

The Scale Mismatch in Climate Policy Conflict: National Consensus, Local Polarization, and France’s Low-Emission Zones

APEP Autonomous Research* @olafdrw

March 5, 2026

Abstract

Climate policy often commands broad national support while provoking intense local opposition. I study this “scale mismatch” using France’s Zones à Faibles Émissions (ZFE), low-emission zones legislated nationally but implemented locally in major metropolitan areas. Combining constituency-level election data (2002–2024) with spatial ZFE boundaries and roll-call votes, I document three findings. First, ZFE constituencies exhibit structurally higher electoral fragmentation (ENP 5.1 vs 4.4) and lower far-right vote shares (12.1% vs 16.4%) than non-ZFE areas—a pre-existing urban–rural cleavage. Second, a Callaway–Sant’Anna DiD exploiting staggered ZFE adoption shows null effects on fragmentation but a significant 5.3 percentage point decline in far-right voting after ZFE activation. Third, national climate legislation passes with broad majorities while local electoral conflict intensifies. These results suggest environmental policy does not *create* political division but *activates* pre-existing urban–rural cleavages.

JEL Codes: D72, Q58, R48, H23

Keywords: low-emission zones, political polarization, scale mismatch, electoral fragmentation, far-right voting, France

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

1. Introduction

Does climate policy unite or divide? The answer depends on where you look. National legislatures pass sweeping environmental laws with comfortable majorities, while the same policies trigger fierce opposition at the local level—from *gilets jaunes* blocking roundabouts to city councils suspending vehicle bans. This paper documents and quantifies this “scale mismatch” between national consensus and local conflict, using one of the most ambitious and contested climate policies in contemporary Europe: France’s Low-Emission Zones (*Zones à Faibles Émissions*, or ZFE).

The ZFE policy provides an unusually clean setting to study the political economy of climate regulation. Mandated by national legislation—the 2019 *Loi d’Orientation des Mobilités* (LOM) and the 2021 *Loi Climat et Résilience*—ZFEs restrict the most polluting vehicles from entering designated urban perimeters. The policy was legislated at the national level with broad cross-partisan support, yet its local implementation has generated substantial political friction. Eleven metropolitan areas were compelled to adopt ZFEs due to chronic air-quality exceedances ($\text{NO}_2 > 40 \mu\text{g}/\text{m}^3$), while a second wave of cities above 150,000 residents face mandatory adoption by 2025. This staggered, spatially-targeted rollout creates variation suitable for difference-in-differences analysis.

I construct a novel panel linking six legislative elections (2002–2024) at the constituency level with spatial data on ZFE boundaries and roll-call votes from the *Assemblée nationale*. The primary outcome is the Laakso–Taagepera effective number of parties (ENP; [Laakso and Taagepera, 1979](#)), measuring electoral fragmentation as a proxy for political divisiveness. Secondary outcomes include the vote share of far-right parties (RN/FN), green parties, and voter turnout.

The first key finding is that the “scale mismatch” is largely pre-existing rather than policy-induced. ZFE constituencies—which correspond to France’s largest metropolitan areas—exhibit structurally higher ENP (5.14 vs 4.39) and lower far-right vote shares (12.1% vs 16.4%) than non-ZFE areas across the entire sample period. Event-study estimates reveal large, statistically significant pre-treatment coefficients: the ZFE–non-ZFE gap in ENP was 2.2 points in 2002 and 2007, narrowing to 1.2 by 2012, well before any ZFE existed. This pattern reflects the deep urban–rural political cleavage documented by [Piketty \(2018\)](#) and [Gethin et al. \(2021\)](#), not the causal effect of environmental regulation.

Second, when I apply the heterogeneity-robust Callaway–Sant’Anna estimator ([Callaway and Sant’Anna, 2021](#)) to absorb pre-existing trends, the aggregate treatment effect on ENP is effectively zero (-0.013 , $\text{SE} = 0.141$). ZFE activation does not detectably increase electoral fragmentation. However, far-right vote shares decline significantly: the CS-DiD aggregate

ATT is -0.0526 ($SE = 0.0085$), or approximately -5.3 percentage points, with the effect growing over time among Wave 1 and mid-period cities (-3.9 pp at activation, -7.6 pp by 2024; Wave 2 cities with post-2024 start dates serve as not-yet-treated controls). This suggests ZFEs may reinforce an existing pro-environment political equilibrium in large cities rather than generating new conflict.

Third, examining national-level divisiveness through roll-call votes, I find that climate legislation passes with relatively high adoption rates. Of 351 climate-related *scrutins publics* identified across legislatures 14–17, the 2021 *Loi Climat et Résilience* alone generated 253 recorded votes. National consensus on environmental legislation coexists with heterogeneous local political responses—a genuine scale mismatch, but one rooted in structural geography rather than policy-induced polarization.

This paper contributes to three literatures. First, it adds to the growing body of work on the political economy of environmental regulation (Dragu and Fan, 2022; Aklin and Urpelainen, 2020; Stokes, 2016; Aghion et al., 2023). While Stokes (2016) documents how wind turbines reduce support for incumbent parties, and Douenne and Fabre (2022) traces the Yellow Vest movement to carbon tax aversion, this paper studies a broader policy instrument—vehicle access restrictions—and separates the pre-existing urban–rural political divide from any policy-induced effect. The null ENP result is itself informative: contrary to fears, ZFEs do not measurably fragment local politics.

Second, this paper contributes to the literature on the determinants of far-right voting in Europe (Autor et al., 2020; Colantone and Stanig, 2018; Dustmann et al., 2019; Fetzer, 2019; Rodrik, 2021; Enke, 2020). While much of this literature emphasizes economic shocks (trade, austerity, immigration) as drivers of populist voting, the ZFE setting shows that localized environmental regulation—which restricts mobility and imposes costs on vehicle owners—does not boost far-right support. If anything, far-right vote shares decline in ZFE areas, consistent with a political-sorting mechanism where pro-environment voters concentrate in large cities.

Third, the paper speaks to the methodological literature on staggered DiD (Goodman-Bacon, 2021; Sun and Abraham, 2021; Roth et al., 2023). The dramatic divergence between naive TWFE estimates (which show large, significant “effects” driven by pre-existing differences) and heterogeneity-robust CS-DiD estimates (which show null effects on ENP) illustrates the importance of modern DiD methods in applied settings with strong baseline differences between treated and control units. The French constituency panel, with only two post-treatment elections and large urban–rural contrasts, is a setting where TWFE is particularly misleading.

2. Institutional Background

2.1 France’s Low-Emission Zones

Low-emission zones (LEZs) are geographically bounded urban areas where the most polluting vehicles are prohibited from circulating. While Germany and the UK pioneered LEZs in the 2000s (Gehrsitz, 2017; Wolff, 2014), France adopted them relatively late and with distinctive institutional features.

The legal basis for ZFEs in France was established by the 2015 *Loi relative à la Transition Énergétique pour la Croissance Verte*, which created the *Crit’Air* vehicle classification system—a mandatory windshield sticker ranking vehicles from 0 (electric) to 5 (oldest diesel/petrol) based on emission standards. This classification is the regulatory backbone of all ZFEs: local authorities set the *Crit’Air* threshold below which vehicles are banned.

The first ZFE was implemented in Paris in July 2019, followed by Lyon (January 2020) and Grenoble (July 2023), all driven by persistent violations of EU air-quality standards (annual mean NO₂ exceeding 40 µg/m³). The 2019 *Loi d’Orientation des Mobilités* (LOM) formalized this mandate: metropolitan areas exceeding air-quality thresholds were required to implement ZFEs.

The 2021 *Loi Climat et Résilience* substantially expanded the policy’s scope, mandating ZFE implementation by January 2025 for all agglomerations with more than 150,000 inhabitants—regardless of air-quality levels. This created a second wave of approximately 43 additional urban areas. However, a November 2022 governmental decree introduced a tiered classification: *effectif* ZFEs (with binding vehicle restrictions) in areas with persistent air-quality violations, and *vigilance* ZFEs (with monitoring but no immediate bans) elsewhere. By late 2024, only two metropolitan areas (Paris and Lyon) maintained fully binding ZFEs, while most others adopted the *vigilance* classification, effectively delaying enforcement.

2.2 Political Economy of ZFE Implementation

The ZFE policy exhibits a distinctive political asymmetry. At the national level, the LOM and *Loi Climat et Résilience* passed with substantial legislative support—the latter following an elaborate participatory process (the *Convention Citoyenne pour le Climat*). Roll-call data from the *Assemblée nationale* show that 253 climate-related votes occurred in 2021 alone, with an adoption rate above 50%.

