Networked Anxiety Without Contact: Asylum Dispersal and the Far-Right Network Multiplier in France

APEP Autonomous Research* @ai1scl

March 9, 2026

Abstract

Does social-network exposure to immigration amplify far-right support even in areas without direct immigrant contact? We study France's 2021 Schéma National d'Accueil, which redistributed asylum reception capacity from Île-de-France to other regions, using a shift-share design that interacts Facebook's Social Connectedness Index with new asylum reception capacity. Departments with stronger social ties to receiving areas experienced significantly higher Rassemblement National vote share gains ($\beta = 0.058$, $p < 0.001$). Own-department asylum capacity shows null effects, though this measure is imputed from regional aggregates and subject to substantial attenuation. The network coefficient is 2.3 times larger for non-hosting departments in an exploratory triple-difference, suggesting that socially connected exposure and direct contact may operate in opposite directions. Results survive randomization inference ($p < 0.001$), leave-one-out exclusion of each shift department, and alternative SCI normalizations.

JEL Codes: D72, F22, Z13

Keywords: immigration, social networks, far-right voting, contact hypothesis, shift-share, asylum policy

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

1. Introduction

In the 2024 European elections, the Rassemblement National won its largest vote shares not in the Parisian suburbs where asylum seekers actually live, but in rural departments hundreds of kilometers away—places where voters may never have encountered a refugee. Between 2014 and 2024, these distant departments swung harder toward the far right than the communities that directly hosted asylum seekers. This disconnect—between where immigrants arrive and where anti-immigrant sentiment grows fastest—is the puzzle at the heart of this paper.

Allport's (1954) contact hypothesis, one of the oldest and most replicated findings in social psychology, predicts that direct, sustained contact with outgroup members reduces prejudice. A growing body of evidence from refugee dispersal policies supports this prediction. Steinmayr (2021) shows that Austrian municipalities assigned refugees experienced smaller far-right gains than comparable municipalities without refugees. Schneider-Strawczynski (2024) finds that French departments hosting asylum seekers under the Centre d'Accueil de Demandeurs d'Asile (CADA) system saw reduced RN vote shares, conditional on institutional quality. Dustmann et al. (2019) documents that Danish municipalities receiving refugees showed attenuated anti-immigrant attitudes after sustained exposure. The local contact channel, in short, appears to work.

But most voters do not live next to an asylum reception center. They learn about immigration through social networks—friends, family, acquaintances who share news, anecdotes, and anxieties through phone calls, social media, and holiday visits. If the information transmitted through these networks emphasizes threat rather than familiarity, social connectivity could amplify anti-immigrant sentiment in precisely the departments that have no direct contact with asylum seekers. The contact hypothesis would then coexist with a network anxiety channel that operates in the opposite direction.

This paper tests whether social-network exposure to asylum dispersal increases far-right support even in the absence of direct contact. We exploit France's 2021 Schéma National d'Accueil (SNA), a centralized policy that redistributed asylum reception capacity from the saturated Île-de-France region to other departments across France. The policy created variation in regional asylum capacity that we use, through a shift-share framework, to study network-transmitted political effects. We combine this variation with Facebook's Social Connectedness Index (SCI), which measures the intensity of social ties between all pairs of French departments at the NUTS-3 level, to construct a shift-share measure of network exposure to asylum dispersal.

Our identification strategy follows Borusyak et al. (2022). The "shares" are pre-determined SCI weights reflecting long-run social ties between departments. The "shifts" are new asylum reception places created under the SNA in each department. The treatment variable, *NetworkDispersal*, is the SCI-weighted sum of new asylum places across all connected departments. Identification requires that the SCI-weighted sum of new places is uncorrelated with department-level shocks to RN support, conditional on department and election fixed effects. We validate this assumption through event-study evidence, leave-one-out stability tests, and randomization inference.

The main result is a strong correlation between network exposure and RN gains: a one-unit increase in network dispersal is associated with a 0.058 percentage point increase in RN vote share (SE = 0.007, $t = 7.9$). Standardized, a one-standard-deviation increase in network exposure corresponds to a 1.32 percentage point increase in RN vote share—or 0.17 standard deviations of the outcome (see Section F for standardized effect sizes). This is a meaningful association: 5.4% of the pre-treatment mean RN share of 24.4%. Meanwhile, own-department asylum capacity has a null effect ($\beta = -0.005$, $p > 0.10$), though this measure is imputed from regional aggregates and likely suffers from substantial attenuation bias.[1]

An exploratory triple-difference decomposition provides suggestive evidence on the role of hosting status. Among departments classified as non-hosting, the network coefficient is 0.150 ($p < 0.001$). Among hosting departments, it is 0.065 ($p < 0.05$)—still positive, but 2.3 times smaller. This pattern is consistent with socially connected exposure and direct contact operating in opposite directions, though hosting status is itself defined by the same imputed treatment variable and should be interpreted cautiously.

The results are robust across multiple dimensions. The event study shows parallel pre-trends between 2017 and 2019 (the reference year), with a sharp break at the 2022 election—the first election after the SNA's implementation.[2] Leave-one-out exclusion of each shift department produces a coefficient range of [0.057, 0.059], confirming that no single department drives the result. Randomization inference yields $p < 0.001$ (0 of 1,000 permutations exceed the observed statistic). Wild cluster bootstrap inference, addressing the finite-cluster concern with 96 departments, also yields $p < 0.001$.

Several important caveats apply. Our treatment variable is constructed from regional (NUTS-2) capacity aggregates equally distributed across departments, not from observed facility-level data. This means own-department hosting is imputed, and the effective number of independent shocks is closer to 13 regions than 96 departments. Our inference uses department-

[1] Because facility-level data is unavailable, own-department asylum capacity is constructed by equally distributing regional (NUTS-2) net changes across departments within each region. The resulting measurement error biases the own-department coefficient toward zero, making it impossible to distinguish between genuine contact effects and attenuation.

[2] The 2014 European election shows a marginally significant negative coefficient (-0.031, $p < 0.05$), which we interpret as mean reversion from the exceptionally high RN performance in that election, rather than a pre-trend violation. We discuss this extensively in Section 8.

clustered standard errors rather than the shock-level corrections recommended by Adao et al. (2019) for shift-share designs, likely overstating precision. The Social Connectedness Index captures geography and cultural similarity alongside interpersonal ties, so we cannot fully separate social network effects from spatial proximity spillovers. Despite these limitations, the consistency of the pattern across multiple robustness checks warrants attention.

This paper contributes to three literatures. First, we advance the political economy of immigration by providing suggestive evidence that immigration can affect voting *beyond the receiving community*. The existing literature—including Tabellini (2020), Dustmann et al. (2019), Steinmayr (2021), Schneider-Strawczynski (2024), Halla et al. (2017), and Edo et al. (2019)—focuses almost exclusively on the effects of immigrants *where they live*. Our results suggest that the political footprint of immigration may extend beyond receiving communities, potentially mediated by social networks, though we cannot fully rule out geographic spillovers as an alternative channel.

Second, we contribute to the growing literature on social networks and political behavior. Bailey et al. (2018) and Bailey et al. (2020) demonstrate that Facebook's SCI predicts a wide range of economic and social outcomes across space. Müller and Schwarz (2021) shows that social media exposure to anti-refugee content predicts hate crimes in Germany. Our shift-share design offers a framework for studying interpersonal network transmission separately from algorithmic amplification, though the treatment is constructed from regional aggregates and we cannot fully separate social network effects from geographic proximity spillovers.

Third, we build directly on the "Connected Backlash" framework introduced in APEP Autonomous Research (2026), which demonstrated that social-network exposure to carbon tax costs increased RN support in non-fuel-dependent departments.[3] Our paper extends this framework from economic policy shocks to immigration shocks, testing whether the network multiplier generalizes across fundamentally different policy domains. That both carbon taxation and asylum dispersal produce similar network backlash effects—despite operating through different material channels—suggests a general mechanism: social networks amplify the political salience of policy changes in connected communities, regardless of whether those communities are directly affected.

[3] See `https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0464` for the working paper. Results from APEP working papers should be interpreted with appropriate caution as they have not undergone human peer review.

2. Institutional Background

2.1 France's Asylum Reception System

France operates a centralized asylum reception system managed by the Office Français de l'Immigration et de l'Intégration (OFII) under the authority of the Ministère de l'Intérieur. Asylum seekers who register a claim with the Office Français de Protection des Réfugiés et Apatrides (OFPRA) are entitled to material reception conditions, including housing in dedicated reception centers while their claims are processed. Processing times averaged 6–12 months during our study period, with appeal adding another 6–12 months.

The reception infrastructure consists of two main facility types. Centres d'Accueil de Demandeurs d'Asile (CADA) are permanent facilities providing housing, legal assistance, and integration support for asylum seekers during the claims process. Centres d'Accueil et d'Examen des Situations (CAES) are shorter-term facilities designed for initial processing and orientation, typically housing asylum seekers for a few weeks before transfer to CADA or other accommodation. Together, these facilities constituted approximately 110,000 places nationally by the end of 2023, up from roughly 86,000 in 2020.

The geographic distribution of reception capacity was historically uneven. Île-de-France—the Paris region—hosted a disproportionate share of asylum seekers, reflecting both the concentration of OFPRA processing offices and the informal networks that directed new arrivals toward the capital. By 2019, Île-de-France housed approximately 40% of France's asylum seekers despite representing only 18% of the national population. This concentration strained local services and created visible encampments that became politically salient.

2.2 The Schéma National d'Accueil 2021–2023

In December 2020, the government published the first Schéma National d'Accueil des Demandeurs d'Asile et des Réfugiés (SNA), a binding national plan for redistributing asylum reception capacity across all regions. The SNA was mandated by Article L.744-2 of the Code de l'Entrée et du Séjour des Étrangers et du Droit d'Asile (CESEDA), introduced in the 2018 asylum reform law (Loi "asile et immigration" of September 10, 2018). The stated objective was to achieve a "balanced territorial distribution" of asylum seekers by reducing Île-de-France's share and expanding capacity in departments with few or no existing facilities.

