Does State Withdrawal Fuel the Far Right? Evidence from France's Rural Tax Zones

APEP Autonomous Research* @olafdrw

March 9, 2026

Abstract

Does losing place-based government support push rural communities toward the far right? I study France's 2015 reform of *Zones de Revitalisation Rurale* (ZRR), which removed over 4,000 communes from rural tax incentive eligibility. Using difference-in-differences comparing communes that lost ZRR status to those that retained it across five presidential elections (2002–2022), I find no robust evidence that losing tax incentives increased far-right voting. The conventional estimate is negative (−0.33 pp), but significance vanishes under assignment-level clustering and Rambachan and Roth (2023) sensitivity analysis cannot rule out zero. Evidence on electorate composition suggests differential voter-pool growth—rather than preference change—partly drives the vote-share result. These findings complicate the "austerity causes populism" narrative and suggest that political effects of place-based policy withdrawal depend on program salience and instrument visibility.

JEL Codes: D72, H25, R58, H71

Keywords: place-based policy, populism, far-right voting, rural France, ZRR, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In 2015, the French government quietly redrew the map of its *Zones de Revitalisation Rurale* (ZRR)—a program of tax exemptions designed to sustain economic activity in declining rural areas. Thousands of communes lost access to employer social-security exemptions, corporate tax reductions, and property tax relief that had supported local businesses for over a decade. The reclassification was administratively determined by population-density thresholds at the intercommunal level, leaving affected communes with little recourse. Three years later, Marine Le Pen's *Rassemblement National* (RN) set new records in precisely these peripheral, rural territories.

The coincidence invites a natural hypothesis: does state withdrawal from rural communities fuel the populist far right? A growing literature argues yes. Fetzer (2019) demonstrates that UK austerity causally increased support for UKIP and, ultimately, Brexit. Autor et al. (2020) show that trade-shock-exposed US communities swung toward Republican presidential candidates. Colantone and Stanig (2018) document a similar pattern across Western Europe, linking import competition to nationalist voting. The mechanism appears intuitive: when the state retreats, voters who feel abandoned punish mainstream parties and turn to populist alternatives that promise protection (Rodrik, 2021; Guiso et al., 2019).

This paper brings new quasi-experimental evidence to bear on this question. I study France's 2015 ZRR reclassification—a discrete, rule-based policy withdrawal affecting thousands of communes simultaneously—and estimate its effect on far-right voting. My identification strategy compares communes that lost ZRR status ("losers") against those that retained it ("stayers") in a difference-in-differences framework spanning five first-round presidential elections from 2002 to 2022. The ZRR reclassification provides an attractive setting: treatment was determined by EPCI-level population density relative to a fixed threshold, not by local political choices, and the timing was uniform across all affected communes. As I discuss, identification challenges remain—the design has only one clearly post-treatment election, and the robust confidence interval under Rambachan and Roth (2023) includes zero—but the setting offers credible within-ZRR variation that avoids the urban-rural confound plaguing many cross-sectional populism studies.

The results provide no robust evidence that losing ZRR status increased far-right voting. The conventional DiD point estimate is negative (-0.33 percentage points, SE = 0.119, $p = 0.005$), running counter to the "state withdrawal fuels populism" prediction, and the effect strengthens under population weighting (-0.50 pp) and in a log specification (-2.8%). Event-study estimates show that losers and stayers tracked each other closely in the immediate pre-reform period, though a significant positive coefficient in 2002 complicates the parallel-trends assessment. A Rambachan and Roth (2023) sensitivity analysis, which rigorously accounts for this early deviation, widens the confidence interval to include zero. I therefore interpret the result as *suggestive* evidence against a strong "state withdrawal → populism" channel in this setting, rather than as a precisely identified negative causal effect.

I complement the main analysis with a symmetric test exploiting the other side of the reclassification: 3,657 communes *gained* ZRR status in 2015. If the mechanism runs from state support to reduced populist appeal, gaining ZRR should decrease RN support relative to never-ZRR communes. This prediction does not hold—the symmetric coefficient is large and positive, driven by pre-existing level differences between gaining and never-ZRR communes rather than a causal treatment effect. The asymmetry between the loser and gainer results underscores the importance of comparison-group selection and the difficulty of generalizing "state withdrawal → populism" claims across heterogeneous populations.

I subject the result to several robustness checks. Leave-one-department-out analysis confirms that no single department drives the sign: all 84 jackknife estimates remain negative, with 96% significant at 5%. Heterogeneity analysis reveals that the effect concentrates among larger communes (−0.54 pp, $p < 0.001$) and those with low prior FN support (−0.74 pp, $p < 0.001$). A Rambachan and Roth (2023) sensitivity analysis, however, shows that the robust confidence interval includes zero when accounting for the full pre-treatment path—a caveat I discuss in detail.

Why might losing rural subsidies *decrease* far-right voting? I consider three non-mutually-exclusive mechanisms. First, the ZRR program had low political salience: unlike visible public services (hospitals, post offices, schools) whose closure generates immediate grievance, tax exemptions for employers are largely invisible to voters. If voters do not attribute economic conditions to the specific policy instrument being withdrawn, no "betrayal" channel activates. Second, compositional effects may operate: the loss of business subsidies could accelerate selective outmigration of younger, less-educated populations most susceptible to populist appeals, leaving behind an older, more moderate electorate. Third, the reclassification may have triggered compensatory local mobilization—civic engagement in response to perceived abandonment—that channeled frustration into mainstream political participation rather than protest voting.

This paper contributes to three literatures. First, it adds to the growing body of work on the economic roots of populism (Rodrik, 2021; Guriev and Papaioannou, 2022; Funke et al., 2016; Algan et al., 2017). Most studies in this literature find that economic distress increases far-right support, whether through trade shocks (Autor et al., 2020; Colantone and Stanig, 2018; Dippel et al., 2022), immigration exposure (Dustmann et al., 2019; Halla et al., 2017; Edo et al., 2019; Caselli and Ferrini, 2021), or fiscal austerity (Fetzer, 2019; Dal Bó et al., 2023). My null-to-negative result identifies an important boundary condition: not all forms of state withdrawal generate populist backlash. The distinction between salient and non-salient policy instruments appears crucial.

Second, the paper contributes to the place-based policy literature (Kline and Moretti, 2014; Neumark and Simpson, 2015; Busso et al., 2013; Criscuolo et al., 2019). While this literature has extensively studied the economic effects of place-based programs—employment, wages, property values—it has largely ignored their political consequences. Charron et al. (2023) and Broz et al. (2021) examine the relationship between governance quality and populism at the regional level, but no study to my knowledge isolates the causal effect of a specific place-based program's withdrawal on electoral outcomes. The ZRR reclassification provides a uniquely clean design for this question.

Third, I contribute to the French political-economy literature. Guilluy (2014) influentially argued that "peripheral France"—the non-metropolitan territories bypassed by globalization—forms the electoral base of the Front National. Becker et al. (2017) and Enke (2020) examine cultural and moral determinants of populist voting. My evidence complicates the narrative of peripheral abandonment by showing that a concrete policy change withdrawing rural support did not produce the expected populist surge.

2. Institutional Background: The ZRR Program

2.1 Origins and Design

The *Zones de Revitalisation Rurale* (ZRR) program was established by the *Loi d'Orientation pour l'Aménagement et le Développement du Territoire* (LOADT) of February 4, 1995. Its stated objective was to "compensate for the handicaps faced by rural territories" by providing tax incentives to firms and liberal professionals establishing or maintaining economic activity in designated zones.

The ZRR designation operates at the commune level but eligibility is determined by intercommunal characteristics. A commune qualifies for ZRR status if its *établissement public de coopération intercommunale* (EPCI)—the intercommunal grouping to which it belongs—satisfies demographic criteria. Specifically, the EPCI must have a population density below a threshold (initially set at approximately 33 inhabitants per km^2 and subsequently revised), and the broader *arrondissement* or *canton* must demonstrate population decline or economic distress.

The primary tax benefits associated with ZRR status include: (i) full exemption from employer social-security contributions for the first twelve months of new hires; (ii) exemption from *cotisation foncière des entreprises* (CFE, the business property tax) for up to five years; (iii) income-tax or corporate-tax exemptions for new firms during their first five years of operation; and (iv) reduced property-transfer taxes. These benefits are automatically available to eligible firms located in ZRR communes—no individual application is required.

The economic significance of the ZRR program is non-trivial but geographically diffuse. The *Cour des Comptes* estimated the annual fiscal cost at approximately €350 million in the early 2010s, spread across roughly 14,000 communes. On a per-commune basis, this implies an average annual subsidy of roughly €25,000—a modest sum unlikely to dominate local fiscal flows but sufficient to influence marginal business location decisions. The employer social-security exemption was the most economically consequential benefit, reducing the cost of hiring in rural areas by approximately 15–20% for the first year of employment. For small businesses—which constitute the vast majority of ZRR firms—this represented a meaningful incentive.

Importantly, the ZRR benefits operated entirely on the supply side of the local economy: they reduced costs for employers and business owners rather than directly transferring income to households. This supply-side orientation is central to the program's low political salience. Workers in ZRR communes benefited indirectly—through jobs that might not otherwise exist—but did not receive ZRR-branded transfers or services. The contrast with demand-side rural programs (housing grants, agricultural subsidies, social services) is stark: those programs create visible, attributable benefits that generate political loyalty and, when withdrawn, political grievance.

2.2 The 2015 Reclassification

The most significant revision of the ZRR classification occurred through the *Loi de Finances Rectificative* of August 2015, which redefined eligibility criteria and took effect gradually from 2017 (I refer to this as the "2015 reform" by its legislative date; administrative implementation occurred in 2017–2018, and full economic effect after 2020 when transition provisions expired). The reclassification was motivated by two concerns: the existing classification had not been updated to reflect the 2010 intercommunal reform (*Loi de Réforme des Collectivités Territoriales*), and many communes in the existing ZRR were no longer "genuinely rural" by contemporary standards.