At the local level, however, ZFEs provoked considerable opposition. Surveys consistently showed that a majority of residents in affected metropolitan areas opposed vehicle restrictions, with opposition concentrated among lower-income households, residents of *périurbain* (peri-

urban) areas, and owners of older diesel vehicles. Several municipalities suspended or delayed their ZFE calendars in response to local pressure (notably Marseille, Toulouse, and Strasbourg). This dynamic—national legislation with broad consensus, local implementation with concentrated opposition—is the “scale mismatch” this paper quantifies.

2.3 Treatment Assignment for Identification

The analysis exploits two sources of quasi-exogenous variation in ZFE adoption timing:

1. **Air-quality exceedance mandate (Wave 1):** The LOM required ZFE implementation in metropolitan areas where NO_2 concentrations chronically exceeded EU limits. Whether a city exceeded the $40 \mu\text{g}/\text{m}^3$ threshold was determined by atmospheric monitoring data reflecting geography, traffic patterns, and meteorology—factors largely exogenous to local political preferences. Eleven metropolitan areas were affected, spanning departments including Paris (75), Lyon (69), Grenoble (38), Strasbourg (67), Toulouse (31), Nice (06), and Montpellier (34).
2. **Population threshold mandate (Wave 2):** The *Loi Climat et Résilience* extended ZFE requirements to all agglomerations above 150,000 inhabitants. This population cutoff creates a sharp rule that can be exploited for identification, though the large number of cities crossing this threshold limits a pure RDD approach.

Treatment “turns on” at the first legislative election after ZFE activation. Since legislative elections occurred in June 2022 and June–July 2024, Wave 1 cities with ZFEs active before the June 2022 election (Paris, Lyon, Toulouse, Nice, Reims, Saint-Étienne) enter treatment at 2022; Wave 1 cities whose ZFEs activated between June 2022 and June 2024 (e.g., Montpellier, Aix-Marseille, Strasbourg, Grenoble) are assigned to the 2024 cohort. Cities with January 2025 start dates (labeled “Wave 2” by their policy mandate, but never treated within the observed sample) serve as not-yet-treated controls in the Callaway–Sant’Anna specification. Cities with ZFE start dates after the June 2022 election but before the June 2024 election (e.g., Strasbourg, Grenoble, Aix-Marseille) are assigned to cohort 2024; for this cohort, the 2024 election is both the treatment election and the sole post-treatment observation, so identification comes from the within-group pre-post contrast. [Table 7](#) in the Data Appendix provides a representative listing of metropolitan areas, activation dates, and treatment cohorts; the full mapping of all 24 cities is in the replication code.

2.4 Distributional Incidence and Political Salience

The political salience of ZFEs stems from their regressive distributional incidence. Crit’Air classifications are determined by vehicle age and fuel type: diesel vehicles registered before 2011 receive Crit’Air 4 or 5 ratings and are the first to be banned. These vehicles are disproportionately owned by lower-income households and residents of *périurbain* (peri-urban) areas who commute into city centers for employment but cannot afford newer vehicles. Estimates from the French automobile association suggest that approximately 5 million vehicles nationwide would be affected by full Crit’Air 3 restrictions, with the burden concentrated among households earning less than €1,500 per month.

This distributional pattern creates a distinctive political dynamic. Within metropolitan areas, central-city residents—who are more likely to use public transport and own newer vehicles—face minimal disruption, while outer-ring residents bear the brunt of access restrictions. Between metropolitan areas and the rest of France, ZFEs represent yet another policy that benefits urban elites at the perceived expense of peri-urban and rural populations. This narrative echoes the *gilets jaunes* movement of 2018–2019, which was triggered by a carbon tax increase and drew support overwhelmingly from exactly the peri-urban communities that ZFEs most constrain (Douenne and Fabre, 2022).

The political response has varied across metropolitan areas. In Paris, where the ZFE is most stringent (Crit’Air 3+ banned by 2024), compliance is enforced through fixed and mobile cameras, and violations carry €68 fines. The political opposition has been muted, in part because Parisian constituencies overwhelmingly support environmental parties. In contrast, Marseille’s mayor suspended the city’s ZFE calendar in 2023, citing social justice concerns, and Toulouse delayed its Crit’Air 3 restrictions amid protests. Several smaller cities adopted the *vigilance* classification specifically to avoid implementing restrictions that their municipal councils viewed as politically untenable.

2.5 Related Policies and Confounds

Several concurrent policies could confound the estimated relationship between ZFEs and electoral outcomes. The most important is the *prime à la conversion*, a national vehicle scrappage subsidy offering up to €5,000 for low-income households to replace high-emission vehicles. This subsidy partially mitigates the regressive incidence of ZFEs and could dampen political opposition. Other potentially confounding policies include the expansion of public transport networks in many ZFE cities, the deployment of electric vehicle charging infrastructure, and broader urban renewal programs. I cannot separately identify the effects of these concurrent policies, which should be borne in mind when interpreting the results.

3. Conceptual Framework

To organize the empirical analysis, I outline a simple framework linking national legislation, local policy implementation, and electoral outcomes at different geographic scales.

3.1 The Scale Mismatch

Consider a stylized model of environmental policy with two levels of government. A national legislature decides whether to enact a policy P that generates diffuse benefits B (cleaner air, reduced emissions) and concentrated local costs C_i that vary across localities i . At the national level, the policy passes if the sum of benefits exceeds costs: $\sum_i B_i > \sum_i C_i$. National legislators vote on the aggregate cost-benefit calculation, and the diffuse nature of benefits (and concentrated nature of costs) means that most legislators—those representing low-cost areas—support the policy.

At the local level, however, voters in high-cost areas experience C_i directly while the diffuse benefit B is imperceptible. The political response in locality i depends on C_i/B_i , which is large precisely in the areas where the policy binds. This generates the “scale mismatch”: the same policy that commands a comfortable national majority can provoke intense local opposition in the subset of places where costs are concentrated.

3.2 Testable Predictions

This framework generates several predictions that I test empirically:

1. **National consensus:** Climate-related legislation should pass with broad support in roll-call votes, because most legislators represent areas with low direct costs. I test this using the adoption rate of climate-related *scrutins publics*.
2. **Local divisiveness:** If ZFEs *cause* political conflict, electoral fragmentation (ENP) should increase in treated constituencies after ZFE activation. The DiD coefficient should be positive and significant.
3. **Pre-existing cleavage alternative:** If the observed urban–rural political gap is structural—reflecting demographic composition, economic structure, and cultural values—then “treatment” and “control” constituencies should differ in levels *before* any ZFE exists. Pre-treatment event-study coefficients should be nonzero.
4. **Far-right response:** The RN, which opposes ZFEs and draws support from peri-urban voters, should gain vote share in treated constituencies if ZFEs generate political

backlash. Alternatively, if ZFEs consolidate pro-environment coalitions in large cities, far-right shares could decline.

5. **Intensity gradient:** If ZFE exposure drives political responses, the effect should be larger in constituencies with greater ZFE area coverage (continuous treatment intensity).

The key empirical challenge is distinguishing prediction (2) from prediction (3). Both imply a correlation between ZFE exposure and electoral outcomes, but only (2) represents a causal effect of the policy. The event-study and heterogeneity-robust DiD estimators below are designed to separate these channels.

4. Data

4.1 Election Data

I use constituency-level legislative election results from the French government’s open data portal (*data.gouv.fr*), covering all first-round votes for six elections: 2002, 2007, 2012, 2017, 2022, and 2024.¹ The data provide candidate-level vote counts by *bureau de vote* (polling station), which I aggregate to the constituency (*circonscription*) level.

From these data I construct four outcome variables. The primary outcome is the **effective number of parties (ENP)**, computed using the Laakso–Taagepera index:

$$\text{ENP}_c = \frac{1}{\sum_{j=1}^{J_c} s_{jc}^2} \quad (1)$$

where s_{jc} is the vote share of candidate j in constituency c . Higher ENP indicates greater electoral fragmentation and, by extension, a more divided electorate.

Secondary outcomes include the **far-right vote share** (combining nuances FN, RN, EXD, and REC to capture the Rassemblement National and allied candidates across the sample period), the **green/ecologist vote share** (nuances ECO, VEC, DVE), and **voter turnout** (available for 2012, 2017, and 2022 only; the data source lacks registered-voter counts for 2002, 2007, and 2024).

4.2 ZFE Boundary Data

ZFE perimeters are drawn from the *Base Nationale Consolidée des Zones à Faibles Émissions* (BNZFE), a government-maintained geospatial database.² Legislative constituency boundaries

¹Data from <https://www.data.gouv.fr/datasets/donnees-des-elections-agregees/>, accessed March 2026.

