Across the Channel: Social Networks and the Cross-Border Housing Effects of Brexit*

APEP Autonomous Research

Contributor: @SocialCatalystLab

February 2026

Abstract

How do political shocks propagate across borders through social networks? We study the effect of Brexit on French housing markets using a continuous-treatment difference-in-differences design that exploits variation in social connectedness between French départements and the United Kingdom. We address identification challenges through pre-determined census stocks, residualized exposure, and a triple-difference comparing houses to apartments. Cluster bootstrap inference confirms that cluster-robust standard errors are reliable despite only 96 clusters. A four-country placebo battery (Belgium, Netherlands, Italy, Spain) measured at the GADM1 level, combined with the existing German placebo, provides multi-country diagnostics; a horse-race triple-difference specification with all countries simultaneously renders every placebo insignificant while the UK coefficient remains positive. Rambachan and Roth (2023) sensitivity analysis quantifies robustness to pre-trend non-linearities. Commune-level estimation with approximately 50 times more observations yields more precise triple-difference estimates.

JEL: F22, R31, R23, F15

Keywords: Social networks, Brexit, housing prices, cross-border spillovers, triple-difference, cluster bootstrap

*This paper is a revision of APEP Working Paper 0460 v3. We thank the editors and three anonymous referees for detailed comments that substantially improved inference and identification. Housing transaction data are from the French *Demandes de Valeurs Foncières* (DVF), available at `data.gouv.fr`. Social connectedness data are from Meta (Facebook SCI). Census data are from INSEE *Recensement de la Population* 2016 (reference date January 1, 2016). Replication materials: https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0460.

1 Introduction

In the rolling countryside of the Dordogne, British retirees have for decades purchased stone farmhouses, transforming quiet communes into anglophone enclaves. The June 2016 Brexit referendum—and the subsequent 10% depreciation of sterling within a single quarter—threatened this decades-old pattern. For French property markets closely linked to British demand, Brexit represented an exogenous shock transmitted through pre-existing social and economic networks. This paper asks: did départements with stronger pre-existing ties to the UK experience differential housing price dynamics after the Brexit vote?

The answer is not straightforward. A naïve continuous-treatment difference-in-differences (DiD) design—interacting a département's Social Connectedness Index (SCI; Bailey et al., 2018) with a post-referendum indicator—yields a positive and significant coefficient. But three identification failures undermine this result. First, the SCI is measured in 2021, five years *after* Brexit, raising post-treatment contamination concerns. Second, the coefficient collapses when controlling for baseline housing price levels, suggesting the treatment variable proxies for metropolitan status rather than UK-specific connectivity. Third, and most damning, an identical specification using German SCI yields a *larger* coefficient than the UK specification, despite Germany experiencing no analogous political shock.

These failures are instructive. They reflect a general challenge in shift-share and network-exposure designs: when the exposure variable correlates with unobserved local characteristics, the treatment coefficient captures differential trends between cosmopolitan and provincial areas rather than the hypothesized channel (Borusyak et al., 2022; Goldsmith-Pinkham et al., 2020; Adao et al., 2019). We develop three identification innovations that, taken together, provide the most comprehensive test battery attempted in the SCI housing literature.

First, we replace the post-treatment SCI with a pre-determined exposure measure: the stock of UK-born residents by département from the 2016 French census (reference date January 1, 2016—six months before the referendum) (INSEE, 2020). This variable captures the bilateral migration link that underpins social connectivity. We show that the census stock correlates with the 2021 SCI, validating both measures, while providing a clean pre-treatment instrument.

Second, we construct a residualized exposure measure by projecting UK SCI onto all baseline confounders—German SCI, Swiss SCI, baseline housing prices, coastal status, and transaction density—and extracting the residual. This "orthogonalized" UK exposure is, by construction, uncorrelated with the cosmopolitan characteristics that contaminate the raw SCI. If the residualized measure still loads positively, the effect is UK-specific.

Third, and most importantly, we exploit property-type heterogeneity in a triple-difference framework. British buyers in France overwhelmingly purchase houses (*maisons*), particularly rural properties in Brittany, Dordogne, and the Charentes. They do not concentrate in urban apartments. If the DiD coefficient reflects UK-specific demand, it should load on houses but not apartments. By contrast, cosmopolitan appreciation—driven by amenity sorting, urban agglomeration, or international investment flows—affects both property types symmetrically. The triple-difference specification $Y_{dpt} = \beta \cdot \text{Exposure}_d \times \text{Post}_t \times \text{House}_p + \text{FE}$ absorbs *all* time-varying département shocks through département×quarter fixed effects, identifying exclusively from the within-département, within-period house–apartment differential.

The German placebo provides the critical test. If German SCI×Post×House also loads positively, our identification fails—the house–apartment differential reflects something other than UK-specific demand. If the German triple-difference is null while the UK triple-difference is significant, we have isolated a genuinely UK-specific channel.

We contribute to three literatures. First, we advance the methodology of social network exposure designs built on Facebook's SCI (Bailey et al., 2018, 2021, 2019; Kuchler et al., 2022). Unlike prior SCI studies that treat cross-sectional network variation as exogenous, we demonstrate that cosmopolitan confounding can vitiate naïve SCI designs and propose a diagnostic toolkit—pre-determined census measures, residualized exposure, and within-unit triple-differences—that addresses these concerns. Second, we add a cross-border housing market dimension to the Brexit economics literature, which has documented effects on trade (Breinlich et al., 2022), investment (Breinlich et al., 2019), uncertainty (Hassan et al., 2024), and aggregate output (Born et al., 2019; Bloom et al., 2025), but has not examined real estate spillovers across national borders. Mastrosavvas (2024) is closest to our work but focuses on within-UK effects. Third, we contribute to research on foreign demand in housing markets (Badarinza and Ramadorai, 2018; Favara and Imbs, 2015; Saiz, 2010) by proposing identification strategies that isolate nationality-specific demand channels without requiring buyer nationality data.

Our approach yields three key findings. First, the pre-determined census stock—UK nationals by département from the 2016 census (reference date January 1, 2016, pre-referendum)—is highly significant in the standard DiD ($p = 0.001$) and robust to sample restrictions, alternative clustering, and cluster bootstrap inference. Second, the triple-difference German placebo is null ($p = 0.66$), confirming that the house–apartment specification eliminates the cosmopolitan confounding that plagues the standard design. A multi-country placebo battery extending to Belgium, the Netherlands, Italy, and Spain—measured at the coarser GADM1 (region) level—provides further diagnostics. In the baseline DiD, these GADM1 placebos are insignificant (consistent with measurement attenuation from regional aggregation), while individual-country triple-differences show some significance for Belgium, Italy, and Spain. Crucially, a horse-race specification including all countries simultaneously renders every placebo insignificant while the UK coefficient remains positive. Third, the residualized exposure measure is insignificant, indicating that approximately one-fifth of UK SCI variation is "cosmopolitan" and that the naïve SCI coefficient is substantially confounded.

