.5)	(7.1)	(8.0)
AER	+0.05	+0.73*	+0.54	+0.66	+0.78*	−34.4	+12.4*	+10.1	+11.5	+16.3*
	(0.92)	(0.39)	(0.40)	(0.49)	(0.46)	(43.1)	(6.6)	(8.1)	(8.4)	(9.8)
PC	+1.11	−0.91**	+0.46	−0.15	+0.57	−4.3	+0.2	+12.7*	+10.8	+16.9*
	(0.93)	(0.41)	(0.48)	(0.65)	(0.49)	(42.9)	(7.0)	(7.7)	(8.3)	(9.2)
Panel B: Full welfare (W_c^{++}) — Core plus statutorily anchored items										
Right-wing	−0.17	−0.25	−0.71**	−0.80	−0.68	+3.6	−1.2	−5.7	−1.5	−1.8
	(0.81)	(0.30)	(0.29)	(0.55)	(0.46)	(8.8)	(3.4)	(3.7)	(6.2)	(5.0)
Populism	+0.45	−0.14	+0.44	+0.57	+0.45	−21.7	+6.4	+9.7	+11.4	+14.2*
	(1.01)	(0.46)	(0.37)	(0.48)	(0.41)	(47.9)	(7.6)	(8.4)	(7.5)	(8.2)
AER	+0.13	+0.62	+0.41	+0.60	+0.74	−33.9	+11.3*	+9.2	+11.8	+16.6*
	(0.91)	(0.41)	(0.42)	(0.53)	(0.48)	(43.0)	(6.6)	(8.2)	(8.8)	(10.0)
PC	+1.06	−0.76*	+0.22	+0.45	+0.88*	−5.8	+1.0	+13.1*	+12.0	+18.1*
	(0.94)	(0.43)	(0.43)	(0.62)	(0.47)	(42.6)	(7.0)	(7.6)	(8.3)	(9.3)

Note: Each cell reports the bias-corrected gradient PCRD estimate; robust standard errors in parentheses. Outcome: means-tested welfare expenditure, share % (% of total municipal expenditure) and EUR per capita (2018 population). Sample: not-term-limited mayors under the L. 35/2022 classification (5,000-inhabitant threshold against the 2021 legal population), $n = 758$ (excluding the three smallest regions). Covariates: region FE, log municipal population, and 2018 baseline of the same outcome. Bandwidth held fixed at the MSE-optimal value computed on the Full sample. The right-wing share-form 2023 cells inherit the EUR-per-capita bandwidth to eliminate a column-name sensitivity in rdrobust specific to those cells. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

N Universalist Welfare Expenditure

Table 45 reports the gradient PCRD estimates on the universalist welfare measure (U_c), aggregating municipal expenditure on childcare, school meals, and school transport contracts. The 2018 placebo column is significantly positive under the Populism design (+1.26 percentage points, $p < 0.05$ for the share form; +11.60 EUR per capita, $p < 0.10$ for the EUR-per-capita form), indicating pre-treatment sorting on universalist service spending. The post-treatment PopT cells are also negative and significant in the late-mandate windows (24M3 and 24M6), but the placebo failure precludes a causal interpretation. The placebo is null under the right-wing, AER, and people-centrism designs. The right-wing post-treatment cells are

mixed in sign and statistically null, in contrast to the negative pattern documented in the means-tested measures. This confirms that right-wing welfare retrenchment is concentrated on the targeted, means-tested margin and does not extend to broadly available services. AER and PC post-treatment cells are uniformly small and (with the exception of PC's U_c share at 24M3, -1.53 percentage points, $p < 0.10$) statistically null.

Table 45: Effect of treatment on universalist welfare expenditure (Gradient PCRD).

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Universalist welfare (U_c) — childcare (asili nido) + school meals (mense scolastiche) + school transport										
Right-wing	+0.98	+0.12	−0.22	−0.33	−0.31	+8.2	+7.5*	−0.1	+5.3	+4.7
	(0.66)	(0.32)	(0.52)	(0.63)	(0.50)	(5.5)	(4.3)	(5.1)	(8.9)	(7.8)
Populism	+1.26**	−0.44	−0.36	−1.15**	−0.83*	+11.6*	−6.7*	−8.7*	−15.7**	−17.5***
	(0.56)	(0.30)	(0.36)	(0.53)	(0.44)	(6.0)	(4.0)	(4.9)	(6.8)	(6.7)
AER	+0.10	+0.08	+0.02	−0.15	+0.09	+2.1	+0.7	+0.5	+4.0	+4.5
	(0.56)	(0.28)	(0.35)	(0.55)	(0.45)	(6.4)	(3.7)	(4.4)	(7.2)	(6.7)
PC	+0.74	−0.40	−0.49	−1.53*	−0.66	+9.0	−5.7	−5.7	−8.3	−11.0
	(0.57)	(0.29)	(0.47)	(0.91)	(0.68)	(6.2)	(4.0)	(5.1)	(7.7)	(7.2)

Note: Each cell reports the bias-corrected gradient PCRD estimate; robust standard errors in parentheses. Outcome: universalist welfare expenditure (childcare contracts .15.010, school meals contracts .15.006, school transport contracts .15.002), in budget-share form (% of total municipal expenditure) and EUR per capita (2018 population). The 2018 placebo column is significantly positive under the Populism (PopT) design, indicating pre-treatment sorting on universalist service spending and breaking RDD identification for that design; PopT estimates are reported but not interpreted causally. Sample: $n = 1065$ close-race municipalities (excluding the three smallest regions — Molise, Sardegna, Basilicata — for numerical stability). The right-wing share-form 2023 cell inherits the EUR-per-capita MSE-optimal bandwidth to eliminate a column-name sensitivity in rdrobust specific to that cell; all other cells use their own MSE-optimal bandwidth. Covariates: region fixed effects, log municipal population, and the 2018 baseline of the same outcome. The 24M3 and 24M6 windows use the cumulative monthly cash flows through March and June 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

O Term-Limit Decomposition

Tables 46–47 replicate the term-limit decomposition of Figures 6–7 (Tables 43–44) using CT-PCRD (fuzzy RD, 2SLS) instead of gradient PCRD (sharp RD). Because CT-PCRD divides the cutoff-induced jump in the outcome by the cutoff-induced jump in the treatment intensity, the coefficients are on the per-unit-treatment-score scale and are therefore numerically much larger than the gradient PCRD estimates. For any cell, the CT-PCRD point estimate is approximately the gradient PCRD estimate divided by the first-stage coefficient (around 0.20 for the right-wing design, smaller for populism). The reported magnitudes should therefore be read as the effect of moving the winning candidate's design-specific score by one full unit of intensity (a counterfactual much larger than what the data actually deliver), rather than

as the treatment effect of a marginal close-race victory. The signs and significance of the load-bearing cells (right-wing W_c and W_c^{++} EUR-per-capita and share at 2023, $24M3$, and $24M6$) are qualitatively consistent across the two estimators. Cells in the populism designs (PopT, AER, PC) and in the term-limited U_c rows are too noisy at the per-unit-score scale to support causal interpretation under CT-PCRD.

Table 46: Effect of treatment on means-tested welfare expenditure (CT-PCRD) — Term-limited mayors.

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Panel A: Core welfare (W_c) — social-assistance anchor										
Right-wing	−14.64	−9.92**	−18.32***	−27.83***	−21.44	−194.8	−193.0**	−237.5***	−291.9***	−333.9**
	(21.23)	(5.04)	(5.20)	(9.83)	(14.84)	(299.9)	(82.6)	(76.3)	(91.4)	(145.2)
Populism	−25.72	−4.26	−0.60	+4.38	+18.77	−239.4	+40.4	+19.7	+22.8	+142.6
	(35.74)	(12.95)	(12.00)	(17.93)	(25.97)	(354.7)	(182.3)	(130.4)	(140.1)	(239.3)
AER	−60.94	−3.42	+4.59	+8.06	+41.32	−662.7	+117.1	+10.4	−169.7	+189.2
	(52.66)	(26.17)	(21.70)	(36.64)	(56.63)	(643.0)	(433.2)	(316.6)	(320.6)	(565.0)
PC	−44.22	+5.88	+6.64	+25.20	+52.32	−396.7	+36.9	+44.9	+41.3	+226.9
	(44.23)	(14.89)	(16.80)	(31.98)	(45.79)	(508.5)	(188.7)	(161.0)	(181.0)	(273.1)
Panel B: Full welfare (W_c^{++}) — Core plus statutorily anchored items										
Right-wing	−14.78	−11.99**	−19.72***	−35.88***	−23.31	−184.4	−200.0**	−257.3***	−355.8***	−392.1**
	(21.33)	(5.56)	(5.77)	(11.89)	(15.78)	(302.0)	(91.9)	(91.6)	(129.6)	(165.0)
Populism	−25.02	−2.99	−2.15	+1.85	+17.63	−230.9	+2.1	+2.0	−32.4	+121.4
	(35.78)	(13.82)	(11.78)	(16.69)	(26.08)	(357.0)	(202.5)	(144.3)	(146.4)	(239.7)
AER	−63.36	−1.51	+4.65	+5.46	+43.68	−680.8	+148.2	−36.2	−221.4	+179.8
	(53.09)	(30.58)	(22.91)	(37.60)	(57.49)	(645.5)	(499.1)	(342.6)	(367.9)	(585.0)
PC	−45.81	+4.69	−3.88	+24.53	+56.13	−373.0	+9.0	+15.0	−62.1	+185.4
	(45.33)	(16.33)	(13.74)	(33.90)	(50.14)	(503.8)	(211.4)	(183.1)	(210.5)	(282.4)

Note: CT-PCRD (fuzzy RD, 2SLS) replication of the gradient PCRD term-limit decomposition (Tables 43–44). The treatment instrument is the $\mathbf{1}\{M_i > 0\}$ cutoff dummy and the endogenous treatment intensity is the winning candidate's design-specific score. Sample: term-limited mayors under the L. 35/2022 classification (5,000-inhabitant threshold against the 2021 legal population), $n = 291$ (excluding the three smallest regions). CT-PCRD uses the full close-elections sample (no caliper); the 5% |diff| trim applies only to the gradient PCRD. Covariates: region FE, log municipal population, and 2018 baseline of the same outcome. Bandwidth held fixed at the MSE-optimal value computed on the Full sample. The right-wing share-form 2023 cells inherit the EUR-per-capita bandwidth. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 47: Effect of treatment on means-tested welfare expenditure (CT-PCRD) — Not-term-limited mayors.