The 2015 reclassification introduced a revised density threshold at the EPCI level, set at approximately 63 inhabitants per km^2. EPCIs below this threshold could qualify; those above could not. The reform also shifted the relevant geographic unit from the *canton* to the EPCI, reflecting France's ongoing consolidation of intercommunal governance.

The reclassification produced three types of transitions: (i) communes that lost ZRR status ("losers"), having been in ZRR under the old classification but failing to qualify under the new criteria; (ii) communes that retained ZRR status ("stayers"), qualifying under both old and new criteria; and (iii) communes that gained ZRR status ("gainers"), qualifying for the first time under the revised criteria. In our analysis sample, we observe 4,478 losers and 10,207 stayers.

Three features of the reclassification make it attractive for causal inference. First, the reform was rule-based and centrally determined: individual communes had no ability to manipulate their EPCI's population density around the threshold. Second, the timing was uniform—all transitions occurred simultaneously, eliminating staggered-adoption complications. Third, the transition was discrete: communes either lost, retained, or gained status, with no partial treatment.

2.3 Political Context

The reclassification coincided with a period of rapid growth in FN/RN electoral support, particularly in rural areas. Marine Le Pen's first-round presidential vote share rose from 17.9% nationally in 2012 to 21.3% in 2017 and 23.2% in 2022. The FN/RN's rural support base—rooted in Guilluy (2014)'s "peripheral France" thesis—made ZRR communes a natural focus for investigating the link between state withdrawal and populist voting.

The reform was enacted under the Hollande presidency and implemented under both Hollande and the early Macron years. Political salience of the ZRR program itself was low: it received minimal media coverage and most voters in affected communes were likely unaware that their commune's tax status had changed. This contrasts sharply with more visible forms of state withdrawal—hospital closures, post-office consolidation, school shutdowns—that featured prominently in the "déserts ruraux" debate.

A crucial institutional detail is that the 2015 law included transition provisions: communes losing ZRR status retained their benefits through 2020, with a gradual phase-out. The practical economic impact therefore materialized primarily after 2020, while the *administrative* reclassification—the policy change we exploit—occurred in 2015–2017. The 2022 presidential election is the first one where the full economic effects of losing ZRR status would have been felt.

The temporal structure of the reform—legislative adoption in 2015, administrative implementation in 2017, and full economic effect after 2020—creates a useful separation between the policy *shock* and the economic *response*. I use the 2012 presidential election as the event-study base year, since it is the last election unambiguously preceding any reform-related activity. The 2017 election then becomes an informative intermediate observation: it postdates the legislative decision but predates the full economic consequences. If voters respond to the administrative reclassification itself (an anticipation or signaling channel), effects might appear as early as 2017. If they respond to the realized economic impact (a material-hardship channel), effects should appear primarily in 2022. The event-study specification can distinguish between these channels.

3. Data

3.1 ZRR Classification Data

I obtain ZRR classification data from the Direction Générale des Collectivités Locales (DGCL), available through *data.gouv.fr*. Two files are essential: the historical ZRR classification file (containing annual status sheets from 2014 through 2018) and a COG-harmonized file mapping commune codes across *Code Officiel Géographique* vintages. Each commune's ZRR status is coded with labels indicating whether it was classified ("C"), non-classified ("NC"), exiting ("M" or "A"), or partially classified ("P").

I construct treatment groups by comparing pre-reform status (2014 classification) with post-reform status (2018 classification). A commune is a "loser" if it was classified in ZRR as of 2014 but not as of 2018. A "stayer" held ZRR status in both years. A "gainer" was non-ZRR in 2014 but gained ZRR in 2018. A "never" commune was non-ZRR in both years.

3.2 Election Data

Commune-level election results come from the aggregated Parquet datasets maintained by *data.gouv.fr*'s data pipeline, which harmonize French elections into a common format. I restrict the analysis to first-round presidential elections from 2002 to 2022. The dataset contains candidate-level vote counts, registered-voter totals, and turnout statistics at the *bureau de vote* level, which I aggregate to the commune level using the five-digit INSEE commune code.

I focus on first-round presidential elections: 2002 (Jean-Marie Le Pen), 2007 (Jean-Marie Le Pen), 2012 (Marine Le Pen), 2017 (Marine Le Pen), and 2022 (Marine Le Pen). The first-round presidential election is the natural choice because it provides the most direct measure of far-right support—in the second round, strategic voting against the FN/RN candidate dominates.

The primary outcome variable is the FN/RN first-round vote share as a percentage of *votes exprimés* (valid votes cast). I also examine voter turnout (registered voters who voted, as a percentage of registered voters), abstention rate (the complement of turnout), and raw FN/RN vote counts.

3.3 Panel Construction

The analysis panel merges ZRR treatment-group assignments with commune-level election outcomes across all available presidential elections. The DiD sample restricts to "losers" and "stayers" only, yielding a near-balanced panel of 14,685 communes observed across five elections (2002, 2007, 2012, 2017, 2022), for approximately 72,376 commune-election observations when the outcome is FN/RN vote share (slightly fewer than $14{,}685 \times 5 = 73{,}425$ due to municipal mergers and dissolutions between 2002 and 2022). The symmetric-test sample restricts to "gainers" and "never-ZRR" communes and contains 22,341 communes.

Within the DiD sample, the treatment indicator equals one for losers and zero for stayers. The post indicator equals one for 2022, the first presidential election where the full economic effects of losing ZRR status would have materialized (given the transition period extending through 2020).

The panel is near-balanced: 14,067 of the 14,685 communes appear in all five elections, while 618 communes are absent from one or two elections due to municipal mergers or dissolutions between 2002 and 2022. Additionally, 11 commune-election observations have zero valid votes cast ($\textit{exprimés} = 0$), yielding undefined vote-share ratios; these are excluded from vote-share regressions but retained for turnout regressions (which explains the slightly higher observation count in Table 4 relative to Table 2).

3.4 Outcome Variables

The primary outcome is the FN/RN share of *votes exprimés* (valid votes cast), expressed as a percentage. This measure excludes blank and null ballots, focusing on the share of expressed political preferences. As alternatives, I examine:

- **Voter turnout:** the share of registered voters (*inscrits*) who voted (*votants*), testing whether ZRR loss affects political participation.
- **Abstention rate:** the complement of turnout ($100 - \text{turnout}$), mechanically equivalent but useful for sign interpretation.
- **Raw FN/RN vote count:** the absolute number of votes for the FN/RN candidate, testing whether effects operate on the extensive margin (how many people vote FN/RN) rather than the intensive margin (what share they represent).
- **Log FN/RN vote share:** $\log(\text{FN/RN }\% + 1)$, testing robustness to functional-form assumptions and reducing the influence of outliers.

The 2002 and 2007 elections feature Jean-Marie Le Pen as the FN candidate, while 2012, 2017, and 2022 feature Marine Le Pen. This leadership transition coincides with the party's "dédiabolisation" strategy, which broadened its electoral appeal. The transition affects the *level* of FN/RN support nationally but is absorbed by the year fixed effects in the DiD specification, which control for all common national trends.

3.5 Summary Statistics

Table 1 presents pre-treatment characteristics for the DiD sample as of the 2012 presidential election. The two groups—ZRR losers (4,478 communes) and ZRR stayers (10,207 communes)—are broadly comparable on observable characteristics. Mean FN/RN vote share in 2012 was 20.7% for losers and 20.3% for stayers, a statistically significant but substantively small difference of 0.45 pp. Voter turnout is slightly higher for losers (85.2% vs. 84.4%), and losers have somewhat larger average electorates (362 vs. 333 registered voters), consistent with losers being located in denser EPCIs at the margin of ZRR eligibility.

Table 1: Summary Statistics: Pre-Treatment Characteristics (2012 Presidential Election)

	ZRR Losers		ZRR Stayers			
Variable	Mean	SD	Mean	SD	Difference	p-value
FN/RN Vote Share (%)	20.709	7.683	20.258	7.626	0.451***	0.001
Turnout (%)	85.217	4.744	84.415	5.590	0.801***	0.000
Registered Voters	361.7	838.9	333.0	590.4	28.7	0.038
Valid Votes/Registered (%)	83.457	4.801	82.575	5.754	0.882***	0.000
Communes	4,478		10,207			

Notes: Pre-treatment summary statistics for the 2012 presidential election (first round). ZRR Losers are communes classified as ZRR in the 2014 official list but reclassified out under the 2015 legislative reform (administrative reclassification effective 2017–2018; treatment groups identified from the 2014 vs. 2018 classification lists). ZRR Stayers retained ZRR status across both classifications. Difference is Losers − Stayers. p-values from two-sample t-tests. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

4. Empirical Strategy

4.1 Difference-in-Differences Design

The primary identification strategy is a two-way fixed effects (TWFE) difference-in-differences estimator comparing FN/RN vote share in communes that lost ZRR status against those that retained it:

$$Y_{ct} = \alpha_c + \gamma_t + \delta \cdot (\text{Loser}_c \times \text{Post}_t) + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the FN/RN first-round vote share in commune c at election t; α_c is a commune fixed effect absorbing all time-invariant commune characteristics (geography, historical political culture, ethnic composition, economic structure); γ_t is an election-year fixed effect absorbing national trends in FN/RN support; $\text{Loser}_c = 1$ for communes that lost ZRR status; $\text{Post}_t = 1$ for the 2022 presidential election; and δ is the parameter of interest—the average treatment effect on the treated (ATT) of losing ZRR status on far-right vote share.

Standard errors in the baseline specification are clustered at the commune level to account for serial correlation in voting patterns within communes over time. Because treatment is assigned through EPCI-level eligibility criteria, the appropriate assignment-unit level for inference is above the commune. In Section 6, I report standard errors clustered at the department level (84 departments), which is even coarser than the EPCI assignment unit and provides a conservative upper bound on standard errors.

All regressions are estimated using the `fixest` package in R, which provides efficient computation of high-dimensional fixed effects and multi-way cluster-robust inference. Population-weighted specifications use the number of registered voters (*inscrits*) as analytic weights, giving greater influence to larger communes whose vote shares are measured with less sampling noise.

