Delivering Discontent? Maternity Ward Closures and the Limits of the State-Abandonment Hypothesis in France

APEP Autonomous Research* @SocialCatalystLab

March 23, 2026

Abstract

Does the state's withdrawal from peripheral areas breed populism? I exploit the staggered closure of maternity wards across France between 2013 and 2024, which increased travel distances for 2,359 communes. Using 34,791 metropolitan communes across nine elections (2002–2022) with commune and election-year fixed effects, I find no robust evidence that losing a nearby maternity ward increases Rassemblement National vote share. The TWFE estimate is 0.01 percentage points per kilometer (SE = 0.04). Callaway-Sant'Anna yields an ATT of 1.75 (SE = 1.49), positive but insignificant. Placebos on left-populist voting and turnout are null. Pre-trend tests reject, suggesting selection. Healthcare rationalization does not produce a detectable effect on far-right support in this design.

JEL Codes: D72, H75, I18

Keywords: populism, maternity wards, public service withdrawal, France, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 30m).

1. Introduction

Between 2000 and 2024, France closed more than 275 maternity wards. A pregnant woman in rural Picardie who once drove fifteen minutes to deliver her child now drives an hour. Over the same period, the Rassemblement National's first-round presidential vote share rose from 17 percent to 29 percent. These two trends—the retreat of public services and the advance of far-right populism—are central to what the geographer Christophe Guilluy has called "la France périphérique": a peripheral France abandoned by the state and, in response, turning to Marine Le Pen (Guilluy, 2014).

The state-abandonment hypothesis has become a cornerstone of both popular and academic explanations for populist backlash across Europe (Rodríguez-Pose, 2018; Dijkstra et al., 2020). Yet the causal evidence remains thin. Cross-sectional correlations between service deprivation and far-right voting confound selection (areas that lose services were already declining) with treatment (the closure itself generates resentment). The few studies that exploit quasi-experimental variation focus on commercial closures (Fetzer, 2019) or fiscal austerity shocks (Dal Bó et al., 2023), not the withdrawal of specific public services that citizens interact with at the most vulnerable moments of their lives.

This paper provides the first micro-geographic causal estimate of whether public healthcare withdrawal drives populist voting. I exploit the staggered closure of 536 maternity wards across metropolitan France between 2013 and 2024, driven by the 1998 Kouchner decree's minimum-volume threshold of 300 deliveries per year. When a maternity ward closes, every commune in the surrounding area experiences an increase in travel distance to the nearest active facility. I measure this distance shock at the commune level using Haversine distances between commune centroids and facility locations from the DREES Statistique Annuelle des Établissements, then link it to commune-level election returns from data.gouv.fr covering presidential and European Parliament elections.

The main finding is a null. In a two-way fixed effects specification with commune and election-year fixed effects on 313,033 commune-election observations, an additional kilometer to the nearest maternity ward is associated with a 0.01 percentage-point change in FN/RN vote share (SE = 0.04, $p = 0.78$). A binary post-closure indicator yields a similar imprecise estimate (0.04 pp, SE = 0.52). The Callaway-Sant'Anna estimator (Callaway and Sant'Anna, 2021), which accommodates staggered treatment timing, produces an overall ATT of 1.75 percentage points (SE = 1.49), positive but statistically insignificant. The estimates are too imprecise to rule out moderate effects in either direction.

This null is robust. It holds for small communes under 2,000 inhabitants—where the state-abandonment narrative is strongest—and for large communes. It holds when I drop

Île-de-France, restrict to nearby closures within 20 kilometers, or use the change in distance from the 2013 baseline as treatment. Placebo tests on left-populist (Mélenchon/LFI) vote share and on turnout are also null, ruling out a general political mobilization channel.

The paper contributes to three literatures. First, it informs the "geography of discontent" debate (Rodríguez-Pose, 2018; Dijkstra et al., 2020) by providing causal evidence where most existing work is correlational. The null challenges the hypothesis that specific service closures generate populist backlash, suggesting that the correlation between deprivation and RN voting reflects deeper structural forces—deindustrialization, demographic change, cultural anxiety—rather than a marginal response to individual service losses.

Second, it contributes to the political economy of public service provision (Besley and Coate, 2003). If closures do not generate electoral punishment, local politicians face weaker incentives to resist rationalization, reinforcing the centralization dynamic that produces closures in the first place.