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Panel A: Core welfare (W_c) — social-assistance anchor										
Right-wing	+1.97	−1.67	−4.00**	−4.08*	−3.78*	+136.9**	−26.3	−44.2*	−28.1	−28.7
	(3.90)	(1.37)	(1.64)	(2.34)	(1.96)	(66.7)	(21.5)	(23.6)	(24.9)	(24.6)
Populism	+5.24	−1.37	+5.56	+6.04	+4.69	−187.8	+42.1	+101.4	+115.5	+137.7*
	(8.75)	(4.24)	(3.84)	(4.99)	(4.19)	(424.6)	(70.3)	(88.0)	(75.9)	(82.2)
AER	+2.51	+15.54*	+12.90	+16.96	+18.63*	−802.8	+265.6	+242.0	+278.6	+368.8
	(24.36)	(9.23)	(10.32)	(11.74)	(11.14)	(1167.0)	(169.2)	(202.8)	(200.0)	(242.6)
PC	+15.64	−12.88**	+6.61	−3.08	+5.79	−48.9	−3.4	+162.9	+128.5	+214.7*
	(11.33)	(5.52)	(6.96)	(8.70)	(8.33)	(555.1)	(91.7)	(104.3)	(109.5)	(127.9)
Panel B: Full welfare (W_c^{++}) — Core plus statutorily anchored items										
Right-wing	+1.88	−1.92	−4.47***	−4.22*	−4.21*	+135.0**	−29.3	−47.8**	−30.3	−33.5
	(3.94)	(1.38)	(1.67)	(2.47)	(2.16)	(67.0)	(20.6)	(23.8)	(25.3)	(24.9)
Populism	+4.97	−0.67	+4.54	+5.37	+4.25	−191.0	+64.5	+93.7	+109.3	+131.8
	(8.76)	(4.22)	(4.26)	(5.32)	(4.41)	(424.8)	(69.5)	(85.1)	(78.7)	(84.5)
AER	+4.57	+13.36	+10.14	+15.72	+18.01	−791.5	+240.1	+222.4	+277.8	+373.6
	(24.04)	(9.68)	(10.69)	(12.78)	(11.76)	(1163.9)	(169.1)	(206.1)	(208.7)	(249.5)
PC	+15.04	−10.66*	+3.10	+4.79	+9.72	−70.0	+8.2	+168.0	+145.8	+230.7*
	(11.70)	(5.74)	(5.63)	(8.76)	(8.33)	(550.5)	(91.0)	(102.9)	(109.9)	(128.9)

Note: CT-PCRD (fuzzy RD, 2SLS) replication of the gradient PCRD term-limit decomposition (Tables 43–44). The treatment instrument is the $\mathbf{1}\{M_i > 0\}$ cutoff dummy and the endogenous treatment intensity is the winning candidate's design-specific score. Sample: not-term-limited mayors under the L. 35/2022 classification (5,000-inhabitant threshold against the 2021 legal population), $n = 758$ (excluding the three smallest regions). CT-PCRD uses the full close-elections sample (no caliper); the 5% |diff| trim applies only to the gradient PCRD. Covariates: region FE, log municipal population, and 2018 baseline of the same outcome. Bandwidth held fixed at the MSE-optimal value computed on the Full sample. The right-wing share-form 2023 cells inherit the EUR-per-capita bandwidth. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

P Term-Limit Classification / Delrio Law

Tables 48–49 replicate the term-limit decomposition (Tables 43–44) using the Delrio Law (Legge 56/2014) classification: the three-term exemption applies to municipalities below 3,000 inhabitants, evaluated against the 2011 legal population. This was the law in effect at the time of the 2019 election and governed the first three years of the mandate. The main specification uses Law n. 35/2022 (5,000-inhabitant threshold against the 2021 legal population), reflecting the legal regime that applied for the bulk of the budget windows analyzed. The 5,000 reform took effect in April 2022, three years into the mandate. The

reclassification affects 66 municipalities: second-term mayors in the 3,001–5,000 population bracket who move from term-limited (under the 3,000 threshold, the Delrio classification) to not-term-limited (under the 5,000 threshold, the main classification).

The pattern of changes across the two classifications is consistent with the incentive mechanism. The 66 reclassified mayors are second-term incumbents in the 3,001–5,000 bracket who were elected in 2019 under a two-term limit and began their mandate as lame ducks. Law n. 35/2022 restored their re-election eligibility only in April 2022. These mayors sit in different groups depending on the threshold: not-term-limited under the 5,000 (main) classification, term-limited under the 3,000 (Delrio) classification. Each classification therefore has one clean group and one group contaminated by these 66 ambiguous mayors.

Under the 5,000 (main) classification, the term-limited group is clean, containing only mayors who were term-limited under both laws. The not-term-limited group, by contrast, contains the 66 mayors who entered office expecting to be lame ducks. Under the 3,000 (Delrio) classification, the not-term-limited group contains only mayors who always had re-election incentives, while the term-limited group is contaminated by the 66 mayors who regained eligibility mid-term. For the right-wing retrenchment channel, which operates through the term-limited group, the Delrio classification therefore attenuates the estimates relative to the main specification: the 66 late-reprieve mayors enter the active group and dilute the signal. For the populism expansion channel, which operates through the not-term-limited group, the two classifications differ only by the 66 reprieved mayors, a small share of the not-term-limited universe. The populism cells therefore move very little across the two specifications. The empirical results conform to this prediction.

Table 48: Effect of treatment on means-tested welfare expenditure — Term-limited mayors, Delrio Law classification (Gradient PCRD).

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Panel A: Core welfare (W_c) — social-assistance anchor										
Right-wing	−5.17	−0.12	−0.51	−0.44	−2.56	−19.4	−11.6	−11.8	−16.6	−23.9
	(3.50)	(1.23)	(1.14)	(1.61)	(2.48)	(56.5)	(13.3)	(12.1)	(12.7)	(23.3)
Populism	−2.15	−0.12	−0.36	−0.34	+1.92	−17.0	+5.3	−3.2	−14.4	+1.4
	(2.80)	(1.28)	(1.44)	(2.58)	(2.13)	(30.0)	(18.0)	(23.3)	(34.2)	(34.9)
AER	−2.52	+0.01	−0.09	−0.41	+1.67	−32.0	+10.8	−8.0	−15.4	−0.4
	(2.09)	(1.20)	(1.38)	(2.42)	(2.60)	(25.9)	(20.0)	(24.6)	(29.5)	(36.7)
PC	−1.84	+0.18	+0.94	+2.16	+3.90	−16.1	+1.7	−6.4	−13.2	+2.0
	(2.97)	(1.28)	(1.30)	(2.12)	(2.73)	(34.8)	(16.0)	(28.2)	(38.3)	(38.1)
Panel B: Full welfare (W_c^{++}) — Core plus statutorily anchored items										
Right-wing	−5.13	−0.28	−1.86	−1.23	−1.93	−17.4	−10.8	−36.6*	−24.4	−31.7
	(3.51)	(1.21)	(1.35)	(1.55)	(2.89)	(57.9)	(13.8)	(21.6)	(16.4)	(27.4)
Populism	−2.04	−0.01	−0.43	−0.47	+1.95	−14.7	+3.2	−5.3	−16.6	+2.4
	(2.88)	(1.30)	(1.56)	(2.57)	(2.16)	(30.3)	(17.4)	(25.5)	(34.0)	(34.8)
AER	−2.66	+0.32	+0.13	+0.03	+1.92	−33.5	+14.7	−7.8	−5.8	+4.5
	(2.10)	(1.39)	(1.41)	(2.39)	(2.67)	(26.0)	(22.4)	(25.0)	(23.7)	(35.6)
PC	−1.89	+0.20	+0.23	+2.30	+4.38	−13.4	+1.4	−6.2	−16.5	+2.5
	(3.07)	(1.39)	(1.17)	(2.37)	(3.10)	(34.7)	(16.8)	(28.4)	(38.2)	(37.5)

Note: Each cell reports the bias-corrected gradient PCRD estimate; robust standard errors in parentheses. Outcome: means-tested welfare expenditure, share % (% of total municipal expenditure) and EUR per capita (2018 population). Sample: term-limited mayors under the Delrio Law (Legge 56/2014) classification (3,000-inhabitant threshold against the 2011 legal population), $n = 357$ (excluding the three smallest regions). Covariates: region FE, log municipal population, and 2018 baseline of the same outcome. Bandwidth held fixed at the MSE-optimal value computed on the Full sample. The right-wing share-form 2023 cells inherit the EUR-per-capita bandwidth to eliminate a column-name sensitivity in rdrobust specific to those cells. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 49: Effect of treatment on means-tested welfare expenditure — Not-term-limited mayors, Delrio Law classification (Gradient PCRD).