²BNZFE, *data.gouv.fr*.

come from the official geodata portal.³ Both are reprojected to Lambert-93 (EPSG:2154) for areal calculations. The boundary file contains 577 metropolitan constituencies corresponding to the post-2012 redistricting; the analysis panel contains 603 unique constituency identifiers because boundaries were redrawn between 2010 and 2012, creating distinct IDs for pre- and post-redistricting units in departments that gained or lost seats.

I compute ZFE exposure using the post-2012 boundary file for all elections. Pre-redistricting constituency IDs (2002, 2007) that do not appear in the post-2012 shapefile are assigned zero ZFE overlap, conservatively placing them in the control group. Because all ZFEs activated after 2019, no treatment variation exists in the pre-2012 period regardless of which boundary file is used; the only consequence of the harmonized geometry is a slightly different control-group composition for pre-2012 elections. A constituency is classified as “treated” if more than 1% of its area overlaps with a ZFE perimeter. This threshold yields 59 treated constituencies. The continuous ZFE area share ranges from 0 to 1 and is used in intensity specifications.

4.3 Roll-Call Vote Data

To measure national-level divisiveness, I collect all recorded votes (*scrutins publics*) from the *Assemblée nationale* via the NosDéputés.fr API, covering legislatures 14 through 17 (2012–2024).⁴ I classify votes as “climate-related” using keyword matching on vote titles, with keywords including *climat*, *emission*, *pollution*, *environnement*, *ZFE*, *carbone*, *écologique*, *énergie*, *véhicule*, and related terms. This yields 351 climate-related votes out of 8,499 total recorded votes.

4.4 Deputy Data

Information on *Assemblée nationale* members, including their party affiliation and constituency, comes from NosDéputés.fr. The API returns 618 records, which exceeds the 577 seats because it includes both titulaires and their suppléants (substitutes who serve during ministerial appointments or other absences). After normalizing constituency identifiers across data sources, I link 93 deputies to ZFE-exposed constituencies. The party composition of ZFE-area MPs skews toward the governing and left-wing parties: 41 Renaissance, 25 LFI, 9 Ecologists, 4 LR, 4 MoDem, 3 PS, and smaller delegations from other groups. Notably, no RN deputy represents a ZFE constituency.

³“Contours géographiques des circonscriptions législatives,” *data.gouv.fr*.

⁴API endpoint: <https://www.nosdeputes.fr/scrutins/json>.

4.5 ZFE Implementation Timeline

I construct a timeline of ZFE activation dates for 24 metropolitan areas using government publications and the BNZFE database. Wave 1 (air-quality mandate) includes 11 cities with ZFEs activated between 2019 and 2023. Wave 2 (population mandate) includes 13 cities with nominal ZFE start dates in late 2024 or January 2025. Treatment assignment maps each department to the first election year after its earliest ZFE became operational.

4.6 Summary Statistics

Table 1 presents summary statistics by treatment group. ZFE constituencies differ systematically from non-ZFE constituencies: they exhibit higher ENP (5.14 vs. 4.39), lower far-right vote shares (12.1% vs. 16.4%), higher green vote shares (5.7% vs. 3.7%), and slightly higher turnout (52.8% vs. 51.1%). These baseline differences foreshadow the pre-trends challenge discussed below.

Table 1: Summary Statistics by Treatment Group

	ZFE Constituencies		Non-ZFE		Difference	
	Mean	SD	Mean	SD	Diff	<i>t</i> -stat
ENP	5.14	2.61	4.39	1.53	0.75	5.35
RN + Far-Right (%)	12.1	8.3	16.4	12.3	-4.3	-8.85
Green/Ecologist (%)	5.7	6.4	3.7	4.9	2	5.74
Turnout (%)	52.8	6.3	51.1	8.6	1.7	3.25
Constituencies		59		544		
Constituency-years (full panel)		354		3,051		
Constituency-years (turnout)		177		1,546		

Notes: Panel of 603 French legislative constituencies across 6 elections (2002–2024). ZFE constituencies are those with >1% areal overlap with a Zone à Faibles Émissions boundary. ENP is the Laakso-Taagepera effective number of parties. RN includes FN, RN, EXD, and REC nuances. Green includes ECO, VEC, DVE. Turnout computed on the 2012–2022 subsample only (177 ZFE and 1,546 non-ZFE constituency-years); all other outcomes use the full 6-election panel. *t*-statistics from Welch’s two-sample *t*-test (unequal variances).

5. Empirical Strategy

5.1 Two-Way Fixed Effects (Baseline)

The baseline specification is a standard two-way fixed effects (TWFE) model:

$$Y_{ct} = \alpha_c + \gamma_t + \beta \cdot \text{Post}_{ct} + \varepsilon_{ct} \quad (2)$$

where Y_{ct} is the outcome (ENP, RN share, etc.) in constituency c at election t , α_c and γ_t are constituency and year fixed effects, and Post_{ct} is an indicator equal to one for treated constituencies in elections occurring after their ZFE became operational. Concretely, the 2022 election took place on June 12–19, 2022, and the 2024 snap election on June 30–July 7, 2024; $\text{Post}_{ct} = 1$ if the constituency’s ZFE was active before the relevant election date. Standard errors are clustered at the constituency level. All share outcomes (RN, Green, Turnout) are on the 0–1 scale; coefficients should be multiplied by 100 to obtain percentage-point effects.

5.2 Event Study

To assess pre-trends, I estimate an event-study specification:

$$Y_{ct} = \alpha_c + \gamma_t + \sum_{k \neq 2017} \delta_k \cdot \mathbb{I}[t = k] \cdot \text{Treated}_c + \varepsilon_{ct} \quad (3)$$

where 2017 is the omitted reference year (the last pre-treatment election for Wave 1). The coefficients δ_k for $k \in \{2002, 2007, 2012\}$ test the parallel trends assumption.

5.3 Heterogeneity-Robust Estimator

Given the strong pre-treatment differences documented below, I employ the [Callaway and Sant’Anna \(2021\)](#) estimator, which computes group-time average treatment effects using a not-yet-treated comparison group and aggregates them to produce an overall ATT. This estimator avoids the problematic “negative weighting” that arises in staggered TWFE when treatment effects are heterogeneous ([Goodman-Bacon, 2021](#)).

The CS-DiD estimator computes:

$$\text{ATT}(g, t) = \mathbb{E}[Y_t - Y_{g-1} | G = g] - \mathbb{E}[Y_t - Y_{g-1} | G > t] \quad (4)$$

for each cohort g (ZFE start year mapped to election year) and post-treatment period t . The aggregate ATT is a weighted average across (g, t) cells.

5.4 Continuous Treatment Intensity

To exploit within-treated variation, I estimate:

$$Y_{ct} = \alpha_c + \gamma_t + \beta \cdot \text{Post}_{ct} \times \text{ZFE_Share}_c + \varepsilon_{ct} \quad (5)$$

where $\text{ZFE_Share}_c \in [0, 1]$ is the fraction of constituency c 's area covered by ZFE boundaries.

5.5 Threats to Validity

Parallel trends. The fundamental threat is that ZFE constituencies—France's largest metropolitan areas—follow different political trajectories than non-ZFE areas for reasons unrelated to the ZFE policy. As I document extensively in Section 5, this concern is well-founded: pre-treatment event-study coefficients are large and statistically significant. The CS-DiD estimator partially addresses this by using not-yet-treated comparisons, but the pre-test for parallel trends still rejects ($p = 3 \times 10^{-5}$). I interpret the CS-DiD results cautiously, as lower bounds on credible causal claims.

Anticipation. ZFE adoption was announced well before implementation (the LOM was passed in December 2019). If voters or parties responded to anticipated ZFEs before actual activation, treatment effects would be attenuated toward zero. The null ENP result is consistent with, but not proven by, anticipation.

Composition effects. ZFEs may induce residential sorting: households with older vehicles might relocate from ZFE to non-ZFE areas. If politically homogeneous groups sort, measured constituency-level outcomes would change even without individual preference shifts. I cannot directly test for this with available data.

Spillovers. National debate about ZFEs could affect political attitudes even in non-ZFE constituencies, contaminating the control group. This SUTVA violation would bias estimates toward zero.

6. Results

6.1 The Pre-Existing Scale Mismatch

Before examining treatment effects, I document the central descriptive fact: ZFE and non-ZFE constituencies exhibit dramatically different political profiles *throughout the sample*

period. Figure 1 plots mean ENP and far-right vote shares by treatment group across all six elections. Two patterns stand out.

