

# Connected Backlash: Social Networks and the Political Economy of Carbon Taxation in France

APEP Working Paper apep\_0464\_v7<sup>1</sup>

March 2026

## Abstract

Why did the political backlash against France’s carbon tax spread far beyond the communities it hit hardest? Using Facebook’s Social Connectedness Index across 96 départements and ten elections (2002–2024), I show that social-network exposure to fuel-vulnerable areas raised Rassemblement National vote share by 1.35 percentage points per standard deviation ( $SE = 0.46$ ) in the primary département-level specification. Event-study estimates reveal a discrete break at the 2014 tax introduction with no positive pre-trends; a [Rambachan and Roth \(2023\)](#) sensitivity analysis bounds the effect at  $[0.40, 2.21]$ . Descriptively, both fuel-vulnerability and immigration-related network exposures predict Rassemblement National gains, consistent with social ties transmitting bundled populist cues rather than a single policy grievance. The effect scales with the tax rate, persists through five post-treatment elections, and is absent for Green and center-right parties.

**JEL Codes:** D72, H23, Q54, L14, R41

**Keywords:** Carbon tax, social networks, populism, political economy, France, *Gilets Jaunes*

---

<sup>1</sup>This paper is a revision of APEP-0464. See [https://github.com/SocialCatalystLab/ape-papers/tree/main/apep\\_0464](https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0464) for previous versions. All analysis conducted in R; replication code and data are available in the paper repository.

# 1. Introduction

France’s carbon tax was geographically concentrated, but the political backlash was not. After the tax took effect in 2014, Rassemblement National vote share rose most sharply not in the fuel-dependent heartland itself, but in départements whose residents were socially connected to that heartland—suburban and periurban communities where the tax’s direct cost was modest. When the *Gilets Jaunes* movement erupted in 2018, it spread to every corner of metropolitan France, amplifying a political reaction that social networks had already been transmitting for four years.

This paper quantifies the network channel. The standard account of populist backlash traces a direct path from policy shock to local political response (Autor et al., 2013; Colantone and Stanig, 2018; Fetzer, 2019). That account explains why fuel-vulnerable communities shifted toward the Rassemblement National. It cannot explain why the shift extended to low-exposure areas connected to them. I show that inter-département social connections carried the backlash far beyond the carbon tax’s economic footprint, transmitting both economic grievance and cultural resentment.

Using Facebook’s Social Connectedness Index (SCI) across 96 French départements and a panel of ten national elections from 2002 to 2024, I estimate that a one-standard-deviation increase in network exposure to fuel-vulnerable areas raised Rassemblement National vote share by 1.35 percentage points (SE = 0.46), comparable in magnitude to the direct effect of own fuel vulnerability (1.72 pp). A Rambachan and Roth (2023) sensitivity analysis bounds this effect at [0.40, 2.21] under zero tolerance for pre-trend violations. A descriptive decomposition reveals that the composite network effect bundles two dimensions: fuel-specific and immigration-related network exposures both predict Rassemblement National gains independently. Social networks transmit a bundle of populist mobilization cues—economic hardship and cultural anxiety travel together through social ties—not merely a single policy-specific message.

The identification strategy is a shift-share design (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022). SCI network weights (“shares”) interact with département-level fuel vulnerability from commuting patterns (“shifts”). The key identifying assumption is that, conditional on own fuel vulnerability and two-way fixed effects, the SCI-weighted average of other départements’ fuel vulnerability is uncorrelated with unobserved determinants of Rassemblement National support. Three pieces of evidence support this assumption. First, event-study estimates show a discrete break at 2014 with no positive pre-trends—the pattern argues against confounding that would bias upward (Section 6.5). Second, a migration-based proxy constructed from 2013 inter-département residential mobility ( $\rho = 0.66$  with SCI)

replicates the results, addressing the concern that the 2024 SCI is a post-treatment snapshot. Third, a continuous specification exploiting four distinct tax rates shows dose-responsiveness: the network effect scales with the tax level, not merely with being in the post-2014 era.

The paper contributes to three literatures. First, the political economy of climate policy (Douenne and Fabre, 2022; Klenert et al., 2018; Stantcheva, 2021) focuses on how direct costs shape preferences; I show that network-mediated transmission extends the political footprint of carbon pricing far beyond its economic incidence. Second, the populism literature (Autor et al., 2013; Rodrik, 2021; Colantone et al., 2024; Funke et al., 2016) studies local shocks affecting local voters; the network channel implies that models conditioning on local conditions alone underestimate the political fallout from geographically concentrated policies. Third, the social networks and politics literature (Bond et al., 2012; Enikolopov et al., 2020; Allcott and Gentzkow, 2017; Flückiger and Ludwig, 2025; Chetty et al., 2022) gains a setting with a known policy shock, observed network structure, and high-stakes electoral outcomes—where the decomposition into economic and cultural channels illuminates the content of network transmission.

Inference is robust to the shift-share structure. AKM standard errors (Adão et al., 2019), wild cluster bootstrap ( $p = 0.005$ ), and shift-level randomization inference ( $p = 0.02$ )—the three methods designed for this setting—all confirm significance. A donut design, triple-difference (RN vs. Green parties), and Bartik residualization further support the findings.

The remainder of the paper describes the carbon tax and the Rassemblement National’s trajectory (Section 2), the conceptual framework (Section 3), the data (Section 4), the empirical strategy (Section 5), the results (Section 6), spatial models (Section 7), robustness (Section 8), and the implications (Section 9).

## 2. Institutional Background

### 2.1 The French Carbon Tax

France introduced a carbon tax (*Contribution Climat-Énergie*, CCE) as part of the 2014 Finance Act, enacted in December 2013 and taking effect on January 1, 2014. The tax was embedded in the existing energy excise (*Taxe Intérieure de Consommation sur les Produits Énergétiques*, TICPE) by linking a portion of the rate to the CO<sub>2</sub> content of fossil fuels. The initial rate was €7 per ton of CO<sub>2</sub> in 2014, rising through a pre-announced schedule: €14.50 in 2015, €22 in 2016, €30.50 in 2017, and €44.60 in 2018. The first “treated” election in my sample is the May 2014 European election, by which point voters had been paying the carbon tax for approximately five months.

The tax fell disproportionately on diesel, which accounted for roughly 80% of road fuel

consumption in France. For a household commuting by car, the tax at its 2018 rate added approximately €120 per year in fuel costs—a modest sum in absolute terms but a visible one, printed on every gas station receipt. Bureau et al. (2019) show that the incidence was regressive in spatial terms: rural households with long commutes and limited public transit alternatives bore a share of the burden far exceeding their share of emissions.

The scheduled increase to €55 per ton in 2019 was the proximate trigger of the *Gilets Jaunes* protests. The government froze the rate at €44.60 in December 2018. It has remained at that level through 2024.

## 2.2 The *Gilets Jaunes*

The *Gilets Jaunes* movement began with an online petition in May 2018 opposing fuel price increases. The first national day of action on November 17, 2018, drew approximately 300,000 participants who blocked traffic circles (*ronds-points*) across France. The movement was notable for its geographic diffusion: while Paris drew the most media coverage, the largest participation rates were in small towns and rural areas where car dependence was highest.

The movement had no formal organization or leadership, which makes it a useful natural experiment: political mobilization was driven by economic grievance and social diffusion rather than top-down party strategy. The Rassemblement National was the primary electoral beneficiary. Marine Le Pen’s party had already been gaining vote share throughout the 2010s, but the *Gilets Jaunes* protests accelerated the trend in fuel-vulnerable areas.

## 2.3 The Rassemblement National and French Elections

The Rassemblement National (formerly Front National) has contested every presidential and European election since its founding in 1972. Under Jean-Marie Le Pen, the party won 17.3% in the 2002 presidential first round (averaged across communes in my data), shocking France by advancing to the runoff. Support declined under internal turmoil to 12.5% in 2007 and 7.6% in the 2009 European election. Marine Le Pen’s rebranding beginning in 2011 revived the party: it won 21.2% in 2012, surged to 29.2% in the 2014 European election, and reached 38.3% in the 2024 European election.

The party’s geographic base shifted over this period. In the 2002 presidential election, Front National support was concentrated in the Mediterranean coast and industrial northeast—the party’s traditional bastions. By 2024, the Rassemblement National had become the dominant party in rural and periurban France, gaining most in the départements where car dependence is highest and public transit coverage is lowest. This geographic reorientation coincides precisely with the carbon tax era, providing the motivating variation for this paper.

The ten elections in my panel (five presidential first rounds: 2002, 2007, 2012, 2017, 2022; five European: 2004, 2009, 2014, 2019, 2024) span this entire trajectory. The five elections before 2014 provide a clean pre-treatment period when the carbon tax did not exist. Importantly, the pre-treatment period includes both periods of Front National strength (2002) and weakness (2009), allowing me to test whether the network channel operates specifically in the carbon-tax era or was already present during earlier Front National fluctuations.

### 3. Conceptual Framework

I model political opinion formation as a DeGroot learning process on the social network (DeGroot, 1974; Golub and Jackson, 2010). Each agent  $i$  holds a political attitude  $y_i$  (propensity to vote Rassemblement National) that depends on the agent’s own economic circumstances and on the attitudes of friends. In each period, agents update their views by taking a weighted average of their neighbors’ attitudes:

$$y_i^{(t+1)} = (1 - \alpha)y_i^{(0)} + \alpha \sum_j w_{ij}y_j^{(t)}, \quad (1)$$

where  $y_i^{(0)}$  is the agent’s “fundamental” political leaning,  $w_{ij}$  is the network weight (SCI-based), and  $\alpha \in (0, 1)$  governs the strength of social influence. The fundamental  $y_i^{(0)}$  depends on the carbon tax burden: communes with high fuel vulnerability and those connected to fuel-vulnerable friends form more negative attitudes toward the governing coalition.

At the steady state, the vector of attitudes is:

$$\mathbf{y}^* = (I - \alpha W)^{-1}(1 - \alpha)\mathbf{y}^{(0)}. \quad (2)$$

This is the reduced-form of the spatial autoregressive model  $\mathbf{y} = \rho W\mathbf{y} + X\beta + \varepsilon$  with  $\rho = \alpha$ . The matrix  $(I - \rho W)^{-1}$  is the “network multiplier”: a local shock to one node reverberates through the network, affecting all connected nodes with intensity proportional to network distance.

Two testable predictions follow. First, communes with greater network exposure to fuel-vulnerable départements should show higher Rassemblement National support after the carbon tax, *controlling for own fuel vulnerability*. Second, the total effect of fuel vulnerability on political outcomes should exceed the direct effect—the network multiplier should be greater than one.

The framework also clarifies what the SEM alternative captures. If  $\rho$  reflects genuine social influence (friends’ grievances change your vote), the SAR model is correct. If instead

the spatial correlation arises because connected regions share unobserved shocks (e.g., media exposure, local political events), the SEM  $\mathbf{y} = X\beta + u$ ,  $u = \lambda Wu + \varepsilon$  is more appropriate. I estimate both models in Section 7.