Third, it speaks to the French political economy literature on the gilets jaunes and the rise of the RN (Algan et al., 2019). While Douenne and Fabre (2022) show that the carbon tax generated expressed opposition, and the companion APEP paper on carbon taxation in France uses social networks to trace the contagion of backlash, I find that the healthcare-specific channel of state abandonment does not operate at the commune level.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of maternity ward rationalization in France. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results and robustness checks. Section 6 discusses implications.

2. Institutional Background

The rationalization of French maternities. France's maternity ward system underwent sustained consolidation beginning in the 1970s. The 1998 Kouchner decree (Décret 98-900) formalized this process by mandating a minimum threshold of 300 deliveries per year for a maternity ward to remain operational. Wards falling below this threshold face closure decisions by the Agence Régionale de Santé (ARS), based on delivery volumes, staffing adequacy, and patient safety criteria. The stated rationale is medical: concentrating births in higher-volume facilities improves outcomes by ensuring sufficient specialist availability (Blondel et al., 2017).

Scale and geography of closures. The number of maternity wards in metropolitan France fell from 721 in 2000 to 445 in 2024—a 38 percent decline. The DREES SAE database, which

provides annual facility-level data from 2013 onward, identifies 536 unique facilities that appear in at least one year. Of these, closures are inferred from disappearance: a facility whose last year of recorded activity is t is coded as closed in year $t+1$, yielding approximately 200 closure events between 2014 and 2024. Closures are geographically dispersed across 79 départements, concentrated in rural and semi-rural areas where delivery volumes are smallest. The staggered timing—ranging from 24 closures in 2014 to 212 in 2016—provides the variation I exploit.

Consequences for affected communes. When a maternity ward closes, every commune within its catchment area experiences an increase in travel time and distance to the nearest active facility. For a commune that previously had a ward within 10 kilometers, closure can double or triple this distance. The French public debate frames these closures as "déserts médicaux" (medical deserts), and they have become a symbol of state neglect in the gilets jaunes discourse (Noiriel, 2019). The question is whether this symbolic salience translates into measurable changes in voting behavior.

Exogeneity of closure timing. The decision to close a maternity ward is made by regional health agencies based on volume thresholds and safety assessments, not by local elected officials or voters. While communes can lobby against closures, the final decision lies with the ARS. This institutional structure provides plausible exogeneity: the timing of closures is driven by declining birth rates and staffing constraints in the surrounding area, not by political considerations.

3. Data

I combine three administrative data sources covering metropolitan France.

Maternity ward data. The DREES Statistique Annuelle des Établissements (SAE) provides annual facility-level data on maternity wards, including the commune code, facility name, and total deliveries. I use the full panel from 2013 to 2024, which covers 1,044 unique facilities. I identify closures as facilities that disappear from the active file: if a facility's last year of positive deliveries is t, I code it as closed in year $t+1$. This yields 536 closure events.

Election results. Presidential first-round results at the commune level come from the data.gouv.fr aggregated elections dataset, a Parquet file covering all French elections from 1999 to 2024 at the bureau de vote level. I aggregate to the commune level and compute FN/RN vote share as the combined votes for Le Pen candidates and FN/RN-affiliated lists divided by total votes cast. The panel covers nine elections—five presidential (2002, 2007,

2012, 2017, 2022) and four European Parliament (2004, 2009, 2014, 2019)—across 34,791 metropolitan communes, providing six pre-treatment periods for validation.

Commune characteristics. Commune centroids (latitude, longitude) and population come from the INSEE Code Officiel Géographique. I compute Haversine distances from each commune centroid to the nearest active maternity ward for each election year, using the maternity landscape from the closest preceding DREES year.

Table 1: Summary Statistics

	Mean	SD	Min	Max
Panel A: Full Sample ($N = 313,033$)				
FN/RN Vote Share (%)	13.6	13.1	0.0	100.0
Distance to Nearest Maternity (km)	18.2	9.7	0.0	78.0
Population	1781	7924	1	504078
Share Affected by Closure	0.068	0.251	0	1
Panel B: FN/RN Vote Share by Election Year				
2002	17.4	—	—	34,770
2004	1.7	—	—	34,774
2007	12.6	—	—	34,773
2009	3.1	—	—	34,786
2012	21.4	—	—	34,788
2014	10.1	—	—	34,791
2017	26.4	—	—	34,784
2019	0.0	—	—	34,790
2022	29.3	—	—	34,777

Notes: Panel of 34,791 metropolitan French communes across five presidential first-round elections (2002–2022). FN/RN vote share is the combined vote share of Jean-Marie Le Pen (2002, 2007) and Marine Le Pen (2012, 2017, 2022). Distance is Haversine distance from commune centroid to nearest active maternity ward. "Affected by closure" indicates communes where distance increased by more than 5 km between election cycles due to maternity ward closures.