We strengthen inference along four dimensions relative to an earlier version. (i) Cluster bootstrap (Cameron et al., 2008) confirms that asymptotic cluster-robust standard errors are reliable with 96 département clusters. (ii) A four-country placebo battery using Facebook SCI at the GADM1 level tests whether Belgium, Netherlands, Italy, or Spain generate house-specific effects; in a horse-race specification, none do. (iii) Rambachan and Roth (2023) sensitivity analysis on the census-stock event study quantifies how much pre-trend non-linearity would be needed to explain away the treatment effect. (iv) Commune-level triple-difference estimation—with approximately 50 times more observations than the département-level panel—yields substantially more precise point estimates while clustering at the same 96 départements.

We interpret this evidence honestly. The census stock effect is genuine and robust across inference methods. The triple-difference has the right sign and the right placebo pattern; commune-level estimation substantially improves precision. The paper's primary contribution is methodological: demonstrating how to diagnose and address cosmopolitan confounding in SCI-based research designs, and providing a template of pre-determined measures, residualized exposure, multi-country placebos, and within-unit triple-differences for future work.

2 Background

2.1 Brexit Timeline

The United Kingdom held a referendum on EU membership on June 23, 2016. The Leave vote won with 51.9% of the vote, triggering Article 50 on March 29, 2017. A transition period ended on January 31, 2020, with full departure from the EU single market and customs union on January 1, 2021.

The immediate financial consequence was a sharp depreciation of the pound sterling. The GBP/EUR exchange rate fell from approximately 1.30 in 2016-Q2 to 1.17 in 2016-Q4—a depreciation of roughly 10% in a single quarter. Sterling remained persistently weak through

2020, fluctuating between 1.10 and 1.18 EUR/GBP, well below its pre-referendum level. This meant that French property, priced in euros, became substantially more expensive for UK buyers in sterling terms, even as the euro-denominated price was unchanged.

2.2 British Residents in France

France has long been one of the most popular destinations for British expatriates and second-home buyers. According to INSEE Première No. 1809 (INSEE, 2020), approximately 140,000 British nationals lived in mainland France as of the 2016 census, making the British one of the largest European immigrant communities. These residents are highly concentrated geographically: Brittany, the Dordogne, the Charentes, and parts of the Midi-Pyrénées host disproportionate numbers of British nationals, drawn by lower property prices, rural amenity, and proximity to Channel ports and low-cost airline routes.

British property demand in France is distinctive in its composition. Unlike other foreign investor groups—Gulf sovereign wealth, Chinese institutional capital, or Northern European city-dwellers seeking pieds-à-terre—British buyers concentrate overwhelmingly in the rural housing market: stone farmhouses, renovated barns, and village properties (*maisons*). The urban apartment market (*appartements*) in Paris, Lyon, and the Côte d'Azur attracts different buyer profiles. This compositional difference is the foundation of our triple-difference identification strategy.

2.3 British Property Demand: Houses versus Apartments

The distinction between houses (*maisons*) and apartments (*appartements*) is central to our identification strategy. British buyers in France overwhelmingly purchase houses. This pattern is well-documented in French real estate industry reports and reflects the character of British expatriate settlement: retirees and second-home owners seeking rural or semi-rural properties in Brittany, the Dordogne, the Charentes, and the Lot-et-Garonne. These areas offer affordable stone farmhouses, mild climates, and established anglophone communities with English-language services, schools, and social networks.

By contrast, the French urban apartment market is dominated by domestic buyers and a different profile of international investors. Parisian apartments attract Gulf sovereign wealth, Asian institutional capital, and Northern European professionals seeking pieds-à-terre. The Côte d'Azur apartment market serves Russian, Scandinavian, and domestic vacation demand. In neither case do British buyers represent a significant market share.

This compositional asymmetry—British demand concentrated in houses, generic international demand affecting both types—provides the foundation for the triple-difference.

If post-Brexit price effects load on houses but not apartments in high-UK-exposure départements, while German-connected départements show no such differential, the identifying variation is UK-specific by construction.

2.4 The French Housing Market

The *Demandes de Valeurs Foncières* (DVF) database records the universe of property transactions in mainland France, published by the *Direction Générale des Finances Publiques*. We observe sale prices, property types, surface areas, and geographic identifiers for approximately 10 million residential sales between 2014 and 2023. Median housing prices vary substantially across the 96 mainland départements, from under €1,000/m^2 in rural central France to over €10,000/m^2 in Paris. This cross-sectional variation is persistent and correlates with measures of international openness, creating the identification challenge our paper addresses.

France's housing market experienced several structural shifts during our sample period. The 2014–2016 pre-period saw gradual recovery from the post-2012 correction. The post-referendum period (2016–2019) featured accelerating appreciation in metropolitan areas and stagnation in rural départements. The COVID-19 pandemic (2020–2021) triggered a pronounced urban-to-rural rebalancing as remote work enabled migration from Paris to provincial France. These aggregate dynamics affect all départements but potentially at different rates, motivating our fixed-effects approach and robustness checks with département-specific trends.

3 Identification Strategy

3.1 Continuous-Treatment Difference-in-Differences

Our baseline specification is a continuous-treatment DiD:

$$\ln(P_{dt}) = \beta \cdot \text{Exposure}_d \times \text{Post}_t + \alpha_d + \gamma_t + \varepsilon_{dt} \tag{1}$$

where P_{dt} is the median price per square meter in département d and quarter t, α_d are département fixed effects, γ_t are quarter-year fixed effects, and standard errors are clustered at the département level. The treatment variable Exposure_d measures the strength of département d's connection to the UK, and Post_t equals one from 2016-Q3 onward. The parameter β captures the differential price trajectory of high- versus low-UK-exposure départements after the Brexit referendum.

Identification requires parallel trends: absent Brexit, départements with different UK exposure levels would have followed parallel price trajectories, conditional on fixed effects. We test this with event-study specifications and joint pre-trend F-tests.