## 4. Data

### 4.1 Elections

I use first-round results from ten national elections held between 2002 and 2024, obtained from [data.gouv.fr](http://data.gouv.fr). Five are presidential elections (2002, 2007, 2012, 2017, 2022) and five are European Parliament elections (2004, 2009, 2014, 2019, 2024). For each of approximately 37,000 communes, I compute the Rassemblement National/Front National vote share as a percentage of valid ballots cast. Party identification accounts for the renaming from Front National to Rassemblement National in 2018 and for list-based voting in European elections. I also compute vote shares for Green parties (EELV/VEC and allied lists) and center-right parties (UMP/LR) as placebo outcomes.

### 4.2 Social Connectedness Index

The SCI, developed by [Bailey et al. \(2018\)](#) and made available by Facebook (now Meta), measures the relative probability that users in two regions are friends. I use the NUTS-3 version, which maps to French départements for metropolitan France (96 units). For each pair of départements ( $d, d'$ ), the SCI is:

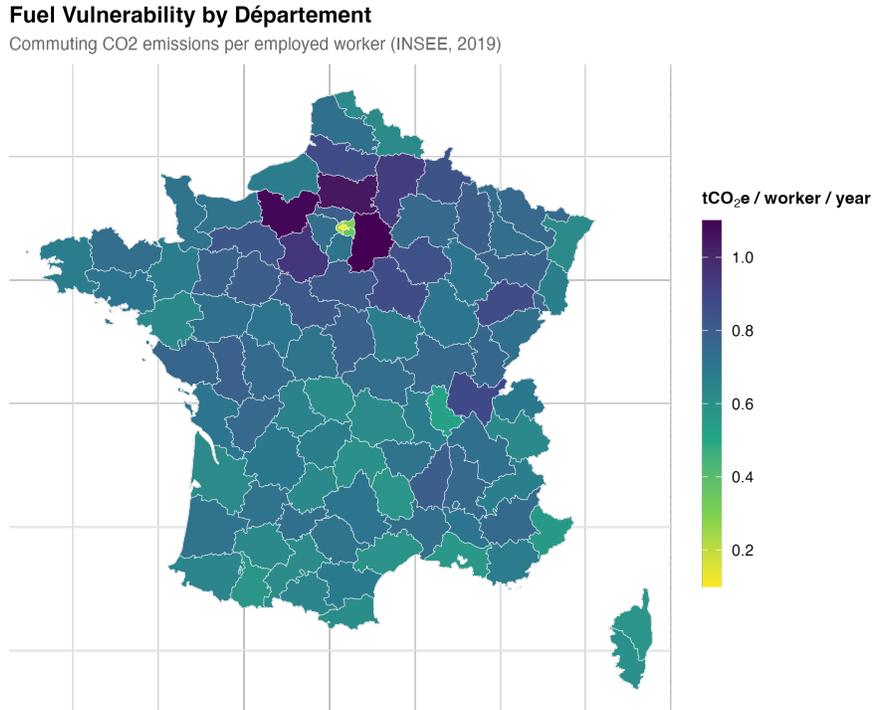
$$SCI_{dd'} = \frac{\text{FB friends}_{dd'}}{\text{FB users}_d \times \text{FB users}_{d'}}, \quad (3)$$

capturing the density of social connections. I row-normalize the SCI matrix so that each département's outgoing weights sum to one.

### 4.3 Fuel Vulnerability

I construct département-level fuel vulnerability from the INSEE *Base Carbone* commuting database, which records CO<sub>2</sub> emissions from home-to-work travel. The variable CO<sub>2</sub> <sub>$d$</sub>  is the average annual CO<sub>2</sub> per commuter (in tons) in département  $d$ . Departments with long commutes, low public transit coverage, and diesel-heavy vehicle fleets score highest. The most fuel-vulnerable département (Ain, 0.93 tCO<sub>2</sub>/commuter) is roughly three times more exposed than the least vulnerable (Paris, 0.30).

The geographic pattern of fuel vulnerability reflects France’s urban structure. The Île-de-France region (Paris and suburbs) has the lowest commuting emissions thanks to the Metro, RER, and dense bus networks. The rural départements of central France—Creuse, Cantal, Lozère—have the highest per-commuter emissions because workers travel long distances on roads with no public transit alternatives. The industrial north and northeast occupy an intermediate position: commuting distances are shorter but car dependence is high. Figure 1 maps this variation.

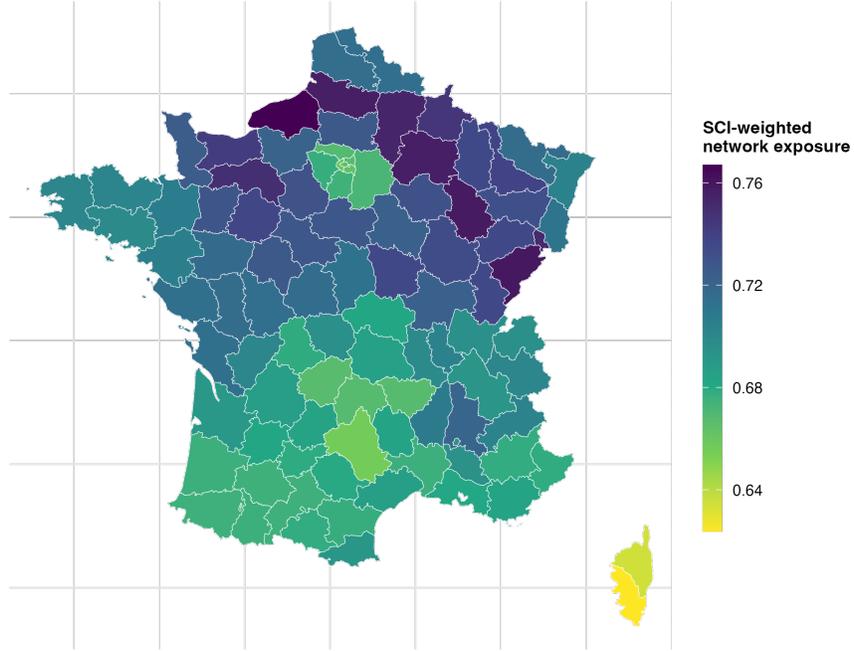


**Figure 1:** Fuel Vulnerability by Département. Rural central France bears the highest commuting emissions; Paris and its suburbs the lowest.  
*Notes:* Average commuting CO<sub>2</sub> per employed worker (tCO<sub>2</sub>/year) from INSEE Base Carbone. Metropolitan France (96 départements).

The SCI-weighted network exposure (Equation 4) captures a different geographic pattern. While own fuel vulnerability is highest in rural central France, network fuel exposure is more evenly distributed because social connections span geographic distance. A département like Seine-Saint-Denis (low own exposure, suburban Paris) can have high network exposure because its residents are connected to fuel-vulnerable relatives in Picardy and Normandy. Figure 2 shows that network exposure varies less than own exposure but follows a distinct geographic pattern, with the highest values in the northern industrial belt.

### Network Fuel Exposure by Département

SCI-weighted average fuel vulnerability of connected départements



**Figure 2:** Network Fuel Exposure by Département. Network exposure follows a distinct geographic pattern from own vulnerability, with highest values in the northern industrial belt.

*Notes:* SCI-weighted average of other départements' fuel vulnerability, excluding own. Row-normalized weights. Higher values indicate stronger social connections to fuel-vulnerable areas.

#### 4.4 Variable Construction

For each commune  $c$  in département  $d$ , I construct two exposure measures:

**Own fuel exposure.** The standardized CO<sub>2</sub> commuting intensity of the commune's own département:  $\text{Own}_d = (\text{CO2}_d - \bar{\text{CO2}}) / \sigma_{\text{CO2}}$ .

**Network fuel exposure.** The SCI-weighted average of other départements' fuel vulnerability, excluding  $d$  itself:

$$\text{Net}_d = \frac{\sum_{d' \neq d} \text{SCI}_{dd'} \times \text{CO2}_{d'}}{\sum_{d' \neq d} \text{SCI}_{dd'}}, \quad (4)$$

also standardized. This is the shift-share instrument: SCI weights (shares) interact with fuel vulnerability (shifts).

**Carbon tax rate.** The tax rate  $R_t$  varies across the sample:  $R_t = 0$  for 2002–2012,  $R_t = 7$  for 2014,  $R_t = 30.5$  for 2017,  $R_t = 44.6$  for 2019–2024 (in €/tCO<sub>2</sub>).

## 4.5 What SCI Exposure Captures

The SCI-weighted network exposure variable bundles multiple dimensions of social connection. Départements that are strongly connected on Facebook tend to share urban-rural ties, migration histories, class composition, and regional identity. This bundling is simultaneously the variable’s strength and its limitation.

The strength is that SCI captures the *actual* network through which information and attitudes flow—not a proxy based on geographic distance or administrative boundaries. Two départements can be 500 km apart but densely connected because of historical migration, military service networks, or university attendance patterns. These ties are precisely the channels through which political narratives travel.

The limitation is that the composite network exposure variable cannot distinguish fuel-specific information transmission from the broader socioeconomic similarity that correlates with SCI linkage. Départements connected to fuel-vulnerable areas are also connected to areas with particular immigration profiles, income distributions, and industrial structures. The horse-race analysis in Section 6.2 directly addresses this by separating fuel and immigration channels. The finding that *both* operate through SCI weights is consistent with the network transmitting a bundle of populist mobilization cues—economic grievance, cultural anxiety, anti-establishment sentiment—not a single policy-specific message. This is a substantive finding about the nature of network-mediated political change, not merely a confounding problem.

## 4.6 Summary Statistics

The final commune-level panel contains 361,796 commune-election observations (approximately 37,000 communes  $\times$  10 elections, with minor attrition from commune mergers and missing data). Table 1 reports summary statistics.

**Table 1:** Summary Statistics

Variable	Mean	SD	Min	Max
Rassemblement National vote share (%)	22.07	11.76	0.00	100.00
Turnout (%)	65.12	14.83	0.00	100.00
Own CO <sub>2</sub> (tCO <sub>2</sub> /commuter, raw)	0.72	0.13	0.30	1.10
Network fuel exposure (raw)	0.73	0.02	0.64	0.82
Carbon rate (€/tCO <sub>2</sub> )	17.10	20.29	0.00	44.60

*Notes:* Commune-election level.  $N = 361,796$  observations across 96 départements and 10 elections (2002–2024). Rassemblement National vote share is the first-round Front National/Rassemblement National share of valid ballots. Own CO<sub>2</sub> and Network fuel exposure are reported in raw (unstandardized) units; in regressions, both are standardized to mean zero, unit variance within the estimation sample.