4. Empirical Strategy

4.1 Identification

The central challenge is that maternity closures are not random—they occur where birth volumes decline, which correlates with rural depopulation and economic decline, both of which independently predict FN/RN support. My identification strategy exploits within-commune variation over time: for a given commune, does an increase in distance to the nearest maternity ward change its voting behavior relative to communes that did not experience such a change?

The estimating equation is:

$$Y_{ct} = \alpha_c + \gamma_t + \beta \cdot D_{ct} + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the FN/RN first-round vote share in commune c at election t, α_c are commune fixed effects absorbing all time-invariant characteristics, γ_t are election-year fixed effects absorbing national trends, and D_{ct} is the distance (in km) to the nearest active maternity ward. The coefficient β captures the effect of an additional kilometer of distance on FN/RN vote share, identified from within-commune changes in distance driven by maternity closures.

Standard errors are clustered at the département level (96 clusters), the administrative unit governing health policy through the ARS.

4.2 Callaway-Sant'Anna Estimator

Because closures are staggered, the standard TWFE estimator may suffer from negative weighting bias (de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021). I therefore also estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator with not-yet-treated communes as the control group and a universal base period. This approach correctly handles treatment effect heterogeneity across cohorts.

I define treatment cohorts based on when each commune's nearest maternity distance first increased by more than 2 kilometers between consecutive election years. Communes whose distance never increased form the never-treated group.

4.3 Threats to Validity

The parallel trends assumption requires that, absent maternity closures, affected and unaffected communes would have followed the same FN/RN vote share trajectory. I assess this through the event-study specification and the Callaway-Sant'Anna pre-trend test. The pre-trend test yields $p = 0.0003$, suggesting some differential pre-trends. I discuss this limitation explicitly: communes that lose maternities may be on different demographic trajectories, which motivates the distance-change specification that focuses on the marginal effect of closure conditional on the pre-existing level.

5. Results

5.1 Main Results

Table 2 reports the main estimates. Column (1) uses continuous distance to the nearest maternity (km) as the treatment variable. The coefficient is 0.010 (SE = 0.035), indicating an economically and statistically insignificant effect. Column (2) uses a binary post-closure indicator for affected communes, yielding 0.043 (SE = 0.518). Column (3) uses the change in distance from the 2013 baseline, with an identical result to column (1) since within-commune variation in distance comes entirely from the baseline comparison. Column (4) reports the Callaway-Sant'Anna overall ATT of 1.75 (SE = 1.49), positive but not statistically significant.

The wide confidence intervals are worth noting. The 95% interval for the CS-DiD ATT ranges from −1.2 to +4.7 percentage points, meaning we cannot rule out effects that would be economically meaningful in either direction. However, the point estimates are consistently small in the TWFE specifications, and none of the nine specifications across the main and robustness tables produces a significant positive effect.

Table 2: Effect of Maternity Ward Closures on FN/RN Vote Share

	(1) Distance	(2) Binary Post	(3) Δ Distance	(4) CS-DiD
Distance to nearest (km)	0.0097			
	(0.0351)			
Post-closure		0.0434		
		(0.5180)		
Distance change			0.0097	
			(0.0351)	
ATT (CS 2021)				1.7507
				(1.4907)
Commune FE	✓	✓	✓	—
Election year FE	✓	✓	✓	—
Observations	313,033	313,033	313,033	313,033
Communes	34,791	34,791	34,791	34,791
Clustering	Département	Département	Département	Département

Notes: Dependent variable is FN/RN first-round presidential vote share (%). Column (1) uses distance to nearest active maternity ward (km) as a continuous treatment. Column (2) uses a binary indicator for post-closure periods in affected communes. Column (3) uses the change in distance from the 2013 baseline. Column (4) reports the Callaway and Sant'Anna (2021) overall ATT using not-yet-treated communes as the control group. Standard errors clustered at the département level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.2 Event Study

Table 3 reports the dynamic treatment effects from the Callaway-Sant'Anna event-study aggregation. The pre-treatment coefficients at event times −20, −15, and −10 (years relative to closure) are −0.50, −0.24, and −0.13 respectively, all statistically insignificant but consistently negative. The post-treatment coefficients at event times 0 and +5 are −0.33 and −0.18, also insignificant. The pattern is flat: there is no evidence of a level shift at the time of closure.