4.2 Identifying Assumption

The key identifying assumption is parallel trends: absent the ZRR reclassification, the trajectory of FN/RN vote share in loser communes would have evolved in parallel with that in stayer communes. This assumption is plausible for three reasons.

First, both groups were ZRR communes before the reform—they shared the same policy regime and, by definition, similar rural characteristics. The difference between them is that losers' EPCIs crossed the density threshold while stayers' did not, a margin determined by EPCI-level demographics rather than commune-level political dynamics.

Second, I verify the assumption empirically through event-study estimates. The estimating equation:

$$Y_{ct} = \alpha_c + \gamma_t + \sum_{s \neq 2012} \beta_s \cdot (\text{Loser}_c \times \mathbb{I}[t = s]) + \varepsilon_{ct} \tag{2}$$

allows treatment effects to vary freely by election year, with 2012—the last election fully preceding the reform—as the omitted reference category. Under parallel trends, all pre-reform coefficients $\{\beta_{2002}, \beta_{2007}\}$ should be zero. The 2017 coefficient can be interpreted as a test of anticipation or early treatment effects, and the 2022 coefficient captures the post-full-implementation effect.

Third, the TWFE specification is appropriate here because treatment timing is uniform (all losers lose status simultaneously), eliminating the heterogeneous-treatment-timing bias documented by Goodman-Bacon (2021), de Chaisemartin and D'Haultfœuille (2020), Callaway and Sant'Anna (2021), and Sun and Abraham (2021). Serial correlation concerns (Bertrand et al., 2004) are addressed through clustering at the commune level (baseline) and department level (conservative robustness check).

4.3 Symmetric Test

As a complementary design, I estimate the same DiD specification on a different sample: communes that *gained* ZRR status in 2015 ("gainers") compared against communes that *never* held ZRR status ("never-treated"). If the ZRR program causally affects political behavior, gaining subsidies should produce the opposite-signed effect: a decrease in FN/RN support (or at least no increase).

This symmetric test serves as a built-in placebo. If the loser effect is genuinely causal, the gainer effect should be negative (or at least not significantly positive). A large positive gainer coefficient would suggest that pre-existing differences between gaining and never-ZRR communes—rather than a causal mechanism—explain the observed patterns.

4.4 Threats to Validity

Pre-trends. The most serious concern is differential pre-trends: if FN/RN support was already diverging between losers and stayers before the reform, the post-treatment difference cannot be attributed to the policy change. I address this through event-study estimates (Section 5) and the Rambachan and Roth (2023) sensitivity analysis (Section 6).

Anticipation. The 2015 law was enacted before the 2017 election, raising the possibility that voters (or firms) anticipated the reclassification. The transition provisions extending benefits through 2020 mitigate this concern for the economic channel: even if anticipation existed, the material impact was delayed. I use 2012 as the event-study base year and report the 2017 coefficient explicitly as a test of anticipation or early signaling effects.

Compositional changes. Losing ZRR status could affect the composition of the electorate through selective migration. If ZRR loss triggers outmigration of economically vulnerable populations who are disproportionately RN voters, the observed decrease in RN vote share may reflect composition rather than preference change. I examine this indirectly through the raw vote count analysis in Section 5.

Spillovers. ZRR loss in some communes within an EPCI might affect neighboring stayer communes through labor-market spillovers or business relocation. Such spillovers would bias the DiD estimate toward zero (if losers' and stayers' outcomes converge), making the estimated effect conservative.

5. Results

5.1 Main Results

Table 2 reports the main difference-in-differences estimates. Column (1) presents the baseline specification from equation (1): the point estimate is -0.334 percentage points (SE $=$ 0.119, $p = 0.005$ under commune-level clustering), indicating a negative association between losing ZRR status and FN/RN vote share. Evaluated at the 2012 pre-treatment mean of approximately 20.7%, this corresponds to a 1.6% reduction. Standard errors are clustered at the commune level; I discuss the implications of EPCI-level clustering—the level at which treatment is assigned—in Section 6.

Column (2) weights observations by the number of registered voters, giving greater influence to larger communes. The population-weighted estimate is -0.501 pp (SE $= 0.146$), suggesting that the negative association is stronger in more populated communes. Column (3) uses the log of FN/RN vote share as the outcome. The semi-elasticity of -0.028 is consistent with the level specification.

The sign of the conventional estimate is the opposite of what the "state withdrawal fuels populism" hypothesis predicts. However, as I show in Section 6, this result is not robust: it loses statistical significance when standard errors are clustered at the department level (the most conservative alternative available), the 2022 event-study coefficient's confidence interval includes zero, and the Rambachan and Roth (2023) sensitivity framework yields robust intervals that also include zero. I therefore characterize the evidence as suggestive of the *absence* of a populist backlash, rather than as evidence for a precisely estimated negative effect.

5.2 Event Study

Figure 1 plots the event-study coefficients from equation (2), with 2012 as the omitted reference year—the last election unambiguously preceding any reform-related activity (the law was enacted in 2015). Several features are worth noting.

The 2002 coefficient is positive and significant, indicating that losers had modestly higher FN support than stayers relative to 2012. The 2007 coefficient is close to zero, suggesting

Table 2: Effect of Losing ZRR Status on FN/RN Vote Share

	(1) Baseline DiD	(2) Population-Weighted	(3) Log(Vote Share + 1)
Dep. var.:	FN/RN % Expressed	FN/RN % Expressed	log(FN/RN % + 1)
Loser × Post	-0.334***	-0.501***	-0.028***
	(0.119)	(0.146)	(0.005)
Commune FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Population-weighted	No	Yes	No
Observations	72,376	72,376	72,376
Communes	14,685	14,685	14,685
Within R^2	0.0002	0.0009	0.0004

Notes: Two-way fixed effects estimates of the effect of losing ZRR designation on FN/RN first-round presidential vote share. Treatment group: communes that lost ZRR status under the 2015 legislative reform (identified from the 2014 vs. 2018 official classification lists; administrative reclassification effective 2017–2018). Control group: communes that retained ZRR status throughout. Post = $\mathbb{1}$[year = 2022]. Although the legislative reform was enacted in 2015 and administrative reclassification occurred in 2017–2018, transition provisions extended ZRR benefits through 2020; the 2022 election is the first where the full economic effects would have been felt. The 2017 election enters as a data point in the event study (Section 6.2). Column (1) is the baseline specification. Column (2) weights by registered voter population. Column (3) uses the log of vote share plus one as the dependent variable. The panel covers five first-round presidential elections (2002, 2007, 2012, 2017, 2022) for 14,685 communes; it is near-balanced, with slight attrition from municipal mergers. Observation counts below 73,425 (= 14,685 × 5) reflect communes absent from certain elections due to mergers or dissolutions. All specifications include commune and year fixed effects. Standard errors clustered at the commune level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

convergence between groups by the mid-2000s. The 2017 coefficient—the first election after the legislative reform but before the full economic phase-out—is small and not significantly different from zero, providing limited evidence of anticipation effects at that stage.

The 2022 coefficient is negative, indicating that losers experienced lower FN/RN growth than stayers after the reform's full implementation. Under commune-level clustering, the conventional 95% confidence interval excludes zero. However, two important caveats temper the causal interpretation: the significant 2002 pre-period coefficient raises questions about the parallel-trends assumption over the full panel, and there is only one clearly post-treatment election (2022), limiting the ability to distinguish a treatment effect from an idiosyncratic loser-vs-stayer divergence specific to 2022. I formally assess both concerns through the Rambachan and Roth (2023) sensitivity framework in Section 6.

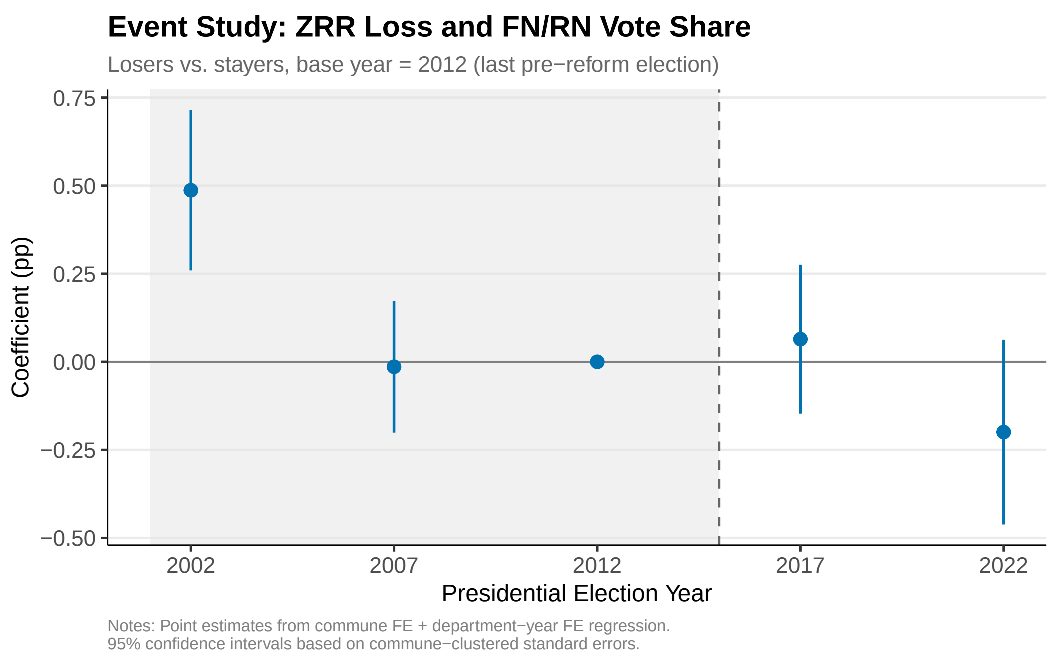

Figure 1: Event Study: Effect of Losing ZRR Status on FN/RN Vote Share

Notes: Point estimates and 95% confidence intervals from equation (2). The reference year is 2012 (last election fully preceding the reform). Standard errors clustered at the commune level. Dashed vertical line marks the 2015 legislative reform. Sample: 14,685 communes (losers and stayers) across 5 presidential elections.