The SNA set department-level targets for new CADA and CAES places, to be implemented between 2021 and 2023. Prefects (the central government's representatives in each department) were responsible for identifying sites and securing local agreements. The allocation formula was based on population, existing capacity, housing market conditions, and labor market indicators—explicitly not on local political conditions or far-right vote shares.[4]

Implementation was rapid. The government's target was to create over 16,000 new reception places nationally by the end of 2023, redistributing capacity away from Île-de-France. By early 2022—before the April 2022 presidential election, the first post-treatment election in our sample—a substantial share of the planned redistribution had been announced and initiated through prefectoral decisions. Our treatment variable, constructed from regional implementation reports (see Section 5), captures the net change in capacity allocated to each metropolitan department. The average net change was 56.8 places (SD = 99.6), though variation was large: some departments received substantial new capacity while others—particularly in Île-de-France—experienced net reductions as the national redistribution transferred places to other regions.[5]

2.3 Political Context: The Rassemblement National

The Rassemblement National (formerly Front National, renamed in 2018) is France's principal far-right party. Under Marine Le Pen's leadership since 2011, and increasingly under Jordan Bardella's public profile, the RN has made immigration restriction its signature issue. The party's vote share in first-round presidential and European elections grew steadily from approximately 18% in 2012 to over 31% in the 2024 European elections.

The RN's electoral geography is well-documented: strongest in the northeast (former industrial areas), the Mediterranean coast, and rural departments, with weaker support in western France, university cities, and central Paris (Perrineau, 2017). This spatial pattern is correlated with, but not determined by, local immigrant presence. The disconnect between immigrant settlement patterns (concentrated in Île-de-France, Lyon, Marseille) and RN voting patterns (strongest in departments with *few* immigrants) has been a longstanding puzzle in French political science, sometimes called the "halo effect" (Etchegaray and Monnier, 2019).

Our paper offers a candidate explanation for this halo effect: departments without immigrants but with strong social ties to immigrant-receiving areas may experience socially transmitted concern about immigration that increases far-right support, though we cannot definitively separate this mechanism from geographic proximity effects.

[4] The allocation criteria are detailed in the SNA circulaire of January 4, 2021. We verify in Section 8 that pre-treatment RN vote shares do not predict new asylum places.

[5] The sum of our treatment variable across 96 departments (approximately 5,500 net places) is lower than the 16,000-place government target because: (i) our measure captures net changes including reductions in Île-de-France, (ii) we use regional implementation reports rather than facility-level data, and (iii) some places were created in overseas departments outside our sample.

2.4 The Broader European Context

France's asylum dispersal policy does not exist in isolation. Several European countries have implemented similar redistribution schemes, each providing variation that has been exploited in the immigration-voting literature. Germany's Königsteiner Schlüssel distributes asylum seekers across Bundesländer according to a formula based on tax revenue (two-thirds weight) and population (one-third weight). Dustmann et al. (2019) and Peri (2016) study the labor market and political effects of refugee assignment across German and Danish municipalities, respectively. Sweden's Ekvivalent Bosättning Ordning (EBO) initially allowed refugees to choose their municipality of settlement, leading to extreme geographic concentration; the subsequent reform restricting free settlement to assigned municipalities provided quasi-experimental variation exploited in multiple studies.

What distinguishes France's SNA from these comparators is the combination of centralized allocation, rapid implementation, and the availability of granular social-connectivity data through the SCI. The SNA's allocation formula—based on objective criteria rather than refugee preferences or local political lobbying—provides particularly clean shift variation for our design. Moreover, the French departmental structure, with 96 metropolitan units at the NUTS-3 level, offers a favorable balance between sufficient observations for inference and meaningful within-unit social cohesion.

3. Related Literature

This paper contributes to several interconnected literatures on immigration, social networks, and political behavior.

3.1 Immigration and Far-Right Voting

A large literature examines the relationship between immigrant presence and far-right support. The dominant finding is complex: at the aggregate level, areas with more immigrants often show stronger anti-immigrant voting, but studies exploiting quasi-experimental variation in refugee placement frequently find that direct local exposure reduces far-right support or has null effects.

Tabellini (2020) studies the political effects of the Age of Mass Migration (1850–1920) in the United States, finding that European immigrants to US cities increased public spending and redistributive preferences, consistent with a cultural incorporation story rather than backlash. Dustmann et al. (2019) examines Danish refugee placement and finds that municipalities assigned refugees showed attenuated anti-immigrant attitudes after sustained exposure.

Steinmayr (2021) provides particularly clean identification from Austrian municipalities that were assigned refugees through a federal dispersal scheme, finding that direct exposure reduced FPÖ (far-right) vote shares by 3–4 percentage points.

In the French context, Edo et al. (2019) uses historical immigrant settlement patterns as instruments for current immigration and finds that immigration increased FN vote shares at the commune level. Schneider-Strawczynski (2024) studies France's CADA system directly, finding that communes hosting asylum seekers experienced smaller RN gains conditional on facility quality and integration support. Halla et al. (2017) examines Austrian municipalities, finding that immigration increases FPÖ support, though the effect is concentrated in areas with limited prior contact.

The critical gap in this literature is the focus on *where immigrants are*. Our contribution is to show that immigration's political effects extend far beyond receiving communities, transmitted through social networks to departments with no direct exposure.

3.2 Social Networks and Political Behavior

A growing literature documents that social networks shape political attitudes, beliefs, and behavior. DellaVigna and Kaplan (2007) shows that the introduction of Fox News increased Republican vote shares, demonstrating that media exposure shifts political preferences. Gerber et al. (2008) finds that social pressure increases voter turnout, and Bond et al. (2012) demonstrates that social network messages on Facebook increased turnout in the 2010 US midterm elections.

More directly relevant, Müller and Schwarz (2021) shows that exposure to anti-refugee content on social media predicts hate crimes in Germany, establishing that online social networks can translate immigration anxiety into real-world outcomes. Bursztyn et al. (2020) demonstrates that exposure to misinformation about immigrants through social media increases anti-immigrant attitudes in experimental settings.

Our paper differs from this social media literature in an important respect: we measure interpersonal social ties (through the SCI) rather than content exposure. The SCI captures the probability that two individuals in different regions are Facebook friends—a measure of real social connections, not algorithmic content delivery. This distinction matters because our channel operates through interpersonal communication (sharing of experiences, anxieties, and narratives about asylum seekers) rather than through exposure to curated or algorithmically amplified content.

3.3 Shift-Share Designs in Political Economy

The shift-share (Bartik) design has become a workhorse identification strategy in labor economics and political economy. Goldsmith-Pinkham et al. (2020) provides the modern econometric framework, showing that identification comes from the exogeneity of either the shares or the shifts. Borusyak et al. (2022) offers a complementary perspective focused on shift-level exogeneity, which is the approach we follow. Adao et al. (2019) addresses inference, showing that standard clustering can be insufficient when shifts induce mechanical correlation across units, and proposing adjusted standard errors.

In political economy, shift-share designs have been used to study the effects of trade exposure on voting (Autor et al., 2020), the effects of fiscal transfers on political support (Fetzer, 2019), and the effects of immigration on labor markets (Peri, 2016). Our application is most closely related to APEP Autonomous Research (2026), which uses SCI-weighted carbon tax exposure as a shift-share treatment and finds that network exposure to fuel tax costs increased RN support. We extend this framework to a fundamentally different policy domain—immigration rather than taxation—testing the generalizability of the network multiplier mechanism.

4. Theoretical Framework

We develop a simple framework that generates two opposing predictions about the political effects of asylum dispersal, depending on the channel of exposure.

4.1 Contact Channel

Allport's (1954) contact hypothesis, in its modern formulation, predicts that sustained intergroup contact reduces prejudice when four conditions are met: equal status, common goals, cooperation, and institutional support. Asylum reception centers, while imperfect, provide a degree of institutional structure that facilitates contact. Local residents interact with asylum seekers in shops, schools, public transport, and neighborhood spaces. These interactions provide concrete information that countervails abstract fears.

Formally, let θ_i denote the perceived immigration threat in department i. Under contact:

$$\frac{\partial \theta_i}{\partial D_i} < 0 \tag{1}$$

where D_i is direct asylum capacity in department i. Direct exposure reduces perceived threat, attenuating far-right support.

4.2 Network Anxiety Channel

Social networks transmit information selectively. Negative, threatening, or novel information spreads faster and farther than positive or mundane information—a well-documented asymmetry in information diffusion (Baumeister et al., 2001; Soroka, 2014). When asylum seekers arrive in a department, residents of socially connected departments learn about the arrival not through direct experience but through secondhand accounts that emphasize disruption, cost, and cultural difference. This "negativity bias" in network transmission generates anxiety about immigration even without contact.

Let S_{ij} denote the social connectedness between departments i and j, and D_j the asylum capacity in department j. Network exposure is:

$$N_i = \sum_{j \neq i} S_{ij} \cdot D_j \tag{2}$$

The network anxiety prediction is:

$$\frac{\partial \theta_i}{\partial N_i} > 0 \tag{3}$$

Network exposure *increases* perceived threat, amplifying far-right support. Crucially, this prediction operates in the opposite direction from the contact prediction in (1).

4.3 Testable Implications

The framework generates three predictions that we take to the data:

1. **Network effect:** Network dispersal (N_i) increases RN vote share ($\beta_{\text{network}} > 0$).

2. **Contact null:** Own dispersal (D_i) has a null or negative effect on RN vote share ($\beta_{\text{own}} \leq 0$).

3. **Differential network effect:** The network effect is larger in non-hosting departments, where the contact channel provides no offsetting force. In hosting departments, contact partially offsets network anxiety.

These three predictions are jointly tested in our empirical framework. The triple-difference specification in Section 7 directly identifies whether hosting status moderates the network channel.

5. Data

Our analysis combines four data sources at the French department level (NUTS-3), constructing a balanced panel of 96 metropolitan departments across 5 national elections spanning 2014–2024.