	Share %					EUR per capita				
	2018 plac.	2022	2023	24M3	24M6	2018 plac.	2022	2023	24M3	24M6
Panel A: Core welfare (W_c) — social-assistance anchor										
Right-wing	−0.12	−0.25	−0.72**	−0.93*	−0.71*	+1.3	−1.5	−5.4	−2.6	−2.6
	(0.83)	(0.31)	(0.30)	(0.52)	(0.43)	(8.3)	(3.5)	(3.5)	(6.0)	(4.8)
Populism	+0.34	−0.24	+0.49	+0.49	+0.37	−24.6	+3.0	+10.2	+12.6	+15.9*
	(1.01)	(0.46)	(0.34)	(0.49)	(0.44)	(49.9)	(6.2)	(8.8)	(7.8)	(9.3)
AER	−0.18	+0.75*	+0.49	+0.53	+0.64	−37.3	+11.6*	+10.1	+10.7	+15.8
	(0.94)	(0.39)	(0.40)	(0.52)	(0.49)	(45.3)	(6.0)	(8.5)	(8.7)	(10.4)
PC	+0.93	−0.94**	+0.46	−0.18	+0.53	−6.3	−1.2	+12.5	+9.6	+16.0
	(0.95)	(0.42)	(0.48)	(0.66)	(0.50)	(44.4)	(5.8)	(8.1)	(8.7)	(9.9)
Panel B: Full welfare (W_c^{++}) — Core plus statutorily anchored items										
Right-wing	−0.18	−0.29	−0.81***	−0.94*	−0.83*	+0.5	−1.7	−5.9	−2.1	−2.8
	(0.85)	(0.31)	(0.30)	(0.55)	(0.46)	(8.5)	(3.5)	(3.7)	(6.3)	(5.1)
Populism	+0.31	−0.18	+0.40	+0.39	+0.30	−25.0	+3.6	+9.3	+11.6	+15.1
	(1.02)	(0.46)	(0.36)	(0.52)	(0.46)	(49.9)	(6.0)	(8.7)	(8.1)	(9.5)
AER	−0.09	+0.62	+0.34	+0.44	+0.57	−36.8	+10.7*	+9.1	+11.1	+15.9
	(0.93)	(0.41)	(0.43)	(0.56)	(0.52)	(45.1)	(6.0)	(8.6)	(9.0)	(10.5)
PC	+0.90	−0.79*	+0.22	+0.41	+0.86*	−8.0	−0.6	+13.0	+11.0	+17.2*
	(0.97)	(0.44)	(0.44)	(0.63)	(0.47)	(44.1)	(5.9)	(8.0)	(8.7)	(9.9)

Note: Each cell reports the bias-corrected gradient PCRD estimate; robust standard errors in parentheses. Outcome: means-tested welfare expenditure, share % (% of total municipal expenditure) and EUR per capita (2018 population). Sample: not-term-limited mayors under the Delrio Law (Legge 56/2014) classification (3,000-inhabitant threshold against the 2011 legal population), $n = 692$ (excluding the three smallest regions). Covariates: region FE, log municipal population, and 2018 baseline of the same outcome. Bandwidth held fixed at the MSE-optimal value computed on the Full sample. The right-wing share-form 2023 cells inherit the EUR-per-capita bandwidth to eliminate a column-name sensitivity in rdrobust specific to those cells. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Tables 48–49 preserve the directional asymmetry across both classifications. The right-wing retrenchment in the term-limited group attenuates noticeably under Delrio. The 66 reprieved mayors are coded as term-limited under the 3,000 threshold, yet for most of the 2022–24 window they had re-election incentives restored by Law n. 35/2022, so they govern as not-term-limited mayors. Their inclusion in the Delrio term-limited group pulls its average toward zero. The not-term-limited populism estimates barely move, as expected from the small-share contamination of that group.

Q Mediation Analysis

Do voters actually observe and react to the welfare cuts documented in Section 7.1? That subsection establishes that right-wing mayors retrench means-tested welfare, but the cuts alone do not show that voters notice them, or that the electoral backlash is connected to the fiscal behavior at all: the backlash could operate through other dimensions of right-wing incumbency entirely (coalition signaling, partisan branding, administrative behavior). A mediation analysis provides suggestive evidence that part of the backlash operates through the welfare channel.

To gauge whether the welfare channel is plausible, a basic Baron–Kenny decomposition (Baron and Kenny, 1986) splits the headline right-wing effect on 2024 vote-weighted ideology into an indirect component running through observed welfare-share changes and a direct component capturing all other channels. Four quantities matter. The *total effect* (c) is the headline RD effect on the 2024 outcome, expected to be negative (a leftward shift). The *first stage* (a) is the RD effect on the welfare-share mediator, expected to be negative (the retrenchment already documented). The *partial effect* (b) is the slope of the 2024 outcome on the mediator given treatment and covariates, expected to be positive: voters reward right-wing mayors who expand welfare rather than cut it. The *direct effect* (c') re-estimates c controlling for the mediator. If voters react to the cuts, c' shrinks in magnitude, and the *mediated share* $(c - c')/c$ is the fraction of the total effect it absorbs. The indirect path is not causally identified: it requires sequential ignorability (Imai, Keele and Tingley, 2010). The estimates are therefore a descriptive decomposition, not an identified mediated effect.

Table 50: Mediation analysis: right-wing → vote-weighted ideology 2024, mediated by welfare share (Baron-Kenny).

Mediator window		Total (c)	A-path (a)	Direct (c')	B-path (b)	Mediated share
Panel A: Core welfare (W_c) share						
Mean(2022, 2023)	Coeff.	-0.066^{***}	-0.011^{**}	-0.051^{**}	$+0.900^{**}$	22.6%
	SE	0.019	0.005	0.022	0.397	$N = 264$
	p-value	0.001	0.030	0.019	0.024	
Mean(2022, 2023, 24M3)	Coeff.	-0.066^{***}	-0.012^{**}	-0.052^{**}	$+0.694^{*}$	21.1%
	SE	0.019	0.006	0.022	0.357	$N = 265$
	p-value	0.001	0.031	0.017	0.053	
Mean(2022, 2023, 24M6)	Coeff.	-0.066^{***}	-0.011^{**}	-0.055^{**}	$+0.526$	16.8%
	SE	0.019	0.005	0.023	0.402	$N = 265$
	p-value	0.001	0.023	0.014	0.192	
Panel B: Full welfare (W_c^{++}) share						
Mean(2022, 2023)	Coeff.	-0.066^{***}	-0.012^{**}	-0.052^{**}	$+0.743^{*}$	21.2%
	SE	0.019	0.005	0.022	0.383	$N = 265$
	p-value	0.001	0.019	0.018	0.053	
Mean(2022, 2023, 24M3)	Coeff.	-0.066^{***}	-0.013^{**}	-0.054^{**}	$+0.547$	19.1%
	SE	0.019	0.006	0.022	0.337	$N = 265$
	p-value	0.001	0.022	0.016	0.106	
Mean(2022, 2023, 24M6)	Coeff.	-0.066^{***}	-0.012^{**}	-0.057^{**}	$+0.354$	14.5%
	SE	0.019	0.005	0.023	0.379	$N = 265$
	p-value	0.001	0.016	0.013	0.350	

Note: Baron–Kenny decomposition of the right-wing → vote-weighted ideology 2024 effect, with mediator = welfare expenditure share. For each mediator window, the three rows report the coefficient (with significance stars), its standard error, and its p-value. The total effect c, A-path a, and direct effect c' are estimated by bias-corrected local-linear RD with a triangular kernel (`rdrobust`), all using the MSE-optimal main bandwidth chosen by `rdbwselect` on the c-path; their standard errors and p-values are robust. The B-path b is the slope of the outcome on the mediator within the same bandwidth, estimated by kernel-weighted OLS, with its OLS standard error and p-value. Mediated share = $100 \times (c - c')/c$; N refers to within-bandwidth observations. Sample: close-race municipalities (excluding the three smallest regions for numerical stability) with finite covariates; the right-wing design restricts to the right side of the spectrum. Covariates: region FE, log municipal population, and 2018 baseline of the same outcome. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 50 reports the decomposition across several mediator windows (the mean of 2022 and 2023, optionally extended to include $24M3$ or $24M6$) for both the Core (W_c) and Full (W_c^{++}) welfare measures. Across specifications, the total effect c is around -0.066 ($p < 0.001$), the first-stage a is in the range -0.011 to -0.013 ($p < 0.05$), the direct effect c' is in the range -0.051 to -0.057 ($p < 0.05$), and the mediated share ranges from approximately 15% to 23%. The partial effect b is statistically significant or marginally significant in the shorter mediator windows but loses significance in the longer ones (the $24M6$ windows and the Full-welfare $24M3$ window). A mediated share close to one-fifth of the total effect is non-trivial and

within the range of magnitudes reported in published political-science mediation analyses.[28]

R Mediation Analysis with EUR-per-capita Mediator

Table 51 reports the Baron–Kenny decomposition with the welfare mediator measured in EUR per capita rather than as a budget share. The total effect (c) and direct effect (c') are unchanged from the share-form specification, because the outcome is the same electoral variable. The A-path (a) is on a different scale, and the B-path (b) is correspondingly rescaled. The mediated share is 14–19% across windows, directionally similar to the share-form specification. The b-path, however, is not statistically significant in any cell, reflecting the lower precision of the EUR-per-capita mediator at the within-bandwidth weighted-OLS step.

[28] For comparison, Imai and Yamamoto (2013) report mediated shares of 14.1–21.3% for the welfare-framing experiments of Slothuus (2008) and 25.3% for the anxiety mechanism in Brader, Valentino and Suhay (2008) when allowing perceived harm to depend on anxiety, magnitudes in the same range as the estimates reported here.

Table 51: Mediation analysis: right-wing → vote-weighted ideology 2024, mediated by welfare EUR per capita (Baron-Kenny).