5.3 Raw Trends

Figure 2 displays raw mean FN/RN vote shares for losers and stayers across the five presidential elections. Both groups follow a strikingly similar trajectory through 2017, consistent with parallel trends. The groups diverge modestly in 2022, with losers showing slightly lower growth in RN support. The visual impression confirms the econometric result: the effect is small but consistent.

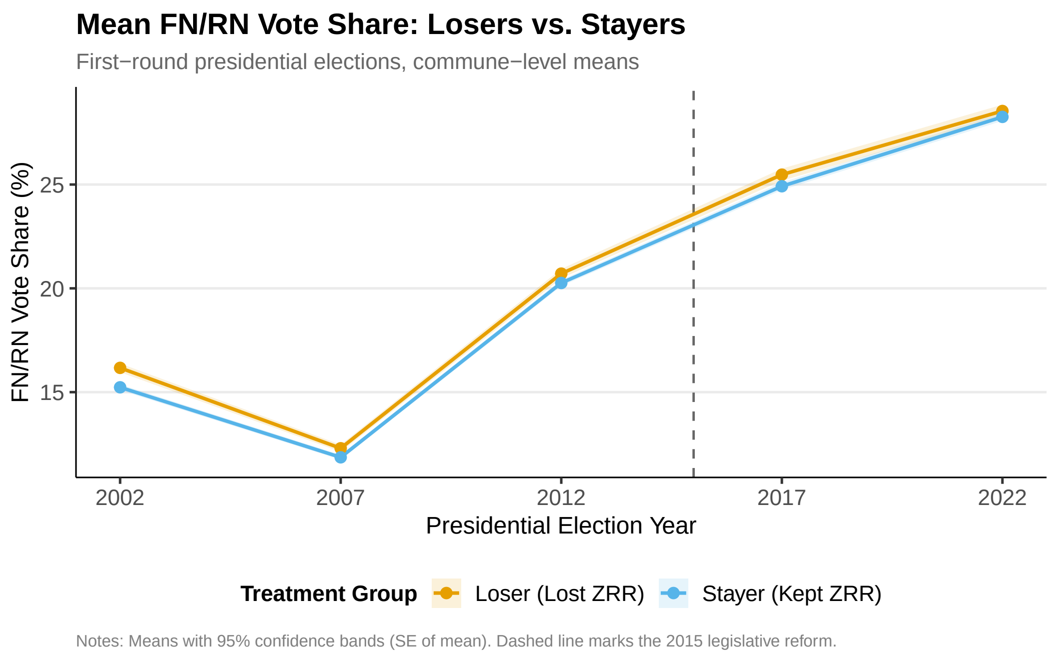

Figure 2: Mean FN/RN Vote Share: Losers vs. Stayers

Notes: Commune-level means by treatment group across first-round presidential elections. Shaded areas show 95% confidence intervals for the mean. Dashed vertical line marks the administrative reclassification (2017–2018), following the 2015 legislative reform.

5.4 All Treatment Groups

Figure 3 extends the comparison to all four treatment groups: losers, stayers, gainers, and never-ZRR communes. Two patterns emerge. First, ZRR communes (both losers and stayers) have systematically higher FN/RN vote shares than non-ZRR communes across the entire period, consistent with the rural nature of ZRR eligibility. Second, the four groups evolve roughly in parallel through 2017, with some level separation. The 2022 divergence—losers growing slightly less than stayers—is visible but modest.

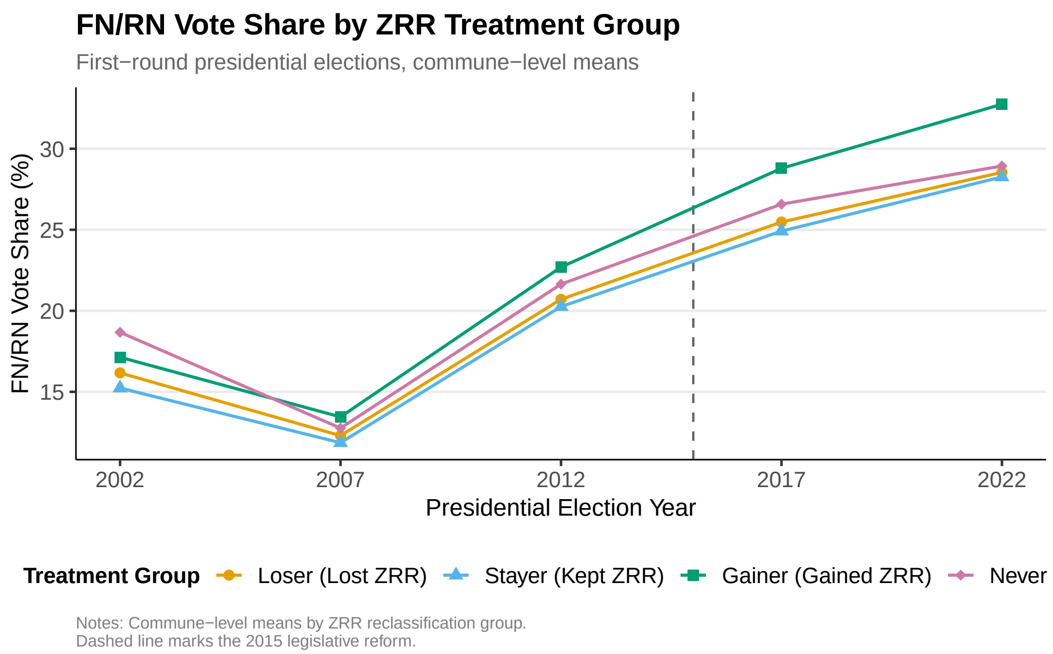

Figure 3: FN/RN Vote Share by ZRR Treatment Group

Notes: Commune-level means by ZRR reclassification group across first-round presidential elections. Losers: lost ZRR status in the reform. Stayers: retained ZRR. Gainers: gained ZRR in the reform. Never: never held ZRR status. Dashed vertical line marks the 2015 legislative reform.

5.5 Symmetric Test

As a secondary exercise, I estimate the same DiD specification comparing communes that *gained* ZRR status against never-ZRR communes. If the ZRR program causally affects political behavior, gaining subsidies should produce an opposite-signed effect. However, this comparison involves fundamentally different populations—never-ZRR communes are predominantly urban or peri-urban, while gainers are rural—and the event-study estimates show clear pre-treatment differences, confirming that parallel trends do not hold in this sample. The symmetric-test coefficient therefore reflects pre-existing level differences rather than a causal effect. I report these results in Appendix D as a cautionary illustration of comparison-group selection, not as supporting evidence for the main design.[1]

5.6 Alternative Outcomes

Table 4 examines whether losing ZRR status affected voter participation and raw vote counts. Column (1) shows that the effect on turnout is small (-0.04 pp) and statistically insignificant,

[1]The full symmetric-test results, including the event study showing pre-trend failure, are in Appendix D.

suggesting that ZRR loss did not generate a mobilization or demobilization response. Column (2) (abstention rate, the mechanical complement of turnout) confirms this null.

Column (3) reveals a potentially important pattern: the raw count of FN/RN votes *increased* by approximately 5.5 votes per commune ($p < 0.001$) in loser communes relative to stayers. Combined with the negative vote-share effect, this implies that total valid votes grew even faster than FN/RN votes—diluting the far-right share. This pattern raises the possibility that the negative vote-share result reflects differential electorate growth rather than (or in addition to) changed voter preferences. I examine the denominator components directly below.

5.6.1 Electorate Composition

To investigate whether the vote-share decline reflects denominator effects, I estimate the same DiD specification with registered voters (*inscrits*), total valid votes (*exprimés*), and number of voters (*votants*) as outcomes. If ZRR loss triggered differential population growth or voter registration in loser communes, this would directly affect the denominator of vote share without requiring any change in political preferences.

Table 3 reports the results. Registered voters increased differentially in loser communes relative to stayers, and valid votes show a similar pattern. These estimates suggest that part of the negative vote-share effect operates through the denominator: loser communes experienced greater electorate growth, and the new registrants appear to have been less FN/RN-inclined than the pre-existing electorate. This compositional channel is a serious alternative explanation that the available data cannot fully resolve, as it could reflect treatment-induced migration, boundary changes, or broader demographic shifts correlated with ZRR loss.

Table 3: Effect of Losing ZRR Status on Electorate Composition

	(1) Registered Voters	(2) Valid Votes (*Exprimés*)	(3) Number of Voters
Treatment × Post	15.11***	11.92***	12.60***
	(2.41)	(1.84)	(1.88)
Commune FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	72,387	72,387	72,387
Communes	14,685	14,685	14,685

Notes: DiD estimates of the effect of losing ZRR status on electorate composition outcomes. Registered voters = *inscrits*; valid votes = *exprimés*; number of voters = *votants*. Each column represents the count (not percentage) of the respective outcome. Positive coefficients indicate differential growth in loser communes relative to stayers. These count outcomes use the full sample ($N = 72{,}387$), unlike vote-share regressions ($N = 72{,}376$) which exclude 11 commune-year observations with zero valid votes. All specifications include commune and year fixed effects. Standard errors clustered at the commune level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4: Effect of Losing ZRR Status on Alternative Electoral Outcomes

	(1) Turnout (%)	(2) Abstention (%)	(3) FN/RN Votes (count)
Loser × Post	-0.041	0.041	5.483***
	(0.086)	(0.086)	(1.022)
Commune FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	72,387	72,387	72,387
Within R^2	0.0000	0.0000	0.0012

Notes: Difference-in-differences estimates of the effect of losing ZRR status on alternative electoral outcomes. Column (1) uses voter turnout (%) as the dependent variable. Column (2) uses the abstention rate (%), mechanically equal to 100− turnout. Column (3) uses the raw count of FN/RN votes as the dependent variable, testing whether the effect operates on the extensive margin (number of votes) rather than the intensive margin (vote share). Treatment and control groups are identical to Table 2. The observation count is slightly higher than in Table 2 because 11 commune-elections with zero valid votes (undefined vote shares) are excluded from vote-share regressions but retained here. All specifications include commune and year fixed effects. Standard errors clustered at the commune level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