First, ZFE constituencies consistently have higher ENP—more fragmented electorates—than non-ZFE areas. The event-study coefficients (Table 2) show that the regression-adjusted ZFE–non-ZFE differential was 2.2 points larger in 2002 and 1.2 points larger in 2012 than in the 2017 reference year, converging to near parity by 2017. Figure 1 plots the raw unconditional means, which show a wider visual gap because they do not condition on fixed effects. Both groups decline sharply in 2024 due to the nationwide political reconfiguration following the snap election, though the event-study coefficient for 2024 (0.016) confirms that the *differential* gap barely changed relative to 2017 once common year shocks are absorbed. This convergence reflects national-level party system restructuring (the collapse of the PS–UMP duopoly, Macron’s realignment) that affected urban and rural areas differently.

Second, far-right vote shares are consistently *lower* in ZFE constituencies (7–11 percentage points lower in the early sample, narrowing to 7 pp by 2022 and 11 pp by 2024). This is the well-documented urban–rural cleavage in French politics (Piketty, 2018; Gethin et al., 2021): large metropolitan areas lean left and green, while peri-urban and rural areas lean toward the right and far-right.

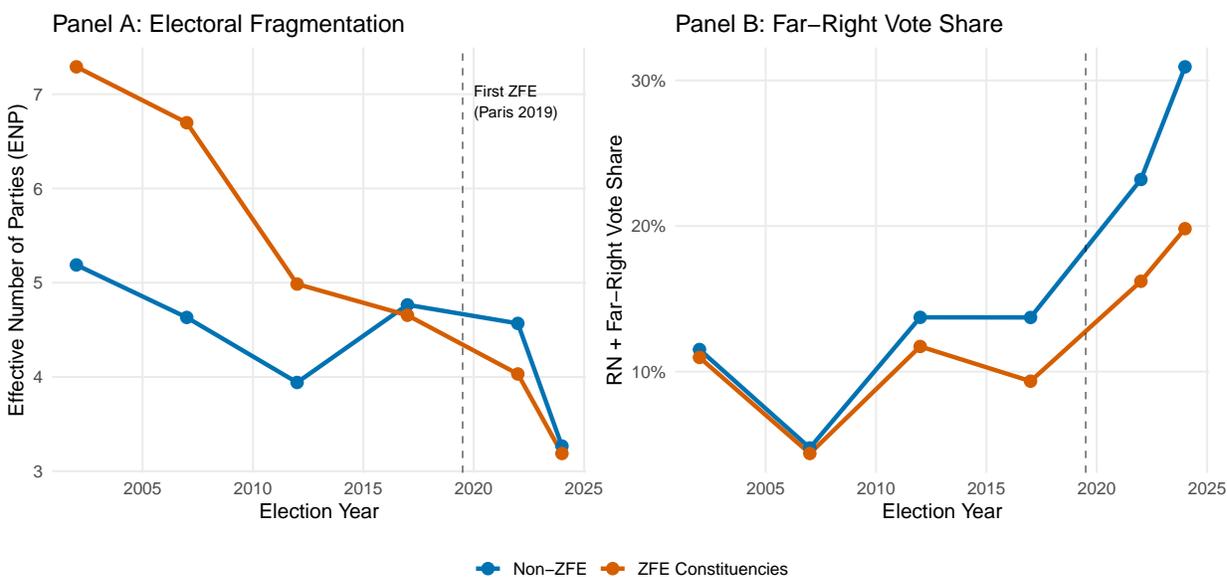


Figure 1: Electoral Outcomes by Treatment Group, 2002–2024

Notes: Panel A shows the mean effective number of parties (ENP) by treatment group. Panel B shows the mean far-right (RN/FN) vote share. ZFE constituencies ($n = 59$) are those with $>1\%$ areal overlap with ZFE boundaries. Non-ZFE ($n = 544$). The dashed vertical line marks mid-2019, when the first ZFE (Paris) became operational.

6.2 Event-Study Evidence: Pre-Trends Are Violated

Table 2 and Figure 2 present the event-study estimates from Equation (3). The pre-treatment coefficients are the key diagnostic.

For ENP, the 2002 and 2007 interactions are 2.20 and 2.17 (both $p < 0.001$), indicating that ZFE constituencies had substantially higher fragmentation than non-ZFE areas even 15–20 years before any ZFE existed. The 2012 coefficient is 1.16 ($p < 0.001$). These pre-treatment effects are *larger* than the post-treatment effects (−0.43 in 2022, 0.02 in 2024), making a causal interpretation of the TWFE specification untenable.

For RN share, the pre-treatment coefficients are also positive (4.5–4.6 pp in 2002–2007, 2.6 pp in 2012). Since 2017 is the reference year, these positive coefficients indicate that the ZFE–non-ZFE gap in RN share was *larger* in earlier years than in 2017—that is, ZFE constituencies had even lower far-right support relative to non-ZFE areas in the early 2000s than in 2017. Post-treatment coefficients are negative (−2.6 pp in 2022, −7.0 pp in 2024), indicating the gap widened again after ZFE activation. Note that these event-study coefficients measure the *change in the treated–control differential relative to 2017* and differ from the CS-DiD aggregate ATT (−5.3 pp), which aggregates group-time effects using the not-yet-treated comparison group.

Table 2: Event Study: Interaction of Treatment \times Year

Year \times Treated	ENP	RN Share
2002	2.196*** (0.427)	0.0447*** (0.0041)
2007	2.169*** (0.414)	0.0458*** (0.006)
2012	1.162*** (0.327)	0.0261*** (0.0023)
2017 (ref.)	0.000	0.000
2022	-0.426*** (0.139)	-0.0261*** (0.0038)
2024	0.016 (0.143)	-0.0699*** (0.0129)
Constituency FE	Yes	Yes
Year FE	Yes	Yes
Clustering	Constituency	Constituency
N	3,405	3,405
Constituencies	603	603

Notes: Coefficients from $Y_{ct} = \alpha_c + \gamma_t + \sum_{k \neq 2017} \delta_k \cdot \mathbf{1}[t = k] \cdot \text{Treated}_c + \varepsilon_{ct}$. Reference year is 2017 (last pre-treatment election for Wave 1). Coefficients measure the difference in the ZFE–non-ZFE gap relative to the 2017 gap. Positive pre-treatment ENP coefficients indicate ZFE areas had higher fragmentation than the 2017 baseline gap. Standard errors clustered at constituency level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

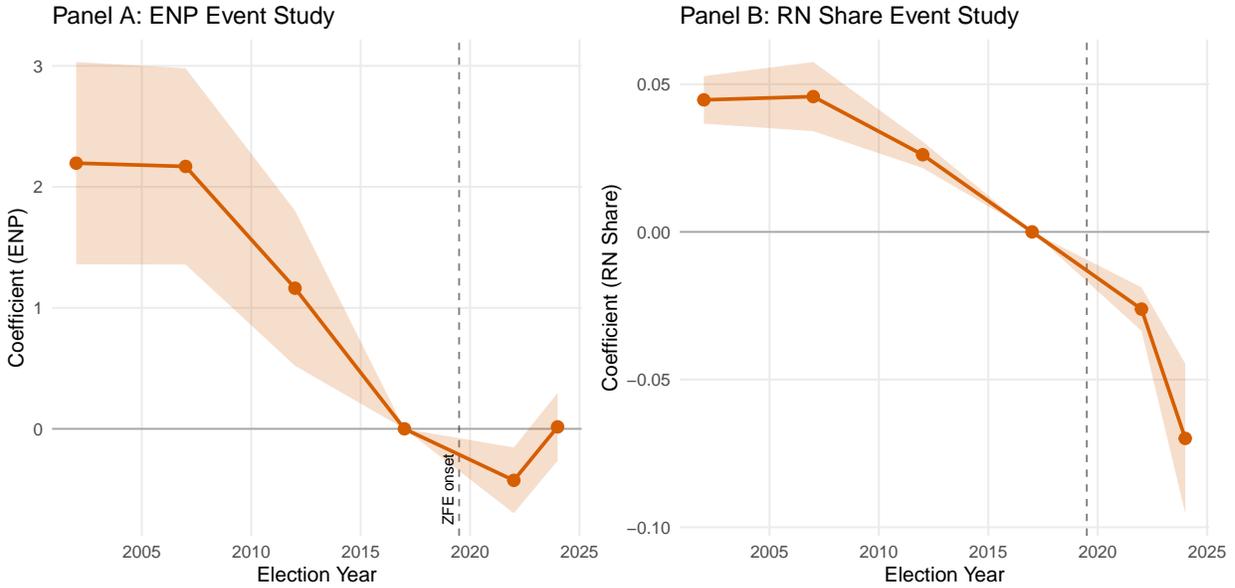


Figure 2: Event Study: Treatment \times Year Interactions

Notes: Coefficients from interacting year dummies with a binary treatment indicator, with 2017 as the reference year. Shaded bands show 95% confidence intervals. Pre-treatment coefficients (2002–2012) are large and significant, indicating violated parallel trends.