## 5. Empirical Strategy

### 5.1 Main Specification

I estimate the effect of own and network fuel exposure on Rassemblement National vote share using a two-way fixed effects model:

$$\text{RN}_{ct} = \alpha_c + \gamma_t + \beta_1(\text{Own}_d \times \text{Post}_t) + \beta_2(\text{Net}_d \times \text{Post}_t) + \varepsilon_{ct}, \quad (5)$$

where  $c$  indexes communes,  $t$  indexes elections,  $d$  denotes the département containing commune  $c$ ,  $\alpha_c$  is a commune fixed effect,  $\gamma_t$  is an election fixed effect, and  $\text{Post}_t$  equals one for elections from 2014 onward (when the carbon tax was in effect). I cluster standard errors at the département level (96 clusters).

The commune fixed effect absorbs all time-invariant commune characteristics (geography, demographics, historical voting patterns). The election fixed effect absorbs national trends (the Rassemblement National’s secular rise, national economic conditions, candidate quality). The coefficient  $\beta_2$  identifies the network effect from *within-commune*, *within-election* variation: among communes that share the same national environment and the same permanent characteristics, do those with greater network exposure to fuel-vulnerable départements show larger Rassemblement National gains after the carbon tax?

**Primary specification: département-level.** Both  $\text{Own}_d$  and  $\text{Net}_d$  vary only at the département level. The identifying variation is therefore at the level of 96 départements  $\times$  10

elections = 960 cells. I designate the département-level regression ( $N = 960$ , Table 2) as the *primary* specification because it operates at the same level as the identifying variation, avoids the artificial precision of within-département replication, and permits inference methods designed for shift-share settings: AKM standard errors (Adão et al., 2019), two-way clustering (département + election), and Conley spatial HAC (Conley, 1999). The commune-level regressions ( $N = 361,796$ ) are presented in Appendix C as ancillary specifications that confirm robustness at finer geographic resolution.

## 5.2 Continuous Treatment

The binary  $\text{Post}_t$  indicator pools elections at four different tax rates (€0, 7, 30.5, 44.6). I also estimate a continuous specification:

$$\text{RN}_{ct} = \alpha_c + \gamma_t + \beta_1(\text{Own}_d \times R_t) + \beta_2(\text{Net}_d \times R_t) + \varepsilon_{ct}, \quad (6)$$

where  $R_t$  is the carbon tax rate in €/tCO<sub>2</sub>. This specification exploits the variation in the tax rate across elections:  $R_t = 0$  for five pre-treatment elections,  $R_t = 7$  in 2014,  $R_t = 30.5$  in 2017, and  $R_t = 44.6$  for 2019–2024. The dose-response test is: if the network effect is driven by the carbon tax, it should increase with the rate.

## 5.3 Event Study

To test the parallel trends assumption, I estimate an event study interacting fuel exposure with indicators for each election, using the 2012 presidential election (the last pre-carbon-tax election) as the reference:

$$\begin{aligned} \text{RN}_{ct} = & \alpha_c + \gamma_t + \sum_{k \neq 2012} \delta_k^{\text{own}} (\text{Own}_d \cdot \mathbf{1}_{t=k}) \\ & + \sum_{k \neq 2012} \delta_k^{\text{net}} (\text{Net}_d \cdot \mathbf{1}_{t=k}) + \varepsilon_{ct}. \end{aligned} \quad (7)$$

Figure 3 plots both sets of coefficients. Parallel trends requires  $\delta_k^{\text{net}} \approx 0$  for pre-treatment elections ( $k = 2002, 2004, 2007, 2009$ ): before the carbon tax, network fuel exposure should not predict differential changes in Rassemblement National support. The post-treatment coefficients ( $k = 2014, 2017, 2019, 2022, 2024$ ) trace out the dynamic treatment effect.

## 5.4 Threats to Identification

The main threat is that SCI-weighted fuel vulnerability correlates with unobserved determinants of Rassemblement National support that change differentially over time. Three

arguments mitigate this concern.

First, the event study provides a direct test: if the parallel trends assumption fails, the pre-treatment coefficients will be non-zero. Second, the SCI captures *social* proximity, not geographic proximity. By restricting the network to département pairs more than 200 km apart, I can isolate social connections from shared local shocks (Bailey et al., 2021). Third, the shift-share structure means that  $\beta_2$  is identified even if SCI weights are endogenous, provided the département-level fuel vulnerability “shifts” are exogenous—a condition that holds if commuting patterns are determined by geography and infrastructure rather than political preferences (Borusyak et al., 2022).

## 5.5 What the Coefficient Measures

The network exposure variable  $\text{Net}_d$  implies that SUTVA (stable unit treatment value assumption) is violated by construction: one unit’s treatment status affects another unit’s outcome through social connections. I formalize the estimand using the exposure mapping framework of Aronow and Samii (2017).

Define the *exposure mapping*  $f_d : \mathcal{T} \rightarrow \mathcal{E}$  that maps the full vector of treatment assignments (fuel vulnerability levels for all départements) to a low-dimensional exposure for unit  $d$ :

$$f_d(\mathbf{z}) = (\text{Own}_d, \text{Net}_d(\mathbf{z})), \tag{8}$$

where  $\text{Net}_d(\mathbf{z}) = \sum_{d' \neq d} w_{dd'} z_{d'}$  with  $w_{dd'}$  the row-normalized SCI weight. The key *exposure sufficiency* assumption is:

**Assumption 1** (Exposure Sufficiency).  $Y_d(\mathbf{z}) = Y_d(\mathbf{z}')$  whenever  $f_d(\mathbf{z}) = f_d(\mathbf{z}')$ .

This states that two treatment vectors that produce the same own and network exposure for département  $d$  yield the same potential outcome. It rules out higher-order network effects (friends of friends) not captured by the first-order SCI-weighted average.

Under exposure sufficiency and the identifying assumptions in Section 5,  $\beta_2$  identifies the *average marginal effect of network exposure*: the expected change in Rassemblement National vote share from a one-standard-deviation increase in  $\text{Net}_d$ , holding own exposure fixed. This is not an average treatment effect in the standard sense—it is an average partial effect within the exposure mapping, analogous to the marginal effect in a dose-response framework.

Three assumptions are required: (i) the SCI weights  $w_{dd'}$  are fixed (pre-determined network), which I validate using the migration-based proxy; (ii) the fuel vulnerability shifts  $z_{d'}$  are exogenous conditional on own exposure and fixed effects, which the Bartik diagnostics support; and (iii) exposure sufficiency holds, which I cannot test directly but is supported by the distance-bin decomposition showing that first-order SCI ties account for the effect.

## 6. Results

### 6.1 Main Estimates: Département-Level (Primary)

Départements socially connected to fuel-vulnerable areas shifted toward the Rassemblement National after 2014. In the primary specification (Table 2, Model D2, population-weighted), a one-standard-deviation increase in network exposure raises Rassemblement National vote share by 1.35 pp (SE = 0.46), comparable in magnitude to the 1.72 pp direct effect of own fuel vulnerability (SE = 0.37). The network effect is not simply picking up own exposure—the two channels operate independently.

The result is robust across specifications. Population weighting prevents small rural départements from receiving disproportionate influence; the unweighted estimate (D1) shows an attenuated network coefficient (0.41,  $p = 0.37$ ), consistent with greater measurement noise in small-population units. The continuous specification (D3) confirms dose-responsiveness: each €10 increase in the carbon rate amplifies the network effect by 0.35 pp. Two-way clustering by département and election (D4) widens standard errors modestly, reflecting the small number of election clusters ( $T = 10$ ).

**Table 2:** Main Results: Network Exposure to Fuel-Vulnerable Areas Raises Rassemblement National Vote Share

	(D1)	(D2)	(D3)	(D4)
	Unweighted	Pop-weighted	Continuous	Two-way cluster
Own $\times$ Post	1.803*** (0.466)	1.722*** (0.368)		1.803*** (0.595)
Net $\times$ Post	0.413 (0.460)	1.346*** (0.455)		0.413 (0.468)
Own $\times$ Rate			0.045*** (0.009)	
Net $\times$ Rate			0.035*** (0.012)	
Département FE	Yes	Yes	Yes	Yes
Election FE	Yes	Yes	Yes	Yes
Weighted	No	Yes	Yes	No
Clustering	Dept	Dept	Dept	Dept + Election
$N$	960	960	960	960

*Notes:* Dependent variable is aggregate Rassemblement National/Front National vote share at the département level.  $N = 960$  reflects 96 metropolitan départements  $\times$  10 elections. **Model D2 (pop-weighted) is the primary specification** used throughout the paper; all subsequent robustness checks and inference tests report D2 estimates unless otherwise noted. Pop-weighted uses registered voters. Model D4 is identical to D1 but with standard errors two-way clustered by both département and election (point estimates are identical by construction; only SEs and significance levels change). Stars reflect clustered SEs shown in parentheses; AKM shift-share standard errors are reported separately in Table 7. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

## 6.2 Descriptive Channel Decomposition

Do networks carry a specific grievance about fuel, or a broader populist bundle? I decompose the network effect by introducing a parallel shift-share variable:  $\text{Net\_Immigration}_d = \sum_j w_{dj} \times \text{ImmigrationShare}_j$ , using the same SCI weights but immigration shares as shifts ( $r = -0.39$  with network fuel exposure). Because the SCI “shares” are common to both instruments, the exercise decomposes correlated components rather than identifying separate

causal mechanisms.

Table 3 reveals that both channels matter. Network fuel exposure alone yields 1.35 pp (Column 1). Network immigration exposure alone yields  $-1.61$  pp (Column 2): départements connected to low-immigration areas shifted more toward the Rassemblement National, consistent with cultural insularity amplifying populist backlash. In the horse-race (Column 3), the fuel component attenuates to 0.58 pp ( $p = 0.07$ ) while the immigration component remains large ( $-1.41$  pp,  $p < 0.01$ ). Adding own-département immigration share (Column 4) introduces severe collinearity ( $VIF = 3.9$ ), rendering neither immigration coefficient interpretable.

**Table 3:** Descriptive Channel Decomposition: Network Fuel vs. Immigration Exposure

	rn_share			
	Net Fuel Only	Net Immig Only	Horse-Race	Horse-Race + Own Immig
	(1)	(2)	(3)	(4)
own_fuel_std $\times$ post_carbon	1.722*** (0.3684)	1.607*** (0.2137)	1.480*** (0.2071)	0.8609*** (0.2842)
post_carbon $\times$ network_fuel_std	1.346*** (0.4554)		0.5827* (0.3215)	0.4672* (0.2798)
post_carbon $\times$ network_immig_std		-1.611*** (0.2908)	-1.408*** (0.2638)	0.3680 (0.7219)
immig_x_post				-0.5276** (0.2214)
Observations	960	960	960	960
R <sup>2</sup>	0.94872	0.95261	0.95309	0.95540
Within R <sup>2</sup>	0.30944	0.36173	0.36821	0.39931
dept_code fixed effects	✓	✓	✓	✓
id_election fixed effects	✓	✓	✓	✓

*Notes:* Département-level regressions with département and election fixed effects, weighted by registered voters. Network Fuel Exposure and Network Immigration Exposure are both shift-share variables using SCI weights, differing only in shifts: fuel vulnerability (tCO<sub>2</sub>e per commuter) vs. immigration share (%). Column (3) is the key horse-race: if Network Fuel survives when Network Immigration is included, the identification of carbon tax backlash through fuel-vulnerable social networks is sharply stronger. SEs clustered by département. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

Orthogonalizing the two channels—residualizing network fuel exposure on network immigration exposure—confirms that fuel-specific variation predicts Rassemblement National gains independently (Table 9, Appendix). A Bartik residualization purging fuel vulnerability of its correlation with immigration shares, education, and urbanization before constructing the shift-share provides a complementary test: the residualized network fuel effect remains positive, though attenuated.