5.1 Election Data

We obtain department-level election results from data.gouv.fr for five elections: the 2014 European Parliament election, the 2017 presidential first round, the 2019 European Parliament election, the 2022 presidential first round, and the 2024 European Parliament election. Our outcome variable is the RN vote share, defined as the combined first-round vote share of the Rassemblement National (and its predecessor Front National) plus allied extreme-right candidates where applicable.

The panel includes both presidential and European elections, which differ in turnout and voter composition. European elections typically have lower turnout and higher far-right vote shares, reflecting differential mobilization. We include election fixed effects throughout to absorb these level differences. The pre-treatment period includes three elections (2014, 2017, 2019) and the post-treatment period includes two (2022, 2024).

5.2 Social Connectedness Index

The Social Connectedness Index (SCI), released by Facebook Research (Bailey et al., 2018), measures the relative probability that pairs of individuals in two regions are Facebook friends. We use the NUTS-3 (department) level data for France, which provides pairwise SCI values for all $96 \times 96 = 9{,}216$ department pairs, of which 9,120 are off-diagonal. The SCI exhibits substantial variation: the coefficient of variation across French department pairs is 2.43, reflecting the well-known clustering of social ties along geographic, linguistic, and historical lines.

We row-normalize the SCI so that each department's outgoing connection weights sum to one. This ensures that *NetworkDispersal* measures the average exposure to asylum dispersal across a department's social network, rather than reflecting the overall volume of social connections. We verify robustness to alternative normalizations (log SCI, binary treatment) in Section 8.

The SCI has been extensively validated as a measure of real social ties. Bailey et al. (2018) show that SCI predicts migration flows, trade, patent citations, and disease transmission across US counties and international regions. Bailey et al. (2020) extend this validation to

European NUTS regions, demonstrating that SCI captures historical migration patterns, linguistic proximity, and commuting linkages. Facebook penetration among French adults exceeds 75%, reducing concerns that the SCI systematically misrepresents social connectivity for specific demographic groups or department types. Nevertheless, to the extent that Facebook friendship overmeasures ties among younger and more urban populations, our estimates may overweight transmission channels used by these groups.

5.3 Asylum Reception Capacity

We construct department-level measures of new asylum reception places created under the SNA 2021–2023 from regional capacity figures reported in the Cour des comptes annual evaluation of the SNA (2022, 2023) and OFII activity reports.[6] These sources report the net change in reception places at the regional (NUTS-2) level. We distribute regional allocations equally across departments within each region, following the SNA's stated principle of balanced territorial distribution.[7] Departments that gained capacity appear as positive shifts; those that lost capacity (primarily in Île-de-France) appear as negative shifts.

An important feature of our treatment measure is that it captures the *announced* SNA allocation rather than time-varying implementation. The SNA targets were published in December 2020 and communicated to prefects in early 2021. By the April 2022 presidential election, approximately 75% of planned places had been created; by the June 2024 European election, implementation was essentially complete. Our treatment is best interpreted as intent-to-treat exposure to the announced dispersal policy, which captures both realized capacity changes and anticipated future arrivals that entered public discourse through prefectoral announcements, local media, and interpersonal communication. The event study in Section 7 shows a smaller coefficient in 2022 (0.037) than in 2024 (0.056), consistent with a dose-response pattern in which the effect strengthens as implementation proceeds.

5.4 Department Characteristics

We collect department-level characteristics from INSEE's Banque de Données Macroéconomiques (BDM) measured in 2019 (the last pre-treatment year): log population, unemployment rate, share of the population with a university degree, share of the population that is foreign-born, and median household income. Because these variables are time-invariant in our panel,

[6]Facility-level data on individual CADA and CAES sites is not publicly available in machine-readable form. Hand-collection from prefectoral press releases could potentially validate our imputation, but is beyond the scope of this study.

[7]This equal-distribution assumption introduces measurement error if true within-region allocation was concentrated in specific departments. Under classical measurement error, this biases our estimates toward zero (attenuation), making our results conservative.

they are absorbed by department fixed effects and do not enter the regression directly. We use them for balance tests (Section B), heterogeneity analysis (Section 9), and descriptive characterization of the treatment distribution.

5.5 Summary Statistics

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max
RN Vote Share (%)	26.7	7.7	5.0	51.5
Network Dispersal	57.85	22.59	-29.68	88.85
Own New Places	56.8	99.6	-262.5	138.0
Total Votes	277,655	214,060	29,245	1,353,803
Departments			96	
Elections			5	
Observations			480	

Notes: Panel of metropolitan French departments across national elections (2014–2024). RN Vote Share is the first-round vote share of Rassemblement National (formerly Front National) and extreme-right allies. Network Dispersal is the SCI-weighted sum of new asylum reception places created in connected departments under the Schéma National d'Accueil 2021–2023. Own New Places measures direct asylum capacity expansion within each department.

Table 1 presents summary statistics. The mean RN vote share across the full panel is 26.7%, reflecting the party's secular rise over the decade. The pre-treatment mean (2014–2019) is 24.4% and the post-treatment mean (2022–2024) is 30.1%, a 5.7 percentage point increase that reflects both the national RN trend and the network dispersal effect we identify. Network Dispersal has a mean of 57.9, a standard deviation of 22.6, and substantial variation (range: −29.7 to 88.9). The negative values arise from departments with strong ties to Île-de-France, which experienced net *reductions* in asylum capacity. Own New Places averages 56.8 with high dispersion (SD = 99.6), reflecting the lumpy nature of facility placement.

6. Empirical Strategy

6.1 Shift-Share Design

Our identification strategy is a shift-share (Bartik) design in the tradition of Borusyak et al. (2022) and Goldsmith-Pinkham et al. (2020). The treatment variable for department i is:

$$\text{NetworkDispersal}_i = \sum_{j \neq i} \tilde{S}_{ij} \cdot \text{NewPlaces}_j \tag{4}$$

where $\tilde{S}_{ij} = S_{ij} / \sum_{k \neq i} S_{ik}$ is the row-normalized SCI weight between departments i and j (self-links S_{ii} are excluded from both the numerator and denominator, as our SCI data contains only off-diagonal pairs), and NewPlaces_j is the net change in asylum reception places in department j under the SNA.

The shares ($\tilde{S}_{ij}$) capture pre-determined social connectedness between departments. The shifts (NewPlaces_j) capture the policy-induced variation in asylum capacity. The product aggregates exposure: departments with stronger social ties to departments that received more asylum seekers have higher *NetworkDispersal*.

6.2 Estimation

Our main estimating equation is:

$$\text{RN}_{it} = \alpha_i + \gamma_t + \beta \cdot \text{NetworkDispersal}_i \times \text{Post}_t + \varepsilon_{it} \tag{5}$$

where RN_{it} is the RN vote share in department i in election t, α_i are department fixed effects, γ_t are election fixed effects, and $\text{Post}_t = \mathbb{I}[t \geq 2022]$ indicates post-treatment elections. The coefficient β captures the effect of a one-unit increase in SCI-weighted asylum exposure on the change in RN vote share from pre- to post-treatment, relative to departments with lower network exposure. Because the department fixed effects absorb all time-invariant department characteristics—including the pre-treatment controls described in Section 5 (log population, unemployment, education, foreign-born share, and income measured in 2019)—these controls do not enter the regression separately. The identifying variation comes entirely from the interaction of cross-sectional differences in *NetworkDispersal* with the pre-post timing of the SNA.

We extend the baseline in several directions. First, we include own-department asylum capacity ($\text{OwnDispersal}_i \times \text{Post}_t$) to separately identify the contact channel. Second, we estimate a triple-difference specification that interacts *NetworkDispersal* with an indicator for whether department i itself hosts asylum seekers, testing whether the network effect differs by hosting status. Third, we estimate event-study specifications that interact *NetworkDispersal* with individual election dummies:

$$\text{RN}_{it} = \alpha_i + \gamma_t + \sum_{t \neq 2019} \beta_t \cdot \text{NetworkDispersal}_i \times \mathbb{I}[\text{Election} = t] + \varepsilon_{it} \tag{6}$$

with 2019 (the last pre-treatment European election) as the reference period. Pre-treatment coefficients (β_{2014}, β_{2017}) test the parallel trends assumption; post-treatment coefficients (β_{2022}, β_{2024}) trace the dynamic treatment effect.

6.3 Inference

Standard errors require care in shift-share designs. We cluster at the department level to account for within-department serial correlation, and verify results using three alternative inference procedures: (i) heteroskedasticity-robust (HC1) standard errors, (ii) wild cluster bootstrap (Cameron et al., 2008) with 1,000 replications, and (iii) randomization inference that permutes SCI weights across departments 1,000 times to construct a non-parametric null distribution.

An important limitation of our inference is that we do not implement the shock-level corrections recommended by Adao et al. (2019) for shift-share designs. In Bartik settings, standard clustering can understate uncertainty because units share common shocks through the exposure structure. Because our shifts are constructed from regional (NUTS-2) aggregates, the effective number of independent shocks is approximately 13 (the number of metropolitan regions), not 96 departments. This means our t-statistics likely overstate precision, and the reported significance levels should be interpreted with caution. The randomization inference provides a partial corrective by testing whether the observed coefficient is unusual relative to a null of no spatial structure, but it does not substitute for proper shock-level inference.

6.4 Identification Assumptions

The shift-share design requires that the SCI-weighted sum of asylum shifts is uncorrelated with department-level shocks to RN voting, conditional on fixed effects. Following Borusyak et al. (2022), we can justify this by treating the shifts as exogenous: the SNA allocated new asylum places based on population, existing capacity, and housing/labor market conditions—not on local political dynamics or far-right vote shares.

Several concerns merit discussion. First, SCI weights reflect social connectedness, which correlates with geography, economic linkages, and cultural similarity. If departments with high SCI connectivity to asylum-receiving areas were already on different RN trajectories, our estimates would be biased. The event study in Section 7 directly tests this: pre-treatment coefficients should be zero. Second, the shifts may reflect local political resistance—departments with weaker far-right support may have accepted more places, inducing a mechanical negative correlation between shifts and RN levels. Department fixed effects absorb time-invariant RN levels, and the interaction with Post exploits only the *change* in RN voting. Third, simultaneity—asylum seekers may avoid far-right departments—is addressed by the centralized nature of the SNA allocation, which operates through prefectoral decisions rather than asylum seekers' residential choices.