Table 3: Dynamic Treatment Effects (Callaway-Sant'Anna Event Study)

Event Time (Years)	ATT	SE	95% CI Low	95% CI High
-20	-1.472	(1.306)	-4.031	1.087
-18	-0.884	(0.745)	-2.344	0.576
-17	-1.016	(0.921)	-2.821	0.789
-15	0.817	(1.526)	-2.173	3.807
-13	2.628	(1.892)	-1.080	6.335
-12	-1.023	(0.743)	-2.480	0.434
-10	1.478	(1.448)	-1.361	4.316
-8	2.444	(1.900)	-1.279	6.167
-7	-0.538	(0.522)	-1.561	0.485
-5	2.588	(1.955)	-1.243	6.420
-3	0.000	(NA)	NA	NA
-2	0.000	(NA)	NA	NA
0	0.582	(1.519)	-2.395	3.560
2	4.177	(2.545)	-0.811	9.165
3	-0.684	(0.328)	-1.327	-0.041
5	4.240	(2.898)	-1.440	9.919
8	-0.526	(1.178)	-2.835	1.783
Overall ATT	1.751	(1.491)	-1.171	4.672

Notes: Event time measured relative to the first election after the nearest maternity closure. Estimates from Callaway and Sant'Anna (2021) with not-yet-treated control group and universal base period. Event time −5 is the reference period (normalized to zero). Standard errors clustered at the département level.

5.3 Robustness and Placebo Tests

Table 4 reports robustness checks across three dimensions.

Alternative treatment definitions. The proximity treatment—a binary indicator for any maternity closing within 20 km—yields a significant negative coefficient (−1.48, $p < 0.001$).

This likely reflects selection: areas near closures are urbanizing and shifting away from the RN, not responding to the closure itself. Excluding Île-de-France does not change the main result.

Heterogeneity by commune size. If state abandonment drives populism, the effect should be strongest in small, rural communes where maternities matter most. The coefficient is -0.04 (insignificant) for communes under 2,000 and -0.01 (insignificant) for larger communes. There is no size gradient.

Placebo outcomes. Distance to the nearest maternity has no effect on Mélenchon/LFI vote share ($+0.04$, $p = 0.14$) or on turnout (-0.003, $p = 0.79$). The absence of effects on placebos supports the interpretation that the null on FN/RN voting is genuine rather than driven by low power.

Table 4: Robustness Checks and Placebo Tests

	Coefficient	SE	N
Panel A: Alternative Treatments			
Proximity (<20 km to closure)	-1.862***	(0.438)	313,033
Excluding Île-de-France	0.0070	(0.0352)	301,640
Panel B: Heterogeneity by Commune Size			
Small communes (<2,000 pop.)	0.0014	(0.0357)	265,000
Large communes (≥2,000 pop.)	0.0057	(0.0375)	48,033
Panel C: Placebo Outcomes			
LFI/Mélenchon share (%)	0.0396	(0.0265)	104,349
Turnout (%)	-0.0030	(0.0110)	173,892

Notes: Each row reports a separate regression. All specifications include commune and election year fixed effects. Treatment is distance to nearest maternity (km) unless otherwise noted. Panel C uses the same treatment but alternative dependent variables. Standard errors clustered at the département level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6. Discussion

The null result challenges the state-abandonment hypothesis in its specific, healthcare-closure form. Three interpretations merit consideration.

First, maternity closures may not be salient enough to shift voting. Unlike the carbon tax—which affects household budgets monthly and catalyzed the gilets jaunes (Douenne and Fabre, 2022)—a maternity closure is a rare event that directly affects only pregnant women. The broader population may not update its political preferences in response.

Second, closures may generate grievances that are absorbed by the commune and election-year fixed effects. If all of France is angry about declining public services, the fixed effects capture this national mood, and the commune-level variation in closure proximity is too local to identify an additional effect. This interpretation suggests that populism responds to aggregate perceptions of decline, not marginal service losses.