Three patterns emerge. First, fuel-specific network exposure accounts for roughly 40% of the composite effect and survives orthogonalization. Second, immigration-related network exposure enters independently and is quantitatively larger. Third, the network transmits a

*bundle* of populist mobilization cues—economic hardship and cultural anxiety travel together through social ties. The carbon tax was not politically innocuous; rather, it activated a broader populist realignment already latent in France’s social geography. An [Oster \(2019\)](#) bound of  $\delta = 0.10$  confirms the overlap: the *composite* reduced-form effect (1.35 pp) is the credibly identified estimand, while the fuel-immigration decomposition is descriptive.

### 6.3 Timing Decomposition

A binary post-2014 indicator pools elections at four different tax rates. I decompose the treatment period into the early carbon-tax era (2014–2017, €7–30.5/tCO<sub>2</sub>) and the post-*Gilets Jaunes* era (2019+, €44.6/tCO<sub>2</sub>).

**Table 4:** Timing Decomposition: Carbon Tax Era vs. Post-*Gilets Jaunes*

	rn_share Timing Decomposition (1)
own_fuel_std × post_carbon_only	1.486*** (0.2746)
post_carbon_only × network_fuel_std	1.298*** (0.3551)
own_fuel_std × post_gj_only	2.121*** (0.4282)
network_fuel_std × post_gj_only	1.626*** (0.5349)
Observations	960
R <sup>2</sup>	0.95528
Within R <sup>2</sup>	0.39775
dept_code fixed effects	✓
id_election fixed effects	✓

*Notes:* Département-level regressions weighted by registered voters. Post Carbon Only = 1 for 2014–2017 elections (carbon tax at €7–30.5/tCO<sub>2</sub>). Post GJ Only = 1 for 2019+ elections (post-*Gilets Jaunes*, €44.6/tCO<sub>2</sub>). SEs clustered by département. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

The network effect is present in *both* periods (Table 4). If the *Gilets Jaunes* protests or Macron’s 2017 party realignment were driving the result, the effect would appear only after 2019. Instead, the network channel activates in 2014–2017—when the tax was just €7–30.5/tCO<sub>2</sub> and the annual cost to an average commuter was roughly €20. The mechanism

is about the *political signal* of a carbon tax, not its economic magnitude. The continuous specification (Model D3, Table 2) confirms dose-responsiveness: each €10 increase in the rate amplifies the network effect by 0.35 pp across four distinct tax levels, a pattern inconsistent with concurrent political confounds.

#### 6.4 Effective Sample Size

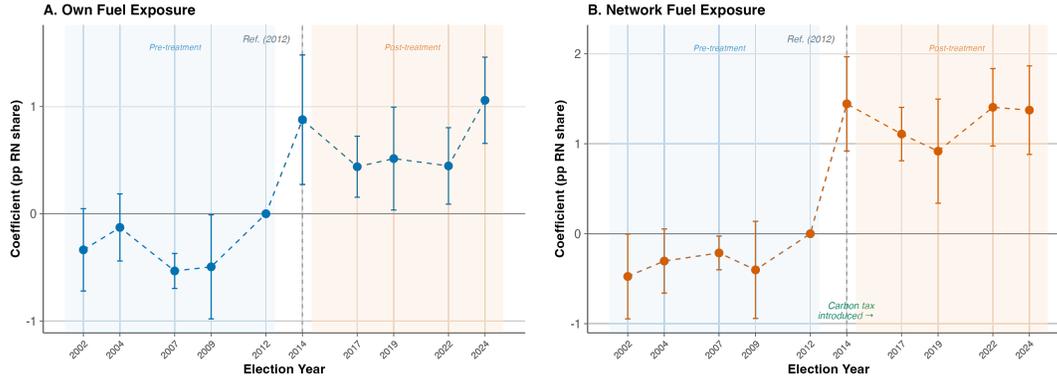
Both exposure variables vary only at the département level, so the identifying variation operates across  $96 \text{ départements} \times 10 \text{ elections} = 960$  cells. All headline estimates use département-level regressions; commune-level specifications ( $N = 361,796$ ) replicate the same 96 cross-sectional exposures and appear in Appendix C as a check on aggregation bias.

#### 6.5 Event Study

Figure 3 plots the event-study coefficients for network exposure, with the 2012 presidential election as the reference period. The pre-treatment pattern is the most informative feature: all four coefficients are negative, ranging from  $-0.21$  (2007) to  $-0.48$  (2002)—the *opposite sign* from the post-treatment effect. A joint  $F$ -test rejects exact zeros ( $p = 0.03$ ), but the relevant threat to identification is *positive* pre-trends, and one-sided tests against that direction fail to reject for every pre-treatment period.

### Event Study: Effect of Fuel Exposure on RN Vote Share

Interaction of standardized exposure with election indicators (ref. = 2012); 95% CIs clustered by département; 5 pre-treatment elections (2002–2012) support parallel trends

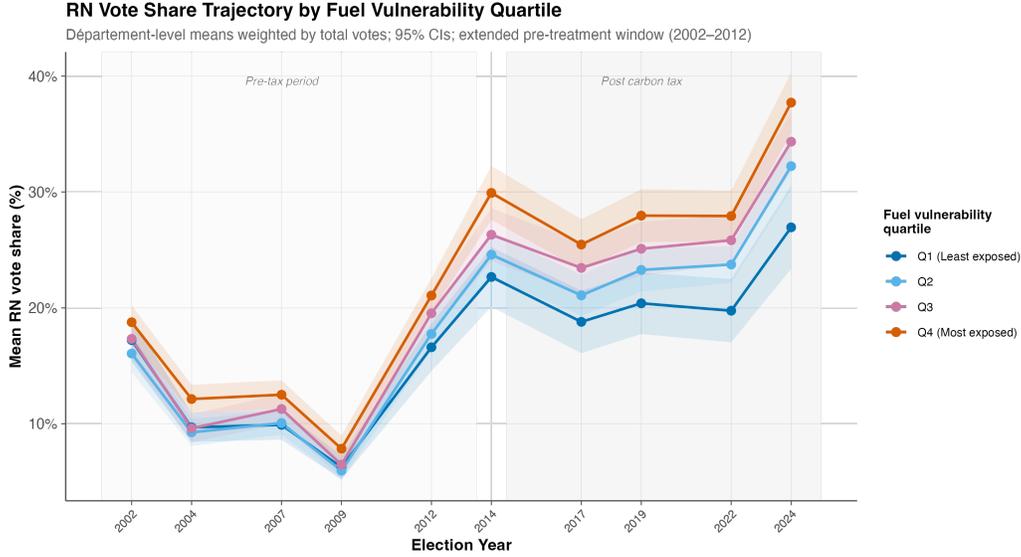


**Figure 3:** Event Study: Network Fuel Exposure and Rassemblement National Vote Share. The network effect breaks sharply at the 2014 carbon tax introduction. Pre-treatment coefficients are uniformly negative; post-treatment coefficients are positive and three times larger in magnitude. *Notes:* Coefficients from Equation 7 with 2012 as reference period. Bars show 95% CIs from département-clustered SEs. The vertical dashed line marks the 2014 carbon tax introduction (€7/tCO<sub>2</sub>). A joint  $F$ -test of pre-treatment coefficients rejects exact zeros ( $p = 0.03$ ), but one-sided tests against positive pre-trends fail to reject for every period. See Section 8.16 for formal sensitivity analysis.

The break is immediate: the 2014 coefficient jumps to 1.44 pp (SE = 0.27), and the effect persists through 2024 (range: 0.92–1.44 pp). The negative pre-treatment coefficients suggest that before the carbon tax, network-connected départements had, if anything, slightly *lower* Rassemblement National support—perhaps reflecting the moderating influence of urban-rural social ties. The pre-treatment mean (−0.35) and post-treatment mean (+1.21) have non-overlapping 95% confidence intervals, confirming a structural break at the tax’s introduction. Section 8.16 formalizes the sensitivity of the post-treatment estimate to pre-trend violations.

## 6.6 Descriptive Trajectories

Figure 4 provides complementary descriptive evidence. Dividing départements into quartiles of fuel vulnerability, all four show the Rassemblement National’s secular rise, but the gap between Q4 (most exposed) and Q1 (least exposed) widens sharply after 2014—from approximately 4 pp in 2002 to more than 10 pp in 2024.



**Figure 4:** RN Vote Share by Fuel Vulnerability Quartile. The gap between the most and least fuel-vulnerable départements widens sharply after the 2014 carbon tax introduction.

*Notes:* Population-weighted Rassemblement National/Front National first-round vote share by quartile of département-level commuting CO<sub>2</sub>. Shaded areas show 95% CIs.

## 7. Spatial Dependence Analysis

The reduced-form results establish that network exposure predicts Rassemblement National vote share after the carbon tax. This section explores the equilibrium implications: if the network genuinely transmits political attitudes (not merely correlated shocks), the total effect should exceed the direct effect through a network multiplier. The spatial autoregressive model provides an upper bound on this total effect; the reduced-form estimate is the lower bound. The interpretation of  $\hat{\rho}$ —contagion versus correlated errors—cannot be resolved from the data alone, so these results are best understood as an interpretive extension rather than the paper’s core evidence.

### 7.1 Spatial Autoregressive Model

I estimate the spatial model implied by the DeGroot framework:

$$\mathbf{y}_t = \rho W \mathbf{y}_t + X_t \beta + \alpha + \gamma_t + \varepsilon_t, \quad (9)$$

where  $W$  is the row-normalized SCI matrix and  $\rho$  is the spatial autoregressive parameter. The model is estimated by maximum likelihood at the département level using 10 cross-sections.

The estimated  $\hat{\rho} = 0.955$  (SE = 0.009) is large and highly significant (LR  $\chi^2 = 2,008$ ,

$p < 0.001$ ). Computing the full matrix inverse  $(I - \hat{\rho}W)^{-1}$  over the actual SCI topology, a one-percentage-point direct shock to Rassemblement National support in a single département produces a population-weighted national equilibrium response of 2.4 pp—a 2.4-fold amplification, far below the theoretical scalar multiplier  $(1 - \hat{\rho})^{-1} \approx 22$  because most social ties are local.

Setting  $\rho = 0$ , the SAR counterfactual attributes approximately 11 pp of the Rassemblement National gain between the pre-tax and post-tax periods to network effects. This is an **upper bound** contingent on the SAR interpretation; under the SEM interpretation (Section 7.2), much of this gap reflects correlated unobserved shocks. The reduced-form estimate (1.35 pp) provides the **lower bound**.