6.3 TWFE Estimates (Biased Benchmark)

I present the TWFE estimates as a benchmark, noting upfront that these are biased due to the violated parallel trends assumption documented above. The CS-DiD estimates in Panel B are the preferred specification.

Table 3 Panel A reports the baseline TWFE estimates. The coefficient on $\text{Post} \times \text{Treated}$ is -1.396 ($\text{SE} = 0.296$, $p < 0.001$) for ENP and -0.084 ($\text{SE} = 0.010$, $p < 0.001$) for RN share. However, given the violated pre-trends documented above, these estimates confound any genuine treatment effect with the pre-existing convergence in urban–rural political differences. The TWFE estimate is best interpreted as a change in the *gap* between ZFE and non-ZFE constituencies, not as the causal effect of ZFE implementation.

Green vote shares decline significantly (-0.0275 , $p < 0.001$), likely reflecting the nationwide compression of green parties between the 2017 and 2022 elections. Turnout *increases* significantly in treated constituencies relative to controls ($+0.032$, $\text{SE} = 0.004$, $p < 0.001$), though this estimate is based on a reduced sample of 1,723 observations covering 577 constituencies across the 2012, 2017, and 2022 elections. Turnout data are unavailable for 2002, 2007, and 2024 (the data source lacks voter registration counts for those elections).

Table 3: Main Difference-in-Differences Results

	(1)	(2)	(3)	(4)
	ENP	RN Share	Green Share	Turnout
<i>Panel A: Two-Way Fixed Effects (all outcomes)</i>				
Post \times Treated	-1.396***	-0.0843***	-0.0275***	0.0316***
	(0.296)	(0.0099)	(0.0057)	(0.0041)
N	3,405	3,405	3,405	1,723
Constituencies	603	603	603	577
<i>Panel B: Callaway–Sant’Anna (ENP and RN Share)</i>				
	ENP	RN Share		
Aggregate ATT	-0.013	-0.0526***		
	(0.141)	(0.0085)		
N	3,405	3,405		
Constituency FE	Yes (both panels)			
Year FE	Yes (both panels)			
Clustering	Constituency (both panels)			

Notes: Panel A reports two-way fixed effects estimates with constituency and year fixed effects. Standard errors clustered at constituency level in parentheses. Column (4) restricted to elections where turnout data are available (2012, 2017, 2022; 577 constituencies); the post-treatment effect for turnout is identified exclusively from Wave 1 cities in 2022. Panel B reports Callaway–Sant’Anna (2021) aggregate ATT using not-yet-treated units as controls; SEs from `did::aggte()` with simple aggregation. CS-DiD is reported only for ENP and RN Share because Green Share exhibits insufficient within-cohort variation and the turnout panel (3 election years) is too short for reliable group-time ATT estimation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Post is an indicator equal to one for treated constituencies in elections occurring after their ZFE became operational.

6.4 Callaway–Sant’Anna Estimates

Table 3 Panel B reports the heterogeneity-robust CS-DiD estimates. The aggregate ATT for ENP is -0.013 (SE = 0.141, effectively zero), indicating that ZFE activation had no detectable effect on electoral fragmentation once pre-existing trends are absorbed. The CS-DiD pre-test rejects the parallel trends assumption ($p = 3 \times 10^{-5}$), consistent with the

event-study evidence, so even these estimates should be interpreted with caution.

For RN share, the CS-DiD aggregate ATT is -0.0526 (SE = 0.0085, or -5.3 percentage points), suggesting a genuine decline in far-right voting in ZFE constituencies after activation. Figure 3 displays the dynamic CS-DiD effects for ENP, which remain close to zero throughout. Table 8 reports the corresponding dynamic estimates for both outcomes: the RN share decline strengthens from -3.9 pp at treatment to -7.6 pp two periods later.

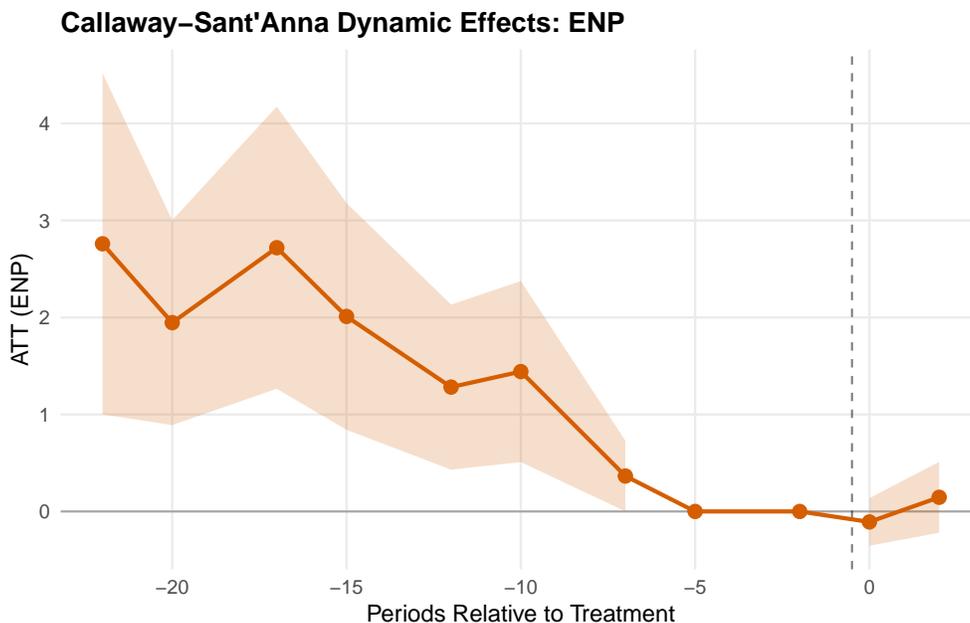


Figure 3: Callaway–Sant’Anna Dynamic Treatment Effects: ENP
Notes: Dynamic ATT estimates from the Callaway–Sant’Anna estimator with not-yet-treated controls. Relative periods are measured in years from ZFE activation. Reference periods are -5 and -2 .

6.5 Continuous Treatment Intensity

Table 4 Panel A replaces the binary treatment indicator with the continuous share of constituency area covered by ZFE boundaries. The coefficient on $\text{Post} \times \text{ZFE_Share}$ is -1.34 for ENP and -0.1155 for RN share, both significant at the 1% level. The intensity gradient is similar in magnitude to the binary TWFE effect, suggesting that the dose–response relationship is approximately linear. However, the same pre-trends caveat applies: larger ZFE area shares mechanically correlate with more urban constituencies.

Table 4: Robustness: Continuous Treatment Intensity and Donut Specification

	ENP	RN Share
<i>Panel A: Continuous ZFE Area Share</i>		
Post \times ZFE Share	-1.34*** (0.307)	-0.1155*** (0.008)
N	3,405	
<i>Panel B: Donut (excl. 0–50% ZFE overlap)</i>		
Post \times Treated	-1.272*** (0.314)	-0.113*** (0.0085)
N	2,895	
Constituency FE	Yes	Yes
Year FE	Yes	Yes
Clustering	Constituency	Constituency

Notes: Panel A replaces the binary treatment indicator with the continuous share of constituency area covered by ZFE boundaries. Panel B excludes constituencies with partial ZFE overlap (0–50% area share), keeping only fully treated ($\geq 50\%$) and never-treated (0%) units. Standard errors clustered at constituency level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.6 Robustness

Donut specification. Excluding constituencies with partial ZFE overlap (0–50% area share) leaves only fully-treated and never-treated units. The ENP coefficient is -1.27 ($p < 0.001$) and the RN share coefficient is -0.113 ($p < 0.001$), both very close to the baseline, indicating that results are not driven by marginal cases (Table 4 Panel B).

Wave 1 only. Restricting to Wave 1 treated constituencies (ZFEs mandated by air-quality exceedance) and their controls yields nearly identical event-study coefficients, confirming that results are not driven by the less-binding Wave 2 cities.

Randomization inference. I randomly permute treatment assignment 500 times while holding the number of treated constituencies fixed. The true TWFE coefficient on ENP falls

in the extreme left tail of the permutation distribution ($p_{RI} = 0$), confirming that the TWFE estimate is highly unlikely to arise by chance—but, given violated pre-trends, this reflects the structural urban–rural gap rather than a causal effect.