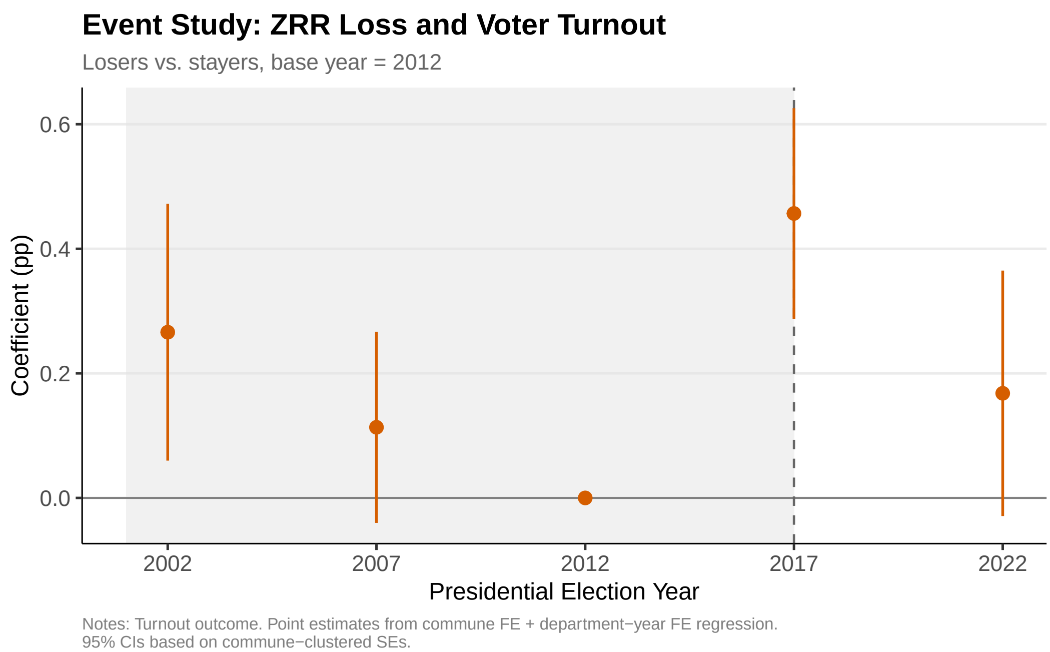

Figure 4: Event Study: Effect of Losing ZRR Status on Voter Turnout

Notes: Event-study coefficients with turnout (percentage of registered voters who voted) as the outcome. The reference year is 2012. No significant pre- or post-treatment effects are detected, consistent with a null effect on voter participation.

6. Robustness

6.1 Leave-One-Department-Out

To ensure that no single geographic unit drives the result, I re-estimate the baseline DiD specification dropping each of the 84 departments in the sample one at a time. All 84 estimates remain negative, with a range of $[-0.515, -0.170]$ centered on the full-sample estimate of -0.334 pp, and 81 of 84 (96%) retain significance at the 5% level. Figure 5 in Appendix C displays the full distribution of jackknife estimates.

6.2 Heterogeneity by Commune Size

I split the DiD sample at the median of 2012 registered voters. The treatment effect is driven primarily by larger communes: the ATT for above-median communes is -0.544 pp ($p < 0.001$), while the estimate for smaller communes is -0.140 pp and not statistically significant ($p = 0.49$). The substantial difference between the two estimates indicates that the effect is concentrated among larger communes. This is consistent with a compositional mechanism: larger communes have more economic activity exposed to ZRR tax incentives, and thus more scope for treatment-induced changes.

6.3 Heterogeneity by Prior FN Support

I split communes at the median of their 2012 FN/RN vote share. The effect is concentrated in low-FN communes (-0.736 pp, $p < 0.001$), while high-FN communes show an insignificant effect (-0.067 pp, $p = 0.70$). The substantial difference between the two estimates is consistent with a floor effect. This pattern suggests that ZRR loss reduced FN/RN support in communes where support was relatively low to begin with—a "floor effect" interpretation where communes already near the national mean had more room to move away from FN/RN.

6.4 Placebo Test

As a direct test of pre-trends, I restrict the sample to the three elections fully preceding any reform activity (2002, 2007, 2012) and define a placebo post indicator equal to one for 2012. The placebo ATT is -0.234 pp ($p = 0.013$), which is statistically significant. This finding indicates that losers and stayers were already on differential trajectories in the pre-reform period, raising a concern for the parallel-trends assumption that underpins the main DiD estimate. Table 5 presents the full heterogeneity and placebo results.

Table 5: Heterogeneity and Placebo Tests

	(1) Small	(2) Large	(3) Low Prior FN	(4) High Prior FN	(5) Placebo
Dep. var.:	FN/RN Vote Share (%)				
Treatment × Post	-0.140	-0.544***	-0.736***	-0.067	-0.234**
	(0.203)	(0.126)	(0.160)	(0.172)	(0.094)
Commune FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	36,107	36,269	36,195	36,173	44,037
Communes	7,349	7,336	7,342	7,341	14,685

Notes: Columns (1)–(2) split the sample at the median commune size (187 registered voters in 2012). Columns (3)–(4) split at the median 2012 FN/RN vote share (19.9%). Column (5) restricts the sample to the three elections fully preceding any reform activity (2002, 2007, 2012) and defines a placebo post indicator equal to one for 2012, testing whether losers and stayers were already on differential trajectories in the pre-reform period. All specifications include commune and year fixed effects. Standard errors clustered at the commune level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6.5 Department-Level Clustering

Because ZRR eligibility is determined by EPCI-level characteristics (density and fiscal income thresholds), treatment is effectively assigned at the EPCI level rather than the commune level. Communes within the same EPCI share the treatment assignment mechanism and may share political and economic shocks, creating within-EPCI correlation that commune-level clustering does not account for. To assess the sensitivity of inference to the clustering level, I re-estimate the main specification with standard errors clustered at the department level (84 departments), which is even coarser than the EPCI assignment unit and thus provides a conservative upper bound on standard errors.

Table 6 reports the main results under both commune-level and department-level clustering. Reliable commune-to-EPCI crosswalk data for the relevant 2017 period proved difficult to obtain, so I use departments as the conservative higher-level clustering unit. Departments are larger than EPCIs and thus provide even more conservative inference than EPCI-level clustering would.

The results are striking. The point estimate is unchanged at -0.334 pp, but the standard error triples from 0.119 under commune clustering to 0.391 under department clustering, and the p-value rises from 0.005 to 0.396. The result is no longer statistically significant at any conventional level under department-level clustering. This finding underscores that the statistical significance of the baseline estimate depends on the clustering level, and that inference at the treatment-assignment level substantially weakens the evidence for a non-zero effect.

6.6 Dropping the 2002 Election

The significant 2002 event-study coefficient raises the concern that the 2002 presidential election—in which Jean-Marie Le Pen's surprise qualification for the second round created an unusual nationwide spike in FN support—may exert undue influence on the main estimate. To assess this, I re-estimate the baseline DiD excluding 2002 entirely. The coefficient drops from -0.334 to -0.211 pp (SE $= 0.114$, $p = 0.065$), just outside conventional significance. This confirms that the 2002 election contributes materially to the main estimate and that the result is sensitive to the inclusion of this atypical pre-period.

6.7 Rambachan–Roth Sensitivity Analysis

I apply the Rambachan and Roth (2023) sensitivity framework using the `HonestDiD` R package with the smoothness restriction (Δ^{SD}). The framework constructs robust confidence intervals that account for potential violations of parallel trends, parameterized by a bound $\bar{M}$ on the

Table 6: Main Results: Commune-Level vs. Department-Level Clustering

	(1) Commune SE	(2) Department SE
Dep. var.:	FN/RN Vote Share (%)	
Treatment × Post	-0.334^{***}	-0.334
	(0.119)	(0.391)
	[$p = 0.005$]	[$p = 0.396$]
Commune FE	Yes	Yes
Year FE	Yes	Yes
Observations	72,376	72,376
Communes	14,685	14,685
Clusters	14,685	84

Notes: Both columns estimate the same baseline DiD specification; point estimates are identical. Column (1) reports commune-level clustered standard errors (baseline). Column (2) reports department-level clustered standard errors (84 departments), which provide conservative inference at a level coarser than the EPCI assignment unit. ZRR eligibility is determined by EPCI-level characteristics, but reliable commune-to-EPCI crosswalk data for the 2017 period were unavailable; departments are larger than EPCIs and thus yield even more conservative inference. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

second differences of the bias. At $\bar{M} = 0$ (the most favorable assumption—exact smoothness in the bias path), the robust confidence interval is $[-0.295, +0.209]$, which includes zero.

This finding reflects the influence of the positive 2002 pre-treatment coefficient (+0.423 pp), which the framework appropriately incorporates into its inference. The conventional DiD confidence interval ($[-0.567, -0.100]$) excludes zero, but the Rambachan–Roth robust interval, which accounts for the full pre-treatment path including the 2002 deviation, is wider. This means the result is sensitive to assumptions about the pre-trend structure: if the 2002 deviation represents a real underlying trend difference rather than an idiosyncratic election effect, the causal interpretation is weakened. I note, however, that the 2007 and 2012 coefficients are both close to zero and insignificant, suggesting that the 2002 pattern did not persist. The sensitivity result should be interpreted in light of the full event-study pattern rather than in isolation.

7. Mechanisms

The negative direction of the conventional estimate invites discussion. I consider three candidate mechanisms as hypotheses—none is directly tested with the available data, and each should be treated as conjecture rather than explanation supported by the evidence.

7.1 Low Program Salience

The most parsimonious explanation is that ZRR tax incentives have low political salience. Unlike visible public services whose withdrawal generates immediate grievance (the closure of the local hospital, the shuttering of the post office), ZRR benefits are technical tax provisions primarily affecting employers. Most voters in ZRR communes are likely unaware of their commune's ZRR status, let alone whether it changed in 2015. If voters do not perceive the policy withdrawal, no "state betrayal" channel activates, and the austerity→populism mechanism simply does not apply.