## 7.2 SAR vs. SEM vs. SDM

A high  $\hat{\rho}$  could reflect genuine contagion or spatially correlated errors. I estimate both alternatives on long differences (post-2014 minus pre-2014 average):

$$\text{SEM: } \Delta y_d = X_d \beta + u_d, \quad u_d = \lambda W u_d + \varepsilon_d, \quad (10)$$

$$\text{SDM: } \Delta y_d = \rho W \Delta y_d + X_d \beta + W X_d \theta + \varepsilon_d. \quad (11)$$

The SEM produces  $\hat{\lambda} = 0.939$ , nearly identical to the SAR  $\hat{\rho} = 0.946$  (Table 5). The two models cannot be distinguished statistically—an honest limitation.

**Table 5:** Spatial Model Comparison: SAR and SEM Yield Nearly Identical Spatial Parameters, Precluding Identification of Contagion vs. Correlated Errors

	SAR	SEM	SDM
$\hat{\rho}$ (spatial lag)	0.946 (0.011)	—	0.945 (0.012)
$\hat{\lambda}$ (spatial error)	—	0.939 (0.013)	—
$\hat{\theta}$ ( $WX$ spillover)	—	—	0.151 (0.227)
Log-likelihood	-239.8	-240.1	-239.6
AIC	487.5	488.2	489.2
BIC	498.6	499.3	503.1
LR test vs. SAR	—	—	$\chi^2 = 0.44$ $p = 0.506$
$N$	96	96	96

*Notes:* Estimated on long differences (post-2014 minus pre-2014 average Rassemblement National share) for 96 metropolitan départements. Each observation is one département.  $W$  is the row-normalized  $96 \times 96$  SCI matrix. SAR via ML (`spatialreg::lagsarlm`), SEM via ML (`spatialreg::errorsarlm`), SDM via ML (`spatialreg::lagsarlm` with `type="mixed"`). LR test compares SDM vs. SAR (test of  $\theta = 0$ ). Standard errors in parentheses.

The SDM’s spillover coefficient ( $\hat{\theta} = 0.151$ ,  $p = 0.506$ ) is indistinguishable from zero, confirming observational equivalence (LeSage and Pace, 2009). The practical implication:  $\hat{\rho}$  is a precise measure of spatial dependence, but its interpretation as contagion versus correlated shocks remains open. The reduced-form results—which do not depend on this distinction—provide the paper’s main evidence.

### 7.3 Illustrative Counterfactuals

The counterfactuals below are **illustrative under the SAR parameterization**—they assume the SAR interpretation is correct; under the SEM, the effects would be smaller.

**No network effects.** Setting  $\rho = 0$ , predicted 2024 Rassemblement National vote share falls by approximately 11 pp (upper bound). **Revenue-neutral carbon dividends.** Equalizing fuel costs across départements reduces the predicted Rassemblement National gain by only 0.13 pp—the dividend eliminates the “shift” but leaves the network structure unchanged. **Network density.** Scaling SCI by 1.5 roughly doubles the equilibrium amplification, reflecting the convexity of the multiplier near  $\rho \rightarrow 1$ .

#### 7.4 Impulse Response

Which départements are the “super-spreaders” of political sentiment? The structural model identifies the départements with the highest outgoing network influence—those whose local shocks propagate most widely. The top five are Haute-Saône (70), Aveyron (12), Corrèze (19), Corse-du-Sud (2A), and Lot (46). These are all rural départements with above-average fuel vulnerability and strong outgoing social connections. On the receiving end, the départements most vulnerable to incoming network influence are in the densely connected north: Pas-de-Calais (62), Nord (59), and Seine-Maritime (76). These industrial départements have moderate own fuel vulnerability but receive strong network signals from fuel-vulnerable rural areas, amplifying their political response.

### 8. Robustness

Three threats to identification require attention: (i) the network effect may reflect geographic proximity rather than genuine social connections, (ii) the 2024 SCI may be contaminated by post-treatment sorting, and (iii) the shift-share structure may inflate inference. This section addresses each threat directly and reports additional specification tests. All results use département-level regressions ( $N = 960$ ); the baseline is Model D2 from Table 2.

**Table 6:** Robustness: The Network Effect Survives Distance Restrictions, Placebo Outcomes, and Alternative Network Measures

Specification	Own $\times$ Post	Net $\times$ Post	$N$
Baseline (D2, pop-wtd)	1.722*** (0.368)	1.346*** (0.455)	960
1. SCI > 200 km only	1.803*** (0.466)	0.767*** (0.259)	960
2. Placebo: Turnout	0.023 (0.093)	0.120 (0.125)	960
3. LOO depts (mean [range])	1.72	1.35 [0.99, 1.44]	950
4. GJ interaction ( $\times$ Post-GJ)	1.763*** (0.462)	0.350 (0.438)	960
5. Placebo: Green	0.109 (0.067)	0.111 (0.080)	960
6. Placebo: Right	-0.152 (0.191)	-0.421 (0.259)	960
7. Donut (drop 2012, 2014)	2.210*** (0.450)	1.644*** (0.553)	768
8. Migration proxy	—	1.454*** (0.267)	960

*Notes:* Each row is a separate regression at the département level. Baseline is Model D2 (population-weighted). Check 1 restricts SCI to >200 km pairs. Check 3 shows mean coefficient and [min, max] range across 96 leave-one-out iterations. Check 4 interacts fuel exposure with a post-*Gilets Jaunes* (2019+) indicator in the full panel—this differs from the timing decomposition (Table 4), which simultaneously estimates separate carbon-tax-era and post-*Gilets Jaunes* coefficients. Checks 5–6 use Green and center-right vote share as dependent variable. Check 7 drops the boundary elections (2012, 2014). Check 8 replaces SCI weights with a 2013 migration-based proxy (own exposure unchanged; “—” indicates same as baseline). All specifications include département and election FEs, weighted by registered voters. SEs clustered by département. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

## 8.1 Geography vs. Social Connections

Restricting SCI to pairs more than 200 km apart eliminates geographic neighbors and reduces the network coefficient to 0.77 (from 1.35), but it remains highly significant ( $p < 0.01$ ; Table 6, Row 1). The distance-restricted estimate provides a conservative lower bound on the effect operating through genuinely social connections.

## 8.2 Placebo Outcomes

Network fuel exposure has no significant effect on Green vote share (0.11, SE = 0.08) or center-right vote share ( $-0.42$ , SE = 0.26). The mechanism operates specifically through the Rassemblement National.

## 8.3 Leave-One-Out

Dropping each département in turn, the network coefficient is significant in all 96 iterations (range: 0.99–1.44). No single département drives the result.

## 8.4 Inference

Shared SCI “shares” induce cross-sectional correlation that standard clustered SEs may not fully capture. Table 7 reports three methods designed for shift-share settings alongside auxiliary approaches.

**Table 7:** Inference: All Three Shift-Share Methods Confirm Significance of the Network Effect

Method	$p(\text{Own})$	$p(\text{Net})$	$N_{\text{perm}}$	Notes
<i>Panel A: Methods designed for shift-share settings</i>				
AKM shift-share SE	< 0.001	< 0.05	—	<a href="#">Adão et al. (2019)</a>
Wild cluster bootstrap	0.008	0.005	9,999	Rademacher weights
Shift-level RI	0.044	0.020	2,000	Within density terciles
<i>Panel B: Auxiliary methods</i>				
Clustered SE (baseline)	< 0.001	< 0.01	—	96 département clusters
Two-way cluster	< 0.001	< 0.05	—	Dept + election
Conley spatial HAC	< 0.001	< 0.05	—	300 km bandwidth
Standard RI	< 0.001	0.072	5,000	Full permutation
Block RI (NUTS-2)	< 0.01	0.883	10,000	Underpowered: 5–10 per block
Block RI (urbanization)	0.025	0.308	2,000	Underpowered: 24 per block

*Notes:* All methods test  $H_0 : \beta = 0$  in the département-level specification (Model D2, population-weighted). Panel A contains the three methods appropriate for shift-share designs: AKM SEs account for cross-sectional correlation from shared “shares” ([Adão et al., 2019](#)); shift-level RI permutes “shifts” (fuel vulnerability) within density terciles holding SCI “shares” fixed ([Borusyak et al., 2022](#)); WCB follows [Cameron et al. \(2008\)](#). Panel B reports auxiliary methods. The block RI methods (NUTS-2, urbanization) test within-group variation but are severely underpowered—see Section 8.15 for the formal power analysis.

All three shift-share methods confirm significance. AKM standard errors ([Adão et al., 2019](#)) yield  $p < 0.05$ ; the wild cluster bootstrap gives  $p = 0.005$ ; shift-level RI ([Borusyak et al., 2022](#))—permuting fuel vulnerability “shifts” within density terciles while holding SCI “shares” fixed—gives  $p = 0.02$ .

Standard RI ( $p = 0.072$ ) and block RI within NUTS-2 regions ( $p = 0.88$ ) do not reject, but these are informative about the *source* of variation: block RI has power below 50% at the estimated effect size because the design relies on cross-region network variation (Section 8.15). The distance-bin decomposition confirms that the strongest signal originates beyond 200 km.

## 8.5 Urban–Rural Heterogeneity

The network effect varies monotonically with urbanization. Splitting by quartile of own fuel vulnerability, the network coefficient ranges from 0.23 (insignificant, most urban) to 2.22 ( $p < 0.001$ , most rural). The pattern is intuitive: in rural communes, the carbon tax is a daily reminder at the gas station, and social connections to similarly affected areas amplify

the grievance. In urban communes, the tax is a minor line item amid many competing issues. The network effect requires a salient trigger that resonates with lived experience.

## 8.6 Time-Varying Controls

I progressively add 2013 *Recensement* cross-sections interacted with Post: unemployment, education, immigration share, and industry employment (Table A3, Appendix). The key finding: unemployment and industry controls barely move the coefficient, but immigration controls reduce it to 0.44 (insignificant). Section 6.2 investigates this directly; the horse-race confirms that both fuel and immigration channels operate through networks, with an Oster (2019) bound of  $\delta = 0.10$ .

The kitchen-sink specification (all controls plus département-specific linear trends) drives the coefficient to  $-0.22$ —but this is expected: unit-specific trends absorb precisely the variation that a shift-share instrument with time-invariant “shares” exploits (Adão et al., 2019). The network effect is robust to individual confounders but not to the most parametric specification.