Mediator window		Total (c)	A-path (a)	Direct (c')	B-path (b)	Mediated share
Panel A: Core welfare (W_c) EUR per capita						
Mean(2022, 2023)	Coeff.	-0.066^{***}	-12.38^{**}	-0.054^{**}	$+0.0005$	18.4%
	SE	0.019	5.83	0.022	0.0003	$N = 265$
	p-value	0.001	0.034	0.016	0.135	
Mean(2022, 2023, 24M3)	Coeff.	-0.066^{***}	-11.12^{*}	-0.056^{**}	$+0.0004$	15.5%
	SE	0.019	5.81	0.022	0.0003	$N = 266$
	p-value	0.001	0.055	0.013	0.217	
Mean(2022, 2023, 24M6)	Coeff.	-0.066^{***}	-12.28^{**}	-0.056^{**}	$+0.0004$	15.6%
	SE	0.019	6.11	0.023	0.0003	$N = 266$
	p-value	0.001	0.044	0.014	0.264	
Panel B: Full welfare (W_c^{++}) EUR per capita						
Mean(2022, 2023)	Coeff.	-0.066^{***}	-13.70^{**}	-0.055^{**}	$+0.0004$	16.6%
	SE	0.019	6.11	0.023	0.0003	$N = 266$
	p-value	0.001	0.025	0.015	0.246	
Mean(2022, 2023, 24M3)	Coeff.	-0.066^{***}	-12.24^{**}	-0.057^{**}	$+0.0003$	13.8%
	SE	0.019	6.01	0.023	0.0003	$N = 266$
	p-value	0.001	0.042	0.012	0.378	
Mean(2022, 2023, 24M6)	Coeff.	-0.066^{***}	-13.51^{**}	-0.057^{**}	$+0.0002$	13.7%
	SE	0.019	6.34	0.023	0.0003	$N = 266$
	p-value	0.001	0.033	0.012	0.453	

Note: Baron–Kenny decomposition of the right-wing → vote-weighted ideology 2024 effect, with mediator = welfare expenditure in EUR per capita. For each mediator window, the three rows report the coefficient (with significance stars), its standard error, and its p-value. The total effect c, A-path a, and direct effect c' are estimated by bias-corrected local-linear RD with a triangular kernel (`rdrobust`), all using the MSE-optimal main bandwidth chosen by `rdbwselect` on the c-path; their standard errors and p-values are robust. The B-path b is the slope of the outcome on the mediator within the same bandwidth, estimated by kernel-weighted OLS, with its OLS standard error and p-value. Mediated share = $100 \times (c - c')/c$; N refers to within-bandwidth observations. Sample: close-race municipalities (excluding the three smallest regions for numerical stability) with finite covariates; the right-wing design restricts to the right side of the spectrum. Covariates: region FE, log municipal population, and 2018 baseline of the same outcome (in EUR per capita). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

S Codebooks and Labeling Examples

S.1 Ideology

S.2 What Are We Measuring?

Ideology measures the **left-right political positioning** of a text, ranging from extreme left (public ownership, workers' power, antifascist identity) through centre/technocratic language

to extreme right (exclusionary nationalism, anti-immigration, economic sovereignty framing).

This is **NOT** about:

- Populism (that's a separate dimension)
- People-centrism (that's a separate dimension)
- General quality of governance

S.3 The Five Codes

Code	Meaning
Extreme Left	Public ownership/nationalization; radical redistribution; anti-capitalism; workers' power; explicit antifascist identity; systemic anti-neoliberal program; welfare/health maximalism; harm-reduction drug policy; civil-rights maximalism
Left	Progressive welfare; labor protections; environmental/social justice; inclusive immigration
Centre	Managerial/technocratic; pragmatic delivery; no clear left/right thrust
Right	Tax cuts; deregulation; private-sector solutions; tradition/family/national pride (secular/cultural) without out-group blame or threat framing
Extreme Right	Exclusionary nationalism; anti-immigration; us-vs-them narratives; ethno-religious supremacy; anti-LGBTQ+ rights; "natural family" exclusivity; law & order as civilizational threat; economic/monetary sovereignty framing

S.4 Default Rule

- **If there's ANY readable text** ⇒ Assign an Ideology label (EL, L, C, R, or ER)
- **OCR_Noise** ⇒ Only for completely unreadable OCR junk

S.5 Coding Decision Tree

S.5.1 Step 0: Is the text readable?

- **If unreadable OCR junk** ⇒ Code as `OCR_Noise`
- **If readable** ⇒ Continue to Step 1

S.5.2 Step 1: Check for Extreme Left (EL) triggers

If you find ANY of these, code as EL immediately:

Antifascist/Partisan Identity

- Explicit **ANPI** (Associazione Nazionale Partigiani d'Italia)
- *Resistenza/partigiani* (Resistance/partisans)
- *Lotta antifascista* (anti-fascist struggle)

Systemic Anti-Neoliberal Program

- Condemnation of *liberismo* (economic liberalism), *privatizzazioni* (privatizations), *austerità* (austerity), *vincoli di bilancio* (balanced-budget constraints)
- With **rollback/reversal** or class/system framing
- **UNLESS** explicitly tied to **national-interest/sovereignty** ⇒ ER

Welfare/Health Maximalism + Emotive Framing

- Abolitionist/moral language vs hazards (e.g., *amianto* [asbestos], *CVM* [vinyl chloride monomer])
- **Zero patient/relative costs regardless of origin**
- Anti-marginalization guarantees

Harm Reduction (Drug Policy)

- Safe-consumption rooms
- Syringe/needle exchange
- Heroin-assisted treatment
- Decriminalization/legalization framed as public health/rights with unconditional public provision

Civil-Rights Maximalism

- **Full LGBTQ+ equality** including *omogenitorialità* (same-sex parenting), *adozione* (adoption), *matrimonio egualitario* (marriage equality)
- Hate-crime enforcement/expansion (e.g., *Legge Mancino* strengthening)
- End-of-life self-determination (euthanasia/*DAT* [advance healthcare directives]) framed as unconditional individual freedom and rights

Immigrant-Rights Maximalism

- **Unconditional access/services/guarantees** "*a prescindere da cittadinanza/status*" (regardless of citizenship/status)

- *Iscrizione anagrafica* (civil registry), *sanità/istruzione/alloggio* (healthcare/education/-housing) *senza requisiti di status* (without status requirements)
- *Corridoi umanitari* (humanitarian corridors)
- *Ius soli / ius scholae* (citizenship by birth/by schooling)
- *Tutela legale gratuita* (free legal aid)

Solidarity with Migrants (Maximalist)

- Explicit *solidarietà/partecipazione diretta* (solidarity/direct participation) for *migranti/stranieri* (migrants/foreigners)
- **Including non-citizens** ("anche senza/ancora senza cittadinanza" [even without/not yet with citizenship])
- With rights or unconditional inclusion framing

Pro-Immigration as Human-Rights Framing

- Explicit *tutela dei diritti umani / dignità / dovere d'asilo* (human rights protection / dignity / duty of asylum)
- *Accoglienza* (reception/welcome) as *diritto* (right) or *obbligo morale/valoriale* (moral/value-based duty)
- Cues: "diritti umani", "dignità della persona", "nessuno è illegale" (no one is illegal), "diritto d'asilo", "protezione internazionale", "accoglienza come dovere costituzionale/etico"

Social-Welfare Maximalism (Non-Health)

- *Gratuità universale* (universal free provision) or *blocco/azzeramento generalizzato delle tariffe* (generalized freeze/zeroing of fees)
- For essential public services (nurseries, school meals, transport, housing) **without** means test
- *Esenzioni totali* (total exemptions) or *fondi di solidarietà universali/automatici* (universal/automatic solidarity funds) as *diritto sociale* (social right)

Distributed Reception (SPRAR/SAI) with Normative Commitment

- Explicit commitments to *gestione diretta comunale* (direct municipal management) of SPRAR/SAI/*accoglienza diffusa* (refugee reception systems/distributed reception)
- With *quota/target proporzionale agli abitanti* (quota/target proportional to residents)
- Solidarity/rights framing ("comunità accogliente" [welcoming community], "nessuno è illegale", "diritto d'asilo")

- With *percorsi di integrazione/formazione/lavori socialmente utili* (integration/training/socially useful work paths)

Self-Managed Social Centers (Positive Endorsement)

- Promoting/supporting/creating *centri sociali autogestiti* (self-managed social centers) as community/cultural hubs
- Neutral/administrative mentions without endorsement remain L/C

Income Support (Universal/Unconditional)

- *Reddito di Cittadinanza* (RdC), *reddito minimo* (minimum income), UBI as **universal/unconditional**
- *Diritto* (right) for *tutti* (everyone), **even without citizenship/status**

If ANY EL trigger is present ⇒ Code as **EL**
If no EL triggers ⇒ Continue to Step 2

S.5.3 Step 2: Check for Extreme Right (ER) triggers

If you find ANY of these, code as ER immediately:

Religious Rhetoric

- *Ispirazione/radici cristiane* (Christian inspiration/roots)
- Faith-based governance, religious duties
- *Sussidiarietà* (subsidiarity) framed religiously

Family Model Exclusivity

- **"Famiglia naturale"** (natural family) as **man–woman only**
- Denial of recognition/rights to same-sex families

Economic/Industrial National Interest

- Explicit *interesse nazionale* (national interest), *sovranità economica* (economic sovereignty), *settori strategici nazionali* (national strategic sectors)
- To justify industrial policy, nationalization/state control, or fiscal/debt rule changes (e.g., *pareggio di bilancio* [balanced budget], *debito* [debt])