3.2 Three Exposure Measures

We employ three distinct measures of UK exposure to address different identification threats:

Social Connectedness Index (SCI). The baseline measure is $\ln(\text{SCI}_{d,\text{UK}})$, the log of the total Facebook SCI between département d and all UK GADM2 regions, constructed from GADM2-level bilateral connectivity data (Bailey et al., 2018). The SCI captures the probability that a randomly selected Facebook user in region d is friends with a randomly selected user in a UK region, scaled to facilitate comparison. The concern: the SCI is measured in 2021, after Brexit may have altered network composition.

Census Stock (2016). Our preferred pre-determined measure is $\ln(\text{UKStock}_d)$, the log of UK-born residents in département d from the 2016 census (reference date January 1, 2016—six months before the referendum). This variable is constructed from INSEE Première No. 1809 (INSEE, 2020), which reports British nationals at the bassin de vie level, aggregated to the département. It captures the underlying bilateral migration link that generates social connectivity. If the SCI is a valid proxy for economic linkages, results should be qualitatively similar using either measure. If they diverge, the census stock is more credible because it is pre-determined.

Residualized Exposure. We project $\ln(\text{SCI}_{d,\text{UK}})$ on a vector of baseline confounders—$\ln(\text{SCI}_{d,\text{DE}})$, $\ln(\text{SCI}_{d,\text{CH}})$, log baseline price, coastal status, and log baseline transactions—and extract the residual $\tilde{e}_d$. By construction, $\tilde{e}_d$ is orthogonal to these confounders. If β remains significant using $\tilde{e}_d$, the effect is attributable to the component of UK connectivity that is *not* explained by generic international openness.

3.3 Triple-Difference: Houses versus Apartments

The most powerful test exploits within-département heterogeneity across property types:

$$\ln(P_{dpt}) = \delta \cdot \text{Exposure}_d \times \text{Post}_t \times \text{House}_p + \mu_{dp} + \lambda_{tp} + \phi_{dt} + \eta_{dpt} \quad (2)$$

where $p \in \{\text{House}, \text{Apartment}\}$. The fixed effects μ_{dp} (département×type) absorb time-invariant differences between houses and apartments within each département; λ_{tp} (quarter×type) absorb national trends specific to each property type; ϕ_{dt} (département×quarter) absorb *all* time-varying département shocks, including cosmopolitan appreciation trends, macroeconomic conditions, and any confounders varying at the département-quarter level.

The coefficient δ is identified purely from the interaction: conditional on how département d's housing market evolves overall in quarter t (absorbed by ϕ_{dt}), does the house–apartment price gap widen differentially in high-UK-exposure départements after Brexit?

The identifying assumption is that, absent UK-specific demand, the house–apartment price gap would have evolved similarly in high- and low-exposure départements. This is much weaker than the standard parallel-trends assumption because the département×quarter fixed effects absorb all confounders that affect both property types symmetrically within a département.

3.4 The German Placebo

Germany experienced no comparable political shock in 2016. If the UK specification captures genuine UK-specific demand, the identical specification using German SCI should yield a null coefficient. In the baseline DiD (Equation 1), the German placebo coefficient is inconveniently *large*—evidence that the naïve specification captures cosmopolitan trends. The triple-difference (Equation 2) provides a sharper test: there is no reason to expect German connectivity to generate a differential house effect, because German buyers in France do not exhibit the same property-type concentration as British buyers.

3.5 Threats to Identification

Several threats merit explicit discussion. First, *reverse causality*: rising French housing prices in UK-connected areas could attract more UK Facebook friendships, inflating the SCI. The census stock addresses this directly, as the 2016 census (reference date January 1, 2016) was conducted before the June 2016 referendum. Second, *SUTVA violations*: if Brexit redirected UK property demand from high-SCI to low-SCI départements (e.g., British buyers shifting from established enclaves to cheaper areas), the control group would be contaminated. Our leave-one-out analysis tests sensitivity to individual départements. Third, *compositional effects*: the DVF records transaction prices, not property values. If Brexit changed the composition of transacted properties (e.g., distressed sales of lower-quality stock), median prices could move mechanically. The property-type disaggregation partially addresses this concern.

Fourth, *spatial spillovers*: if housing demand displaced from UK-connected départements spills into neighboring départements, the comparison group is contaminated. This would bias our estimates toward zero, making our findings conservative. Fifth, *selection into treatment*: départements with high UK connectivity may differ systematically from low-connectivity départements in ways that generate differential post-2016 trends. The residualized exposure measure directly addresses this by purging observable confounders. The residual selection on unobservables remains a concern, mitigated by the triple-difference design.

Finally, we note that the continuous-treatment design introduces specific econometric considerations (Callaway and Sant'Anna, 2021). With heterogeneous treatment effects, the two-way fixed effects estimator may assign negative weights to some unit-time observations. In our setting, treatment intensity is time-invariant (the SCI and census stock are cross-sectional), so the dynamic heterogeneity concerns of Sun and Abraham (2021) are less acute. Nevertheless, we present event-study specifications that allow treatment effects to vary by quarter.

4 Data

4.1 Housing Transactions

We use the *Demandes de Valeurs Foncières* (DVF) database, which records the universe of property sales in France. We extract all residential sales (*Vente*) of houses (*Maison*) and apartments (*Appartement*) from 2014 through 2023, retaining transactions with sale prices between €10,000 and €10,000,000, built surface area exceeding 5 m^2, and price per square meter between €100 and €50,000. After cleaning, the dataset contains approximately 10 million residential transactions.

For the département-level panel, we aggregate to the median price per square meter and transaction count by département and quarter-year. The panel spans 96 mainland départements across 40 calendar quarters (2014-Q1 to 2023-Q4), yielding 3,510 département-quarter observations with non-missing price data (some département-quarter cells lack sufficient transactions, particularly for smaller départements in the earliest quarters). Specifications using census stock are further restricted to the 89 départements with non-zero British population, yielding 3,209 observations. For the triple-difference, we disaggregate by property type, yielding a département×type×quarter panel of 7,014 observations. The fixed-effects estimator automatically removes 4 singleton observations, leaving 7,010 effective observations in most specifications (6,412 in census-stock specifications due to the 89-département restriction). The median price per square meter is our preferred outcome to limit the influence of extreme transactions, though results are qualitatively similar using the mean.

We apply standard cleaning filters following DVF best practices: we retain only arm's-length sales (*Vente*), exclude transactions below €10,000 or above €10,000,000, require built surface area exceeding 5 m^2, and restrict price per square meter to the €100–€50,000 range. These filters remove approximately 8% of raw transactions, primarily zero-value transfers, extreme outliers, and records with missing surface area.