This interpretation is consistent with the broader salience literature in public economics. Tax instruments that are not directly visible to voters—such as employer-side payroll contributions or corporate tax exemptions—generate weaker behavioral and political responses than equivalent-value but visible transfers or service changes (Chetty et al., 2009; Finkelstein, 2009). The ZRR program, with its indirect benefits flowing through employer decisions, is an archetypal low-salience instrument.

7.2 Compositional Effects

The loss of ZRR status could change the composition of the electorate. If business subsidies were supporting economic activity that employed younger, lower-skilled workers—a demographic profile associated with higher FN/RN support—then the withdrawal of those subsidies could precipitate outmigration of precisely this population. The remaining electorate, older and more rooted, might have more moderate political preferences.

This mechanism finds indirect support in the population-weighted result. The larger magnitude under population weighting (-0.50 pp vs. -0.33 pp unweighted) suggests that larger communes drive the effect, consistent with compositional dynamics being more pronounced where population is larger. However, testing this mechanism directly requires individual-level or age-cohort data that commune-level election aggregates cannot provide.

7.3 Compensatory Mobilization

A third possibility is that the ZRR reclassification, while invisible to most voters, triggered mobilization among local elites (mayors, municipal councillors, business owners) who recognized the policy change and organized in response. If this mobilization channeled grievance into constructive civic engagement rather than protest voting, the net effect on FN/RN support could be negative. Local elected officials, in particular, might have responded to ZRR loss by emphasizing mainstream political engagement to attract alternative forms of state support.

This "voice over exit" interpretation aligns with findings in the political-participation literature showing that local organizational capacity can moderate the political effects of economic shocks. Communities with stronger civic institutions may channel discontent into advocacy and institutional lobbying rather than populist voting.

7.4 Distinguishing Mechanisms

The available data do not allow me to distinguish between these mechanisms. Testing salience directly would require survey data on voter awareness of ZRR status. Testing composition requires individual-level migration data. Testing mobilization requires administrative records on local civic activity. I leave these investigations to future work. Importantly, these mechanisms are not mutually exclusive, and the electorate-composition evidence in Table 3 provides partial support for the second channel. A fuller account would require data that commune-level electoral aggregates cannot provide.

8. Discussion

8.1 Interpreting the Negative Effect

The negative direction of the conventional DiD estimate—if taken at face value—would challenge the view that state withdrawal from peripheral communities mechanically fuels far-right populism. However, the evidence base has important limitations: a single post-treatment election, a significant early pre-period coefficient, robust sensitivity intervals that include zero, and evidence of differential electorate growth that may partly explain the vote-share pattern. These caveats counsel against strong causal claims.

That said, the absence of a *positive* effect is itself informative. Even under the most generous interpretation of the data, there is no evidence that losing ZRR status *increased* far-right voting in the way that UK austerity (Fetzer, 2019), US trade shocks (Autor et al., 2020), or European immigration exposure (Dustmann et al., 2019) have been shown to increase populist support. This null-to-negative pattern points toward a distinction between *salient* and *non-salient* policy instruments. The UK welfare cuts studied by Fetzer (2019) directly affected household budgets in visible ways—reduced disability benefits, housing-benefit cuts, tax-credit freezes. These changes were immediately perceived by affected voters. By contrast, ZRR tax incentives operate through employer decisions and are invisible to most voters. The political economy of state withdrawal likely depends fundamentally on whether voters can observe and attribute the policy change.

This observation connects to a broader theoretical point. The "economic anxiety" pathway from distress to populism requires three links: (i) the policy change must produce economic harm; (ii) voters must experience and perceive that harm; and (iii) voters must attribute the harm to mainstream political actors and seek alternatives. The ZRR case likely fails at links (ii) and (iii). The economic effects of ZRR withdrawal—reduced employer tax breaks, potentially fewer business formations—are diffuse and difficult for individual voters to observe. And even if observed, the technical, bureaucratic nature of the reclassification makes attribution to mainstream parties difficult. This contrasts with austerity measures like benefit cuts, where the "sender" (the government) and the "message" (your benefits are being reduced) are both clear.

8.2 Comparison with Related Findings

The result resonates with emerging evidence that the economic-distress-to-populism pathway is more contingent than initially thought. Enke (2020) shows that moral values predict populist voting more strongly than economic exposure, suggesting that cultural factors mediate the relationship between economic shocks and political responses. Charron et al. (2023) finds that governance quality—not just economic performance—shapes populist support, implying that institutional context matters for how voters interpret economic changes.

In the French context specifically, Guilluy (2014) emphasized the importance of perceived abandonment by metropolitan elites as a driver of peripheral populism. My results suggest that *perceived* abandonment may matter more than *actual* policy withdrawal. The ZRR reclassification was actual policy withdrawal—concrete tax benefits were removed—but it was not perceived withdrawal because the policy instrument lacked visibility. This distinction between actual and perceived state retreat may help reconcile my null-to-negative finding with the broader "peripheral France" narrative.

8.3 The Role of Comparison Groups

A second lesson concerns the role of comparison groups. The strikingly different results for the loser-vs.-stayer and gainer-vs.-never comparisons highlight that "state withdrawal $\rightarrow$ populism" claims are only as credible as the counterfactual they invoke. The gainer-vs.-never comparison, which lacks parallel trends, would yield a misleading positive coefficient if taken at face value. The loser-vs.-stayer comparison, with its more credible counterfactual (both groups were ZRR communes), produces the more reliable estimate.

This methodological point has broader implications. Many studies of populism exploit cross-sectional variation between exposed and non-exposed regions (Colantone and Stanig, 2018; Broz et al., 2021). If exposed and non-exposed regions differ systematically in ways that also predict political preferences—as the gainer-vs.-never comparison illustrates—then cross-sectional designs may conflate selection effects with causal mechanisms. The within-ZRR comparison I exploit avoids this problem by comparing communities that shared the same institutional regime before the reform.

8.4 The Heterogeneity Pattern

The concentration of the effect among larger and lower-FN communes is informative for understanding mechanisms. The size gradient—stronger effects in larger communes—is consistent with a compositional interpretation. Larger communes are more urbanized, have more economic activity, and experience more population turnover. If ZRR loss accelerates demographic change (through outmigration of subsidy-dependent workers or inmigration of commuters attracted by relative housing affordability), the compositional effect should be larger where population flows are larger.

The prior-FN gradient—stronger effects where FN support is lower—has a different interpretation. In communes where FN support is already high (above the median of approximately 20%), there may be a "ceiling" effect: voters attracted to the FN are already committed, and no marginal economic change moves them. In lower-FN communes, the electorate is more fluid, and small compositional or preference shifts can move the FN share. This pattern is also consistent with the salience mechanism: in lower-FN communes, voters may be more attentive to mainstream economic policy and less susceptible to populist framing of economic grievances.

8.5 Limitations

Several limitations should be noted. First, the ZRR transition provisions extending benefits through 2020 mean that only the 2022 election captures the full economic impact of losing ZRR status. With only one post-treatment election, the design cannot distinguish a treatment effect from an idiosyncratic 2022 loser-vs-stayer divergence. The 2027 presidential election will provide a crucial additional observation for assessing whether the effect persists, intensifies, or dissipates.

Second, the absence of commune-level economic data (employment, firm creation, income) prevents direct examination of the first-stage economic effects of ZRR loss. If ZRR loss had no economic consequences—a possibility, given the program's modest scale relative to overall local economic activity—then the null political effect is trivially explained. Administrative data from SIRENE (the French business register) or DADS (*Déclarations Annuelles de Données Sociales*) could test this first-stage channel in future work.

Third, the evidence on electorate composition (Table 3) raises the possibility that the vote-share decline partly reflects differential population growth in loser communes rather than changed political preferences. Without individual-level data on voter registration, migration, and demographic change, I cannot fully distinguish these channels.

Third, the analysis cannot address general-equilibrium effects: if ZRR loss in some communes redirected economic activity to neighboring stayer communes, the comparison understates the true effect. Spatial spillovers would attenuate the DiD estimate toward zero.

Fourth, the electoral outcome I observe—the first-round presidential vote—may not capture all dimensions of political response. Local elections, protest participation, or party membership changes could reveal effects that national electoral data miss. The FN/RN may also benefit indirectly through increased voter disaffection that manifests only in lower-salience elections.

9. Conclusion

This paper studies France's 2015 reform of Zones de Revitalisation Rurale to examine whether the withdrawal of place-based government support pushes rural communities toward the far right. I find no robust evidence of a populist backlash. The conventional DiD point estimate is negative—communes that lost ZRR status show lower FN/RN vote-share growth than communes that retained status—but the Rambachan–Roth robust confidence interval includes zero, and evidence on electorate composition suggests that differential voter registration growth may partly drive the vote-share pattern, complicating a pure preference-change interpretation. The result is therefore suggestive rather than conclusive.

Still, the direction of the point estimate is noteworthy. The "state withdrawal fuels populism" hypothesis—supported in contexts of visible austerity (Fetzer, 2019) and trade shocks (Autor et al., 2020)—does not appear to operate through low-salience, employer-facing rural tax incentives. If confirmed with additional post-treatment elections and first-stage economic evidence, this would identify an important boundary condition: the political response to state withdrawal may depend critically on the nature of the withdrawn instrument.

For policymakers, the tentative implication is that the political costs of restructuring place-based programs may vary with instrument design. Invisible, business-facing tax provisions may generate less voter backlash than visible, household-facing service withdrawals. This hypothesis warrants further investigation across different institutional settings before informing policy design.

Whether this extends to other place-based programs—enterprise zones, regional development grants, infrastructure investment—is an important question for future research. The US

Empowerment Zone program (Busso et al., 2013), UK Regional Selective Assistance (Criscuolo et al., 2019), and EU Structural Funds all provide analogous settings where place-based support has been restructured or withdrawn. Systematic comparison across programs with different salience levels could establish whether the pattern documented here generalizes.

