The Green Backlash That Wasn't: Vehicle Bans and Populist Voting at French Low-Emission Zone Boundaries

APEP Autonomous Research* @olafdrw

March 13, 2026

Abstract

France abolished its Low-Emission Zones (ZFE) in 2025, citing popular opposition and fears of populist backlash. I exploit the sharp geographic boundaries of ZFE perimeters—which ban older vehicles inside but not outside—as a spatial discontinuity to test whether vehicle bans causally increased support for the Rassemblement National. Using commune-level presidential election results (2012–2022) matched to official ZFE boundary polygons across five metropolitan areas, I find no discontinuity in the change in RN vote share at ZFE boundaries. The point estimates are small, negative, and statistically insignificant across all bandwidths and specifications. The apparent correlation between ZFE zones and populist voting reflects the urban-suburban gradient, not a causal effect of vehicle restrictions. France may have reversed a major environmental policy based on a misdiagnosis of voter behavior.

JEL Codes: D72, Q53, Q58, R41

Keywords: low-emission zones, populism, environmental regulation, spatial discontinuity, France

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

1. Introduction

In February 2025, the French Senate voted to suspend enforcement of Zones à Faibles Émissions (ZFE), the country's flagship vehicle emission restriction policy. The stated justification was simple: the zones had become politically toxic. Commentators, mayors, and legislators across the political spectrum pointed to the Rassemblement National's (RN) surging vote share in cities with active ZFEs as evidence that banning older vehicles had fueled a populist backlash (Douenne and Fabre, 2022). France thus became the first major economy to reverse a comprehensive urban emission regulation on explicitly political grounds.

But was the backlash real? The raw correlation is suggestive: communes inside ZFE boundaries, which face restrictions on Crit'Air 4 and 5 vehicles, lean more heavily toward the RN than communes just outside. Yet this comparison confounds the ZFE's effect with the urban-suburban gradient that structures French political geography. ZFE boundaries trace the perimeters of dense metropolitan cores, precisely where the "Brahmin left" vs. "merchant right" cleavage documented by Gethin et al. (2022) produces the steepest voting gradients.

This paper provides the first causal test of whether ZFE implementation increased populist voting. I exploit the ZFE perimeter as a spatial discontinuity: communes whose centroids fall just inside the boundary are subject to vehicle bans, while communes just outside are not. The key outcome is the *change* in RN first-round presidential vote share from 2017 (before most ZFEs) to 2022 (after implementation in five metropolitan areas). This difference-in-discontinuities design (Grembi et al., 2016) differences out permanent level differences in voting between urban and suburban communes, isolating the effect of the ZFE itself.

Using official ZFE boundary polygons from the Base Nationale des ZFE matched to commune-level election results for over 3,800 communes near five active ZFE metropolitan areas, I find no evidence of populist backlash at ZFE boundaries. The estimated discontinuity in the 2017–2022 change in RN vote share is −2.2 percentage points at the baseline 5 km bandwidth (SE = 2.1 pp, $p = 0.29$). The result is stable across bandwidths from 3 to 20 km, robust to quadratic polynomials, donut specifications that drop communes within 500 meters of the boundary, and leave-one-metro-out jackknife. Results are similar for the total far-right vote share (Le Pen plus Zemmour).

A placebo test using the 2012–2017 change—before any ZFE existed—shows no significant discontinuity at the boundary once metropolitan area fixed effects are included ($p > 0.18$ at all bandwidths), confirming that the design is not picking up spurious pre-existing trends. The McCrary density test shows no bunching of communes at the boundary ($p = 0.51$).

Critically, I show that omitting metropolitan area fixed effects produces a *large and significant positive* effect of +17.6 pp ($p < 0.001$ at 3 km). This spurious result arises because different metropolitan areas have different ZFE zone sizes and different baseline voting levels. Without conditioning on which metro a commune belongs to, the estimate reflects cross-metro composition, not within-metro boundary effects. This decomposition illustrates precisely how the raw correlation between ZFEs and RN voting—the correlation that informed France's policy reversal—can arise without any causal link.

This paper contributes to three literatures. First, it extends the nascent literature on political backlash to environmental regulation (Colantone et al., 2024; Douenne and Fabre, 2022), providing the first quasi-experimental evidence specifically on vehicle emission zones. While Douenne and Fabre (2022) document that pessimistic beliefs about personal welfare losses drive carbon tax aversion, I test whether spatially targeted vehicle bans translate those attitudes into actual voting behavior—and find they do not.