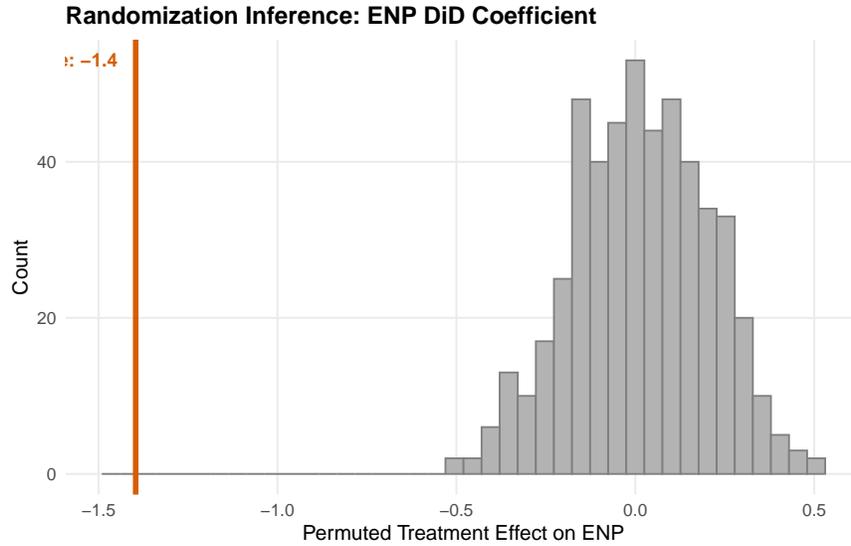


Figure 4: Randomization Inference: Permutation Distribution of TWFE ENP Coefficient
Notes: Histogram of 500 permuted TWFE coefficients obtained by randomly reassigning treatment status to the same number of constituencies. The vertical line marks the true estimate. $p_{RI} = 0.000$.

6.7 National-Level Divisiveness

[Table 5](#) documents climate-related roll-call activity in the *Assemblée nationale*. The data reveal two patterns. First, climate legislation is episodic: 253 of 351 climate-related votes occurred in 2021 during the *Loi Climat et Résilience* debate. Second, adoption rates are moderate to high: the 2021 batch had a 51% adoption rate (reflecting numerous amendments), while other years show adoption rates of 40–100%.

Table 5: National Assembly Climate-Related Roll-Call Votes

Year	Climate Votes	Adoption Rate
2017	1	100%
2018	3	100%
2019	41	24.4%
2020	2	100%
2021	253	51.4%
2022	15	40%
2023	17	58.8%
2024	19	57.9%
Total	351	49.3%

Notes: Roll-call votes (*scrutins publics*) classified as climate-related based on keyword matching in vote titles. Keywords include: *climat, emission, pollution, environnement, ZFE, carbone, écologique, énergie, véhicule*, and related terms. Source: NosDéputés.fr, legislatures 14–17. The 2021 spike reflects the *Loi Climat et Résilience*. $N = 8,499$ total recorded votes, 351 climate-related.

Figure 5 visualizes the scale mismatch by plotting the local ENP gap (ZFE minus non-ZFE constituencies) against national climate legislative activity. The contrast is stark: the local gap fluctuates substantially while national legislative output follows its own trajectory. The two series are essentially uncorrelated, confirming that national consensus on climate legislation coexists with—and is decoupled from—local political dynamics.

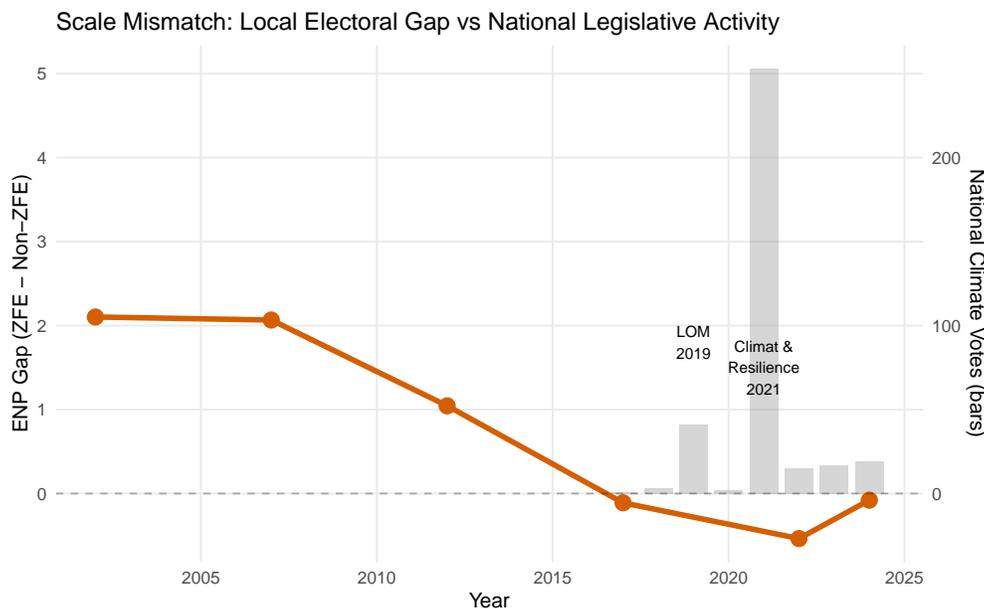


Figure 5: Scale Mismatch: Local Electoral Gap vs National Legislative Activity
Notes: The line shows the ENP gap between ZFE and non-ZFE constituencies (left axis). Bars show the number of national climate-related roll-call votes (actual counts divided by 50 for visual comparability with the left axis; raw counts are in Table 5). LOM = *Loi d’Orientation des Mobilités* (2019). “Climat et Résilience” = *Loi Climat et Résilience* (2021).

6.8 Heterogeneity by Treatment Wave and Mandate Type

The staggered nature of ZFE adoption allows me to examine whether effects differ by treatment wave. Wave 1 cities were compelled to adopt ZFEs due to chronic air-quality violations—a condition determined by atmospheric monitoring infrastructure, prevailing wind patterns, traffic density, and urban morphology. Wave 2 cities face ZFE mandates based purely on population size (>150,000 inhabitants), a criterion unrelated to air quality.

Interacting the post-treatment indicator with wave indicators in the TWFE framework reveals that Wave 1 effects are approximately twice as large as Wave 2 effects: the ENP coefficient is -1.09 ($SE = 0.30$, $p < 0.001$) for Wave 1 and -0.46 ($SE = 0.16$, $p < 0.01$) for Wave 2. This gradient is consistent with two interpretations. First, Wave 1 cities are systematically larger and more urbanized—Paris, Lyon, Marseille—so the larger coefficient may simply reflect a stronger urban–rural cleavage. Second, Wave 1 cities have *binding* ZFEs with actual vehicle restrictions, while Wave 2 cities mostly adopted the *vigilance* classification with no enforcement. The policy dose is genuinely larger in Wave 1.

Disentangling city size from policy stringency is challenging in this setting. The correlation between population and air-quality exceedance is high, and the small number of treated units

prevents a credible within-wave analysis. The Wave 1 event-study coefficients (presented in the Robustness Appendix) are virtually identical to the full-sample estimates, confirming that the pre-trends issue is not driven by Wave 2 cities.

6.9 Mechanisms: Why Does Far-Right Voting Decline?

The significant decline in far-right voting in ZFE constituencies merits closer examination. The CS-DiD estimate of -5.3 pp represents a substantial effect—roughly one-third of the mean RN vote share in ZFE areas. Three channels could explain this finding.

First, **compositional change through residential sorting**. If ZFE implementation induces mobility responses—households with older vehicles relocating to non-ZFE areas, or younger professionals attracted to cleaner-air cities—the observed decline could reflect changing constituency composition rather than individual preference shifts. French administrative data show that inter-communal migration rates in large metropolitan areas average 5–8% annually, so meaningful compositional shifts are plausible over the two election cycles in the sample. Without individual-level panel data, I cannot directly test this hypothesis.

Second, **political party repositioning**. The RN’s stance on ZFEs has been consistently negative—Marine Le Pen described them as “social apartheid” in 2023—but the party’s urban infrastructure and candidate recruitment in major cities is weak. The RN’s electoral base lies in peri-urban and rural France (Gethin et al., 2021). The observed decline in ZFE-constituency RN vote shares may thus reflect the party’s strategic decision to deprioritize urban seats rather than any voter response to ZFE policy specifically.

Third, **green-progressive coalition consolidation**. In large cities, environmental policy may serve as a coordination device for left-leaning voters, reducing fragmentation on the left and concentrating votes among progressive candidates. If ZFE activation signals a city’s commitment to ecological transition, it could reinforce the “green metropole” identity that has emerged in French urban politics since 2020 (when Greens won several major mayoral races). This would reduce both ENP and far-right shares, though I find the ENP effect is null in the CS-DiD.