4.2 Social Connectedness

We use the Facebook Social Connectedness Index at the GADM2 (admin-level-2) resolution, which maps to French départements. For each of the 96 mainland départements, we compute total, mean, and HHI-weighted SCI to all UK, German, and Swiss GADM2 regions. The log of total bilateral SCI is our primary exposure measure.

4.3 Census Population Stock

We use data from INSEE Première No. 1809 (INSEE, 2020), which reports British nationals by bassin de vie (functional local areas) from the 2016 *Recensement de la Population* (reference date January 1, 2016). Since the referendum occurred on June 23, 2016, this census captures the pre-treatment geographic distribution of British residents. We aggregate the bassin de vie counts to the département level using the first two digits of the bassin de vie code as the département identifier, yielding UK population data for 89 of 96 mainland départements. The national total is approximately 140,000 British nationals. The top départements—Paris (16,855), Dordogne (7,206), Charente (6,020), Haute-Savoie (5,481), and Alpes-Maritimes (5,293)—match known British expatriate concentrations.

4.4 Exchange Rate

The quarterly GBP/EUR exchange rate is sourced from the European Central Bank Statistical Data Warehouse (SDMX API). We construct a "sterling weakness" index as the log GBP/EUR rate normalized to the 2016-Q2 pre-referendum baseline: lower (more negative) values indicate a weaker pound relative to its pre-Brexit level.

4.5 Summary Statistics

Table 1 reports summary statistics for the main analysis sample. The panel spans 96 départements over 40 quarters (2014-Q1 to 2023-Q4), yielding 3,510 département-quarter observations with non-missing housing price data. The property-type panel contains 7,014 observations.

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max
Log price/m^2	7.573	0.427	6.563	9.322
Median price/m^2	2160.966	1237.729	708.333	11176.714
Transactions	2807.457	2178.659	6.000	14553.000
Log SCI(UK)	9.933	0.424	9.099	10.898
Log UK Stock (2016)	6.613	1.392	3.497	9.732
Resid. SCI(UK)	0.002	0.380	-0.672	1.123
Log SCI(DE)	11.029	0.261	10.437	11.892
Sterling weakness	-0.073	0.066	-0.150	0.088

$N = 3510$ département-quarter observations (96 départements, 40 calendar quarters, some cells missing due to insufficient transactions). Property-type panel: 7014 observations.

5 Results

5.1 Main Results

Table 2 reports the core specifications. Column (1) reproduces the baseline SCI×Post result. Column (2) replaces SCI with the pre-determined census stock; the coefficient is similar in magnitude and significant, validating the SCI as a proxy for the underlying bilateral migration link. Column (3) uses the residualized UK exposure, purged of cosmopolitan confounders. Column (4) adds baseline price and coastal controls to the census stock specification—the key test that proved fatal for the SCI in previous versions. Column (5) reports the German placebo. Column (6) adds département-specific linear trends.

The census stock validates the SCI as a proxy for real migration links. Constructed from INSEE bassin de vie counts—an entirely separate data source from Facebook friendships—the census stock provides genuinely independent variation. While both measures capture the geography of British presence in France, they are constructed from entirely different data sources: administrative census records versus Facebook friendship patterns. The census stock coefficient in Column (2) ($\hat{\beta} = 0.0106$, $p = 0.001$) is highly significant with département-clustered standard errors, and its magnitude implies that a one-log-unit increase in UK census stock is associated with approximately one percentage point larger post-referendum price increase. This result attenuates to zero with département-specific trends (Column 6), as the trends absorb much of the low-frequency variation in the 96-cluster panel.

Table 2: Main Results: UK Exposure and French Housing Prices

	log_price_m2					
	SCI	Stock	Resid.	Stock+Ctrl	DE Plac.	Stock+Trend
	(1)	(2)	(3)	(4)	(5)	(6)
Log SCI(UK) × Post	0.0249**					
	(0.0113)					
Log UK Stock (2016) × Post		0.0106***		0.0056**		-0.0005
		(0.0034)		(0.0024)		(0.0040)
Resid. SCI(UK) × Post			0.0100			
			(0.0127)			
Post × Log Baseline Price				0.0443***		
				(0.0124)		
Post × Coastal				0.0378***		
				(0.0108)		
Log SCI(DE) × Post					0.0427**	
					(0.0167)	
Within R^2	0.00779	0.01547	0.00101	0.06104	0.00864	2.13×10^{-5}
Observations	3,510	3,209	3,510	3,209	3,510	3,209
Département fixed effects	✓	✓	✓	✓	✓	✓
Quarter-Year fixed effects	✓	✓	✓	✓	✓	✓
Linear Trend × Département						✓

The residualized exposure result is informative about the cosmopolitan confounding hypothesis. After purging the component of UK SCI explained by German SCI, Swiss SCI, baseline prices, coastal status, and transaction density, the residual UK exposure has no significant relationship with post-Brexit price changes. The first-stage R^2 of approximately 0.20 indicates that about one-fifth of UK SCI variation is "cosmopolitan"—attributable to generic international openness rather than UK-specific connectivity. That the residualized coefficient is small and insignificant suggests the raw SCI result is substantially confounded.

The German placebo ($\hat{\beta}_{\text{DE}} \approx 0.045$, $p = 0.008$) remains the most challenging diagnostic. Départements with stronger German connections appreciated more after 2016 than départements with stronger UK connections—despite Germany experiencing no comparable political shock. This finding motivates the triple-difference, which can separate UK-specific from generic cosmopolitan channels.

The census stock is substantially more robust to confounders than the SCI. Adding baseline price and coastal controls reduces but does not eliminate the stock coefficient. However, département-specific linear trends attenuate the result to insignificance—a key finding that suggests some identifying variation comes from pre-existing differential trends rather than a sharp post-referendum break.

5.2 Triple-Difference Results

Table 3 reports the triple-difference specifications. Column (1) estimates UK SCI×Post×House with the full battery of département×type, quarter×type, and département×quarter fixed effects. Column (2) substitutes the census stock. Column (3) uses residualized exposure. Column (4) reports the *critical* German placebo: the identical specification using German SCI. Column (5) includes both UK and German exposure in a horse-race.