Monetary Sovereignty / Central-Bank Nationalism

- *Sovranità monetaria* (monetary sovereignty)
- *Signoraggio* (seigniorage)
- *Banca d'Italia* (Bank of Italy) nationalization or political control
- *Uscita dall'euro/indipendenza da BCE* (exit from euro/independence from ECB)
- *Debito/pareggio di bilancio* (debt/balanced budget) framed as national-sovereignty identity

Exclusionary Nationalism

- Us-vs-them or "replacement" narratives
- Ethno-religious supremacy
- Anti-immigration with threat/civilizational framing

Anti-LGBTQ+ Rights

- Opposition to LGBTQ+ equality
- "Natural family" (man–woman) exclusivity

Law & Order as Civilizational Threat

- Security linked to specific out-groups
- Existential/civilizational threat framing
- **Note**: Routine safety/crime ⇒ R

Income Support (Restrictions/Exclusions)

- Abolish/restrict *Reddito di Cittadinanza*
- **ER** if with exclusion for *stranieri* (foreigners) or stigmatizing/us-vs-them rhetoric
- **R** if general restrictions without out-group targeting

If ANY ER trigger is present ⇒ Code as **ER**
If no ER triggers ⇒ Continue to Step 3

S.5.4 Step 3: Evaluate Left (L) vs Right (R) vs Centre (C)

Code as LEFT (L) if:

- Progressive welfare measures
- Labor protections

- Environmental/social justice
- Inclusive immigration (with administrative or non-maximalist framing)
- Income support with means test/conditionality (ISEE, PUC, *Patto per il Lavoro*)
- SPRAR/SAI as organizational tool without quota/rights/solidarity commitments
- Non-discrimination/recognition of cohabitation **without** equating to marriage
- Isolated critique of privatization (not systemic rollback)

Code as RIGHT (R) if:

- Tax cuts
- Deregulation
- Private-sector solutions
- Tradition/family/national pride (secular/cultural) **without** out-group blame or threat framing
- Family valorization without exclusionary language
- Routine law & order (not linked to out-groups or civilizational threat)

Code as Centre (C) if:

- Managerial/technocratic language
- Pragmatic service delivery
- No clear left/right thrust
- Local technical management (service desks, coordination, monitoring)
- Administrative/community measures (transparency, service desks, "baratto amministrativo" [administrative barter]) by default
- Technical/regulatory public-health framing

When Torn Between R and ER

Prefer **ER** if there's **any** exclusionary, sovereignty-as-threat, or national-interest defense language.

S.6 Key Disambiguation Rules

S.6.1 Privatization & Neoliberal Critiques

- Descriptive/isolated critique ⇒ **L**
- Rollback/anti-capitalist framing ⇒ **EL**
- Rollback framed as defending national interest/sovereignty ⇒ **ER**

S.6.2 Nationalization / State Control

- Redistribution/class framing ⇒ **EL**
- Defending strategic sectors/national sovereignty/*interesse nazionale* ⇒ **ER**

S.6.3 Monetary Sovereignty & Banking Sector

- Measures involving *Banca d'Italia*, BCE/euro, *signoraggio*, central-bank nationalization ⇒ **ER** (sovereignty/identity framing)
- Nationalization argued on redistribution/anti-privatization class grounds **without** sovereignty framing ⇒ **EL**

S.6.4 Immigration/Integration

- **EL** ⇒ Explicit pro-immigration defense grounded in human rights/dignity/duty of asylum; expansion of rights/services regardless of citizenship/status; humanitarian corridors; *ius soli/ius scholae*; free legal protection
- **L** ⇒ Integration projects/services, cultural mediation, funding, participation with administrative or non-maximalist inclusive framing (language courses, mediators, service desks)
- **R/ER** ⇒ Restrictions/priority to citizens; **ER** if there's exclusion, blame, "ethnic replacement", civilizational threat, or security as identity against out-group

S.6.5 SPRAR/SAI

- Administrative only (management/financial): Reference to SPRAR/SAI/*accoglienza diffusa* as organizational tool without quota/rights/solidarity commitments ⇒ **L**
- With normative commitment (see EL triggers above) ⇒ **EL**

S.6.6 Cohabitation and Civil Rights

- **EL** ⇒ *Piena equiparazione* (full equalization): marriage equality, adoption/same-sex parenting, explicit extension of all rights
- **L** ⇒ Non-discrimination/recognition of cohabitation **without** equating to marriage
- **ER** ⇒ "Natural family man–woman" exclusivity / denial of rights

S.6.7 Income Support (RdC / Minimum Income / UBI)

- **EL** ⇒ Universal/unconditional; right for everyone, even without citizenship/status
- **L** ⇒ Means test/conditionality (ISEE, PUC, *Patto per il Lavoro*), administrative defense/improvement

- **C** ⇒ Local technical management (service desks, INPS coordination, monitoring)
- **R/ER** ⇒ Abolish/restrict; **ER** if with exclusion for foreigners or stigmatizing/us-vs-them rhetoric

S.6.8 Public-Health Bans/Regulation

- Technical/regulatory framing ⇒ **C/L**
- Abolitionist/moral/emotive with universal guarantees ⇒ **EL**

S.6.9 Memory/Heritage

- Neutral civic remembrance ⇒ **L/C**
- Activist/identity antifascism ⇒ **EL**

S.6.10 Identity (Non-Religious)

- Cultural pride/heritage ⇒ **R**
- Exclusion/denial of rights to minorities/immigrants/LGBTQ+ ⇒ **ER**

S.6.11 Administrative/Community Measures

- Transparency, service desks, "baratto amministrativo" ⇒ **C** by default
- Shifts to **L** only if there's clear redistributive expansion
- Not **EL** unless explicit universality/unconditionality

S.7 Quick Reference Table

If you see...	Then...
ANPI / Resistenza / lotta antifascista	EL
Systemic anti-neoliberal rollback (no sovereignty)	EL
Welfare/health maximalism + emotive framing	EL
Harm reduction (drug policy)	EL
Full LGBTQ+ equality (marriage/adoption)	EL
Immigrant rights regardless of status	EL
Pro-immigration as human rights	EL
Universal free services (no means test)	EL
SPRAR/SAI with quota/solidarity framing	EL
Self-managed social centers (endorsed)	EL
Universal/unconditional income support	EL
Christian inspiration / religious governance	ER
"Natural family" (man–woman only)	ER
National interest / economic sovereignty	ER
Monetary sovereignty / Banca d'Italia	ER
Exclusionary nationalism / replacement	ER
Anti-LGBTQ+ rights	ER
Law & order as civilizational threat	ER
Income support restricted to citizens	ER
Progressive welfare (means-tested)	L
Labor protections	L
Environmental/social justice	L
Inclusive immigration (non-maximalist)	L
SPRAR/SAI (administrative only)	L
Cohabitation recognition (not marriage)	L
Tax cuts / deregulation	R
Private-sector solutions	R
Cultural pride (no out-group blame)	R
Family valorization (no exclusion)	R
Routine law & order	R
Managerial/technocratic language	C
Service delivery (neutral)	C
Administrative measures	C

S.8 Coding Workflow Summary

1. **Check readability** ⇒ OCR noise? ⇒ `OCR_Noise`
2. **Check EL triggers** ⇒ Any present? ⇒ **EL**
3. **Check ER triggers** ⇒ Any present? ⇒ **ER**
4. **Evaluate L vs R vs C** ⇒ Based on content

S.9 Examples in Context

S.9.1 Extreme Left (EL) Examples

- "Difendiamo i valori della Resistenza e della lotta antifascista" (defend values of Resistance and anti-fascist struggle)
- "Nazionalizziamo i servizi pubblici per redistribuzione, non per interesse nazionale" (nationalize public services for redistribution, not for national interest)
- "Diritto d'asilo e accoglienza come dovere costituzionale" (right of asylum and reception as constitutional duty)
- "Servizi pubblici gratuiti per tutti, a prescindere dalla cittadinanza" (free public services for all, regardless of citizenship)
- "Piena equiparazione: matrimonio egualitario e omogenitorialità" (full equalization: marriage equality and same-sex parenting)

S.9.2 Left (L) Examples

- "Progetti di integrazione e mediazione culturale" (integration projects and cultural mediation)
- "Reddito di Cittadinanza con Patto per il Lavoro" (citizenship income with job pact)
- "Tutela dei diritti dei lavoratori e contratti collettivi" (protection of workers' rights and collective contracts)
- "Riconoscimento delle convivenze senza equiparazione al matrimonio" (recognition of cohabitation without equating to marriage)

S.9.3 Centre (C) Examples

- "Sportello unico per servizi amministrativi" (single desk for administrative services)
- "Monitoraggio KPI e coordinamento con INPS" (KPI monitoring and coordination with social security)
- "Trasparenza e pubblicazione dati online" (transparency and online data publication)

S.9.4 Right (R) Examples

- "Riduzione tasse locali e semplificazione burocratica" (reduction of local taxes and bureaucratic simplification)
- "Valorizzazione della famiglia come rete di valori" (valorization of family as network of values)
- "Orgoglio per le nostre tradizioni e cultura locale" (pride in our traditions and local culture)
- "Sicurezza urbana attraverso controllo del territorio" (urban security through territorial control)

S.9.5 Extreme Right (ER) Examples

- "Radici cristiane della nostra comunità" (Christian roots of our community)
- "La famiglia naturale è solo quella tra uomo e donna" (natural family is only between man and woman)
- "Sovranità monetaria e nazionalizzazione della Banca d'Italia" (monetary sovereignty and nationalization of Bank of Italy)
- "Interesse nazionale nei settori strategici" (national interest in strategic sectors)
- "Prima gli italiani: restrizioni per gli stranieri" (Italians first: restrictions for foreigners)

S.10 Populism

S.11 What Are We Measuring?