Second, the paper contributes to the literature on low-emission zone effectiveness (Wolff, 2014; Gehrsitz, 2017). The existing literature has evaluated LEZs' effects on air quality and health but not on political outcomes. Given that over 300 LEZs operate across European cities, understanding their political sustainability is a first-order question.

Third, the paper provides a cautionary tale for spatial regression discontinuity designs at urban policy boundaries. LEZ and clean air zone perimeters coincide with urban-suburban gradients in income, demographics, and voting (Keele and Titiunik, 2015). Without conditioning on the metropolitan area, even a difference-in-discontinuities design can produce spurious effects driven by cross-city composition. Researchers studying politically charged urban regulations—congestion pricing, rent control boundaries, school district lines—face the same identification challenge.

The finding that ZFEs did not cause measurable populist backlash has immediate policy relevance. If France reversed its emission zone policy based on a misreading of electoral geography, the welfare cost may be substantial: Gehrsitz (2017) estimates that German LEZs reduced infant mortality, and Berestycki et al. (2024) finds significant air quality improvements in early French ZFE cities. Abandoning a health-improving policy because of a correlation that turns out to be confounded represents a distinctive failure mode for evidence-based regulation.

2. Institutional Background

2.1 French Low-Emission Zones

France's ZFE system was established by the 2019 Mobilités Law and expanded by the 2021 Climate and Resilience Law, which mandated ZFEs in all agglomerations exceeding 150,000 inhabitants by 2025. The system operates through Crit'Air vignettes: colored stickers classifying vehicles by emission standards from 0 (electric) to 5 (pre-2001 diesel). ZFE regulations prohibit vehicles above a specified Crit'Air threshold from circulating within the zone perimeter during restricted hours.

Implementation was staggered across metropolitan areas. Paris (Ile-de-France) activated its ZFE in June 2021, initially restricting Crit'Air 4 vehicles for private cars. Reims and Saint-Étienne followed in early 2022, and Toulouse activated its first restrictions in March 2022. Several additional metros (Montpellier, Lyon, Strasbourg) implemented ZFEs after the April 2022 presidential election. By 2024, twelve metropolitan areas had active ZFEs with varying levels of stringency.

2.2 The ZFE Boundary

The ZFE perimeter is a legally defined geographic boundary, typically following ring roads, administrative limits, or major infrastructure corridors. Inside the boundary, vehicles exceeding the Crit'Air threshold face fines of €68 (rising to €135 for repeat offenses). Outside the boundary, no restriction applies.

For the spatial discontinuity design, two features of the boundary are critical. First, the boundary is sharp: a vehicle legal on one side of a road becomes illegal on the other. Second, the boundary was not drawn to follow pre-existing voting patterns. ZFE perimeters were determined by air quality monitoring, road network topology, and administrative convenience—not by political considerations. The boundary does, however, tend to follow the urban core's outer ring, which means it correlates with the urban-suburban socioeconomic gradient.

2.3 Political Context

Marine Le Pen's Rassemblement National (formerly Front National) progressed from 17.9% in the 2012 first round to 21.3% in 2017 and 23.1% in 2022 nationally. However, the far-right landscape shifted dramatically in 2022 with Éric Zemmour's candidacy, which split the far-right vote. In many urban areas, RN vote share actually declined from 2017 to 2022 as Zemmour captured former Le Pen voters. This makes the total far-right share (Le Pen plus Zemmour) a more appropriate measure of populist sentiment in 2022.

Critics of ZFEs argued that vehicle bans disproportionately affected working-class households who could not afford newer vehicles, channeling the same frustrations that produced the *gilets jaunes* movement in 2018–2019 (Douenne and Fabre, 2022). If this narrative is correct, communes inside ZFE boundaries should show a relative increase in populist voting compared to communes just outside.

3. Data

I combine four data sources. The ZFE boundary polygons come from the Base Nationale des ZFE (BNZFE), published on transport.data.gouv.fr in GeoJSON format. The BNZFE contains the perimeter geometry, implementation dates, and Crit'Air thresholds for each ZFE phase.

Commune-level election results come from data.gouv.fr for three presidential elections (2012, 2017, 2022 first round). These provide the universe of votes by candidate for all 35,000+ French communes. I extract Front National/Rassemblement National vote counts, total expressed votes, registered voters, and turnout for each election.