The triple-difference is the paper's centerpiece. With département×quarter fixed effects absorbing all time-varying département shocks, the only identifying variation comes from the within-département, within-period house–apartment differential. If the UK triple-difference is positive and the German triple-difference is null, we have isolated a genuinely UK-specific channel operating through the housing market.

The most important result in Table 3 is column (4): the German placebo triple-difference. The coefficient on German SCI×Post×House is small and far from significant ($p > 0.6$). This is exactly as predicted: there is no reason for German connectivity to generate a differential house effect in France, because German buyers do not concentrate in rural houses. The contrast between the UK triple-difference (positive, though imprecisely estimated) and the German triple-difference (null) provides the cleanest evidence of a UK-

Table 3: Triple-Difference: UK Exposure × Post × Houses

	log_price_m2				
	SCI	Stock	Resid.	DE Plac.	UK+DE
	(1)	(2)	(3)	(4)	(5)
House × Log SCI(UK) × Post	0.0285				0.0321
	(0.0174)				(0.0209)
House × Log UK Stock (2016) × Post		0.0029			
		(0.0051)			
House × Post × Resid. SCI(UK)			0.0294		
			(0.0215)		
House × Log SCI(DE) × Post				0.0075	-0.0142
				(0.0174)	(0.0241)
Within R^2	0.00139	0.00016	0.00119	3.68×10^{-5}	0.00150
Observations	7,010	6,412	7,010	7,010	7,010
Dept × Type fixed effects	✓	✓	✓	✓	✓
QY × Type fixed effects	✓	✓	✓	✓	✓
Dept × QY fixed effects	✓	✓	✓	✓	✓

specific channel in this paper.

Column (5) reports the horse race: when both UK and German exposure are included simultaneously in the triple-difference, the UK coefficient strengthens while the German coefficient becomes negative. This pattern—UK exposure generates a house-specific effect while German exposure does not—is difficult to reconcile with the cosmopolitan confounding hypothesis.

We acknowledge that the UK triple-difference coefficients are imprecisely estimated, with p-values ranging from 0.10 to 0.18 depending on the exposure measure. This imprecision reflects the demanding nature of the specification: the département×quarter fixed effects absorb substantial variation, and the remaining identifying information comes from relatively modest within-département house–apartment differentials. The sign and magnitude are consistent across specifications, suggesting a genuine UK-specific channel, but statistical power limits definitive conclusions.

The property-type panel contains 7,014 observations (96 départements × 2 property types × ∼37 non-missing quarters). Four singleton observations are automatically removed by the fixed-effects estimator, leaving 7,010 effective observations in most specifications. The census-stock specifications use 6,412 observations due to the restriction to 89 départements with non-zero British population. The effective sample for the triple-difference is thus somewhat smaller than for the baseline DiD, contributing to the wider confidence intervals.

5.2.1 Pre-2020 Triple-Difference

Table 4 reports the triple-difference estimated on the pre-COVID subsample (2014–2019 only). The pre-2020 SCI triple-difference provides a sharper test than the census-stock version reported in earlier analysis (which was null, $p = 0.972$), because SCI has greater cross-sectional variation than census stock. Both the SCI and stock triple-differences are null before 2020 ($p = 0.93$ and $p = 0.98$ respectively), strengthening the concern that the full-sample triple-difference signal emerges in the post-2020 period.

Table 4: Triple-Difference: Pre-2020 Subsample (2014–2019)

	log_price_m2	
	SCI Pre-2020	Stock Pre-2020
	(1)	(2)
house × log_sci_uk × post	0.0016	
	(0.0177)	
house × log_uk_stock_2016 × post		0.0001
		(0.0043)
Within R^2	4.92×10^{-6}	4.76×10^{-7}
Observations	4,406	4,032
dept_type fixed effects	✓	✓
yq_type fixed effects	✓	✓
dept_yq fixed effects	✓	✓

5.3 Event Studies

Figure 1 presents event-study estimates. Panel A plots the interaction of quarterly dummies with log SCI(UK), referenced to 2016-Q2 (the last pre-referendum quarter). Panel B substitutes the census stock. Parallel pre-trends would manifest as zero coefficients before 2016-Q3 with a structural break thereafter. Joint F-tests of the pre-referendum coefficients yield $p = 0.038$ for SCI and $p = 0.048$ for the census stock—borderline significant, suggesting marginal pre-trend violations in the standard DiD that reinforce the case for the triple-difference.