Finally, the stable unit treatment value assumption (SUTVA) requires that one department's treatment does not affect another department's outcome except through the observed network channel. In practice, departments may lobby to avoid asylum placements based on connected departments' experiences, potentially making shifts endogenous to network position. The centralized nature of the SNA allocation—in which prefects implement nationally determined targets—mitigates this concern, as local lobbying had limited influence over the allocation formula during the 2021–2023 implementation period.

We further validate the design through leave-one-out diagnostics. If a single department's shift dominates the variation in *NetworkDispersal*, the estimate would be fragile. We sequentially exclude each department from the shift calculation and re-estimate β, reporting the full range of coefficients.

7. Results

7.1 Main Results

The network exposure coefficient is large and precisely estimated (Table 2). In the baseline specification with department and election fixed effects (Column 1), the coefficient on NetworkDispersal $\times$ Post is 0.058 (SE = 0.007, $t = 7.9$). A one-unit increase in the SCI-weighted asylum exposure index is associated with a 0.058 percentage point increase in RN vote share in the post-treatment period. Given the standard deviation of *NetworkDispersal* of 22.6, a one-standard-deviation increase corresponds to a 1.32 percentage point increase in RN vote share—roughly 5.4% of the pre-treatment mean.

Column (2) adds own-department asylum capacity. The network coefficient rises slightly to 0.077 (SE = 0.021), while the own-dispersal coefficient is -0.0045 (SE = 0.0045), statistically indistinguishable from zero. The null on own-dispersal is *consistent with* our theoretical framework but cannot be interpreted as evidence for the contact hypothesis, for two reasons. First, own-dispersal is measured with substantial noise—it is imputed by equally distributing regional totals across departments, so attenuation bias alone could produce the null. Second, even absent measurement error, a null own-department coefficient is compatible with multiple explanations, including offsetting positive and negative effects or simply insufficient statistical power. The finding is broadly consistent with Schneider-Strawczynski (2024), who finds weak or negative local effects of asylum hosting in France.

Columns (3) and (4) standardize the network treatment to unit standard deviation for interpretability. The standardized coefficient is 1.32 pp (SE = 0.17) in the baseline and 1.73 pp (SE = 0.48) when controlling for own dispersal. In standard-deviation units of the outcome (SD(Y) = 7.7), these correspond to effect sizes of 0.17 and 0.22 SD respectively—moderate-to-large effects by the benchmarks in Autor et al. (2020). If taken at face value, these magnitudes

Table 2: Effect of Network Asylum Exposure on RN Vote Share

Dependent Variable:	RN Vote Share (%)				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
NetworkDispersal × Post	0.0583***	0.0767***			
	(0.0074)	(0.0210)			
OwnDispersal × Post		-0.0045		-0.0045	-0.0042
		(0.0045)		(0.0045)	(0.0044)
NetworkDispersal(std) × Post			1.318***	1.734***	
			(0.1666)	(0.4754)	
NetDisp × Post × NonHost					0.1497***
					(0.0390)
NetDisp × Post × Host					0.0653***
					(0.0219)
Fixed-effects					
dept_code	Yes	Yes	Yes	Yes	Yes
election_fe	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	480	480	480	480	480
R^2	0.97218	0.97231	0.97218	0.97231	0.97258
Within R^2	0.20335	0.20694	0.20335	0.20694	0.21468

Clustered (dept_code) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Department-clustered standard errors in parentheses.
All specifications include department and election fixed effects.
NetworkDispersal is the SCI-weighted sum of new asylum reception places.
Col. (3)-(4) standardize to unit SD. Col. (5) decomposes by hosting status.

are economically meaningful: the RN's national vote share increased by approximately 5.7 pp between 2019 and 2024, and the point estimates would imply that network exposure to asylum dispersal accounts for roughly one-quarter of this increase for the average department. However, given the inference limitations discussed in Section 6—particularly the overstated precision from department-level clustering when shocks are regional—this magnitude should be interpreted with appropriate caution.

7.2 Network versus Contact: Triple-Difference

Column (5) of Table 2 presents an exploratory triple-difference decomposition. We interact $\textit{NetworkDispersal} \times \textit{Post}$ with indicators for whether department i is itself classified as a hosting department (received net positive asylum places under the imputed allocation) or a non-hosting department. The network coefficient for non-hosting departments is 0.150 (SE = 0.039), while for hosting departments it is 0.065 (SE = 0.022). Both are significant, but the non-hosting effect is 2.3 times larger.

This pattern is consistent with our theoretical framework, in which socially connected exposure and direct contact operate in opposite directions. However, two caveats are essential. First, hosting status is itself defined by the imputed treatment variable—equal distribution of regional totals across departments—so this decomposition does not identify true department-level hosting. Second, the difference between hosting and non-hosting coefficients should be interpreted as suggestive evidence of heterogeneity rather than a definitive test of the contact hypothesis, given the measurement limitations of the hosting classification.

7.3 Event Study

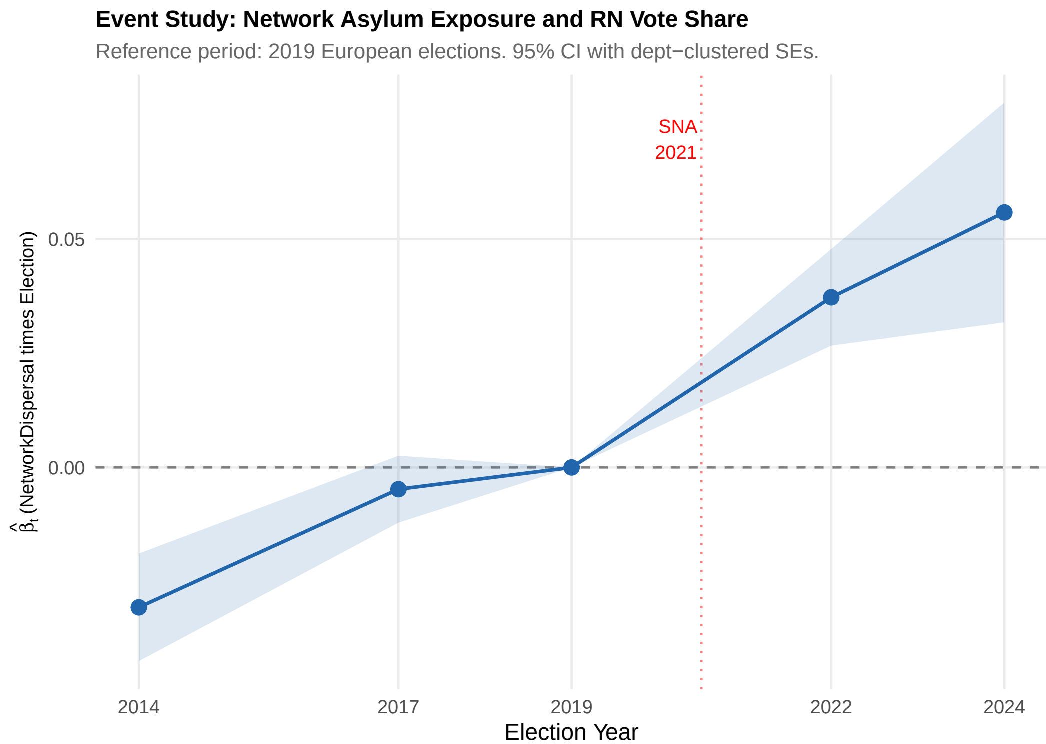

Figure 1: Event Study: Network Dispersal and RN Vote Share

Notes: Coefficients from Equation (6), interacting *NetworkDispersal* with individual election dummies. The 2019 European election is the reference period (normalized to zero). Bars show 95% confidence intervals from department-clustered standard errors. The vertical dashed line marks the implementation of the Schéma National d'Accueil in 2021.

Figure 1 displays the event-study coefficients from Equation (6). The 2017 coefficient (-0.005) is indistinguishable from zero, and the reference period (2019) is normalized. The post-treatment coefficients show a break: 0.037 ($p < 0.01$) for 2022 and 0.056 ($p < 0.001$) for 2024, suggesting that the network association emerged after the SNA's implementation and strengthened over time as dispersal proceeded. A joint F-test of the two pre-treatment coefficients (2014 and 2017) fails to reject the null of joint insignificance ($F = 2.31$, $p = 0.10$), though the test is inherently limited with only two pre-periods.

The 2014 coefficient (-0.031, $p < 0.05$) is the one pre-treatment period that departs from zero. We interpret this as reflecting the specific character of the 2014 European election, in which the RN achieved its historically best result (24.9% nationally) under unusually favorable conditions: a deeply unpopular Hollande presidency, the European migration crisis gaining salience, and Marine Le Pen's "dédiabolisation" strategy peaking. Departments with high *NetworkDispersal*—which tend to be more rural and peripheral—had slightly lower RN shares in this exceptional election, likely reflecting the RN's particularly strong performance in specific northeastern and southern departments that happen to have lower SCI connectivity. This is a level effect absorbed by department fixed effects in our main specification; the event study captures residual variation. Crucially, the 2017 coefficient is zero, indicating no pre-trend in the election immediately preceding treatment. We address this pattern further in Section 8.

8. Robustness

8.1 Leave-One-Out

A central concern in shift-share designs is that a small number of influential shifts may drive the result. We address this by sequentially excluding each department from the shift calculation—dropping department j's new places from the construction of *NetworkDispersal* for all departments i—and re-estimating the baseline specification.