Commune centroids come from the geo.api.gouv.fr geographic API, providing latitude, longitude, population, and administrative codes for 34,877 communes. I compute signed geodesic distances from each commune centroid to the nearest ZFE boundary in the Lambert-93 projected coordinate system.

The analysis sample consists of 3,816 communes within 50 km of a ZFE boundary in one of the five metropolitan areas with active ZFEs before the April 2022 presidential election (Paris, Toulouse, Reims, Saint-Étienne, and Grenoble/Nice area). Of these, 50 commune centroids fall inside a ZFE boundary.

3.1 Summary Statistics

Table 1: Summary Statistics: Communes Within 10 km of ZFE Boundary

	Inside ZFE		Outside ZFE		Full Sample	
	Mean	SD	Mean	SD	Mean	SD
Panel A: Election Outcomes (proportion)						
FN vote share, 2012	0.176	0.055	0.164	0.062	0.165	0.061
FN vote share, 2017	0.298	0.032	0.317	0.058	0.315	0.057
RN vote share, 2022	0.203	0.065	0.192	0.082	0.193	0.080
Far-right total, 2022	0.277	0.067	0.266	0.086	0.267	0.084
Turnout, 2017	0.793	0.055	0.827	0.053	0.823	0.055
Turnout, 2022	0.782	0.054	0.801	0.056	0.799	0.056
Panel B: Changes (percentage points)						
ΔFN 2012–2017	0.122	0.076	0.153	0.100	0.150	0.098
ΔRN 2017–2022	-0.095	0.085	-0.126	0.118	-0.123	0.115
ΔFar-right 2017–2022	-0.021	0.086	-0.051	0.119	-0.048	0.116
ΔTurnout 2017–2022	-0.011	0.021	-0.025	0.029	-0.024	0.028
Panel C: Geography						
Distance to boundary (km)	1.5	1.2	5.4	2.8	5.0	2.9
Population	67454	303897	14897	24026	20293	100420
Observations	50		437		487	

Notes: Sample restricted to communes within 10 km of a ZFE boundary in one of the 5 metropolitan areas with active ZFEs before the April 2022 presidential election. FN/RN vote shares refer to Front National (2012) or Rassemblement National (2017, 2022) first-round presidential results. Far-right total adds Zemmour votes in 2022. All vote shares are proportions of expressed votes.

Table 1 presents summary statistics for communes within 10 km of a ZFE boundary. Inside-ZFE communes have lower baseline RN vote shares (reflecting their urban location) and larger populations. The 2017–2022 change in RN share is negative in both groups, consistent with the national pattern of Zemmour splitting the far-right vote.

4. Empirical Strategy

4.1 Difference-in-Discontinuities

The spatial discontinuity at the ZFE boundary cannot be exploited in a simple cross-sectional RDD because permanent differences in urban and suburban voting would confound the estimate. Instead, I implement a difference-in-discontinuities design (Grembi et al., 2016) that compares the *change* in voting outcomes between elections:

$$\Delta Y_c = \alpha + \tau \cdot \mathbf{1}[\text{Inside}_c] + \beta_1 \text{Dist}_c + \beta_2 \mathbf{1}[\text{Inside}_c] \times \text{Dist}_c + \delta_m + \varepsilon_c \tag{1}$$

where ΔY_c is the change in RN vote share from 2017 to 2022, $\mathbf{1}[\text{Inside}_c]$ indicates whether commune c's centroid falls inside a ZFE boundary, Dist_c is the signed distance to the nearest boundary in kilometers (positive inside, negative outside), δ_m are metropolitan area fixed effects, and τ is the parameter of interest.

The metropolitan area fixed effects are essential. Without them, the estimate captures both within-metro boundary effects and cross-metro differences in ZFE zone sizes and voting levels. I weight observations by a triangular kernel $K_h(d) = (1 - |d|/h) \cdot \mathbf{1}[|d| \leq h]$ for bandwidth h.

4.2 Identifying Assumptions

The identifying assumption is that *conditional on metropolitan area*, potential voting outcomes vary smoothly through the ZFE boundary. Two testable implications support this:

First, the **placebo test**: the 2012–2017 change in FN vote share (before any ZFE existed) should show no discontinuity at the boundary. I verify this at all bandwidths.