Populism measures whether a text frames politics as a **conflict between "the people" and "the elite/system"**, uses **hard accusations** against political elites, or claims **exclusive representation** of popular will. In its **soft** form, populism can also appear as **people-centric orientations *or* weak, non-accusatory anti-elite framing** (e.g., elites/institutions portrayed as distant, neglectful, or unresponsive), even when the text stops short of corruption/capture/betrayal claims and does not assert exclusive legitimacy.

This is **NOT** about:

- Left vs. right ideology (that's a separate dimension)
- Technical policy criticism
- Simple inefficiency complaints
- Attacks on out-groups (that's the ideological axis)

Code	Meaning	Example
None	No people-vs-elite conflict	Technical content, may criticize inefficiency but no hard accusations
Soft	People-centric *or* weak anti-elite	People-centric appeals *or* elite-distance framing (incl. participatory ideals/tools) WITHOUT hard accusations or exclusive claims
Hard	People-vs-elite conflict OR hard accusations OR exclusive claims	Any hard accusation automatically triggers Hard

S.12 The Three Codes

S.13 Coding Decision Tree

S.13.1 Step 0: Is the text readable?

- **If unreadable OCR junk** ⇒ Code as `OCR_Noise`
- **If readable** ⇒ Continue to Step 1

S.13.2 Step 1: Check for hard accusations

If you find ANY of these, code as Hard immediately:

"Casta" (Political Caste)

- Accusatory use of "casta" toward politicians/elites
- Any connection to privileges/abuses/acting against citizens
- **Exception**: Metalinguistic/historical use without accusation ⇒ None

Corruption/Enrichment of the Few

- Accusations that elites enrich themselves at citizens' expense
- "Si intascano i soldi della gente" (they pocket the people's money)
- Corruption benefiting few against the many

Clientelism/Favoritism

- *Mercimonio* (favoritism), "amici degli amici" (friends of friends)
- Rigged contracts, favors for cronies

Betrayal/Rotten System

- Politicians/elites/parties have "betrayed" citizens
- *Sistema marcio* (rotten system)

Capture by Powerful Interests Policies serving **privileged few** AGAINST citizens' interests. Includes:

- Politics "polluted/contaminated/conditioned" by interests serving the few
- Must be "protected from" interests serving the few
- Promises to act "without compromising with *poteri forti*"
- Choices conditioned/polluted by *poteri forti*
- ANY denunciation of "*privilegi di pochi*" / "*vantaggio per pochi*" / "*pochi privilegiati*" in political/administrative context

Important Clarification

Public-private or redistributive conflicts (beach access, parking, fair compensation) WITHOUT elite blame/complicity do NOT trigger this rule ⇒ Evaluate as None unless other hard elements present.
Exception: The formulation "privilegi/vantaggio per pochi" automatically triggers Hard regardless.

Note on *Sfruttamento* (Exploitation)

Generic/rhetorical use (e.g., criticizing fiscal centralism) is NOT automatically a hard trigger. Becomes Hard ONLY if accompanied by:

- Explicit accusations against elites (corruption, clientelism, betrayal), OR
- Reference to *poteri forti*/capture, OR
- Formulations about "privileges/benefits for the few"

Exclusivity of Popular Will Claims to **exclusive/totalizing** representation:

- "We alone represent/interpret the people's will"
- Delegitimization of pluralism or dissent
- **Does NOT trigger** for simple co-identification slogans ("with the people", "we want what the people want")

Total Delegitimization of Expertise/Mediation

- "Technocrats/authorities **rig/suffocate** the popular will"
- Goes beyond calling for citizen input to attacking competence/mediation as illegitimate

If ANY hard accusation is present ⇒ Code as **Hard**
If no hard accusations ⇒ Continue to Step 2

S.13.3 Step 2: Is there soft populism (people-centrism *or* weak anti-elite framing)?

Code as **Soft** if you find **at least one** people-centric *or* weak, non-accusatory anti-elite element below (and **no** hard accusations or exclusive claims).

Voice/Power Transfer Tools "Ridare voce/potere ai cittadini" (restore voice/power to citizens) + **concrete tools**:

- Assemblies, councils, participatory budgets, petitions
- Consultative/confirmatory referendum as procedure
- Transfer of decision power toward citizens (co-decision/binding local mandate)

Diffuse/Continuous Citizen Control Explicit framing of **accountability at any moment**:

- "Sotto il controllo di tutti" (under everyone's control)
- "Valutati in ogni momento dai cittadini" (evaluated at any moment by citizens)
- Periodic public reports to citizens
- **Even without listing specific tools**

Note: Mere "transparency" or "data publication" WITHOUT continuous popular control/evaluation ⇒ None

Anti-Party/Elite "Distance" + Citizen Focus Phrases about distance from elites WITHOUT hard accusations:

- *Nei palazzi* (in the palaces)
- "Riunioni per pochi" (meetings for the few)
- *Zavorre/pesi dei partiti* (burdens/weights of parties)
- *Lontani dalla gente* (distant from the people)
- Combined with citizen-centered approach
- Even without explicit tools

Salary/Compensation Renunciation as Moral Duty Renunciation explicitly presented as **moral duty/honor** in favor of "those most in need":

- Even without hard accusations
- Must be framed within ethical renewal discourse

Bottom-Up Localism + Moral Renewal Anti-old politics framing + bottom-up language:

- *Dal basso* (from below), "con le persone come noi" (with people like us)
- Moral renewal ("gestire bene la *res publica*", "creare sensibilizzazione politica locale")
- Even without tools or diffuse control
- If lacking hard accusations

Comprehensive Participatory Commitments Package of participatory measures:

- Donate portion of salary to community
- Citizen benefits (e.g., citizenship allowance)
- Monthly meetings, statute implementation
- **Even if citing past inaction** ("disatteso" - ignored)
- If lacking hard accusations (no corruption, enrichment, betrayal, capture, casta, rotten system)

"Neglected + Unprecedented Attention" Co-occurrence of institutional neglect + commitment to unprecedented attention:

- "Institutions that neglect/forget" + "unprecedented attention" toward common good/local territory
- Even without tools or diffuse control
- Provided lacking hard accusations and exclusivity
- Narrow exception to None rule for simple/soft criticism

Non-Exclusive Co-Identification Co-identification with "la gente" in administrative tone:

- "Insieme alla Gente" (Together with the People)
- "Vogliamo quello che vuole la gente" (we want what the people want)
- WITHOUT delegitimizing mediation/competence
- WITHOUT hard accusations
- Even without tools

Hyperbolic Delegitimization (Without Hard Accusations) Hyperbolic statements about institutions:

- "The Council meets ONLY to raise taxes / worsen life"
- WITHOUT accusations of corruption/clientelism/capture/privileges for few
- WITHOUT *poteri forti/casta*
- WITHOUT exclusivity of popular will

If any Soft element is present AND no hard accusations ⇒ Code as **Soft**
If no Soft elements and no hard accusations ⇒ Code as **None**

S.13.4 Step 3: Default to None

If none of the above apply, code as **None**.

None includes:

- Technical/managerial text
- Descriptive policy proposals
- May criticize inefficiency/mismanagement
- Simple/soft institutional criticism (neglect, delays, inefficiency, "forgotten by funding") WITHOUT hard accusations or people-vs-elite frame (i.e., purely administrative critique that does not portray elites/institutions as distant/unresponsive to ordinary people)
- Simple reduction or renunciation of salary (without hard/exclusive framing)

S.14 Common Pitfalls and Special Cases

× **Simple Inefficiency Criticism**

- Criticizing delays, bureaucracy, mismanagement
- WITHOUT accusing corruption, betrayal, or elite-people conflict
- ⇒ None

× **Service Delivery**

- "Services for citizens" (neutral language)
- Customer service improvements
- ⇒ None

Referendums & Popular Sovereignty

- Procedural tools (consultative/confirmatory) WITHOUT hard accusations or exclusivity ⇒ Soft
- With hard accusations OR exclusivity ⇒ Hard

Identity Slogans

- "Con/insieme alla gente" WITHOUT exclusivity or hard accusations ⇒ Soft
- Or None if purely ornamental/technical

Public Spending / "People's Money"

- Technical criticism (audits, tenders, KPIs) ⇒ None
- "Soldi della gente" (people's money) + control WITHOUT hard accusations ⇒ Soft
- "Si intascano i soldi della gente" (they pocket the people's money) ⇒ Hard

Administrator Salaries

- Simple reduction or renunciation ⇒ None
- Moral/honor renunciation for the needy ⇒ Soft
- Renunciation + benefit tools + participation ⇒ Soft
- With people-vs-elite frame or enrichment accusations ⇒ Hard

Anti-Party Rhetoric

- "Civic beyond alignments" ⇒ None
- "Parties distant from the people" (elite-distance + citizen focus, even without explicit tools) ⇒ Soft
- "Parties distant from the people" + tools/diffuse control ⇒ Soft
- "Parties = casta/corrupt/traitors" ⇒ Hard

Anti-Expert Rhetoric

- Reviewing opinions/simplifications ⇒ None
- "Enough technicalities that ignore citizens' voices" + tools/diffuse control ⇒ Soft
- "Technocrats rig opinions / block the popular will" ⇒ Hard

Out-Groups

- Attacks on groups = ideological axis (R/ER), not populist per se
- Becomes populism if presented as elite tool against the people ⇒ Hard