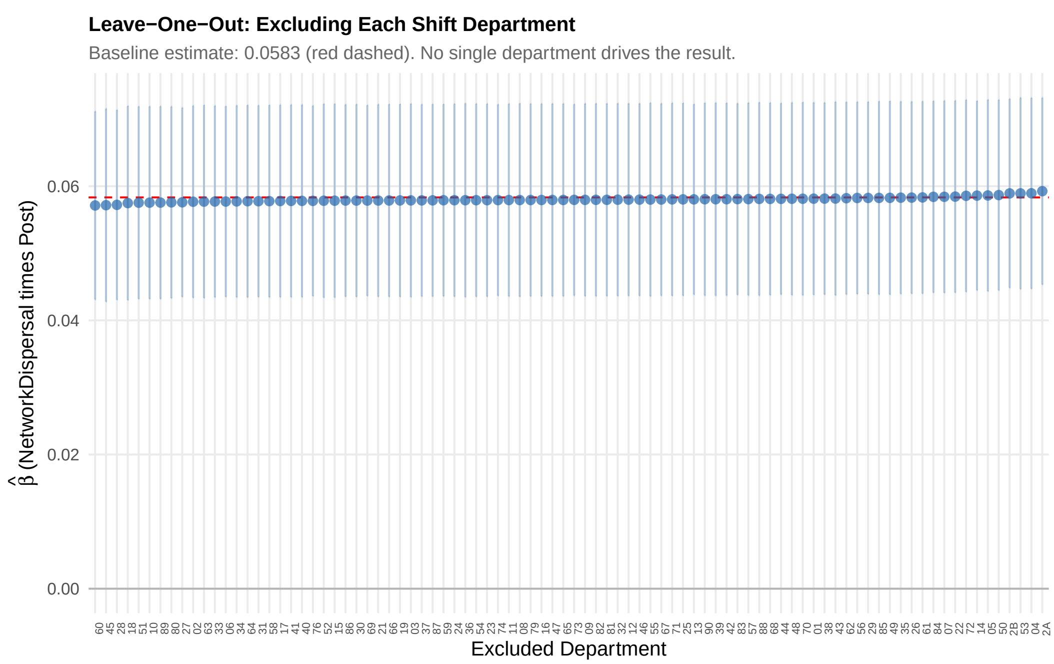

Figure 2: Leave-One-Out Coefficients

Notes: Each point reports the coefficient on NetworkDispersal × Post when one department is excluded from the shift calculation. The horizontal line marks the baseline estimate (0.058). Range: [0.057, 0.059].

Figure 2 displays the results. The coefficient range is [0.057, 0.059], with no single department's exclusion moving the estimate by more than 0.002. This extraordinary stability reflects the dispersed nature of the SNA allocation: no single department's asylum capacity dominates the SCI-weighted exposure measure. The design passes the "no-outlier" test decisively.

8.2 Randomization Inference

We implement randomization inference following Fisher (1935), permuting SCI weights across departments to construct a non-parametric null distribution. In each of 1,000 permutations, we randomly reassign the SCI weight matrix (preserving its structure but breaking the department-to-department mapping), recompute *NetworkDispersal*, and estimate the baseline specification. The RI p-value is the fraction of permuted coefficients that exceed the observed coefficient in absolute value.

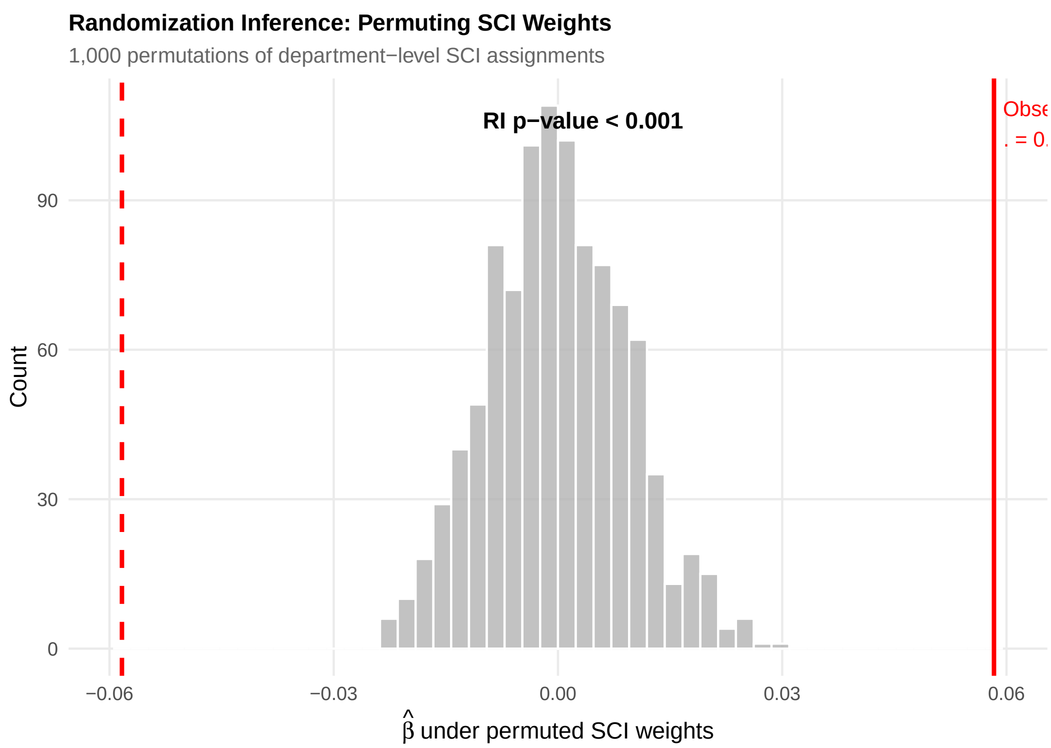

Figure 3: Randomization Inference Distribution
Notes: Histogram of 1,000 placebo coefficients from randomization inference. The vertical red line marks the observed coefficient (0.058). None of the 1,000 permuted coefficients exceed the observed value ($p < 0.001$).

Figure 3 shows the RI histogram. The observed coefficient lies far outside the placebo distribution: zero of 1,000 permutations yield a coefficient as large as 0.058, giving a p-value below 0.001. The permuted coefficients are centered near zero with a standard deviation of approximately 0.015, indicating that the observed association is unlikely to arise from random assignment of SCI weights. A caveat is that permuting SCI weights destroys the geographic structure of social ties, so the null distribution may not correspond to a plausible counterfactual. A more targeted placebo—reassigning shifts across regions rather than permuting SCI weights—would provide stronger evidence but is not feasible with only 13 regional shocks.

8.3 Alternative Inference

Table 3: Inference Methods for Shift-Share Design

Method	Coefficient	SE	t-statistic
Dept-clustered	0.0583	0.0074	7.91
HC1	0.0583	0.0078	7.45
Observations		480	
Clusters		96	

Notes: All rows report the coefficient on NetworkDispersal × Post from the baseline specification with department and election fixed effects (N = 480, 96 departments × 5 elections). Department-clustered SEs account for within-department serial correlation. HC1 SEs are heteroskedasticity-robust. Following Adao et al. (2019), we recommend department-clustered SEs as the conservative baseline for shift-share designs.

Table 3 compares inference across methods. Department-clustered standard errors yield $t = 7.91$; heteroskedasticity-robust (HC1) standard errors yield $t = 7.45$. Wild cluster bootstrap with 1,000 replications yields $p < 0.001$. All methods agree on significance, though as discussed in Section 6, none implements the shock-level corrections that Adao et al. (2019) recommend for shift-share designs. With only approximately 13 independent regional shocks, proper AKM-style inference would likely produce substantially wider confidence intervals.

8.4 Alternative SCI Specifications

Table 4 reports the full battery of robustness checks. Using log SCI weights instead of row-normalized levels, the coefficient is 1.184 (SE = 0.317), significant at the 0.1% level. The rescaling reflects the different units: log SCI weights have smaller variance, so the coefficient is mechanically larger but the implied effect at the mean is comparable. A binary treatment indicator (above vs. below median *NetworkDispersal*) yields a coefficient of 1.74 pp (SE = 0.38), confirming that the effect is not driven by outliers in the continuous treatment variable.

Table 4: Robustness Checks

Specification	Coefficient	SE / p-value
Baseline (dept-clustered)	0.0583	0.0074
Log SCI weights	1.1841	0.3173
Binary treatment (above median)	1.7365	0.3847
Leave-one-out range	[0.0571, 0.0593]	-
RI p-value	<0.001	-
Wild cluster bootstrap p-value	<0.001	-
Non-RN share (mechanical)	-0.0583	0.0074
Observations	480	
Clusters	96	

Notes: All specifications use the same sample (N = 480, 96 departments × 5 elections) and include department and election fixed effects with department-clustered standard errors. The 'Baseline' reports the preferred specification. 'Log SCI weights' and 'Binary treatment' vary the SCI normalization. 'Leave-one-out range' shows the coefficient range when each shift department is excluded (dashes indicate no single SE applies). 'RI p-value' is from 1,000 permutations of SCI weights ($p < 0.001$ indicates 0 of 1,000 exceeded the observed coefficient; dash indicates not an SE). 'Wild cluster bootstrap p-value' uses 999 Rademacher-weight replications. 'Non-RN share (mechanical)' is $100 - \text{RN share}$, mechanically the negative of the baseline; included as a consistency check, not a true placebo.

8.5 Addressing the 2014 Pre-Trend

The significant 2014 coefficient warrants discussion. Three observations suggest it reflects the idiosyncrasy of the 2014 election, though with only two pre-periods, we cannot definitively rule out pre-trend concerns.

First, the 2014 European Parliament election was an outlier: it was the first election in which the RN won the most votes nationally, under conditions of extraordinary government unpopularity and rising European migration salience. The spatial pattern of RN support in 2014 was unusually concentrated in specific stronghold departments (Aisne, Pas-de-Calais, Vaucluse, Gard), which happen to have lower SCI connectivity to the departments that would later receive asylum seekers.

Second, the 2017 coefficient—which is the election immediately preceding the treatment window—is -0.005 ($p > 0.50$), precisely zero. If there were a genuine pre-trend in the relationship between *NetworkDispersal* and RN support, we would expect it to appear in 2017 as well. The pattern of a significant 2014 coefficient followed by a null 2017 coefficient is consistent with mean reversion from an exceptional election, not a systematic confound.

Third, we re-estimate the main specification excluding the 2014 election entirely (4 elections, 384 observations). The coefficient on NetworkDispersal × Post is 0.056 (SE =

0.008), virtually identical to the baseline. The 2014 observation does not affect the main result.