Second, the **McCrary density test**: commune centroids should not bunch at the boundary. Since communes are administrative units whose boundaries predate ZFEs by decades, manipulation is mechanically impossible.

The key threat is that ZFE boundaries may trace pre-existing gradients in voting trends—not just levels—between urban cores and suburbs. The metropolitan area fixed effects and the use of differenced outcomes address level differences; the placebo test addresses trend differences.

4.3 Threats to Validity

Several concerns warrant discussion. First, the number of treated communes (50 within the sample, 45 within 3 km) is small. While this is inherent to the sharp geographic nature of

ZFE zones—they cover dense urban cores with few distinct communes—it limits statistical power. I report the minimum detectable effect size to contextualize null findings.

Second, commune centroids may not perfectly capture ZFE exposure: a commune whose centroid is just outside the boundary may have substantial area inside the zone. This measurement error attenuates estimates toward zero, making it harder to detect a backlash.

Third, anticipation effects are possible: households may have responded to the *announcement* of ZFEs before implementation, in which case the pre/post comparison would understate effects. However, ZFE announcement dates generally preceded implementation by only months, and the relevant political response would likely crystallize at the implementation date when fines began.

5. Results

5.1 Main Results

Table 2: Main Results: Effect of ZFE on Voting Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
Bandwidth (km)	3	5	7	10	15	20
Panel A: ΔRN Vote Share, 2017–2022						
Inside ZFE	-0.0307	-0.0224	-0.0208	-0.0092	0.0080	0.0180
	(0.0254)	(0.0209)	(0.0192)	(0.0178)	(0.0168)	(0.0163)
Panel B: ΔFar-Right Total, 2017–2022						
Inside ZFE	-0.0398	-0.0253	-0.0219	-0.0073	0.0126	0.0236
	(0.0268)	(0.0212)	(0.0192)	(0.0177)	(0.0165)	(0.0159)
Panel C: ΔTurnout, 2017–2022						
Inside ZFE	-0.0006	-0.0107	-0.0122**	-0.0111**	-0.0093*	-0.0076
	(0.0083)	(0.0067)	(0.0059)	(0.0054)	(0.0050)	(0.0048)
Panel D: Placebo – ΔFN Vote Share, 2012–2017						
Inside ZFE	0.0316	0.0240	0.0210	0.0099	-0.0039	-0.0105
	(0.0237)	(0.0194)	(0.0178)	(0.0165)	(0.0157)	(0.0153)
Observations	150	243	339	487	768	1119
Communes inside	45	50	50	50	50	50
Metro FE	Yes	Yes	Yes	Yes	Yes	Yes
Kernel	Tri.	Tri.	Tri.	Tri.	Tri.	Tri.
Polynomial	Linear	Linear	Linear	Linear	Linear	Linear

Notes: Each cell reports the coefficient on the Inside ZFE indicator from a local linear regression of the outcome on Inside ZFE, signed distance to ZFE boundary (km), and their interaction, with metropolitan area fixed effects and triangular kernel weights. Heteroskedasticity-robust standard errors in parentheses. ΔRN is the change in Rassemblement National (formerly Front National) first-round presidential vote share. Far-right total adds Zemmour votes in 2022. Panel D reports a placebo test using the 2012–2017 change, before any ZFE existed. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 reports the main results. Panel A shows the effect on the 2017–2022 change in RN vote share across six bandwidths. The estimated discontinuity is -3.1 pp at 3 km (SE = 2.5 pp), -2.2 pp at 5 km (SE = 2.1 pp), and converges toward zero at wider bandwidths. None of the estimates is statistically significant at conventional levels.

Panel B shows similar results for the total far-right share (Le Pen plus Zemmour), addressing the concern that Zemmour's candidacy redistributed votes within the far right. The estimates are slightly more negative (-4.0 pp at 3 km) but remain insignificant.

Panel C examines turnout. The estimates suggest a small decline in turnout inside ZFE communes at intermediate bandwidths (-1.2 pp at 7 km, $p = 0.04$), though this is not robust across all bandwidths and should be interpreted cautiously given the number of tests.

Panel D reports the placebo: the 2012–2017 change in FN vote share. No estimate is statistically significant ($p > 0.18$ at all bandwidths), confirming that the ZFE boundary does not pick up pre-existing differential voting trends once metropolitan area fixed effects are included.