S.15 Quick Reference Table

If you see...	Then...
"Casta" (accusatory)	Hard
Corruption/enrichment accusations	Hard
Clientelism/"amici degli amici"	Hard
Betrayal/"sistema marcio"	Hard
"Poteri forti" conditioning choices	Hard
"Privilegi/vantaggio per pochi"	Hard
Politics polluted by interests of few	Hard
Exclusive popular will claims	Hard
Technocrats rig/suffocate popular will	Hard
Participatory budget/assemblies (no hard)	Soft
"Under everyone's control"	Soft
"Palazzi distant" + citizen focus (no hard)	Soft
Salary renunciation as moral duty	Soft
"Dal basso" + moral renewal (no hard)	Soft
"Neglected" + "unprecedented attention"	Soft
"Together with the people" (non-exclusive)	Soft
"Council meets only to raise taxes" (no hard)	Soft
Technical criticism of inefficiency	None
"Civic beyond parties"	None
Simple salary reduction	None
Service delivery language	None

S.16 Coding Workflow Summary

1. **Check readability** ⇒ OCR noise? ⇒ `OCR_Noise`
2. **Check hard accusations** ⇒ Any present? ⇒ **Hard**
3. **Check soft populism elements (people-centric *or* weak anti-elite)** ⇒ Any present? ⇒ **Soft**
4. **Default** ⇒ **None**

S.17 Examples in Context

S.17.1 None Examples

- "Snelliremo la burocrazia con sportello unico digitale" (streamline bureaucracy with single digital desk)

- "Lista civica senza appoggio dei partiti; il programma nasce dal confronto con la comunità" (civic list without party support; program from community dialogue)
- "Le Marine hanno molti tratti privati; il Comune negozierà accordi bonari per parcheggi e accesso pubblico" (Marina areas have private sections; Municipality will negotiate fair agreements for parking and public access)

S.17.2 Soft Examples

- "Politica locale dal basso, non la vecchia politica dei partiti nazionali" (local politics from below, not old politics of national parties)
- "Il Palazzo è distante: assemblee e bilancio partecipativo per ridare voce ai cittadini" (Palace is distant: assemblies and participatory budgets to restore citizens' voice)
- "Amministrazione pulita e trasparente, sotto il controllo di tutti; saremo valutati in ogni momento" (clean and transparent administration, under everyone's control; evaluated at every moment)
- "Fiducia senza le zavorre dei partiti chiusi nei palazzi" (trust without burdens of parties closed in palaces)
- "Onorabilità, dovere chiari dei candidati, rinuncia al 40% dei compensi a favore delle persone che più hanno bisogno" (honor, clear duty of candidates, renunciation of 40% of compensation for those most in need)
- "La Laguna maltrattata/dimenticata dalle istituzioni: serve un'attenzione mai vista" (Lagoon mistreated/forgotten by institutions: unprecedented attention needed)
- "Insieme alla Gente... vogliamo quello che vuole la Gente" (Together with the People... we want what the People want)

S.17.3 Hard Examples

- "La casta si è arricchita coi soldi della gente" (the caste enriched itself with the people's money)
- "Mettere al sicuro la politica del Comune da interessi per pochi e da inquinamenti" (protect Municipality's politics from interests for the few and from pollution)
- "Scuola media solo con risorse nostre, senza scendere a compromessi con i poteri forti" (middle school only with our resources, without compromising with strong powers)
- "Decisioni prese per mantenere privilegi a pochi" (decisions made to maintain privileges for the few)

S.18 Anti-Elite Rhetoric (AER)

S.19 What Are We Measuring?

Anti-Elite Rhetoric (AER) measures whether a text frames political elites or institutions as **distant, neglectful, or corrupt**, ranging from neutral technical language to accusations of capture and betrayal.

This is **NOT** about:

- General policy criticism
- Resource constraints
- Technical inefficiencies without political blame

S.20 The Three Codes

Code	Meaning	Example
None	No anti-elite rhetoric	Technical/managerial language, neutral policy, resource issues without political blame
Soft	Soft anti-elite rhetoric	Political distance, institutional neglect, anti-party sentiment, hyperbolic delegitimation WITHOUT hard accusations
Hard	Hard anti-elite rhetoric	Corruption, capture, betrayal, privilege extraction, systematic elite harm

S.21 Coding Decision Tree

S.21.1 Step 0: Is the text readable?

- **If unreadable OCR junk** ⇒ Code as `OCR_Noise`
- **If readable** ⇒ Continue to Step 1

S.21.2 Step 1: Check for hard anti-elite rhetoric

If you find ANY of these, code as Hard immediately:

"Casta"

- *Casta* + privileges/enrichment/abuses
- Any accusatory use toward political elites

Corruption/Enrichment

- *Corruzione*, *arricchimento*
- "Si intascano" (they pocket), "rubano" (they steal)
- "Soldi pubblici nelle tasche" (public money in their pockets)

Clientelism/Nepotism

- *Clientelismo*, *mercimonio*
- "Amici degli amici" (friends of friends)
- "Appalti truccati" (rigged contracts)

Betrayal/System Rot

- *Tradimento*, "tradito i cittadini" (betrayed citizens)
- *Sistema marcio* (rotten system)

Capture/Privilege

- *Poteri forti* (strong powers)
- Politics "*condizionata da*" (conditioned by), "compromessi con *poteri forti*"
- "*Privilegi di/per pochi*" (privileges for the few)
- "*Vantaggi per pochi*" (advantages for the few)
- "*Pochi privilegiati*" (few privileged)
- "*Interessi di pochi*" (interests of the few)
- "Politica *inquinata da*" (politics polluted by)
- "Mettere al sicuro la politica da" (protect politics from)

Accusatory Anti-Expertise

- "Tecnici truccano" (technocrats rig)
- "Esperti soffocano la volontà popolare" (experts suffocate popular will)
- "Pareri manipolati" (manipulated opinions)

Frame: Elite **intentionally harms** citizens through corruption, capture, or privilege extraction.

If ANY hard rhetoric is present ⇒ Code as **Hard**
If no hard rhetoric ⇒ Continue to Step 2

S.21.3 Step 2: Check for soft anti-elite rhetoric

Code as **Soft** if you find **at least one** of these three sub-types:

Anti-Party Distance Lexical markers about political parties:

- *Partiti*, *apparati*, *palazzi*
- *Vecchia politica*, *logiche di partito*
- *Zavorre dei partiti*, *pesi dei partiti*
- "Riunioni per pochi", "senza l'appoggio dei partiti"

Frame: Political parties are distant from citizens, operate in closed spaces.
Examples:

- "Basta logiche di partito nei palazzi"
- "Zavorre dei partiti che bloccano il rinnovamento"
- "Senza l'appoggio dei partiti, direttamente con i cittadini"

Institutional Neglect/Distance Lexical markers about political/institutional actors:

- *Indifferenza*, *assenza*, *lontananza*, *distanza*
- *Dimenticata/o*, *trascurata/o* (when referring to **political actors/institutions**, not just resources)
- *Istituzioni lontane/assenti/indifferenti*

Frame: Political/institutional actors have been absent or indifferent to citizens' needs.
Examples:

- "L'indifferenza di chi ha governato"
- "Istituzioni lontane dai bisogni reali"
- "La Laguna dimenticata per vent'anni"
- "Assenza dell'amministrazione nei quartieri"

Cross-Measure Note

If neglect is framed with **implied people-centrism** (e.g., "istituzioni lontane; noi ripristineremo presenza"), this also triggers **PC Soft** via cross-measure rule.

Hyperbolic Institutional Delegitimation Lexical markers:

- Intensifiers: *solo*, *sempre*, *mai*, *unicamente*, *nulla se non*
- Pattern: [Institution/Council] + intensifier + negative action

Frame: Institutions act in systematically biased ways (hyperbolic claim).
Examples:

- "Il Consiglio si riunisce solo per alzare le tasse"
- "La Giunta non fa mai ciò che serve"
- "Sempre le stesse decisioni a favore dei soliti"

Critical Blocker

If hyperbole appears WITH hard terms (*casta*, *corruzione*, *clientelismo*, "privilegi per pochi", *poteri forti*) ⇒ Escalate to **Hard**.

If any soft anti-elite element is present AND no hard rhetoric ⇒ Code as **Soft**
If no soft or hard anti-elite rhetoric ⇒ Code as **None**

S.21.4 Step 3: Default to None

If none of the above apply, code as **None**.

None includes:

- Technical administration language
- Resource management frames
- Neutral policy proposals
- Problems framed as technical/resource-based, NOT caused by political actors

Semantic Field for None Technical terms and neutral verbs:

- *Sportello*, *KPI*, *cronoprogramma*, *capitolato*, *SUAP*, *VAS*, *gare*
- *Implementare*, *semplificare*, *digitalizzare*, *monitorare*
- *Risorse*, *finanziamenti*, *investimenti* (without agent blame)
- Simple reductions: "Riduciamo l'indennità del 10%" (no moral frame, no instruments)

S.22 Quick Reference Table

If you see...	Then...
"Casta" + privileges/enrichment	Hard
Corruption/"si intascano"/"rubano"	Hard
Clientelism/"amici degli amici"	Hard
Betrayal/"sistema marcio"	Hard
"Poteri forti" conditioning politics	Hard
"Privilegi/vantaggi per pochi"	Hard
"Politics polluted by interests of few"	Hard
"Technocrats rig/manipulate"	Hard
"Logiche di partito nei palazzi"	Soft
"Zavorre dei partiti"	Soft
"Istituzioni lontane/indifferenti"	Soft
"Dimenticata/trascurata" (political blame)	Soft
"Council meets only to raise taxes" (no hard)	Soft
Technical terms (KPI, sportello, gare)	None
"Needs investments/resources"	None
Simple salary reduction (no moral frame)	None
Neutral policy language	None

S.23 Coding Workflow Summary

1. **Check readability** ⇒ OCR noise? ⇒ `OCR_Noise`
2. **Check hard anti-elite rhetoric** ⇒ Any present? ⇒ **Hard**
3. **Check soft anti-elite rhetoric** ⇒ Any present? ⇒ **Soft**
4. **Default** ⇒ **None**