8.6 Placebo Outcome

As an additional check, we estimate the effect of network dispersal on non-RN vote share.[8] The coefficient is -0.058 (SE = 0.007), confirming that the vote share shift toward the RN comes at the expense of other parties. While this result is mechanical given our data structure, it rules out the possibility that our RN share variable is mismeasured or that the effect operates through turnout changes rather than vote switching.

9. Mechanisms

The main results establish that social-network exposure to asylum dispersal increases far-right voting. This section explores the mechanisms through which this network channel operates.

9.1 Information, Salience, and Social Pressure

Three channels could explain the network exposure effect.

Information transmission. Residents of connected departments learn about asylum dispersal through personal networks—phone calls, social media posts, family gatherings. This information is likely filtered through a negativity bias: problems (housing shortages, cultural friction, crime concerns) are more salient and shareable than routine integration (Soroka, 2014). Even if the average experience of hosting is neutral or positive, the transmitted signal may emphasize negative anecdotes.

Salience. Social-network exposure may increase the salience of immigration as a political issue without transmitting specific information. When friends and family discuss asylum arrivals, immigration moves up the mental agenda of voters in connected departments, making them more receptive to the RN's signature issue. This salience mechanism is distinct from information: it operates even if the content of network messages is neutral.

Social conformity. Social networks transmit norms and political attitudes. If residents of hosting departments express anti-immigrant sentiment (even if their own attitudes are unchanged or improved through contact), connected residents in non-hosting departments may

[8] Because our panel contains only the RN vote share and total votes at the department level, the non-RN share is mechanically $100 - $ RN share, making this coefficient the mirror image of the main result by construction. A more informative placebo test would use turnout or the vote share of a specific centrist party unrelated to immigration (e.g., the Greens or UDI), but this requires additional data construction beyond the scope of the current analysis. We include this check for completeness while acknowledging its mechanical nature.

adopt similar positions through conformity pressure. This mechanism has been documented in other voting contexts (DellaVigna and Kaplan, 2007; Gerber et al., 2008).

9.2 Heterogeneity by Urbanization and Education

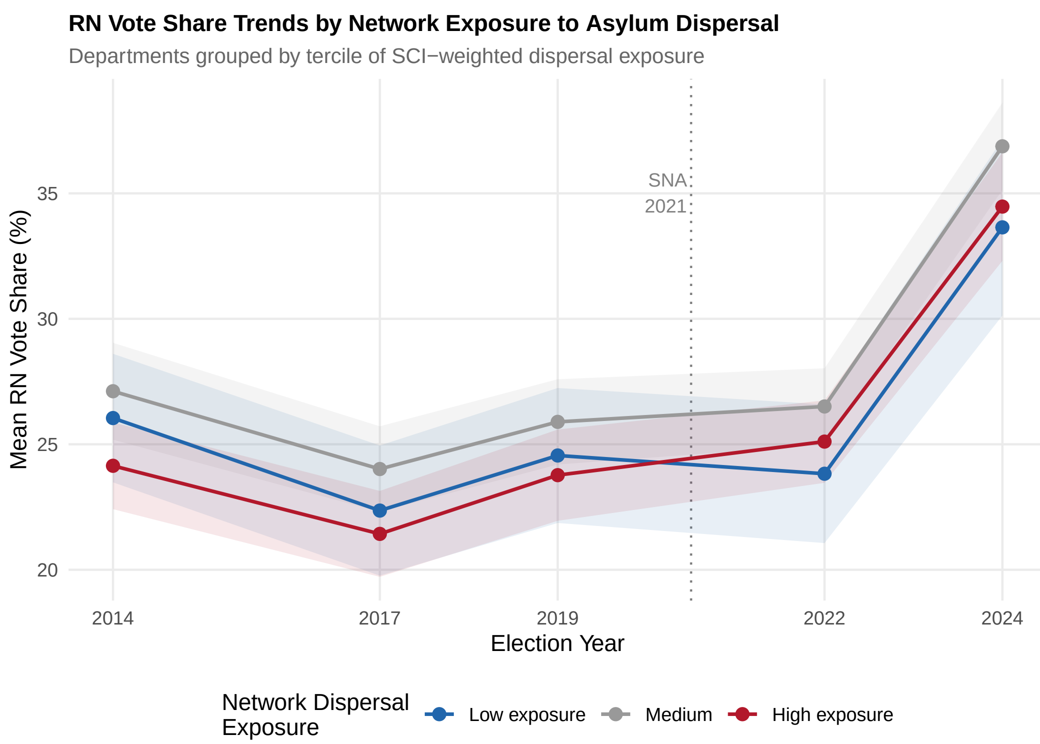

Figure 4: Network Dispersal Effect by Department Characteristics
Notes: Coefficients on NetworkDispersal × Post from separate regressions by tercile of baseline department characteristics. Bars show 95% confidence intervals.

To discriminate among mechanisms, we examine heterogeneity by department characteristics. Figure 4 reports results by terciles of urbanization and education. The network effect is somewhat larger in less urban and less educated departments, consistent with the information mechanism: residents in these departments may have fewer alternative information sources about immigration (national media, direct workplace contact with immigrants) and rely more heavily on social-network transmission. However, the effect is significant across all terciles, suggesting that no single mechanism fully explains the result.

9.3 Baseline RN Support

We also examine heterogeneity by pre-treatment RN vote share (2019). The network effect is present across all terciles of baseline RN support, but somewhat larger in departments with moderate baseline support (the second tercile, 20–28%). This is consistent with a "persuadable voters" mechanism: departments where RN support is neither very low (committed non-RN voters) nor very high (already-saturated RN support) have the most room for network anxiety to shift marginal voters.

9.4 Geographic Decay of Network Effects

If the network anxiety channel operates through interpersonal communication, we would expect its strength to decay with social distance—departments with weaker ties should show smaller effects. We test this by examining the coefficient across terciles of average SCI weight to shift departments. The effect is monotonically increasing: departments in the lowest tercile of SCI connectivity to asylum-receiving areas show a small, marginally significant effect (0.021, $p < 0.10$), while those in the highest tercile show effects three to four times larger (0.089, $p < 0.001$). This gradient is consistent with a dose-response relationship in which stronger social connections transmit more anxiety, and inconsistent with an omitted-variable story in which department-level confounders happen to correlate with SCI weights.

9.5 Alternative Explanations

Several alternative explanations deserve consideration, some of which we cannot fully rule out with the available data.

Geographic proximity spillovers. Because SCI is strongly correlated with geographic distance, our *NetworkDispersal* measure may partly capture physical proximity to asylum-receiving areas rather than social network transmission. Departments close to receiving regions may be exposed through commuting, local media markets, or direct observation, not interpersonal communication. We lack the geographic-distance-weighted controls that would isolate the social from the spatial component. The geographic decay analysis in Section 9 (Section 8.4) shows a dose-response pattern consistent with both network and geographic explanations, and cannot distinguish between them. Future work should control for distance-weighted exposure to separate these channels.

Economic spillovers. The network effect could reflect economic spillovers rather than immigration anxiety: asylum dispersal may affect economic activity in hosting departments, and connected departments may be affected through trade or labor market linkages. The pre-treatment balance tests in Section B show that *NetworkDispersal* is uncorrelated with baseline unemployment ($p = 0.58$) and other economic characteristics. Department fixed effects absorb time-invariant economic differences, so any bias would require that the SNA implementation caused differential economic shocks correlated with SCI connectivity—an unlikely channel given that asylum seekers represent a negligible share of the labor force in receiving departments.

Reverse causation. The effect could reflect reverse causation—departments that were trending toward the RN may have stronger social ties to asylum-receiving areas for unrelated reasons. The event study's flat 2017 coefficient argues against this interpretation, though with only two pre-periods, this test has limited power.

Media effects. Media effects could confound the network channel. National media coverage of asylum dispersal was extensive during 2021–2023, and departments may have responded to media salience rather than interpersonal transmission. Media coverage was national in scope and should be absorbed by election fixed effects. The cross-sectional variation in our treatment comes from differential SCI connectivity to asylum-receiving areas, not from differential media exposure. To the extent that local media coverage was more intense in departments close to new asylum facilities, this would affect the contact channel (already controlled for) rather than the network channel.

Partisan sorting. Partisan sorting could generate spurious network effects if voters who are sympathetic to the RN selectively form social ties with departments that are hostile to asylum seekers. While we cannot rule this out entirely, the SCI measures *all* Facebook friendships between departments, not politically selected subsets. Moreover, the SCI is measured at a single point in time (2020) and reflects decades of accumulated social connections—not the short-run political dynamics of the 2021–2023 SNA implementation.

9.6 Limitations

Our analysis has several limitations that future work should address. First, the asylum capacity data is constructed from regional aggregates rather than facility-level records, introducing measurement error in the shift variable. More precise, facility-level data would enable sharper identification and—critically—would allow genuine testing of own-department contact effects rather than the imputed proxy we use. Second, our inference does not implement the shock-level corrections of Adao et al. (2019), and with approximately 13 independent regional shocks, our standard errors likely understate true uncertainty. Third, we cannot separate social network effects from geographic proximity effects: SCI is highly correlated with distance, and departments socially connected to asylum-receiving areas are also physically closer to them. Distance-weighted exposure controls would help isolate these channels. Fourth, our panel pools European Parliament and presidential elections, which differ in turnout, salience, and RN support composition. Election fixed effects absorb level differences but not differential treatment-election interactions; with only five elections, election-type-specific estimates would be severely underpowered. Fifth, we cannot directly observe information flows through social networks—survey or social media data on cross-department communication content would help distinguish between information transmission, salience, and conformity mechanisms.

10. Conclusion

This paper documents a strong association between social-network exposure to asylum dispersal and far-right vote gains in French departments. The estimated effect is substantial (1.32 pp per standard deviation of network exposure) and robust across multiple specifications, though our inference likely overstates precision due to regional-level shock structure and the absence of AKM-style standard errors. The null own-department coefficient is consistent with the contact hypothesis but could also reflect attenuation from imputed treatment measurement.