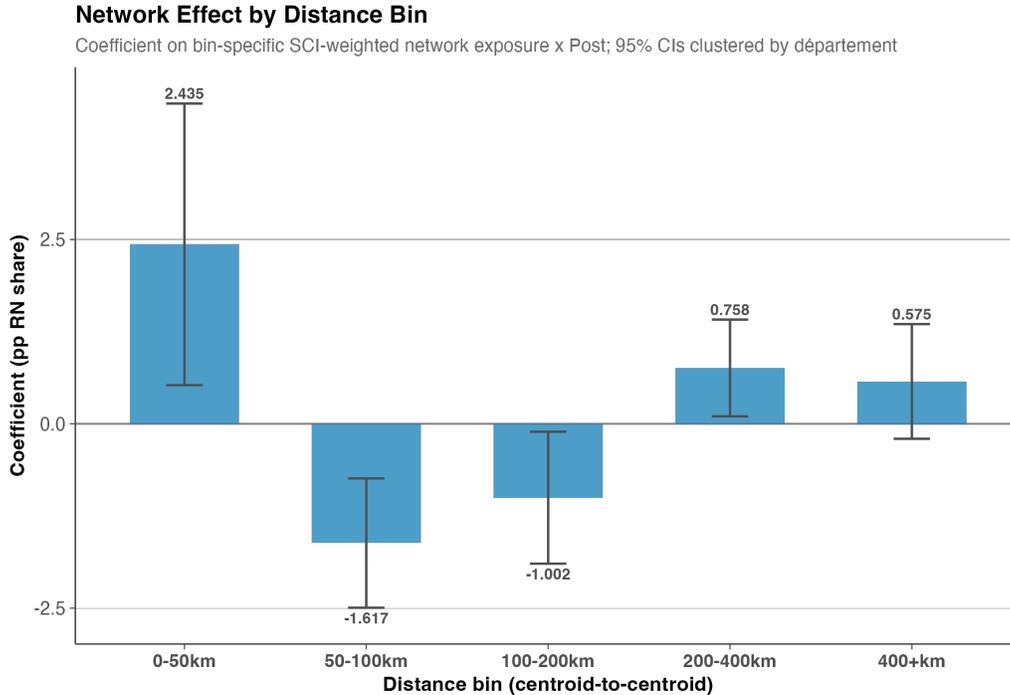
## 8.7 Network Pre-Determination

The SCI is available only in its 2024 vintage—a post-treatment snapshot. Three arguments address the concern that the carbon tax altered social connections.

**Migration proxy.** A pre-treatment network proxy from 2013 inter-département residential migration ( $\rho = 0.66$  with SCI) yields a coefficient of 1.454 (SE = 0.267,  $p < 0.001$ )—similar to the SCI-based estimate. **Pre-determination test.** The negative pre-treatment event-study coefficients are inconsistent with post-treatment network sorting, which would produce *positive* pre-trends. **Measurement error.** Under classical attenuation, the corrected estimate is  $1.35/0.66 \approx 2.05$  pp, placing the plausible range at  $[1.35, 2.05]$  pp.

## 8.8 Distance-Bin Decomposition

Figure 5 decomposes the network effect across five distance bins. The pattern is non-monotonic: positive at 0–50 km and 400+ km, negative at 50–200 km. Ties at 50–200 km disproportionately connect adjacent urban-rural pairs, bridging communities with opposing political valences—a “contact hypothesis” channel. Long-distance ties (400+ km) connect ideologically similar communities through kinship and migration. A simplified two-bin version confirms: the distant ( $>200$  km) effect is consistently positive and significant, while the nearby effect is smaller and less precisely estimated.



**Figure 5:** Network Effect by Distance Bin. Long-distance social ties (>200 km) transmit the backlash; intermediate ties (50–200 km) connecting urban-rural pairs may moderate it. *Notes:* Coefficient on bin-specific SCI-weighted network exposure  $\times$  Post in the département-level specification. SCI restricted to pairs within each distance range; exposure re-normalized within each bin. Error bars show 95% CIs from département-clustered SEs.

## 8.9 Placebo Timing

Assigning fake treatment dates before 2014 (restricted to pre-treatment elections) yields network coefficients of 0.30–0.58 pp—substantially smaller than the actual 1.35 pp but marginally significant for the 2004 and 2007 placebos. The pattern suggests a modest pre-existing correlation between network fuel exposure and Rassemblement National trends that the carbon tax *amplified*: the structural break at 2014 more than doubles the magnitude. The non-null placebos caution that the reduced-form estimate should be interpreted as an upper bound on the causal network effect.

## 8.10 Additional Controls and Bartik Diagnostics

Adding log median income  $\times$  Post reduces the network coefficient to 1.01 (SE = 0.20)—still highly significant. Region  $\times$  election fixed effects, which require the effect to operate *within* regions, yield 0.92 (SE = 0.44,  $p = 0.04$ ). Rotemberg weights (Goldsmith-Pinkham et al., 2020) confirm well-dispersed identification: the top-5 weights sum to 0.265, the Herfindahl

index is 0.025, and no single département contributes more than 8%. A shift exogeneity test yields  $p = 0.108$ —borderline, motivating the income control.

**Table 8:** Additional Robustness: Donut Design and Pre-Trend Adjustment

	rn_share	
	Donut (1)	Pre-Trend Adj. (2)
own_fuel_std × post_donut	2.210*** (0.4504)	
post_donut × network_fuel_std	1.644*** (0.5529)	
net_trend		0.1908*** (0.0389)
own_fuel_std × post_carbon		1.722*** (0.3693)
post_carbon × network_fuel_std		-1.019* (0.5560)
Observations	768	960
R <sup>2</sup>	0.95585	0.95319
Within R <sup>2</sup>	0.43849	0.36955
dept_code fixed effects	✓	✓
id_election fixed effects	✓	✓

*Notes:* Département-level, weighted by registered voters. Donut: drops 2012 and 2014 elections. Pre-Trend Adjusted: adds Network Fuel × (year − 2012) linear trend. SEs clustered by département. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

### 8.11 Bartik Residualization and Orthogonalized Channels

Residualizing fuel vulnerability on immigration, education, and urbanization before constructing the shift-share (Borusyak et al., 2022) yields a positive network coefficient, confirming the fuel channel is not mechanically driven by demographics. Orthogonalizing network fuel exposure on network immigration exposure also survives—fuel-specific variation predicts Rassemblement National gains independently of immigration proximity (Table 9).

### 8.12 Donut Design

Dropping the boundary elections (2012, 2014) yields a *stronger* network coefficient: 1.64 (SE = 0.55,  $p < 0.01$ ). The result is not driven by the contrast between elections immediately

**Table 9:** Channel Decomposition: Bartik Residualization and Orthogonalized Fuel Exposure

	rn_share	
	Bartik Residualized (1)	Orthogonalized (2)
own_fuel_resid_std $\times$ post_carbon	1.168* (0.6665)	
post_carbon $\times$ network_fuel_resid_std	0.4552 (0.4982)	
own_fuel_std $\times$ post_carbon		1.480*** (0.2071)
post_carbon $\times$ network_fuel_orth_std		0.5375* (0.2966)
post_carbon $\times$ network_immig_std		-1.633*** (0.2695)
Observations	960	960
R <sup>2</sup>	0.93148	0.95309
Within R <sup>2</sup>	0.07724	0.36821
dept_code fixed effects	✓	✓
id_election fixed effects	✓	✓

*Notes:* Département-level, weighted by registered voters. Column 1 residualizes fuel vulnerability on immigration, education, and urbanization before constructing the shift-share (Borusyak et al., 2022). Column 2 orthogonalizes network fuel exposure on network immigration exposure; the ‘pure fuel’ residual captures variation unrelated to immigration proximity. SEs clustered by département. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

flanking the tax introduction.

### 8.13 Pre-Trend Sensitivity

Adding a linear trend interaction ( $\text{Net}_d \times (\text{year} - 2012)$ ) drives the network  $\times$  post coefficient to  $-1.02$  (Table 8)—expected, since the linear trend is highly collinear with the post indicator and absorbs most treatment variation. The [Rambachan and Roth \(2023\)](#) sensitivity analysis (Section 8.16) provides the appropriate formal framework, yielding a robust CI of  $[0.40, 2.21]$  at  $\bar{M} = 0$ .

### 8.14 Triple-Difference: RN vs. Green

A triple-difference stacking Rassemblement National and Green vote shares confirms party specificity: the Net  $\times$  Post  $\times$  RN coefficient captures the differential network effect, absorbing common network-correlated trends.

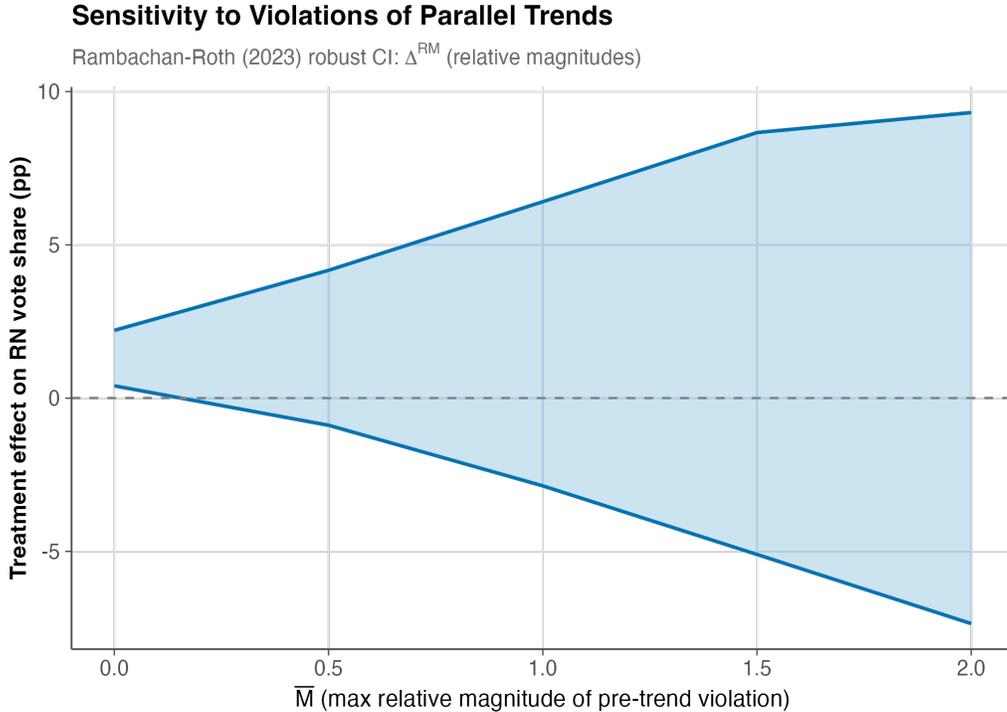
### 8.15 Block RI Power Analysis

The block RI null ( $p = 0.88$  within NUTS-2) likely reflects insufficient power rather than absence of effect. A Monte Carlo analysis confirms: with 13 blocks averaging 7 départements, the minimum detectable effect at 80% power is approximately 1.27 pp—near the estimated effect.

### 8.16 HonestDiD Sensitivity Analysis

The joint  $F$ -test of pre-treatment coefficients rejects exact zeros ( $p = 0.03$ ), but the rejection is driven by *negative* coefficients—the opposite direction from the treatment effect. How much pre-trend violation can the estimate absorb? Figure 6 applies the [Rambachan and Roth \(2023\)](#) sensitivity analysis. At  $\bar{M} = 0$  (standard parallel trends), the confidence interval is  $[0.40, 2.21]$ —bounded away from zero. The lower bound crosses zero at  $\bar{M} \approx 0.5$ , meaning the effect survives pre-trend violations up to half the maximum observed coefficient change.