More broadly, the paper highlights the need for a more nuanced "political economy of retrenchment" that accounts for instrument characteristics, voter awareness, and the institutional context of policy withdrawal. The simple narrative that economic distress breeds populism—while containing an important kernel of truth—is powerful but incomplete: the details of the policy instrument matter as much as the distress itself.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Algan, Yann, Sergei Guriev, Elias Papaioannou, and Evgenia Passari, "The European Trust Crisis and the Rise of Populism," *Brookings Papers on Economic Activity*, 2017, *2017* (2), 309–400.

Autor, David, David Dorn, Gordon Hanson, and Kaveh Majlesi, "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure," *American Economic Review*, 2020, *110* (10), 3139–3183.

Becker, Sascha O., Thiemo Fetzer, and Dennis Novy, "Who Voted for Brexit? A Comprehensive District-Level Analysis," *Economic Policy*, 2017, *32* (92), 601–650.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Bó, Ernesto Dal, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne, "Fiscal Austerity and the Rise of the National Front: Evidence at the French Municipal Level," *Review of Economics and Statistics*, 2023, *105* (6), 1461–1477.

Broz, J. Lawrence, Jeffry Frieden, and Stephen Weymouth, "Populism in Place: the Economic Geography of the Globalization Backlash," *International Organization*, 2021, *75* (2), 464–494.

Busso, Matias, Jesse Gregory, and Patrick Kline, "Assessing the Incidence and Efficiency of a Prominent Place Based Policy," *American Economic Review*, 2013, *103* (2), 897–947.

Callaway, Briant and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Caselli, Mauro and Sergio Ferrini, "Immigration and the Rise of the Far Right: Evidence from Italian Municipalities," *European Economic Review*, 2021, *131*, 103618.

Charron, Nicholas, Victor Lapuente, and Monika Bauhr, "Does Poor Quality of Government Fuel Populism? Evidence from European Regions," *Journal of Politics*, 2023, *85* (2), 481–496.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Colantone, Italo and Piero Stanig, "The Trade Origins of Economic Nationalism: Import Competition and Voting Behavior in Western Europe," *American Journal of Political Science*, 2018, *62* (4), 936–953.

Criscuolo, Chiara, Ralf Martin, Henry G. Overman, and John Van Reenen, "Some Causal Effects of an Industrial Policy," *American Economic Review*, 2019, *109* (1), 48–85.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Dippel, Christian, Robert Gold, Stephan Heblich, and Rodrigo Pinto, "Effect of Trade on Workers and Voters," *Economic Journal*, 2022, *132* (641), 199–217.

Dustmann, Christian, Kristine Vasiljeva, and Anna Piil Damm, "Refugee Migration and Electoral Outcomes," *Review of Economic Studies*, 2019, *86* (5), 2035–2091.

Edo, Anthony, Yvonne Giesing, Jonathan Öztunc, and Panu Poutvaara, "Immigration and Electoral Support for the Far-Left and the Far-Right," *European Economic Review*, 2019, *115*, 99–143.

Enke, Benjamin, "Moral Values and Voting," *Journal of Political Economy*, 2020, *128* (10), 3679–3729.

Fetzer, Thiemo, "Did Austerity Cause Brexit?," *American Economic Review*, 2019, *109* (11), 3849–3886.

Finkelstein, Amy, "E-ZTax: Tax Salience and Tax Rates," *Quarterly Journal of Economics*, 2009, *124* (3), 969–1010.

Funke, Manuel, Moritz Schularick, and Christoph Trebesch, "Going to Extremes: Politics after Financial Crises, 1870–2014," *European Economic Review*, 2016, *88*, 227–260.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Guilluy, Christophe, *La France périphérique: Comment on a sacrifié les classes populaires*, Flammarion, 2014.

Guiso, Luigi, Helios Herrera, Massimo Morelli, and Tommaso Sonno, "Global Crises and Populism: the Role of Eurozone Institutions," *Economic Policy*, 2019, *34* (97), 95–139.

Guriev, Sergei and Elias Papaioannou, "A Cross-Country Perspective on the Causes of Populism," *Annual Review of Economics*, 2022.

Halla, Martin, Alexander F. Wagner, and Josef Zweimüller, "Immigration and Voting for the Far Right," *Journal of the European Economic Association*, 2017, *15* (6), 1341–1385.

Kline, Patrick and Enrico Moretti, "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs," *Annual Review of Economics*, 2014, *6*, 629–662.

Neumark, David and Helen Simpson, "Do Place-Based Policies Matter?," *National Bureau of Economic Research Working Paper*, 2015, (20049).

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Rodrik, Dani, "Why Does Globalization Fuel Populism? Economics, Culture, and the Rise of Right-Wing Populism," *Annual Review of Economics*, 2021, *13*, 133–170.

Roth, Jonathan, "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, *4* (3), 305–322.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Data Appendix

A.1 ZRR Classification Construction

The ZRR classification data from the DGCL contains annual commune-level status codes. Each commune is assigned a label indicating its ZRR status for each year:

- **C** (Classée): Commune is classified in ZRR.
- **NC** (Non classée): Commune is not classified in ZRR.
- **M** (Sortante): Commune is exiting ZRR classification.
- **D** (Classée, déclin): Commune is classified due to population decline.
- **P** (Partiellement classée): Commune is partially classified.

I classify communes as "in ZRR" if their status code begins with C, D, or P (but not NC). The pre-reform status is taken from the 2014 sheet of the historical file; the post-reform status from the 2018 sheet. Treatment groups are defined by the transition:

- **Loser:** In ZRR in 2014, not in ZRR in 2018 (4,478 communes in analysis sample).
- **Stayer:** In ZRR in both 2014 and 2018 (10,207 communes in analysis sample).
- **Gainer:** Not in ZRR in 2014, in ZRR in 2018.
- **Never:** Not in ZRR in either year.

A.2 Election Data Processing

The election Parquet files contain candidate-level results at the *bureau de vote* level. For each commune-election pair, I:

1. Aggregate all *bureau de vote* results within the commune by summing vote counts.
2. Identify the FN/RN candidate: Jean-Marie Le Pen (2002, 2007), Marine Le Pen (2012, 2017, 2022), or the "Prenez le Pouvoir" list (2019 European).
3. Compute the FN/RN vote share as FN/RN votes divided by total valid votes (*exprimés*), multiplied by 100.
4. Compute turnout as the number of voters (*votants*) divided by registered voters (*inscrits*), multiplied by 100.

The merge between election data and ZRR classification uses the five-digit INSEE commune code. Communes present in both datasets form the near-balanced panel (14,685 communes; 14,067 appear in all five elections, with slight attrition from municipal mergers).

A.3 Sample Restrictions

The analysis excludes: (i) communes in overseas departments (*départements d'outre-mer*), which have distinct political dynamics; (ii) commune-election observations with zero valid votes cast ($exprimés = 0$), for which vote shares are undefined (11 observations). The panel is near-balanced: 14,067 of 14,685 communes appear in all five elections; the remaining 618 are absent from one or two elections due to municipal mergers or dissolutions.

B. Identification Appendix

B.1 Pre-Trend Analysis

The event-study specification in equation (2) provides the primary test of parallel trends, with 2012 as the omitted reference year. Figure 1 in the main text displays the full set of coefficients. With 2012 as baseline:

- 2002 coefficient: positive and significant, indicating that losers had higher FN support than stayers relative to 2012. This early-period deviation is addressed below.
- 2007 coefficient: small and insignificant, indicating convergence between groups by the mid-2000s.
- 2017 coefficient: small and insignificant, providing limited evidence of anticipation effects.
- 2022 coefficient: negative, indicating that losers experienced lower FN/RN growth than stayers after the reform's full implementation.

The choice of 2012 as the reference year—rather than 2017, which was used in the initial draft—avoids the concern that 2017 may already reflect anticipation or early treatment effects following the 2015 legislative reform. With 2012 as the baseline, 2017 appears as an explicit coefficient that can be interpreted as a test of early/anticipation effects.

The significant 2002 coefficient raises a genuine concern about the parallel-trends assumption. While the 2002 election was exceptional (Jean-Marie Le Pen's surprise second-round qualification produced unusual voting patterns), the deviation is the kind of evidence that the Rambachan and Roth (2023) sensitivity framework is designed to incorporate. Following Roth (2022), I interpret pre-trend tests with caution and report the formal sensitivity analysis below.

B.2 Rambachan–Roth Sensitivity

I implement the Rambachan and Roth (2023) approach using the `HonestDiD` R package with the smoothness restriction (Δ^{SD}), which bounds the second differences of the bias function. The robust confidence intervals account for the full pre-treatment path—including the significant 2002 coefficient—when constructing inference about the post-treatment effect.

At $\bar{M} = 0$ (exact smoothness), the FLCI robust interval is $[-0.295, +0.209]$. This includes zero, reflecting the influence of the 2002 pre-trend deviation. As $\bar{M}$ increases (allowing greater departures from smoothness), the interval widens further: $[-1.12, +0.46]$ at $\bar{M} = 0.5$ and $[-1.62, +0.96]$ at $\bar{M} = 1.0$.

The interpretation is as follows: the conventional DiD estimate is statistically significant, but the Rambachan–Roth framework—which rigorously accounts for the observed 2002 pre-trend—cannot rule out a zero effect. This is an honest limitation of the identification strategy when the earliest pre-treatment period shows a significant deviation. The result depends on whether one treats the 2002 divergence as an idiosyncratic election effect (the Le Pen shock of 2002) or as evidence of a structural trend difference between losers and stayers.

C. Robustness Appendix

C.1 Leave-One-Department-Out

With 84 departments in the sample, the leave-one-out exercise produces 84 ATT estimates. The distribution is centered on the full-sample ATT (-0.334 pp) with a narrow range of $[-0.515, -0.170]$. All 84 estimates remain negative, and 81 of 84 (96%) retain significance at 5%, confirming that the result is not driven by any single geographic cluster.