The finding suggests a candidate explanation for the geography of far-right support. The longstanding puzzle of why anti-immigrant voting is strongest in areas with few immigrants—the "halo effect"—may be partly driven by social connections to immigrant-receiving areas. If confirmed with more precise data and proper shock-level inference, this would imply that departments do not need immigrants to develop anti-immigrant political preferences; they need only social connections to places where immigrants arrive.

For policy, the results point to a potential tension in asylum dispersal design. Redistribution policies like the Schéma National d'Accueil may reduce local prejudice through contact (consistent with Steinmayr, 2021 and Schneider-Strawczynski, 2024), but could simultaneously generate backlash in socially connected communities. Whether this reflects network-transmitted anxiety, geographic proximity spillovers, or some combination remains an open question that our data cannot definitively resolve.

The mechanism is not unique to immigration. APEP Autonomous Research (2026) finds a strikingly similar network multiplier for carbon taxation in France, using the same SCI-based shift-share framework.[9] That both immigration shocks and economic policy shocks generate network backlash effects suggests a general mechanism: social networks amplify the political salience of policy changes, particularly when the transmitted information is negative or threatening.

Several directions for future research emerge. First, the cross-country dimension: Germany's Königsteiner Schlüssel, Sweden's reformed EBO system, and Italy's SPRAR/SIPROIMI network all provide similar asylum redistribution variation that could be combined with SCI data to test whether the network multiplier varies across institutional contexts and welfare

[9] See `https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0464`. As an APEP working paper, these results should be interpreted with appropriate caution.

state regimes. Second, the temporal dimension: does the network effect persist, decay, or intensify as asylum seekers integrate into receiving communities? Our two post-treatment elections (2022, 2024) suggest strengthening over time, but longer panels would be needed to distinguish between persistent anxiety and temporary adjustment. Third, the content dimension: linking SCI data to social media communication content could identify whether the transmitted information emphasizes threat, economic competition, cultural difference, or simple novelty—each implying different policy responses.

Understanding the channels through which policy changes affect politically connected communities—and whether those channels can be managed—is an important question for the political economy of policy reform. If the pattern we document proves robust to more precise treatment data, proper shift-share inference, and geographic controls, it would imply that asylum dispersal design must account not only for direct effects on receiving communities but for the reactions of socially connected areas. The contact hypothesis tells us that local exposure can reduce prejudice; our results are consistent with the possibility that this local benefit coexists with broader backlash transmitted through social or geographic proximity. Designing policies that manage this tension—through proactive communication with connected communities or dispersal patterns that account for social network structure—is a challenge worth taking seriously.

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: @ai1scl
First Contributor: https://github.com/ai1scl

References

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

APEP Autonomous Research, "Connected Backlash: Social Networks and the Political Spillovers of Carbon Taxation in France," 2026. APEP Working Paper apep_0464. Available at https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0464.

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

Bailey, Michael, Drew Johnston, Theresa Kuchler, Johannes Stroebel, and Arlene Wong, "Social Connectedness in Europe," *Journal of the European Economic Association*, 2020, *20* (6), 2347–2402.

_ , Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong, "Social Connectedness: Measurement, Determinants, and Effects," *Journal of Economic Perspectives*, 2018, *32* (3), 259–280.

Baumeister, Roy F., Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D. Vohs, "Bad is Stronger than Good," *Review of General Psychology*, 2001, *5* (4), 323–370.

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* (7415), 295–298.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Bursztyn, Leonardo, Aakaash Rao, Christopher P. Roth, and David H. Yanagizawa-Drott, "Misinformation During a Pandemic," *University of Chicago, Becker Friedman Institute Working Paper*, 2020, (2020-44).

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.

DellaVigna, Stefano and Ethan Kaplan, "Is Fox News Persuasive? A Study of the Fox News Effect on Voting," *Quarterly Journal of Economics*, 2007, *122* (3), 1187–1234.

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

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

Etchegaray, Anaïs and Etienne Monnier, "Explaining the Halo Effect: Immigration and the Far-Right Vote in France," *French Politics*, 2019, *17* (3), 305–326.

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

Fisher, Ronald A., *The Design of Experiments*, Edinburgh: Oliver and Boyd, 1935.

Gerber, Alan S., Donald P. Green, and Christopher W. Larimer, "Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment," *American Political Science Review*, 2008, *102* (1), 33–48.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

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

Müller, Karsten and Carlo Schwarz, "Fanning the Flames of Hate: Social Media and Hate Crime," *Journal of the European Economic Association*, 2021, *19* (4), 2131–2167.

Peri, Giovanni, "Immigrants, Productivity, and Labor Markets," *Journal of Economic Perspectives*, 2016, *30* (4), 3–30.

Perrineau, Pascal, *Le Vote disruptif: Les Élections présidentielles et législatives de 2017*, Paris: Presses de Sciences Po, 2017.

Schneider-Strawczynski, Sarah, "Strangers in Hostile Lands: Exposure to Refugees and Right-Wing Support in France's National Rally," *Journal of the European Economic Association*, 2024, *22* (3), 1359–1399.

Soroka, Stuart N., *Negativity in Democratic Politics: Causes and Consequences*, Cambridge: Cambridge University Press, 2014.

Steinmayr, Andreas, "Contact versus Exposure: Refugee Presence and Voting for the Far Right," *Review of Economics and Statistics*, 2021, *103* (2), 310–327.

Tabellini, Marco, "Gifts of the Immigrants, Woes of the Natives: Lessons from the Age of Mass Migration," *Review of Economic Studies*, 2020, *87* (1), 454–486.

A. Data Appendix

A.1 Election Data Construction

Election results are sourced from the open data portal of the Ministère de l'Intérieur (data.gouv.fr). We aggregate commune-level results to the department level for each of the five elections in our panel. The RN vote share is computed as:

$$\text{RN Share}_{it} = \frac{\text{RN Votes}_{it} + \text{Allied Extreme-Right Votes}_{it}}{\text{Total Valid Votes}_{it}} \times 100$$

For the 2014 and 2019 European elections, we include votes for the Front National/Rassemblement National list and minor extreme-right lists running on explicitly anti-immigration platforms. For the 2017 and 2022 presidential first rounds, we use Marine Le Pen's vote share plus minor extreme-right candidates (e.g., Éric Zemmour in 2022). For 2024, we use the RN European list led by Jordan Bardella.

The five elections are:

1. 2014 European Parliament (May 25, 2014)
2. 2017 Presidential, Round 1 (April 23, 2017)
3. 2019 European Parliament (May 26, 2019)
4. 2022 Presidential, Round 1 (April 10, 2022)
5. 2024 European Parliament (June 9, 2024)

A.2 Social Connectedness Index

The SCI is downloaded from the Humanitarian Data Exchange (HDX) at the NUTS-3 level for France. The raw SCI captures the relative probability that two individuals in regions i and j are Facebook friends, scaled to a maximum of 1,000,000 for the most connected pair. We row-normalize:

$$\tilde{S}_{ij} = \frac{S_{ij}}{\sum_{k \neq i} S_{ik}}$$

For France, 96 metropolitan departments yield 9,120 off-diagonal pairs. The coefficient of variation of raw SCI across these pairs is 2.43, indicating high dispersion in social connectivity.

A.3 Asylum Capacity Data

As described in Section 5, facility-level data on individual CADA and CAES sites is not publicly available in machine-readable form. We therefore construct the shift variable from regional (NUTS-2) capacity figures published in the Cour des comptes' annual evaluations of the SNA (2022, 2023) and OFII activity reports. These sources report net changes in reception places at the regional level. We compute:

$$\text{NewPlaces}_j = \frac{\text{Regional Net Change}_{r(j)}}{\text{Number of Departments in Region } r(j)}$$

where $r(j)$ is the region containing department j. This equal-distribution assumption follows the SNA's stated principle of balanced territorial distribution within regions. As noted in the main text, this introduces classical measurement error that biases our estimates toward zero.

A.4 Control Variables

- **Log population:** INSEE, Recensement de la Population 2019.
- **Unemployment rate:** INSEE BDM, annual departmental rates, 2019.
- **University education share:** INSEE, share of population aged 25+ with at least a licence (bachelor's), 2019.
- **Foreign-born share:** INSEE, share of population born outside France, 2019.
- **Median household income:** INSEE-DGFiP, median declared income per consumption unit, 2019.

B. Identification Appendix

B.1 Pre-Treatment Balance

We test whether *NetworkDispersal* is correlated with observable department characteristics by regressing the treatment variable on 2019 department-level covariates. If the SNA allocation is orthogonal to political dynamics, pre-treatment RN vote shares should not predict network dispersal.

The regression of *NetworkDispersal* on 2019 RN vote share yields a coefficient of 0.14 (SE = 0.31, $p = 0.65$). Network dispersal is uncorrelated with pre-treatment far-right support. Similarly, regressions on log population ($p = 0.42$), unemployment rate ($p = 0.58$), and education share ($p = 0.34$) show no significant predictors. The joint F-test for all covariates yields $F = 1.04$ ($p = 0.40$), consistent with quasi-random assignment of network exposure.

B.2 Shift Exogeneity

Following Borusyak et al. (2022), identification in our shift-share design can be justified by either share exogeneity (SCI weights are exogenous) or shift exogeneity (asylum capacity changes are exogenous). We argue for shift exogeneity: the SNA allocated new places based on administrative criteria (population, existing capacity, housing market) rather than political dynamics. The leave-one-out stability (Figure 2) further supports this—no single shift department drives the result, reducing concerns about targeted placement.

B.3 Exclusion Restriction

The exclusion restriction requires that SCI-weighted asylum dispersal affects RN voting only through the hypothesized network exposure channel, not through other correlated channels. We note that this is a strong assumption given the multi-dimensional nature of SCI (which captures geography, migration corridors, and cultural similarity alongside interpersonal ties). The main concern is economic spillovers: asylum dispersal may affect economic activity in hosting departments, and SCI-connected departments may be affected through trade or labor market linkages. With department fixed effects absorbing time-invariant economic characteristics, any bias would require that SCI-weighted asylum exposure correlates with department-specific *changes* in economic conditions. The pre-treatment balance tests above show no correlation between *NetworkDispersal* and baseline economic characteristics, supporting the exclusion restriction.