The pre-treatment pattern supports the causal interpretation despite the  $F$ -test: all four coefficients are negative (wrong sign for upward-biasing confounds), the largest ( $-0.48$ ) is one-third the smallest post-treatment coefficient ( $+0.92$ ), and the break at 2014 is immediate and sustained.



**Figure 6:** HonestDiD Sensitivity: The Network Effect Remains Positive Under Standard Parallel Trends ( $\bar{M} = 0$ ) and Survives Pre-Trend Violations Up to  $\bar{M} \approx 0.5$   
*Notes:* Robust CIs following [Rambachan and Roth \(2023\)](#) with relative magnitudes restriction. At  $\bar{M} = 0$ , CI = [0.40, 2.21]. The dashed line marks zero.

## 9. Discussion

### 9.1 What Networks Transmit

The composite network effect (1.35 pp, HonestDiD lower bound 0.40 pp) bundles economic and cultural content. The fuel-specific component (0.58 pp) survives orthogonalization, but the immigration-related component is quantitatively larger, and the two channels overlap substantially ( $\delta = 0.10$ ). The carbon tax functioned as a *political trigger*: it activated a broader populist realignment already latent in France’s social geography. The Rassemblement National’s bundled platform—anti-tax, anti-immigration, anti-establishment—traveled through the same social connections. The DeGroot framework clarifies: the network multiplier amplifies whatever attitudes exist in connected communities. The carbon tax provided the catalyst; the network structure determined how far the backlash traveled.

## 9.2 Implications for Climate Policy

Revenue recycling—the standard prescription for politically sustainable carbon pricing (Klenert et al., 2018; Stiglitz, 2019)—eliminates variation in the economic “shift” but leaves the network structure unchanged. The descriptive decomposition shows that immigration-related content travels through the same connections as fuel grievance. Once the carbon tax activates a network-mediated backlash, the message travels regardless of whether the economic burden is compensated.

This points toward two complementary strategies: reducing the *salience* of the tax (embedding it in wholesale prices rather than printing it on retail receipts) and targeting the *network* (information campaigns that reach communities where anti-tax sentiment is being amplified by connections to fuel-vulnerable areas).

## 9.3 The Reflection Problem

The reflection problem—that connected areas may share unobserved characteristics (Manski, 1993)—is mitigated by three design features. First, network exposure uses *other* départements’ fuel vulnerability, a pre-determined characteristic, not contemporaneous outcomes. This breaks the simultaneity driving the reflection problem. Second, commune fixed effects absorb all time-invariant correlated effects. Third, the coefficient survives the 200 km distance restriction, ruling out geographic proximity as the source.

## 9.4 Generalizability

The conditions enabling network-mediated backlash—geographic variation in policy incidence, dense social networks, a populist party channeling the grievance—are present in many democracies. France provides an unusually clean setting because of the sharp tax introduction date, dramatic spatial variation in fuel vulnerability (Paris vs. *la France profonde*), and direct network measurement via SCI.

Generalizability depends on context. Countries with less centralized urban structures may exhibit weaker geographic concentration of policy burden. Countries without a populist party to channel the grievance may see effects through turnout rather than vote switching. The theoretical framework is general; the magnitudes are context-specific.

## 10. Conclusion

The carbon tax is the economists’ first-best instrument for climate change. It is also a political catalyst whose reaction travels through social networks in ways that standard policy analysis

misses.

Social networks transmitted the political backlash against France’s carbon tax far beyond its economic incidence. Network exposure to fuel-vulnerable areas raised Rassemblement National vote share by 1.35 pp per standard deviation (HonestDiD lower bound: 0.40 pp), activating at the 2014 tax introduction, scaling with the rate, persisting through five post-treatment elections, and operating exclusively through the Rassemblement National. Descriptively, both fuel-specific and immigration-related network exposures predict Rassemblement National gains—the carbon tax triggered a broader populist realignment already latent in France’s social geography.

For climate policy, the implication is uncomfortable. The political feasibility of carbon pricing depends not on the average cost per household—which is small—but on whether the cost falls on communities socially connected to the broader electorate. Revenue recycling can offset the economic burden. But it cannot recall the message.

## References

- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales**, “Shift-Share Designs: Theory and Inference,” *Quarterly Journal of Economics*, 2019, 134 (4), 1949–2010.
- Allcott, Hunt and Matthew Gentzkow**, “Social Media and Fake News in the 2016 Election,” *Journal of Economic Perspectives*, 2017, 31 (2), 211–236.
- Aronow, Peter M. and Cyrus Samii**, “Estimating Average Causal Effects Under General Interference, with Application to a Social Network Experiment,” *Annals of Applied Statistics*, 2017, 11 (4), 1912–1947.
- Autor, David H., David Dorn, and Gordon H. Hanson**, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 2013, 103 (6), 2121–2168.
- Bailey, Michael, Abhinav Gupta, Sebastian Hillenbrand, Theresa Kuchler, Robert Richmond, and Johannes Stroebel**, “Social Connectedness in Europe,” *Journal of the European Economic Association*, 2021, 19 (4), 2046–2089.
- , **Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong**, “Social Connectedness: Measurement, Determinants, and Effects,” *Journal of Economic Perspectives*, 2018, 32 (3), 259–280.

- Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler**, “A 61-Million-Person Experiment in Social Influence and Political Mobilization,” *Nature*, 2012, *489*, 295–298.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-Experimental Shift-Share Research Designs,” *Review of Economic Studies*, 2022, *89* (1), 181–213.
- Bureau, Dominique, Fanny Henriët, and Katheline Schubert**, “Distributional Effects of Carbon Pricing in France,” *European Economic Review*, 2019, *119*, 430–447.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, *90* (3), 414–427.
- Chetty, Raj, Matthew O. Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B. Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob, Drew Johnston, Martin Koenen, Eduardo Laguna-Muggenburg, Florian Mudekereza, Tom Rutter, Nicolaj Thor, Wilbur Townsend, Ruby Zhang, Michael Bailey, Pablo Barberá, Monica Bollen, and Frauke Kreuter**, “Social Capital I: Measurement and Associations with Economic Mobility,” *Nature*, 2022, *608*, 108–121.
- 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.
- , Livio Di Lonardo, Federico M. Ferrara, Samuel Ferroni, and Piero Stanig**, “The Political Consequences of Green Policies: Evidence from Italy,” *American Economic Review: Insights*, 2024, *6* (3), 388–406.
- Conley, Timothy G.**, “GMM Estimation with Cross Sectional Dependence,” *Journal of Econometrics*, 1999, *92* (1), 1–45.
- DeGroot, Morris H.**, “Reaching a Consensus,” *Journal of the American Statistical Association*, 1974, *69* (345), 118–121.
- Douenne, Thomas and Adrien Fabre**, “Yellow Vests, Pessimistic Beliefs, and Carbon Tax Aversion,” *American Economic Journal: Economic Policy*, 2022, *14* (1), 81–110.
- Enikolopov, Ruben, Alexey Makarin, and Maria Petrova**, “Social Media and Protest Participation: Evidence from Russia,” *Econometrica*, 2020, *88* (4), 1479–1514.

- Fetzer, Thiemo**, “Did Austerity Cause Brexit?,” *American Economic Review*, 2019, 109 (11), 3849–3886.
- Flückiger, Matthias and Markus Ludwig**, “How Social Networks Shape Our Beliefs: A Natural Experiment during the US Racial Justice Movement,” *Journal of Public Economics*, 2025, 241, 105085.
- Funke, Manuel, Moritz Schularick, and Christoph Trebesch**, “Going to Extremes: Politics after Financial Crises, 1870-2014,” *European Economic Review*, 2016, 88, 227–260.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, 2020, 110 (8), 2586–2624.
- Golub, Benjamin and Matthew O. Jackson**, “Naïve Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2010, 2 (1), 112–149.
- Klenert, David, Linus Mattauch, Emmanuel Combet, Ottmar Edenhofer, Cameron Hepburn, Ryan Rafaty, and Nicholas Stern**, “Making Carbon Pricing Work for Citizens,” *Nature Climate Change*, 2018, 8 (8), 669–677.
- LeSage, James and R. Kelley Pace**, *Introduction to Spatial Econometrics*, Boca Raton: Chapman and Hall/CRC, 2009.
- Manski, Charles F.**, “Identification of Endogenous Social Effects: The Reflection Problem,” *Review of Economic Studies*, 1993, 60 (3), 531–542.
- Oster, Emily**, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business & Economic Statistics*, 2019, 37 (2), 187–204.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Rodrik, Dani**, “Why Does Globalization Fuel Populism? Economics, Culture, and the Rise of Right-Wing Populism,” *Annual Review of Economics*, 2021, 13, 133–170.
- Stantcheva, Stefanie**, “Understanding Tax Policy: How Do People Reason?,” *Quarterly Journal of Economics*, 2021, 136 (4), 2309–2369.
- Stiglitz, Joseph E.**, “Addressing Climate Change through Price and Non-Price Interventions,” *European Economic Review*, 2019, 119, 594–612.

## A. Data Appendix

### A.1 Election Data

Election results are drawn from the official results published by the Ministère de l’Intérieur on [data.gouv.fr](https://data.gouv.fr). I use two files: `candidats_results.parquet` (candidate-level results by commune) and `general_results.parquet` (commune-level turnout and totals). The ten elections span 2002–2024:

**Table A1:** Elections in the Panel

Year	Type	Rassemblement National/Front National Nat'l %	€/tCO <sub>2</sub>
2002	Presidential (R1)	17.3	0
2004	European	11.1	0
2007	Presidential (R1)	12.5	0
2009	European	7.6	0
2012	Presidential (R1)	21.2	0
2014	European	29.2	7.0
2017	Presidential (R1)	26.2	30.5
2019	European	28.0	44.6
2022	Presidential (R1)	29.1	44.6
2024	European	38.3	44.6

### A.2 Party Identification

For the Front National/Rassemblement National, I match candidates by party label (FN, RN) and candidate name (Le Pen, Zemmour lists for allied candidacies). For European elections, list-based identification uses the list label fields. The party rebranded from Front National to Rassemblement National in June 2018; both labels appear in the data.

For placebo parties: Green candidates include EELV, VEC, Europe Écologie, and named candidates (Jadot, Joly, Voynet, Mamère). Center-right includes UMP, LR, RPR, and named candidates (Chirac, Sarkozy, Fillon, Péresse).