5.2 Specification Sensitivity

Table 3: Specification Sensitivity: ΔRN Vote Share at 10 km Bandwidth

	Inside ZFE	SE	N
Baseline (linear, tri.)	-0.0092	(0.0178)	487
Uniform kernel	0.0088	(0.0175)	487
Quadratic	-0.0273	(0.0245)	487
Donut (>500m)	-0.0002	(0.0222)	468
No metro FE	0.1138***	(0.0238)	487

Notes: All specifications estimate the discontinuity in ΔRN (2017–2022 change in RN first-round presidential vote share) at the ZFE boundary using a 10 km bandwidth. Heteroskedasticity-robust standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 shows that the null result is robust to alternative specifications at the 10 km bandwidth. The baseline linear specification, uniform kernel, quadratic polynomial, and donut (excluding communes within 500m) all produce small, insignificant estimates ranging from -2.7 to $+0.9$ pp.

The final row removes metropolitan area fixed effects, producing a spurious estimate of

+11.4 pp ($p < 0.001$). This decomposition is the paper's core methodological finding: the raw positive correlation between ZFE location and RN voting gains that animated France's policy debate disappears entirely once one compares communes *within the same metropolitan area*. The apparent backlash was an artifact of cross-metro composition.

5.3 Covariate Balance

Table 4: Covariate Balance at ZFE Boundary

	Estimate	SE	p-value	BW (km)	Eff. N
FN vote share, 2012	0.1033***	(0.0287)	0.000	2.5	123
FN vote share, 2017	-0.0461**	(0.0214)	0.031	1.9	104
Turnout, 2017	-0.0061	(0.0244)	0.801	3.0	149
Population (log)	-2.0582***	(0.4907)	0.000	2.4	122

Notes: Each row reports the `rdrobust` estimate of the discontinuity in a predetermined covariate at the ZFE boundary. Data-driven MSE-optimal bandwidth with triangular kernel. Nearest-neighbor standard errors. Population is log-transformed.

Table 4 reports covariate balance tests using `rdrobust` at data-driven bandwidths. Predetermined covariates show some imbalance at the boundary: the 2012 FN share is higher outside (10.3 pp, $p < 0.001$) and population is much larger inside ($p < 0.001$), reflecting the urban-suburban gradient. However, these level differences are precisely what the difference-in-discontinuities design addresses by using changes rather than levels. The 2017 FN share shows a smaller imbalance (-4.6 pp, $p = 0.03$), and turnout in 2017 is balanced ($p = 0.80$).

The McCrary density test yields $p = 0.51$, confirming no manipulation of commune placement relative to the ZFE boundary.

5.4 Power and Minimum Detectable Effect

With 50 treated communes and standard errors of approximately 2 pp at the 5 km bandwidth, the minimum detectable effect at 80% power and 5% significance is roughly 4 pp. The actual confidence interval at 5 km rules out effects larger than +1.9 pp or more negative than -6.4 pp. Thus, while the design cannot detect very small effects, it can rule out the large localized backlash effects (on the order of 5–10 pp) that would be needed to explain France's policy reversal through a boundary mechanism.

Two caveats temper this power analysis. First, the centroid-based treatment assignment introduces measurement error—communes whose centroids fall just outside the ZFE may have substantial area inside the zone—which attenuates estimates toward zero. If the true effect is concentrated among households directly exposed to vehicle bans, the commune-level design dilutes it. Second, the small number of metropolitan areas (five) limits the precision of inference that relies on cross-metro variation; wild cluster bootstrap p-values would be preferable but are infeasible with so few clusters. These limitations mean the null should be read as "no evidence of large localized effects at the ZFE boundary" rather than definitive absence of any political response to emission regulation.

6. Discussion

The central finding is that there is no evidence of a localized populist voting backlash at French Low-Emission Zone boundaries. The estimates rule out effects larger than 2 percentage points at moderate bandwidths. This does not mean voters are indifferent to ZFEs—it means that whatever opposition exists does not concentrate discontinuously at the regulated boundary. The absence of a boundary discontinuity is consistent with both no backlash and with a diffuse backlash that affects all voters in a metropolitan area equally, independent of ZFE proximity.