S.24 Examples in Context

S.24.1 None Examples

- "Attiveremo sportello unico digitale con KPI" (activate single digital desk with KPIs)
- "La rete idrica necessita di investimenti" (water network needs investments)
- "Pubblicazione trimestrale di bilanci in albo online" (quarterly publication of budgets in online register)
- "Riduciamo l'indennità del 15%" (reduce compensation by 15% — neutral, no accusation)

- "Il porto è stato trascurato: piano lavori e gare entro 60 giorni" (port neglected: work plan and tenders within 60 days — resource issue, no political blame)

S.24.2 Soft Examples

Anti-Party Distance:

- "Basta logiche di partito nei palazzi" (enough party logic in the palaces)
- "Zavorre dei partiti che bloccano il rinnovamento" (burdens of parties that block renewal)
- "Senza l'appoggio dei partiti, direttamente con i cittadini" (without party support, directly with citizens)

Institutional Neglect/Distance:

- "L'indifferenza di chi ha governato" (the indifference of those who governed)
- "Istituzioni lontane dai bisogni reali" (institutions distant from real needs)
- "La Laguna dimenticata per vent'anni" (Lagoon forgotten for twenty years)
- "Assenza dell'amministrazione nei quartieri" (absence of administration in neighborhoods)

Hyperbolic Institutional Delegitimation:

- "Il Consiglio si riunisce solo per alzare le tasse" (Council meets only to raise taxes)
- "La Giunta non fa mai ciò che serve" (Executive never does what's needed)
- "Sempre le stesse decisioni a favore dei soliti" (always same decisions in favor of the usual ones)

S.24.3 Hard Examples

- "La casta si è arricchita coi soldi pubblici" (the caste enriched itself with public money)
- "Appalti truccati per gli amici degli amici" (rigged contracts for friends of friends)
- "Hanno tradito i cittadini" (they betrayed the citizens)
- "Politica condizionata dagli interessi di pochi" (politics conditioned by interests of the few)
- "Senza compromessi con i poteri forti" (without compromises with strong powers)
- "Decisioni per privilegi di pochi" (decisions for privileges of the few)

S.25 People-Centrism (PC)

S.26 What Are We Measuring?

People-Centrism (PC) measures whether a text frames **ordinary citizens as political agents** who should influence decisions or exercise oversight over **local/municipal government**.

This is **NOT** about:

- Citizens as beneficiaries ("services for citizens")
- Citizens as customers ("better information at URP offices")
- National/parliamentary politics
- Workplace democracy

S.27 The Three Codes

Code	Meaning	Example
None	No citizen agency in municipal governance	Technical content, service delivery, or no governance connection
Soft	Citizen agency rhetoric OR non-binding consultation	"Citizens will be involved in decisions," public consultations, participation promises
Hard	Strong participatory mechanisms	Participatory budgets, binding referendums, citizen oversight bodies, recall mechanisms

S.28 Coding Decision Tree

S.28.1 Step 0: Is the text readable?

- **If OCR noise/unreadable** ⇒ Code as `OCR_Noise`
- **If readable** ⇒ Continue to Step 1

S.28.2 Step 1: Is this about municipal governance?

Ask: Does the paragraph discuss **local/municipal institutions and decisions**?

✓ YES — Municipal governance includes:

- Municipal bodies: *comune*, *giunta*, *consiglio comunale*, *sindaco*, *assessori*

- Municipal decisions: local budget, public works, urban planning, municipal regulations
- Municipal services and administration when tied to governance/decisions

× **NO — Exclude these topics (always None):**

- **National politics**: parliament, national government, deputies, senators, national reforms
- **Workplace governance**: worker control of firms, union democracy, labor contracts
- **Service delivery only**: URP offices, website information, opening hours

Special Rule for Mixed Content

If a paragraph mentions both national and municipal governance, code **ONLY** the municipal parts and ignore national references.

S.28.3 Step 2: Does it have BOTH required elements?

For any score above None (i.e., Soft or Hard), you need **BOTH** of these:

Required Element A: Agency/Participation Cue The text must show citizens as **actors/agents** using words like:

- *partecipazione*, *coinvolgimento*, *confronto*
- "Citizens must count/matter in decisions"
- "Bring citizens closer to institutions/politics"
- "The city belongs to those who live in it"
- "Sovereignty belongs to citizens" (only with municipal anchor)
- **"Public consultation"** or **"citizen consultation"**

Important Note on *consultazione*

Simple *consultazione* only counts if explicitly public/citizen-focused:

- ✓ "consultazione pubblica", "consultazione dei cittadini"
- × "consultazione con i tecnici" (technical consultation doesn't count)

Required Element B: Municipal Governance Anchor The text must mention **local government institutions or decisions**:

- Institutions: *comune*, *giunta*, *consiglio*, *sindaco*, *assessori*, *uffici comunali*

- Decision objects: municipal budget, public contracts, urban plans, local regulations

If BOTH elements are present ⇒ Continue to Step 3
If either element is missing ⇒ Code as **None**

S.28.4 Step 3: Are there strong participatory mechanisms?

Look for any **Hard trigger** (if found, code as Hard immediately):

Hard Trigger #1: Participatory Budget (Automatic)

- Any mention of "*bilancio partecipativo*" or "*budget partecipativo*" ⇒ **Hard**
- This always overrides everything else

Hard Trigger #2: Binding Citizen Power Words showing citizens have **determinant/obligatory** power:

- *vincolante*, *obbligatorio*, *decisivo*, *determinante*
- "No decision without citizens/referendum/assembly"
- "Only with citizen approval"
- *mandato imperativo*, *revoca*, *recall* (of municipal officials)

Hard Trigger #3: Institutionalized Public Oversight Requires **ALL THREE** components:

1. **Control vocabulary**: *controllo*, *monitoraggio*, *vigilanza*, *verifica*, *audit*, *rendicontazione*, *valutazione*
2. **By the public**: citizens/population/community OR "controllo pubblico/civico/democratico"
3. **Of municipal governance**: budget, spending, contracts, public works, municipal acts/regulations

Don't Confuse With

"Controllo del territorio" (public order/policing) is **NOT** about accountability.

Hard Trigger #4: Transparency for Citizen Control

- **Hard** if: "publication/transparency **to enable citizens to verify/control**"
- **Soft** if: transparency about budget/contracts but without explicit citizen-control purpose
- **None** if: transparency only about services (schedules, contact info)

Hard Trigger #5: Anti-Mediation Language Explicit rejection of intermediation:

- "Only we represent the people"
- "No more mediation... the people decide directly"
- Must be about municipal representation

If ANY Hard trigger is present ⇒ Code as **Hard**
If no Hard triggers ⇒ Code as **Soft**

S.29 Common Pitfalls (Always None)

These do **NOT** count as people-centrism, even if they sound participatory:

× **Service Delivery**

- URP offices, municipal website, information desks
- "Informing/communicating with citizens"
- Customer service improvements

× **Beneficiary Framing**

- "In favor of citizens"
- "Putting citizens at the center"
- "Respecting citizen needs"
- These frame citizens as recipients, not agents

× **Social Volunteering**

- Community projects, volunteering initiatives
- Social *protagonismo* without governance connection
- Co-production of services (unless tied to governance decisions)

× **Corporatist Consultation**

- Consultation with professional categories, firms, expert orders
- Unless explicitly framed as public/citizen consultation

S.30 Special Cases

S.30.1 Subgroup Participation

Participation limited to specific groups (youth, disabled, elderly) is typically **None** UNLESS:

1. It concerns governance decisions/policies affecting that subgroup, AND
2. Explicit agency language is present

Examples:

- ✓ Soft: "Nothing about us without us" + decisions on municipal accessibility policies
- ✓ Soft: Youth council for municipal youth policies
- × None: Youth council for recreational events
- × None: Business association involvement in contracts

S.31 Quick Reference Table

If you see...	**Then...**
Only service delivery language	None
National/parliamentary politics	None
Citizens as beneficiaries only	None
Participation language + municipal anchor	Check for Hard triggers
"Bilancio partecipativo"	Hard automatically
Binding referendum on municipal works	Hard
Citizen control/monitoring of budget	Hard (if all 3 elements present)
Public consultation on urban plan	Soft (if no Hard triggers)
"Involving citizens in decisions"	Soft (if no Hard triggers)
Technical consultation with experts	None

S.32 Coding Workflow Summary

1. **Check readability** ⇒ OCR noise? Stop.
2. **Check domain** ⇒ Not municipal? ⇒ None
3. **Check Step-1 gate** ⇒ Missing agency OR municipal anchor? ⇒ None
4. **Check Hard triggers** ⇒ Any present? ⇒ Hard
5. **Default** ⇒ Soft

S.33 Examples in Context

S.33.1 None Examples

- “Il comune migliorerà i servizi di informazione online” (service delivery only)
- “Chiederemo al parlamento di riformare la giustizia” (national politics)
- “Servizi a favore dei cittadini” (beneficiary framing)

S.33.2 Soft Examples

- “Coinvolgeremo i cittadini nelle decisioni urbanistiche” (participation promise)
- “Consultazione pubblica sul piano regolatore” (non-binding consultation)
- “I cittadini devono contare nelle scelte del comune” (agency rhetoric)

S.33.3 Hard Examples

- “Introdurremo il bilancio partecipativo” (participatory budget)
- “Referendum vincolante sulle opere pubbliche” (binding mechanism)
- “Controllo pubblico della spesa comunale attraverso assemblee cittadine” (institutionalized oversight)
- “Pubblicazione degli appalti per permettere ai cittadini di verificare” (transparency-as-control)