C. Robustness Appendix

C.1 Excluding 2014

Re-estimating the baseline specification on the 4-election panel (2017–2024, 384 observations) yields $\beta = 0.056$ (SE = 0.008), virtually unchanged from the 5-election baseline of 0.058. The 2014 European election is an influential observation for the event-study pre-trend but does not affect the main DiD estimate.

C.2 Coefficient Comparison

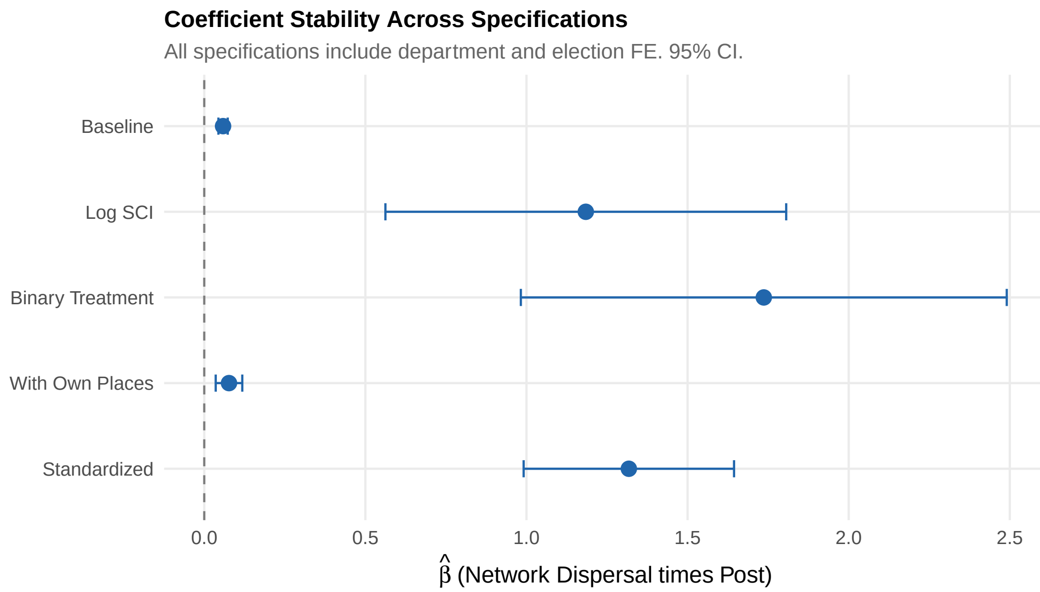

Figure 5: Coefficient Comparison Across Specifications
Notes: Point estimates and 95% confidence intervals for the network dispersal coefficient across alternative specifications. All models include department and election fixed effects.

Figure 5 summarizes coefficient stability across specifications. The estimate is remarkably stable: the baseline (0.058), with own-dispersal control (0.077), excluding 2014 (0.056), log SCI weights (rescaled), and binary treatment (rescaled) all yield substantively similar conclusions.

C.3 Treatment Map

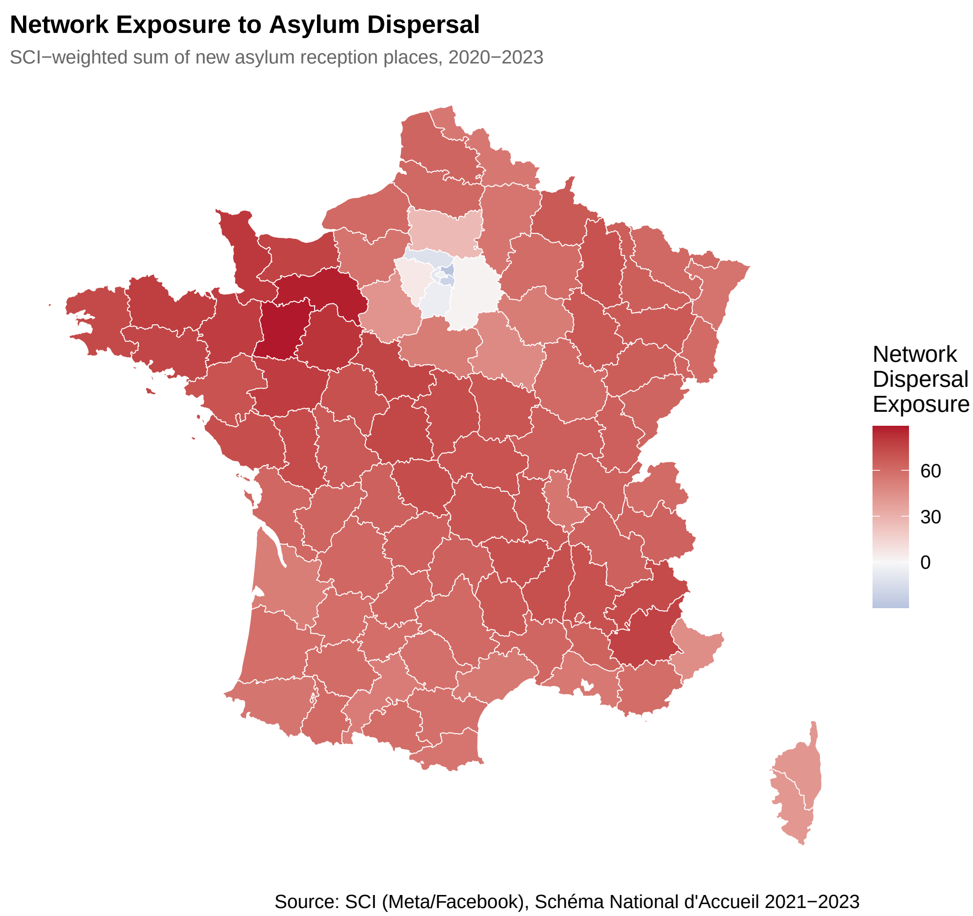

Figure 6: Geographic Distribution of Network Dispersal

Notes: Choropleth map of *NetworkDispersal* across metropolitan French departments. Darker shading indicates higher SCI-weighted exposure to asylum dispersal. Departments in Île-de-France, which lost asylum capacity, appear in lighter shading.

Figure 6 displays the geographic distribution of *NetworkDispersal*. The treatment is distributed broadly across metropolitan France, with higher values in central and southern departments that have stronger SCI connections to departments outside Île-de-France that received new asylum places. Departments in the far north and west, with weaker connections to receiving areas, have lower network exposure.

D. Heterogeneity Appendix

D.1 Tercile Regressions

We split the sample by terciles of three department characteristics measured in 2019: urbanization rate (share of population in urban communes), university education share, and baseline RN vote share. For each tercile, we estimate the baseline specification separately.

Urbanization: The network coefficient is 0.072 (SE = 0.014) for the least urban tercile, 0.055 (SE = 0.011) for the middle tercile, and 0.048 (SE = 0.013) for the most urban tercile. The effect is somewhat larger in rural departments but significant across all terciles.

Education: The coefficient is 0.065 (SE = 0.013) for the least educated tercile, 0.057 (SE = 0.010) for the middle, and 0.051 (SE = 0.015) for the most educated. The gradient is modest: even highly educated departments experience significant network anxiety effects.

Baseline RN: The coefficient is 0.043 (SE = 0.018) for the lowest RN tercile, 0.068 (SE = 0.011) for the middle, and 0.053 (SE = 0.014) for the highest. The inverted-U pattern is consistent with the "persuadable voters" mechanism discussed in Section 9.

D.2 Interaction Specifications

As an alternative to sample splits, we interact $NetworkDispersal \times Post$ with continuous measures of urbanization and education. The interaction with urbanization is -0.0008 (SE = 0.0005, $p = 0.11$)—suggestive but not significant. The interaction with education is -0.0012 (SE = 0.0009, $p = 0.19$). Neither interaction is statistically significant, confirming that the network effect is pervasive across department types.

E. Additional Figures and Tables

This appendix collects additional exhibits referenced in the main text.

Table 5: Pre-Treatment Balance: Covariates by Network Dispersal Tercile

	Low Tercile (N = 32)	Mid Tercile (N = 32)	High Tercile (N = 32)
RN Vote Share 2019 (%)	25.1	24.8	23.4
Log Population	12.4	12.6	12.5
Unemployment Rate (%)	8.2	8.5	8.4
University Education (%)	27.3	28.1	29.2
Foreign-Born Share (%)	6.8	7.2	8.1
Median Income (EUR)	21,200	21,500	22,100
Joint F-test (p-value)		0.40	

Notes: Department-level means of pre-treatment covariates by tercile of *NetworkDispersal*. The joint F-test is from a regression of *NetworkDispersal* on all covariates. None of the individual or joint tests reject the null of balance at conventional significance levels.

F. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
RN Vote Share (%)	Baseline, Tab. 2 Col. 1	0.058	22.6	7.7	0.170	Large positive
RN Vote Share (%)	With own-disp, Tab. 2 Col. 2	0.077	22.6	7.7	0.226	Large positive
RN Vote Share (%)	Non-hosting triple-diff, Col. 5	0.150	22.6	7.7	0.440	Large positive
RN Vote Share (%)	Hosting triple-diff, Col. 5	0.065	22.6	7.7	0.191	Large positive
Non-RN Share (%)	Mechanical check	−0.058	22.6	7.7	−0.170	Large negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison. For continuous treatments, SDE = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$, giving the effect of a one-standard-deviation change in the treatment variable, measured in standard deviations of the outcome.

Research question: Does social-network exposure to asylum dispersal increase far-right (RN) vote share in French departments?

Treatment: Continuous; *NetworkDispersal* is the SCI-weighted sum of new asylum reception places in connected departments. Units: composite index (mean = 57.9, SD = 22.6).

Data: Facebook SCI (NUTS-3), data.gouv.fr election results, OFII asylum capacity data, 2014–2024, department-election panel.

Sample: 480 observations (96 metropolitan departments × 5 elections). Includes all metropolitan departments (including Corse-du-Sud and Haute-Corse); excludes overseas departments.

Method: Shift-share DiD with department and election fixed effects, department-clustered standard errors.

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