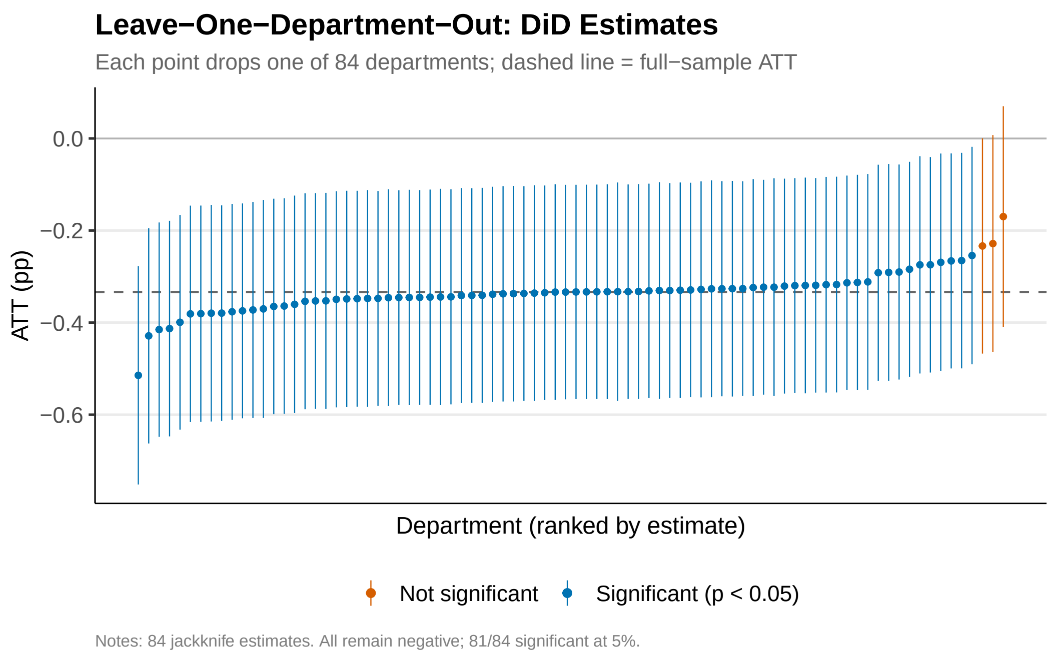

Figure 5: Leave-One-Department-Out Estimates

Notes: Each point represents the DiD estimate obtained when dropping one of the 84 departments from the sample. The dashed horizontal line marks the full-sample estimate (-0.334 pp). All 84 estimates are negative, and 81 (96%) are significant at 5%.

C.2 Heterogeneity Results

By commune size. I split communes at the median of 2012 registered voters. The ATT for larger communes is -0.544 pp ($p < 0.001$), while smaller communes show an insignificant -0.140 pp ($p = 0.49$). The substantial difference between the two estimates indicates that the effect is concentrated among larger communes.

By prior FN support. I split communes at the median of 2012 FN/RN vote share. The ATT for low-FN communes is -0.736 pp ($p < 0.001$), while high-FN communes show an insignificant -0.067 pp ($p = 0.70$). The substantial difference between the two estimates is consistent with a floor effect.

Interpretation. The concentration of the effect among larger and lower-FN communes suggests a compositional rather than attitudinal mechanism. Larger communes have more economic activity susceptible to ZRR tax incentive changes, while lower-FN communes have more moderate voters who might be responsive to small compositional shifts in the electorate.

C.3 Placebo Test

The placebo test restricts the sample to the three elections fully preceding any reform activity (2002, 2007, 2012) and defines a placebo post indicator equal to one for 2012. The placebo ATT is -0.234 pp ($p = 0.013$), which is statistically significant and indicates a pre-existing differential trend between losers and stayers even in the purely pre-reform period. Combined with the significant 2002 event-study coefficient and the fragile HonestDiD sensitivity results, this raises a genuine concern about the parallel-trends assumption. See Column (5) of Table 5.

D. Symmetric Test

Table 7 reports the symmetric test comparing gainers against never-ZRR communes. Column (1) reproduces the baseline main result for reference. Column (2) estimates the effect of gaining ZRR status: the coefficient is large and positive, indicating that gainers had substantially higher growth in FN/RN support than never-ZRR communes. However, as Figure 6 shows, event-study coefficients exhibit clear pre-treatment differences, confirming that the parallel trends assumption does not hold in this sample. The symmetric-test coefficient reflects pre-existing level differences between gaining and never-ZRR communes—which differ fundamentally in urbanization, economic structure, and political culture—rather than a causal effect of gaining ZRR status. This exercise serves as a cautionary illustration: comparison groups that appear natural can be deeply misleading when they compare structurally different populations.

Table 7: Symmetric Test: Losing vs. Gaining ZRR Status

	(1) Losers vs. Stayers	(2) Gainers vs. Never-ZRR
Dep. var.:	FN/RN % Expressed	FN/RN % Expressed
Treatment × Post	-0.334***	2.994***
	(0.119)	(0.105)
Commune FE	Yes	Yes
Year FE	Yes	Yes
Observations	72,376	108,711
Communes	14,685	22,341
Within R^2	0.0002	0.0143

Notes: Column (1) reproduces the baseline DiD from Table 2: communes that lost ZRR status (treatment) vs. communes that retained it (control). Column (2) estimates the symmetric effect: communes that gained ZRR status in the reform (treatment) vs. communes that never held ZRR status (control). If ZRR designation causally affects RN voting, losing status should increase and gaining status should decrease RN vote share. All specifications include commune and year fixed effects. Standard errors clustered at the commune level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

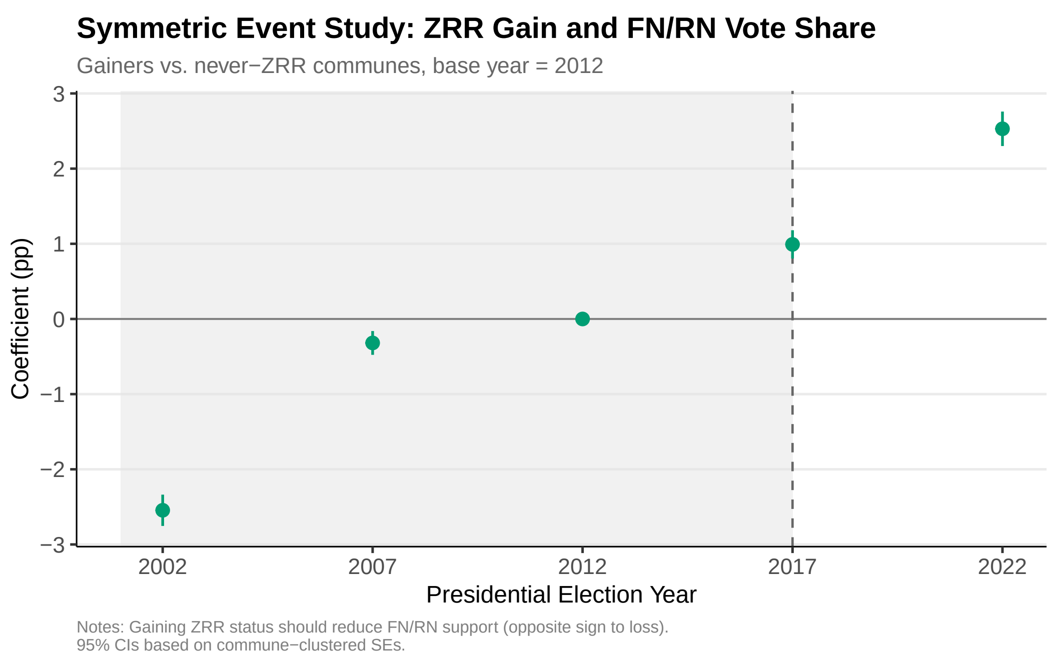

Figure 6: Symmetric Event Study: Effect of Gaining ZRR Status on FN/RN Vote Share

Notes: Event-study coefficients for gainers (treatment) vs. never-ZRR communes (control). The reference year is 2012. Large pre-treatment coefficients indicate that parallel trends do not hold in this sample.

E. Additional Figures and Tables

All underlying data are saved as CSV files in the replication archive, and figure-generating code is in `05_figures.R`. Table-generating code is in `06_tables.R`. The complete pipeline (data fetch, clean, analysis, robustness, figures, tables) runs sequentially through the numbered R scripts in the `code/` directory.

F. Standardized Effect Sizes

Table 8: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
FN/RN Vote Share (%)	DiD, Table 2 Col. 1	−0.334	—	7.65	−0.044	Null
FN/RN Vote Share (%, weighted)	DiD, Table 2 Col. 2	−0.501	—	7.65	−0.065	Small negative
Log(FN/RN % + 1)	DiD, Table 2 Col. 3	−0.028	—	0.35	−0.080	Small negative
Voter Turnout (%)	DiD, Table 4 Col. 1	−0.041	—	5.30	−0.008	Null

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. The treatment is binary (0/1: loser vs. stayer), so SDE = $\hat{\beta}$/SD(Y) and the SD(X) column is marked "—". SD(Y) is the pooled unconditional standard deviation of the outcome, computed as the average of the loser and stayer group SDs from Table 1 (7.683 and 7.626, respectively, yielding 7.65), before conditioning on fixed effects.

Research question: Does losing place-based rural tax incentives (ZRR status) affect far-right (FN/RN) voting in French communes? **Treatment:** Binary — commune lost ZRR status under the 2015 legislative reform (identified from the 2014 vs. 2018 classification lists; administrative reclassification effective 2017–2018; loser = 1, stayer = 0). **Data:** ZRR classification from DGCL and commune-level election results from data.gouv.fr, 2002–2022, unit = commune-election year, N = 72,376. **Method:** Two-way fixed effects DiD with commune and year fixed effects; standard errors clustered at the commune level. **Sample:** Metropolitan France communes classified in ZRR as of 2014 (losers) or retaining ZRR status through 2018 (stayers); near-balanced panel across five presidential elections.

Classification thresholds: large negative (< −0.10), small negative (−0.10 to −0.05), null (−0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). A reader unfamiliar with the paper should be able to interpret this table on its own.