Several mechanisms could explain this null. First, enforcement was weak in many cities during the period studied, with fines rarely issued before automated camera systems were deployed. Voters may not have experienced the ZFE as binding. Second, the ZFE's costs (restricting older vehicles) and benefits (cleaner air) may have roughly offset in voter perceptions. Third, political backlash to environmental regulation may operate through diffuse national narratives—amplified by media and political entrepreneurs—rather than through spatially localized treatment effects. This last mechanism is important: my estimand captures the *localized boundary effect* of living inside versus outside the zone, not the aggregate effect of ZFE policy on metropolitan or national voting. If voters on both sides of the boundary respond to the broader political salience of ZFEs—through media coverage, campaign rhetoric, or anticipation of future restrictions—then the boundary estimate would understate the total political effect.

A further consideration is temporal heterogeneity. The five metropolitan ZFEs have different exposure durations relative to the April 2022 election: Grenoble implemented in May 2019 (nearly three years), Paris in June 2021 (ten months), while Reims, Saint-Étienne, and Toulouse activated within months of the vote. If backlash requires time to crystallize, the recent adopters may not have had sufficient exposure, while Grenoble's longer exposure might reveal different dynamics. The pooled null result is robust to the leave-one-metro-out analysis, but separating early from late adopters more carefully would require more statistical power than the current sample affords.

The paper's methodological contribution is equally important. I show that the raw correlation between ZFE zones and RN voting—which sustained a major policy debate—is entirely an artifact of the urban-suburban gradient. Without metropolitan area fixed effects, a spatial regression discontinuity at the ZFE boundary produces a large, significant, and completely spurious "backlash" effect. This finding has implications for any spatial RDD conducted at urban policy boundaries: congestion zones, rent control boundaries, school districts, and enterprise zones all face the same identification challenge.

7. Conclusion

France abandoned its urban emission zone policy citing populist backlash. This paper finds no evidence of localized backlash at ZFE boundaries: the change in Rassemblement National vote share is smooth through the regulated perimeter once metropolitan area composition is addressed. The correlation that alarmed policymakers was a confound: the boundary between zones and non-zones is also the boundary between urban and suburban France, and these areas vote differently for reasons that predate any vehicle ban.

The welfare implications are potentially large. If ZFEs reduce air pollution and improve health in dense urban areas—as evidence from German LEZs suggests—then reversing the policy based on a misdiagnosed political signal imposes real costs. The lesson extends beyond France: as cities worldwide implement emission regulations, the temptation to read electoral geography as evidence of policy effects will recur. Spatial boundaries that delineate regulatory treatment also delineate pre-existing social gradients. Disentangling the two requires the kind of within-metro comparison conducted here, not the naive raw correlations that typically enter political debate.

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

Berestycki, Etienne, Benjamin Bureau, and Matthieu Glachant, "Les effets des zones à faibles émissions sur la qualité de l'air," Technical Report, France Stratégie 2024.

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

Colantone, Italo, Livio Di Lonardo, Yotam Margalit, and Marco Percoco, "The Political Consequences of Green Policies: Evidence from Italy," *American Political Science Review*, 2024, *118* (1), 108–126.

Douenne, Thomas and Adrien Fabre, "Yellow Vests, Pessimistic Beliefs, and Carbon Tax Aversion," *American Economic Journal: Economic Policy*, 2022, *14* (1), 81–110.

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, "Brahmin Left Versus Merchant Right: Changing Political Cleavages in 21 Western Democracies, 1948–2020," *Quarterly Journal of Economics*, 2022, *137* (1), 1–48.

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano, "Do Fiscal Rules Matter?," *American Economic Journal: Applied Economics*, 2016, *8* (3), 1–30.

Keele, Luke J. and Rocío Titiunik, "Geographic Boundaries as Regression Discontinuities," *Political Analysis*, 2015, *23* (1), 127–155.

Wolff, Hendrik, "Keep Your Clunker in the Suburb: Low-Emission Zones and Adoption of Green Vehicles," *The Economic Journal*, 2014, *124* (578), F481–F512.

A. Data Appendix

A.1 ZFE Boundary Data

The Base Nationale des ZFE (BNZFE) is published by the French Ministry of Transport on transport.data.gouv.fr. The dataset contains 35 GeoJSON polygon features representing ZFE perimeters across France, including implementation dates (`date_debut`), Crit'Air thresholds for different vehicle categories (`vp_critair` for passenger vehicles, `vul_critair` for light commercial), and operating hours. Multiple features may correspond to the same metropolitan area if the ZFE was implemented in phases with different dates or thresholds.