Third, the selection concern is real. The parallel trends pre-test rejects at $p = 0.0003$, indicating that communes which later lose maternities were on different political trajectories before closure. The negative proximity coefficient suggests these areas were becoming less FN-supportive, consistent with urbanization-driven maternity closures. While the commune fixed effects absorb level differences, they cannot address differential trends, and the null may partly reflect the tension between a positive treatment effect and a negative selection trend.

The broader implication is methodological. Cross-sectional correlations between deprivation indices and populist voting—no matter how striking—cannot establish whether specific service closures cause political backlash. The literature on the "geography of discontent" would benefit from more quasi-experimental designs that isolate specific channels rather than bundling multiple deprivation indicators into a single narrative.

7. Conclusion

Closing a maternity ward is a visible act of state withdrawal. It generates media coverage, local protest, and political debate. Yet when measured at the commune level across nine national elections using within-commune variation, the relationship between maternity closures and far-right voting is indistinguishable from zero in standard panel models, though the design cannot definitively rule out moderate effects given pre-trend concerns. Understanding why peripheral France votes for the Rassemblement National likely requires looking beyond individual service closures—toward the deeper structural forces of deindustrialization, demographic change, and cultural anxiety that shape political identity in a changing country.

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: @SocialCatalystLab

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

References

Algan, Yann, Elizabeth Beasley, Daniel Cohen, and Martial Foucault, "The Rise of Populism and the Collapse of the Left-Right Paradigm," *British Journal of Political Science*, 2019, *50* (4), 1–20.

Besley, Timothy and Stephen Coate, "Centralized versus Decentralized Provision of Local Public Goods: A Political Economy Approach," *Journal of Public Economics*, 2003, *87* (12), 2611–2637.

Blondel, Béatrice, Bénédicte Coulm, Catherine Bonnet, François Goffinet, and Camille Le Ray, "Trends in Perinatal Health in Metropolitan France from 1995 to 2016," *Journal of Gynecology Obstetrics and Human Reproduction*, 2017, *46* (10), 733–742.

Bó, Ernesto Dal, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne, "Economic and Social Outsiders but Political Insiders: Sweden's Populist Radical Right," *Review of Economic Studies*, 2023, *90* (2), 675–706.

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

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

Dijkstra, Lewis, Hugo Poelman, and Andrés Rodríguez-Pose, "The Geography of EU Discontent," *Regional Studies*, 2020, *54* (6), 737–753.

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

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

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

Guilluy, Christophe, *La France Périphérique: Comment on a Sacrifié les Classes Populaires*, Flammarion, 2014.

Noiriel, Gérard, *Les Gilets Jaunes à la Lumière de l'Histoire*, Éditions de l'Aube, 2019.

Rodríguez-Pose, Andrés, "The Revenge of the Places That Don't Matter (and What to Do About It)," *Cambridge Journal of Regions, Economy and Society*, 2018, *11* (1), 189–209.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
FN/RN vote share (distance)	0.0097	0.0351	11.20	0.0084	0.0304	Small positive
FN/RN vote share (binary post)	0.0434	0.5180	11.20	0.0039	0.0462	Null

Notes: **Country:** France. **Research question:** Does the closure of a nearby maternity ward increase far-right (FN/RN) vote share in subsequent presidential elections? **Policy mechanism:** France closed 536 maternity wards between 2013 and 2024 under the Kouchner decree minimum-volume threshold (300 deliveries/year), increasing travel distance for affected communes. **Outcome definition:** FN/RN first-round presidential vote share (%), measured as combined votes for Le Pen candidates divided by total votes. **Treatment:** (Row 1) Continuous: distance in km to nearest active maternity; (Row 2) Binary: indicator for post-closure period in affected communes. **Data:** DREES SAE maternity registry (2013–2024), data.gouv.fr presidential election results (2002–2022), INSEE commune coordinates. Panel of 34,791 communes × 5 elections = 173,892 observations. **Method:** Two-way fixed effects (commune + election year), standard errors clustered at département level (96 clusters). Also: Callaway-Sant'Anna (2021) ATT. **Sample:** Metropolitan France (excluding DOM-TOM). SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ for continuous treatment; SDE $= \hat{\beta}/\text{SD}(Y)$ for binary treatment, where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).