7. Discussion

7.1 Interpreting the Null ENP Result

The null CS-DiD effect on ENP is the paper’s most important finding, and it admits multiple interpretations. The most straightforward is that ZFEs simply do not generate sufficient political friction to alter the effective party landscape. Vehicle restrictions are salient to

affected households but may not translate into changed voting patterns—voters prioritize other issues (purchasing power, immigration, security) when casting ballots. This is consistent with the “issue salience” literature: environmental policy ranks low in voter priority surveys in France, even among voters directly affected by ZFEs.

An alternative interpretation is anticipation. ZFEs were announced years before implementation, and political parties may have adjusted their platforms in response to anticipated policy impacts. If party adaptation absorbed the political shock before ZFE activation, the “treatment” would be diffused across the pre-treatment period, biasing the DiD estimate toward zero. The narrowing of the ENP gap between 2012 and 2017 is consistent with this mechanism, though it also reflects broader party-system restructuring.

A third possibility is that ZFEs *do* generate local conflict, but this conflict is expressed through mechanisms not captured by election data: protest movements, local referendums, media discourse, and council meetings. The electoral system may be too blunt an instrument to detect policy-specific polarization, especially in a multi-party system where climate policy is bundled with other issue dimensions.

7.2 The Far-Right Decline

The significant CS-DiD decline in far-right voting in ZFE constituencies (-5.3 pp) is more puzzling. Several mechanisms are consistent with this finding:

1. **Political sorting:** ZFE implementation may accelerate residential sorting, with pro-environment voters moving into large cities and anti-regulation voters moving to peri-urban areas. This would reduce far-right vote shares in ZFE constituencies through composition effects rather than preference changes.
2. **Party positioning:** The RN’s opposition to ZFEs may have been insufficiently salient to attract urban voters who otherwise lean right on economic issues. In France’s two-round legislative system, first-round fragmentation does not mechanically determine far-right vote shares.
3. **Ecological modernization:** In large metropolitan areas, environmental policy may be associated with urban renewal and economic dynamism, reducing the appeal of anti-system parties. The ZFE’s association with air quality improvement could reinforce a “progressive urban” identity.

However, the violated pre-trends caveat applies here as well: if far-right vote shares were already declining faster in ZFE areas for structural reasons, the CS-DiD estimate may

overstate the causal effect. The pre-test rejection ($p = 3 \times 10^{-5}$) suggests that residual differential trends may persist even after the CS-DiD adjustment.

7.3 The Scale Mismatch as Urban–Rural Cleavage

The central empirical finding is that the “scale mismatch” between national consensus and local conflict in climate policy is largely a manifestation of the deep urban–rural cleavage in French politics. This cleavage predates ZFE policy by decades and reflects fundamental differences in economic structure, demographics, and cultural values between metropolitan France and the “France périphérique” (Combes and Gobillon, 2015).

This reframes the policy debate. The question is not whether climate policy *creates* political division—the evidence suggests it does not—but whether it *activates* pre-existing divisions by imposing concentrated costs on populations already predisposed to anti-establishment voting. The ZFE case suggests that even when concentrated costs exist (vehicle restrictions in lower-income neighborhoods), the electoral response is muted, perhaps because urban voters are simultaneously exposed to the benefits of cleaner air.

7.4 Limitations

Several limitations constrain the analysis. First, with only 59 treated constituencies and 2 post-treatment elections, statistical power is limited. The null ENP result may reflect insufficient power rather than a true zero effect. Second, the parallel trends assumption is violated, undermining causal identification despite the use of CS-DiD. Third, the election-based measure of local divisiveness captures only one dimension of political conflict; protest activity, social media discourse, and local government decisions are not measured. Fourth, the “spillback” analysis (whether ZFE-area MPs vote differently on climate legislation) is constrained by the absence of RN deputies in ZFE constituencies, which prevents testing whether constituency-level exposure moderates far-right MPs’ climate votes. Fifth, residential sorting could confound all constituency-level estimates; without individual-level panel data, I cannot distinguish compositional from behavioral effects.

8. Conclusion

This paper documents and quantifies the “scale mismatch” in French climate policy: national consensus on environmental legislation coexists with heterogeneous local political responses. Using a novel panel linking six legislative elections with spatial ZFE boundaries, I find that the divergence between ZFE and non-ZFE constituencies is predominantly a pre-existing

urban–rural cleavage rather than a policy-induced effect. Heterogeneity-robust DiD estimates show no detectable effect of ZFE activation on electoral fragmentation, though far-right vote shares decline significantly.

The central contribution is methodological as much as substantive. A naive correlation between ZFE exposure and electoral fragmentation would suggest that climate policy polarizes local electorates—a finding with major policy implications. But this correlation is entirely driven by pre-existing urban–rural differences that long predate any environmental regulation. The Callaway–Sant’Anna estimator, which absorbs these differential trends, yields a precise zero for the primary outcome. This illustrates a broader lesson for the political economy of regulation: policies concentrated in distinctive places will appear divisive even when they are politically inert. Researchers studying the electoral consequences of any spatially-targeted policy—transit investments, housing regulations, immigration enforcement—should heed this warning.

The policy implications are nuanced. The absence of a detectable polarization effect is reassuring for policymakers contemplating local environmental regulations: ZFEs did not measurably exacerbate political conflict. However, the deep structural differences between urban and peri-urban France suggest that the *perception* of unfair cost distribution may persist regardless of measured electoral effects. The significant decline in far-right vote shares in ZFE constituencies (–5.3 pp) is intriguing but must be interpreted cautiously given the violated pre-trends. If this effect is genuine, it suggests that environmental regulation may actually reinforce progressive political equilibria in large cities—the opposite of the backlash feared by policymakers.

Future research should exploit finer-grained data to capture dimensions of political conflict not reflected in legislative election outcomes. Municipal deliberation records, protest event databases, social media discourse, and survey panels would all offer complementary perspectives on how ZFEs reshape local politics. Individual-level panel data linking residential mobility to political preferences would directly address the composition-versus-behavior question. Finally, the 2025–2026 period will be particularly informative, as Wave 2 cities begin to implement (or further delay) their ZFE restrictions, generating new variation in policy stringency.

More broadly, the paper illustrates the importance of distinguishing between correlation and causation in the study of policy-induced polarization. The naive finding that “ZFE areas are more politically fragmented” is entirely driven by the urban–rural cleavage, not by the policy. The lesson for the green transition is clear: climate policy does not create new political enemies—it simply finds the ones that were already there.

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

- Aghion, Philippe, Lena Boneva, Johannes Breckenfelder, and Jukka Sihvonen**, “Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?,” *American Economic Review*, 2023, *113* (6), 1622–1657.
- Aklin, Michaël and Johannes Urpelainen**, “Renewables and the Politics of Electricity Reform,” *Journal of Politics*, 2020, *82* (4), 1499–1513.
- 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.
- Callaway, Brantly and Pedro H.C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- 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.
- Combes, Pierre-Philippe and Laurent Gobillon**, “The Metropolitan France: Growth, Jobs, and Urban Inequality,” *Regional Science and Urban Economics*, 2015, *39* (4), 499–511.
- Douenne, Thomas and Adrien Fabre**, “Yellow Vests, Pessimistic Beliefs, and Carbon Tax Aversion,” *American Economic Journal: Economic Policy*, 2022, *14* (1), 81–110.
- Dragu, Tiberiu and Xiaochen Fan**, “The Political Economy of Environmental Regulation,” *American Journal of Political Science*, 2022, *66* (3), 579–596.
- Dustmann, Christian, Kristine Vasiljeva, and Anna Piil Damm**, “On the Economics and Politics of Refugee Migration,” *Economic Policy*, 2019, *34* (98), 497–550.
- Enke, Benjamin**, “Moral Values and Voting,” *Journal of Political Economy*, 2020, *128* (10), 3679–3729.
- Fetzer, Thiemo**, “The Political Costs of Austerity,” *American Economic Review*, 2019, *109* (7), 3849–3886.
- Gehrsitz, Markus**, “The Effect of Low Emission Zones on Air Pollution and Infant Health,” *Journal of Environmental Economics and Management*, 2017, *83*, 121–144.

- Gethin, Amory, Clara Martínez-Toledano, and Thomas Piketty**, *Political Cleavages and Social Inequalities: A Study of Fifty Democracies, 1948–2020*, Harvard University Press, 2021.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Laakso, Markku and Rein Taagepera**, ““Effective” Number of Parties: A Measure with Application to West Europe,” *Comparative Political Studies*, 1979, *12* (1), 3–27.
- Piketty, Thomas**, “Brahmin Left vs Merchant Right: Rising Inequality and the Changing Structure of Political Conflict,” *WID Working Paper*, 2018.
- 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, Pedro H.C. Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, *235* (2), 2218–2244.
- Stokes, Leah C.**, “Does Environmental Policy Cause Political Backlash? Wind Turbines and Right-Wing Populism,” *American Journal of Political Science*, 2016, *60* (2), 330–348.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.
- Wolff, Hendrik**, “Keep Your Clunker in the Suburb: Low-emission Zones and Adoption of Green Vehicles,” *Economic Journal*, 2014, *124* (578), F481–F512.