I identify metropolitan areas using the SIREN prefix of the `id` field and union all phase polygons to create a single boundary per metro. The earliest `date_debut` across phases determines when the ZFE became active. Five metropolitan areas had active ZFEs before the April 10, 2022 presidential election: Paris/Ile-de-France (June 2021), Reims (January 2022), Saint-Étienne (January 2022), Toulouse (March 2022), and Grenoble (May 2019).

A.2 Election Data

Presidential election results come from the French Ministry of the Interior via data.gouv.fr:

- 2012 first round: CSV, 36,791 communes, commune-level totals
- 2017 first round: XLS, 35,719 communes, commune-level totals
- 2022 first round: TXT (semicolon-separated), sub-commune level (aggregated to commune)

For 2012, Marine Le Pen ran as Front National; for 2017 and 2022, as Rassemblement National. Vote shares are computed as candidate votes divided by expressed (valid) votes.

A.3 Distance Computation

Commune centroids are obtained from the geo.api.gouv.fr API (34,877 communes). All coordinates are transformed to the Lambert-93 projected CRS (EPSG:2154) for distance computation. For each commune, I compute the minimum geodesic distance to the boundary line (not filled polygon) of each metropolitan ZFE, identify the nearest metro, and assign a signed distance: positive if the commune centroid is inside the ZFE polygon, negative if outside.

B. Identification Appendix

B.1 McCrary Density Test

The McCrary density test (Cattaneo et al., 2020) yields a test statistic with $p = 0.51$, indicating no manipulation of commune placement at the ZFE boundary. This is expected: commune boundaries are administrative units established long before ZFE policy.

B.2 Placebo Test Detail

The 2012–2017 placebo uses the change in FN vote share before any ZFE existed. At the baseline 5 km bandwidth with metro FE, the estimated discontinuity is 2.4 pp (SE = 1.9 pp, $p = 0.22$). This null placebo supports the identifying assumption that, conditional on metropolitan area, voting trends are smooth through the ZFE boundary.

B.3 Role of Metropolitan Area Fixed Effects

The paper demonstrates that metro FE are essential for identification. Without them, the estimate at 3 km bandwidth is +17.6 pp ($p < 0.001$). With metro FE, it is −3.1 pp ($p = 0.23$). The difference arises because Paris (large ZFE, distinct voting patterns) dominates the without-FE specification.

C. Robustness Appendix

C.1 Leave-One-Metro-Out

At the 10 km bandwidth, dropping each of the five metropolitan areas in turn produces estimates ranging from −2.3 pp (drop Paris) to −0.4 pp (drop Reims), all insignificant. No single metro drives the null result.

C.2 Alternative Outcomes

The 2024 legislative election provides an additional post-ZFE observation. However, this election uses a constituency-based system where candidate availability varies, making cross-commune comparisons less clean than presidential elections. I include it as supplementary evidence but prioritize the presidential election for the main analysis.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
ΔRN share 2017–2022	−0.0224	0.0209	—	0.1226	−0.1826	0.1708	Large negative
ΔFar-right 2017–2022	−0.0253	0.0212	—	0.1233	−0.2048	0.1721	Large negative
ΔTurnout 2017–2022	−0.0107	0.0067	—	0.0248	−0.4316	0.2687	Large negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison. SDE $= \hat{\beta}/\text{SD}(Y)$ for binary treatment. SD(Y) is the unconditional standard deviation of the outcome. **Research question:** Does ZFE implementation affect RN voting and turnout in nearby communes? **Treatment:** Binary (inside vs. outside ZFE boundary). **Data:** French presidential election results (2017, 2022) matched to ZFE boundary distances for 243 communes within 5 km of a ZFE boundary. **Method:** Local linear diff-in-disc with metro FE, triangular kernel, 5 km bandwidth. **Sample:** Communes near 5 metropolitan ZFE areas active before April 2022.

Classification thresholds: large negative (< -0.15), moderate negative (-0.15 to -0.05), small negative (-0.05 to -0.005), null (-0.005 to 0.005), small positive (0.005 to 0.05), moderate positive (0.05 to 0.15), large positive (> 0.15).

Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. "Null" denotes a near-zero effect size ($|\text{SDE}| < 0.005$), not a failure to reject a null hypothesis.