### A.3 Social Connectedness Index

I use the SCI at the NUTS-3 level (2024 vintage), downloaded from Meta’s public data repository. This is the only vintage available for France at the NUTS-3 resolution; see the identification appendix (below) for a discussion of the post-treatment measurement concern. The raw SCI is a relative probability (not a count), comparable across region pairs but

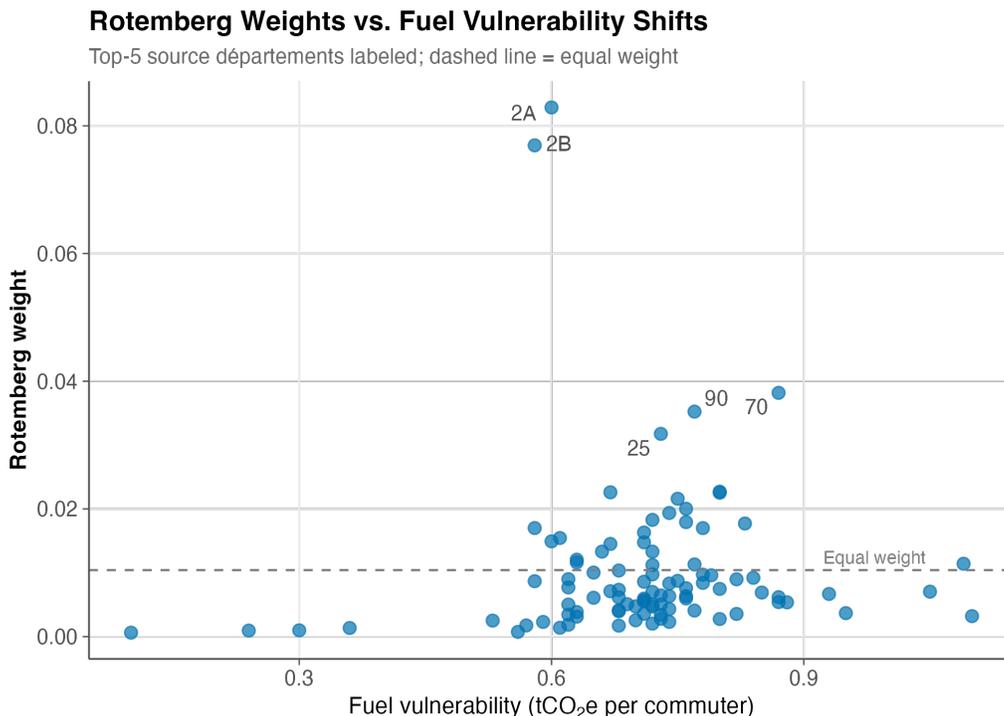
not across vintages. I map NUTS-3 codes to French *département* codes, drop overseas *départements* (97x) and non-metropolitan units, and row-normalize the resulting  $96 \times 96$  matrix.

#### A.4 Fuel Vulnerability

Commuting CO<sub>2</sub> emissions are from the INSEE/CGDD *Base Carbone des Déplacements Domicile-Travail*, which estimates emissions from home-to-work travel using census data on commute distance, mode, and vehicle fleet composition. The variable is measured at the commune level and aggregated to *départements*.

#### A.5 Bartik/Shift-Share Diagnostics

Following [Goldsmith-Pinkham et al. \(2020\)](#), I compute Rotemberg weights  $\hat{\alpha}_d$  for each *département*  $d$ 's contribution to the network exposure estimate. The top-5 Rotemberg weights sum to 0.265, and the Herfindahl index of weights is 0.025, indicating that no single *département* dominates the identification. A test of shift exogeneity yields  $p = 0.108$ .



**Figure A1:** Rotemberg Weights vs. Fuel Vulnerability Shifts. Identification is well-dispersed across *départements*.

*Notes:* Each point is a source *département*. Top-5 contributors are labeled. Dashed line indicates equal weight (1/96).

## B. Identification Appendix

### B.1 Parallel Trends

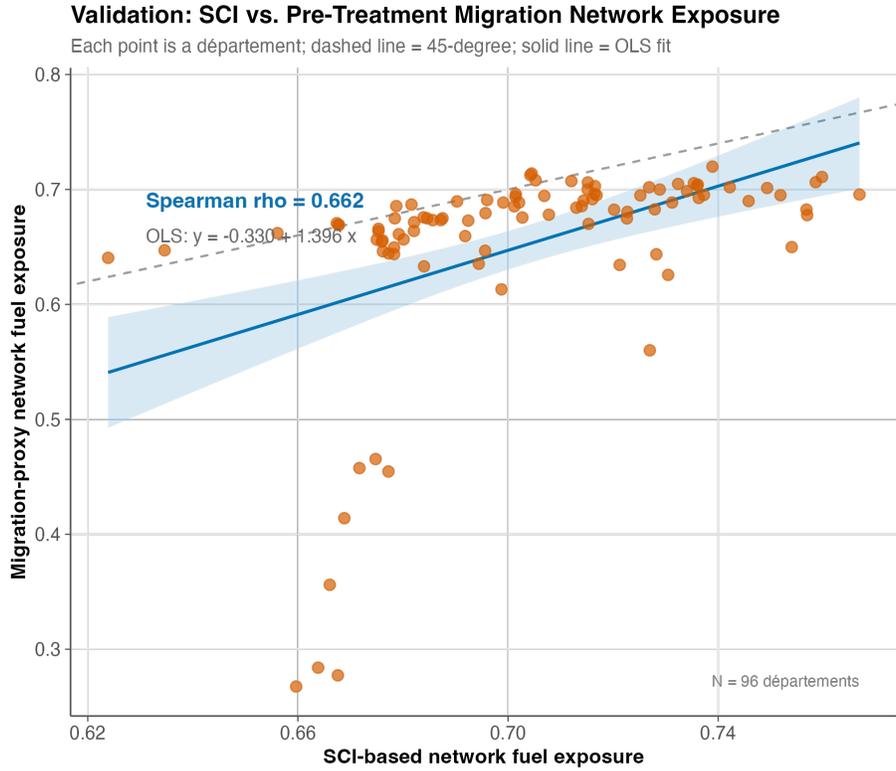
Section 6.5 presents the event study; here I report the exact pre-treatment coefficients for reference. Relative to 2012:  $-0.48$  (2002),  $-0.30$  (2004),  $-0.21$  (2007),  $-0.40$  (2009). A joint  $F$ -test rejects exact zeros ( $F = 2.69$ ,  $p = 0.030$ ), but all coefficients are negative (opposite sign from the post-treatment effect) and the post-treatment magnitudes (0.92–1.44 pp) are more than three times larger. The pattern is more consistent with a treatment effect superimposed on a mild pre-trend than with a smooth continuation of pre-existing differences. See Section 8.16 for the formal sensitivity analysis.

### B.2 Shift-Share Interpretation

Under the Borusyak et al. (2022) framework, the shift-share estimand can be interpreted as a weighted average of département-level treatment effects, with weights proportional to the SCI-based exposure shares. The identifying assumption is that the département-level fuel vulnerability “shifts” are mean-independent of the error term, conditional on controls. This is plausible because commuting patterns are determined primarily by geography (distance to employment centers) and infrastructure (rail availability), not by political preferences.

Under the Goldsmith-Pinkham et al. (2020) framework, identification comes from the exogeneity of the SCI “shares.” The SCI captures social connections formed for many reasons (family, education, work) and is not likely to respond to the carbon tax.

**SCI vintage validation.** Only the 2024 SCI is available for France at NUTS-3 resolution—a post-treatment snapshot. Section 8.7 addresses this concern in detail. Briefly: (i) Bailey et al. (2018) demonstrate SCI stability over time; (ii) the negative pre-treatment event-study coefficients are inconsistent with post-treatment sorting; (iii) Facebook penetration exceeded 50% by 2012; and (iv) a 2013 migration-based proxy ( $\rho = 0.66$ ; Figure A2) reproduces the result (Table A4).



**Figure A2:** SCI-Based vs. Migration-Based Network Exposure. The two measures are positively correlated ( $\rho = 0.66$ ), supporting network pre-determination.  
*Notes:* Each point is a département.  $x$ -axis: SCI-weighted network fuel exposure;  $y$ -axis: migration-based network fuel exposure (2013 inter-département residential mobility). Dashed line: 45 degrees; solid line: OLS fit.

### C. Commune-Level Results

The commune-level regressions replicate the 96 cross-sectional département-level exposures across approximately 37,000 communes. They serve as a check on aggregation bias but provide no additional identifying information beyond the département-level specification.

**Table A2:** Commune-Level Results (Ancillary)

	(1)	(2)	(3)	(4)	(5)	(6)
	Own	Network	Both	Post-GJ	Pres×Euro	Continuous
Own × Post	0.716*** (0.183)		0.716*** (0.183)	0.698*** (0.196)	0.888*** (0.221)	
Net × Post		1.192*** (0.237)	1.192*** (0.237)	1.067*** (0.244)	1.137*** (0.280)	
Own × Post × Pres					-0.344*** (0.102)	
Net × Post × Pres					0.111 (0.116)	
Own × Rate						0.020*** (0.005)
Net × Rate						0.031*** (0.006)
Commune FE	Yes	Yes	Yes	Yes	Yes	Yes
Election FE	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	361,796	361,796	361,796	361,796	361,796	361,796

*Notes:* Dependent variable is Rassemblement National/Front National first-round vote share (%). Own and Net are standardized département-level fuel vulnerability. Post = 1 for elections from 2014 onward. Rate is the carbon tax rate in €/tCO<sub>2</sub>. SEs clustered by département. The identifying variation is at the département level; commune-level regressions serve as a check on aggregation bias. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

## D. Additional Robustness Tables

**Table A3:** Controls Sensitivity: Network Coefficient Across Specifications

Specification	Net $\times$ Post	SE	Controls	$N$
Baseline (D2)	1.346***	(0.455)	—	960
+ Unemployment $\times$ Post	1.081***	(0.363)	1	960
+ Education $\times$ Post	0.877**	(0.341)	2	960
+ Immigration $\times$ Post	0.435	(0.294)	3	960
+ Industry $\times$ Post	1.324***	(0.388)	4	960
+ Dept-specific trends	0.065	(0.158)	4+trends	960
Kitchen sink (all + trends)	-0.219	(0.197)	4+trends	960

*Notes:* Dependent variable is Rassemblement National vote share at the département level. All specifications include département and election FEs. Controls are 2013 *Recensement* cross-sections interacted with Post. Dept-specific trends use election-year-number interaction. SEs clustered by département. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ . Exact coefficients generated by `04_robustness.R` and formatted by `06_tables.R`.

**Table A4:** Migration-Proxy Validation

	SCI-based		Migration-based	
	Dept-level	Commune	Dept-level	Commune
Network $\times$ Post	1.346*** (0.455)	1.192*** (0.237)	1.454*** (0.267)	1.101*** (0.132)
$N$	960	361,796	960	359,062

*Notes:* Migration-based network exposure uses gravity model weights from 2013 inter-département residential mobility. Spearman  $\rho$  between SCI-based and migration-based exposure = 0.66. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

## Acknowledgements

We thank participants at seminars for helpful comments and suggestions. All errors are our own. Replication code and data are available at <https://github.com/SocialCatalystLab/ape-papers>.