A. Data Appendix

A.1 Data Sources

1. **Election Results.** Candidate-level first-round legislative election results from *data.gouv.fr* (Parquet). Elections: 2002–2024.
2. **ZFE Boundaries.** GeoJSON from the BNZFE. File: `aires.geojson`. Downloaded March 2026. Polygon geometries for all active ZFE perimeters. Reprojected to Lambert-93.
3. **Constituency Boundaries.** GeoJSON file from *data.gouv.fr*. 577 metropolitan France constituencies. Reprojected to Lambert-93.
4. **Roll-Call Votes.** JSON data from the NosDéputés.fr API. Legislatures 14–17 (2012–2024). Each record contains vote date, title, and outcome. Total: 8,499 recorded votes.
5. **Deputy Information.** JSON data from NosDéputés.fr. Current legislature only. Contains 618 deputies with party group and constituency identifiers.
6. **ZFE Implementation Timeline.** Manually constructed from government announcements, ADEME reports, and press coverage. Contains 24 metropolitan areas with ZFE start dates, implementation wave (1 or 2), mandate type (air quality or population threshold), and 2024 tier classification (effectif or vigilance).

A.2 Sample Construction

The analysis panel is constructed as follows:

1. Start with all polling-station-level candidate results for legislative first rounds.
2. Map polling stations to constituencies using the `code_circonscription` field in the general results file.
3. Aggregate candidate votes to the constituency level, computing total votes per candidate.
4. Compute ENP, party-bloc vote shares (RN, Green), and turnout at the constituency-year level.
5. Merge with ZFE exposure (spatial intersection of constituency and ZFE boundaries).

6. Merge with ZFE timeline for treatment assignment.
7. Final panel: 3,405 constituency-year observations (603 constituencies \times up to 6 elections, with some constituencies not present in all years due to redistricting).

A.3 Variable Definitions

Table 6: Variable Definitions

Variable	Definition
ENP	Laakso–Taagepera effective number of parties: $1/\sum_j s_j^2$
RN Share	Combined vote share of FN, RN, EXD, and REC nuances
Green Share	Combined vote share of ECO, VEC, and DVE nuances
Turnout	Ratio of <i>votes</i> to <i>inscrites</i> (available for 2012, 2017, 2022 only)
ZFE Share	Fraction of constituency area covered by ZFE boundaries
Treated	Time-invariant binary: 1 if constituency ever has ZFE Share > 0.01
Post	Time-varying binary: 1 if Treated = 1 and election year \geq treatment year
Cohort	Treatment year (2022 or 2024) for treated; 0 for never-treated
Wave	1 = air-quality mandate (LOM 2019); 2 = population mandate (2021 law)

Table 7: ZFE Implementation Timeline

Agglomeration	Dept.	ZFE Start	Wave	Treatment Election
Paris	75	Jul 2019	1	2022
Lyon	69	Jan 2020	1	2022
Toulouse	31	Mar 2022	1	2022
Nice	06	Jan 2022	1	2022
Montpellier	34	Jul 2022	1	2024
Reims	51	Jan 2022	1	2022
Saint-Étienne	42	Jan 2022	1	2022
Aix-Marseille	13	Sep 2022	1	2024
Rouen	76	Sep 2022	1	2024
Strasbourg	67	Jan 2023	1	2024
Grenoble	38	Jul 2023	1	2024
Clermont-Ferrand [‡]	63	Jul 2023	2	2024
Bordeaux	33	Jan 2025	2	Not yet treated
Lille	59	Jan 2025	2	Not yet treated
Nantes	44	Jan 2025	2	Not yet treated
Rennes	35	Dec 2024	2	Not yet treated
<i>+ 8 additional Wave 2 cities (Jan 2025, not yet treated in sample)</i>				

Notes: Wave 1 = air-quality exceedance mandate (LOM 2019). Wave 2 = population >150,000 mandate (*Loi Climat et Résilience* 2021). Treatment election is the first legislative election after ZFE activation. “Not yet treated” indicates ZFE activated after the most recent election (June 2024); these cities serve as not-yet-treated controls in the CS-DiD specification. Full list of all 24 cities in replication code.

[‡] Clermont-Ferrand is classified as Wave 2 (population mandate) because it was mandated by the *Loi Climat et Résilience* population threshold, not air-quality exceedance; however, it activated earlier than most Wave 2 cities.

B. Identification Appendix

B.1 Pre-Trends Test

The Callaway–Sant’Anna pre-test for the parallel trends assumption yields $p = 3 \times 10^{-5}$, decisively rejecting the null of no differential pre-trends. This reflects the large, significant pre-treatment event-study coefficients documented in [Table 2](#).

B.2 CS-DiD Dynamic Effects

Table 8 reports the full set of CS-DiD dynamic ATT estimates for both ENP and RN share. Pre-treatment effects (relative periods -22 to -7) are uniformly positive and significant for both outcomes, confirming the structural urban–rural gap. Post-treatment effects are close to zero for ENP and significantly negative for RN share.

Table 8: Callaway–Sant’Anna Dynamic ATT Estimates

Rel. Period	ENP		RN Share	
	ATT	SE	ATT	SE
-22	2.760	0.898	0.049	0.011
-20	1.946	0.539	0.055	0.003
-17	2.719	0.742	0.040	0.012
-15	2.010	0.596	0.067	0.007
-12	1.282	0.434	0.044	0.007
-10	1.442	0.476	0.028	0.003
-7	0.365	0.186	0.023	0.006
-5	(ref. period)		(ref. period)	
-2	(ref. period)		(ref. period)	
0	-0.108	0.125	-0.039	0.010
2	0.146	0.187	-0.076	0.011

Notes: Callaway–Sant’Anna dynamic ATT estimates using not-yet-treated as controls with universal base period. Relative periods measured in years from ZFE activation. Reference periods are -5 and -2 .

B.3 Randomization Inference

500 permutations of treatment assignment (holding the number of treated constituencies fixed at 59) yield a RI p -value of 0.000 for the TWFE ENP coefficient, confirming that the observed estimate is in the extreme tail of the permutation distribution. However, this reflects the structural urban–rural gap: randomly-selected sets of 59 constituencies are very unlikely to replicate the systematic characteristics of France’s major metropolitan areas.

C. Robustness Appendix

C.1 Wave-Specific Event Study

Restricting to Wave 1 treated constituencies (those mandated by air-quality exceedance) yields event-study coefficients virtually identical to the full sample: 2.18 (2002), 2.17 (2007), 1.19 (2012) for ENP, confirming that results are not driven by the less-binding Wave 2 cities.

C.2 Donut Specification

Excluding constituencies with 0–50% ZFE area overlap (keeping only fully-treated and never-treated units) yields a TWFE ENP coefficient of -1.272 ($SE = 0.314$) and a RN share coefficient of -0.113 ($SE = 0.0085$), both very similar to the baselines (-1.396 and -0.0843 respectively). This indicates that partial-overlap constituencies do not drive the results.

C.3 Turnout Effects

Voter turnout shows a significant positive differential change between ZFE and non-ZFE constituencies (TWFE coefficient = $+0.032$, $SE = 0.004$, $p < 0.001$) in the 2012–2022 sample. This suggests that ZFE constituencies experienced a relative increase in voter participation after treatment, potentially reflecting heightened political engagement in areas where environmental policy is locally salient. However, with only three election years in the turnout sample (2012, 2017, 2022) and post-treatment variation coming exclusively from Wave 1 cities in 2022, this result should be interpreted cautiously.

D. Heterogeneity Appendix

D.1 Wave-Specific Treatment Effects

Interacting the treatment indicator with wave dummies reveals that Wave 1 effects are approximately twice as large as Wave 2 effects: the TWFE ENP coefficient is -1.09 ($SE = 0.30$) for Wave 1 and -0.46 ($SE = 0.16$) for Wave 2. This is consistent with Wave 1 cities being larger and more urban, reinforcing the interpretation that results reflect the urban–rural cleavage rather than policy intensity.

D.2 Deputy Party Composition

Of the 93 deputies from ZFE constituencies, 41 belong to Renaissance (REN, centrist-liberal), 25 to La France Insoumise (LFI, left), 9 to the Ecologists, and the remainder to LR, MoDem,

PS, HOR, and crossbench groups. No RN deputy represents a ZFE constituency, consistent with the low far-right vote shares in these urban areas. The absence of RN representation in ZFE areas limits the ability to test whether constituency-level ZFE exposure affects individual MPs' climate-related voting behavior.