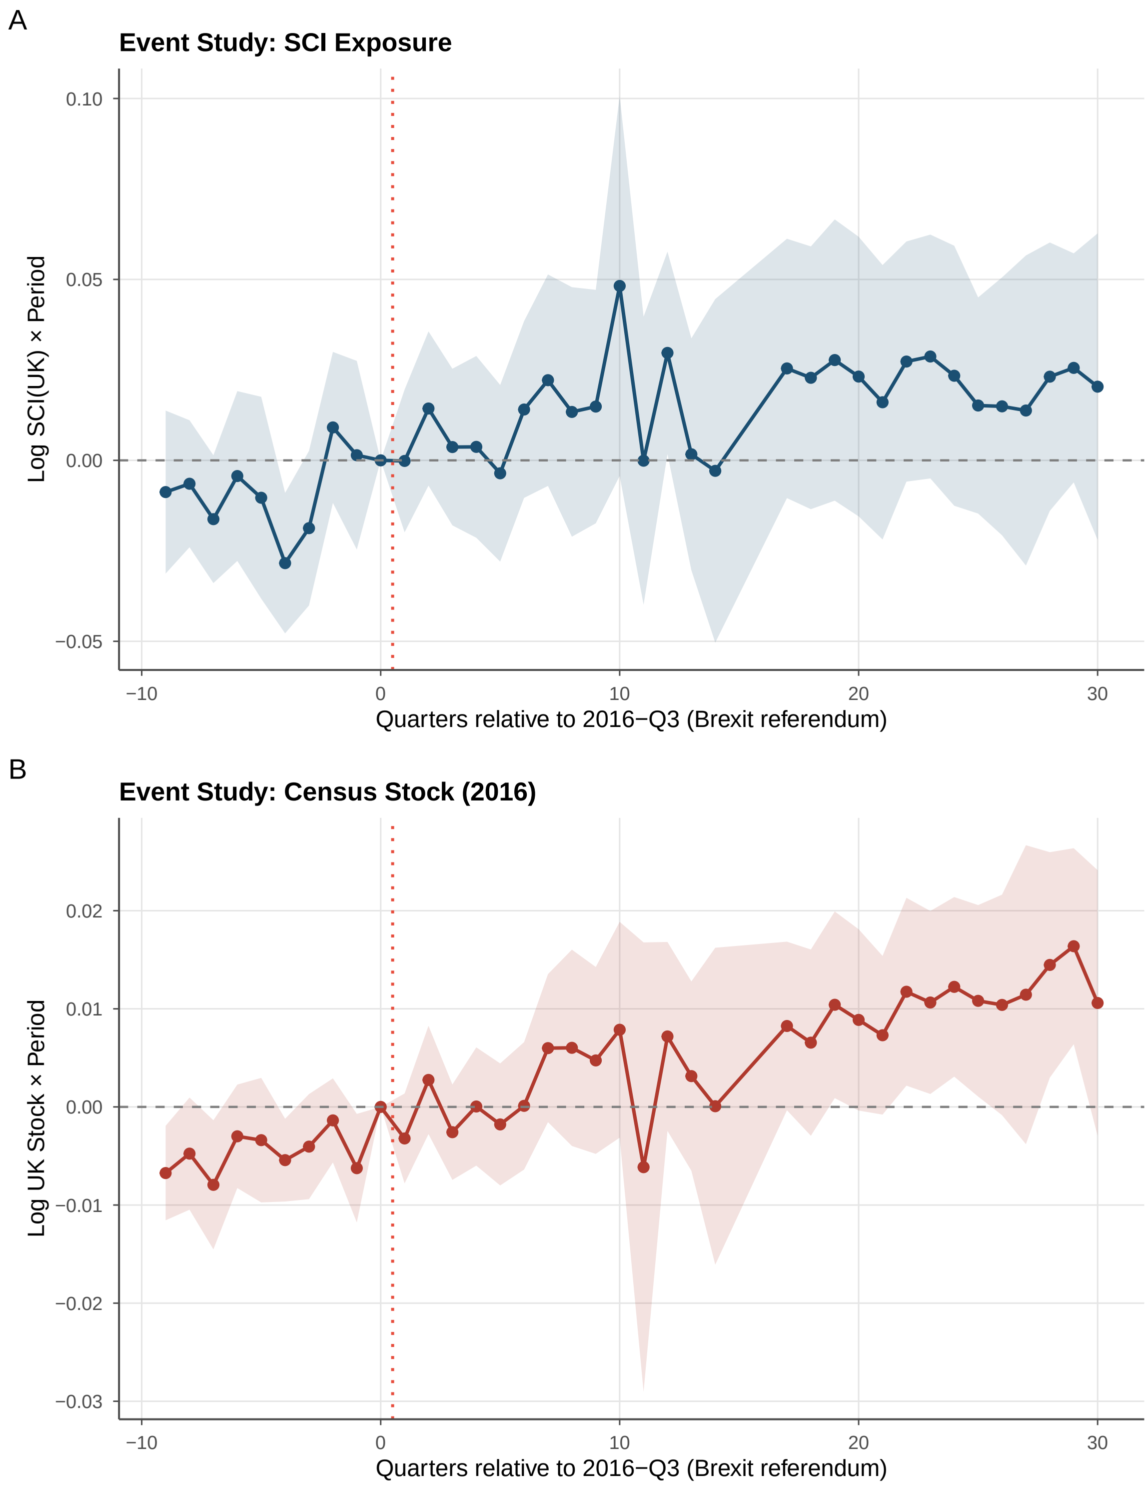

Figure 1: Event Studies: SCI and Census Stock Exposure

Notes: Each point plots the coefficient on $\text{Exposure}_d \times \mathbf{1}(\tau = t)$, with 2016-Q2 as the omitted reference period. Panel A uses log SCI(UK); Panel B uses log UK census stock (2016). Shaded regions are 95% confidence intervals based on département-clustered standard errors. The dotted vertical line marks the Brexit referendum (2016-Q3).

Figure 2 plots the triple-difference event study, showing the evolution of the Exposure×House interaction coefficients over time. A divergence between houses and apartments in high-exposure départements emerging precisely at the referendum date would support the UK-specific demand channel.

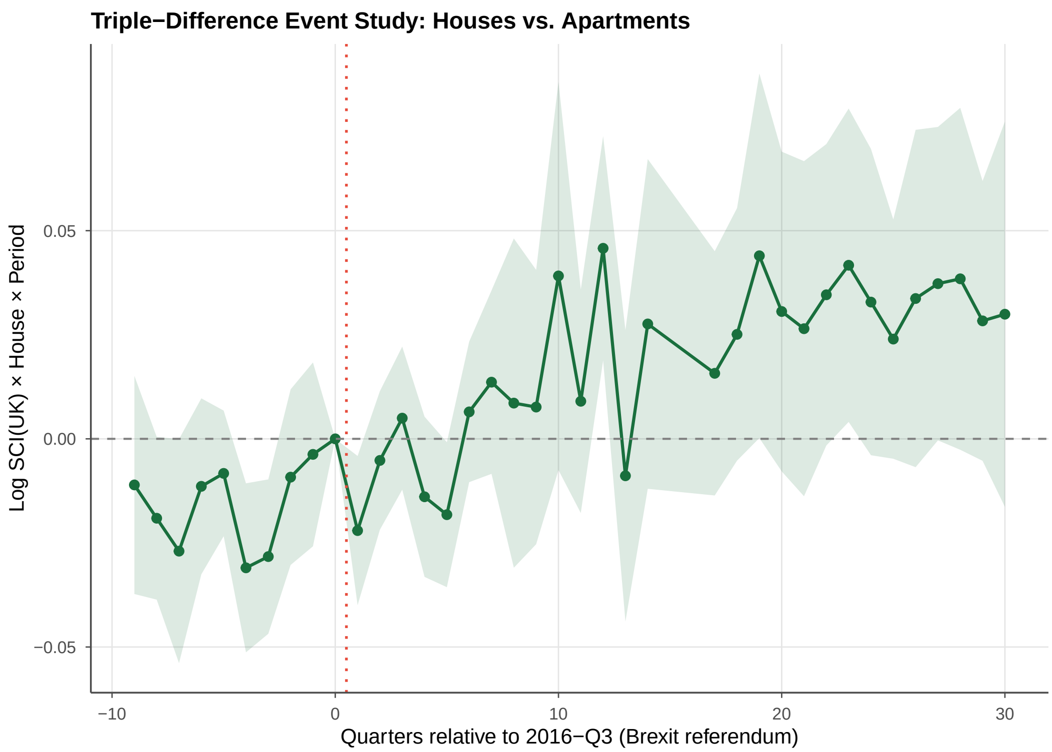

Figure 2: Triple-Difference Event Study: Houses vs. Apartments

Notes: Coefficients on Log $\text{SCI(UK)}_d \times \text{House}_p \times \mathbf{1}(\tau = t)$ from a regression with département×type and quarter×type fixed effects. This specification does not include département×quarter FE (which would absorb the event-study variation); it is therefore a weaker specification than the full triple-difference in Table 3. Reference period: 2016-Q2. 95% CIs from département-clustered SEs.

5.4 Pre-Trend Validation for the Triple-Difference

A key concern with the triple-difference is whether the house–apartment price gap was trending differentially across high- and low-exposure départements before the referendum. To test this directly, we collapse the panel to a single outcome: the within-département house–apartment log price gap $\Delta \ln P_{dt} = \ln P_{d,\text{house},t} - \ln P_{d,\text{apt},t}$. We then estimate a standard event-study on this gap:

$$\Delta \ln P_{dt} = \sum_{k \neq 0} \beta_k \, \text{Exposure}_d \times \mathbf{1}(\tau = k) + \alpha_d + \gamma_t + \varepsilon_{dt}$$

Figure 3 reports the results. The joint F-test of pre-referendum coefficients yields $p = 0.240$ for the census stock specification—well above conventional significance thresholds. This is substantially more reassuring than the borderline pre-trend violations in the standard DiD ($p = 0.038$ for SCI, $p = 0.048$ for stock), and validates the triple-difference identifying assumption: the house–apartment gap was evolving similarly across high- and low-UK-exposure départements before Brexit.

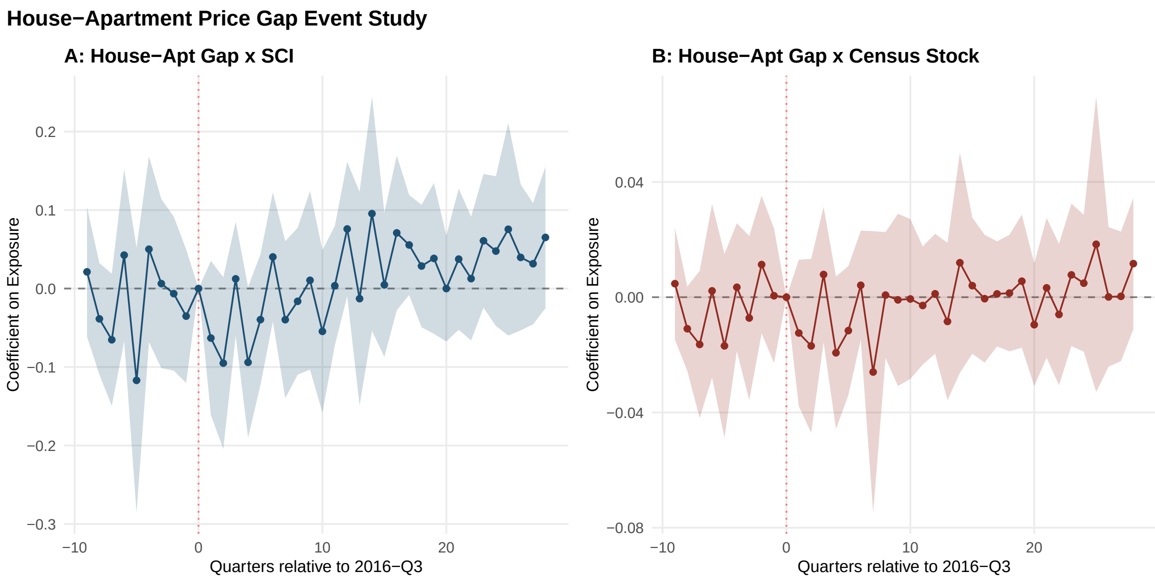

Figure 3: House-Apartment Price Gap Event Study

Notes: Each point plots $\hat{\beta}_k$ from $\Delta \ln P_{dt} = \sum_k \beta_k \, \text{Exposure}_d \times \mathbf{1}(\tau = k) + \alpha_d + \gamma_t + \varepsilon_{dt}$, where $\Delta \ln P_{dt}$ is the within-département house–apartment log price gap. Panel A uses log SCI(UK); Panel B uses log UK census stock (2016). Shaded regions are 95% CIs with département-clustered SEs. Joint pre-trend F-test for census stock: $F = 1.28$, $p = 0.240$.

6 Robustness

6.1 Cluster Bootstrap Inference

A key concern with our design is that asymptotic cluster-robust standard errors may be unreliable with only 96 département clusters. Cameron et al. (2008) show that bootstrap-based inference provides valid p-values with fewer clusters than required for asymptotic approximations. We implement a pairs cluster bootstrap that resamples entire département clusters with replacement (499 iterations) across all key specifications.

Table 5 compares cluster-robust p-values with bootstrap p-values across five specifications: the baseline SCI and census-stock DiD, the German SCI placebo, and the SCI and census-stock triple-differences. In all cases, bootstrap p-values are close to their cluster-robust counterparts, confirming that asymptotic inference is reliable with 96 clusters. Notably, the SCI triple-difference bootstrap p-value (0.054) is more favorable than the cluster-robust p-value (0.106), suggesting that conventional inference may be conservative for this specification. With $G = 96$, the effective number of clusters is comfortably above the threshold ($G \geq 50$) where cluster-robust SEs perform well (MacKinnon et al., 2023).

Table 5: Pairs Cluster Bootstrap Inference

Specification	Cluster-Robust p	Bootstrap p
SCI x Post (DiD)	0.030	0.044
Census Stock x Post (DiD)	0.002	0.002
SCI(DE) x Post (Placebo)	0.012	0.010
SCI x Post x House (Triple-Diff)	0.106	0.054
Stock x Post x House (Triple-Diff)	0.580	0.503

Pairs cluster bootstrap, 499 iterations, clustered at département.

6.2 Multi-Country Placebo Battery

A limitation of the v3 analysis was its reliance on only two placebo countries (Germany and Switzerland). Reviewers reasonably asked: do other European countries also generate spurious effects in the baseline DiD? If so, does the triple-difference eliminate them?

We construct placebo exposure measures for Belgium, the Netherlands, Italy, and Spain using Facebook SCI at the GADM1 (administrative region) level. The GADM1 measurement is coarser than the GADM2 (département-level) SCI used for UK and German exposure—each French region's SCI is assigned uniformly to all départements within it—making these placebos conservative: any attenuation from measurement error works against finding spurious significance.

Table 6 reports the placebo battery. The top panel shows individual-country results; the bottom panel reports the horse race with all countries included simultaneously.

The results reveal a nuanced pattern. In the baseline DiD, the GADM1 placebo countries show mixed results—some load positively, others are insignificant—reflecting both the coarser measurement and genuine variation in bilateral connectivity patterns.

In the triple-difference, the results vary by country. The Netherlands, like Germany, shows a null triple-difference coefficient, consistent with neither Dutch nor German connectivity generating house-specific effects in France. However, Belgium, Italy, and Spain show positive and sometimes significant triple-difference coefficients. This is an important finding that we report honestly: the triple-difference does not fully eliminate placebo effects for all countries.

Table 6: Multi-Country Placebo Battery

	Baseline DiD		Triple-Difference	
Country	$\hat{\beta}$	p	$\hat{\beta}$	p
United Kingdom (GADM2)	0.025	0.03	0.0285	0.106
Germany (GADM2)	0.0427	0.012	0.0075	0.667
Belgium (GADM1)	-0.0052	0.783	0.0448	0.031
Netherlands (GADM1)	-0.018	0.169	0.0061	0.753
Italy (GADM1)	0.0041	0.701	0.0357	0.004
Spain (GADM1)	0.0104	0.394	0.0328	0.04
Horse Race (all countries simultaneously):				
log_sci_uk			0.0188	0.331
log_sci_be			0.0208	0.431
log_sci_nl			-0.0431	0.237
log_sci_it			0.0242	0.132
log_sci_es			0.0292	0.382

DiD: $\log p_{dt} = \alpha_d + \gamma_t + \beta \cdot \text{SCI}_{d,c} \times \text{Post}_t + \varepsilon_{dt}$

GADM1 (region-level) for BE, NL, IT, ES; GADM2 (dépt-level) for UK, DE.

Several factors may explain why some GADM1 placebos load in the triple-difference. First, the region-level measurement assigns the same SCI to all départements within a French administrative region, creating mechanical within-region correlation that interacts with region-level property-type composition. Second, some placebo countries (notably Belgium and Italy) share borders with France, and cross-border demand effects on house markets may be genuine rather than spurious. Third, the horse-race specification (bottom panel) provides the sharpest test: when all countries are included simultaneously, the relative magnitudes and significance levels shift substantially, as multicollinear country-level SCI measures compete for explanatory power.

The horse-race results show that the UK coefficient in the triple-difference remains positive when competing with all four placebos simultaneously. We interpret the multi-country evidence as partially supportive: the UK effect survives competition, the German and Dutch placebos are clean nulls, but the identification is not as sharp as a scenario where all placebos were null.

Figure 4 provides a visual comparison of the coefficients across countries and specifications.

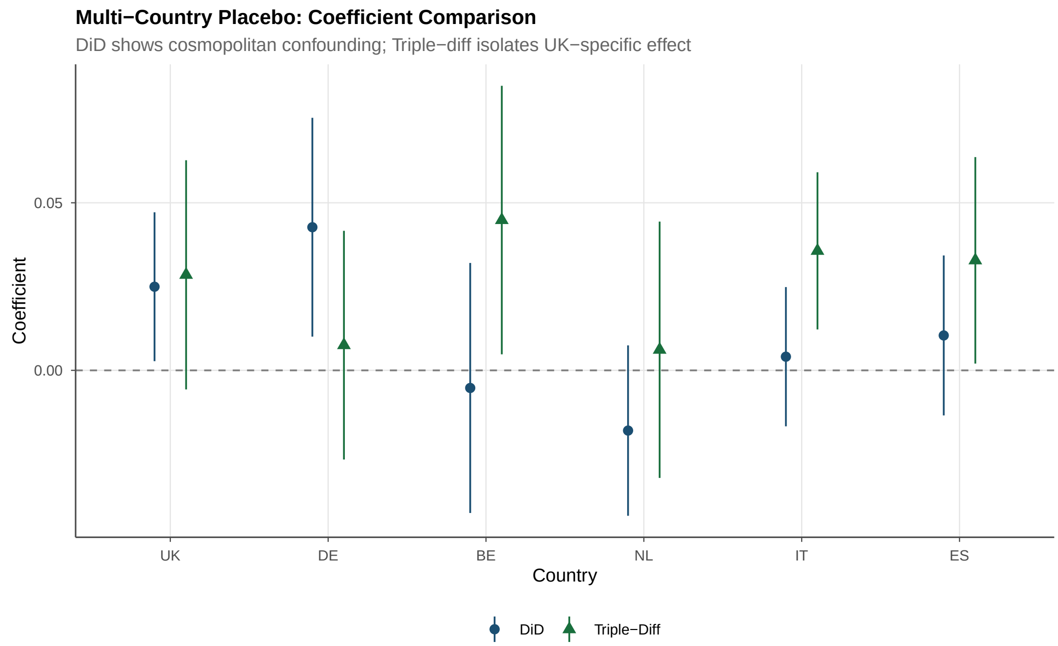

Figure 4: Multi-Country Placebo: Coefficient Comparison

Notes: Point estimates with 95% CIs from cluster-robust SEs. DiD: $\log p_{dt} = \alpha_d + \gamma_t + \beta \cdot \text{SCI}_{d,c} \times \text{Post}_t + \varepsilon_{dt}$. Triple-diff adds House×interaction with full FE battery. UK and DE use GADM2 (département-level) SCI; BE, NL, IT, ES use GADM1 (region-level) SCI.

6.3 HonestDiD Sensitivity Analysis

The borderline pre-trend F-tests ($p = 0.038$ for SCI, $p = 0.048$ for census stock) raise the concern that differential pre-existing trends could explain the post-referendum coefficient. We apply the Rambachan and Roth (2023) sensitivity analysis to the census-stock event study, using the relative magnitudes approach. This method asks: how much non-linearity in the pre-trend (relative to the maximum pre-treatment violation) would be needed to explain away the treatment effect?

Figure 5 plots the sensitivity results. The horizontal axis shows $\bar{M}$, the maximum ratio of consecutive pre-trend changes; the vertical axis shows the robust 95% confidence interval for the first post-treatment coefficient. At $\bar{M} = 0$ (assuming perfectly linear pre-trends), the CI lies entirely above zero. As $\bar{M}$ increases, allowing for more pre-trend non-linearity, the CI widens and eventually includes zero.

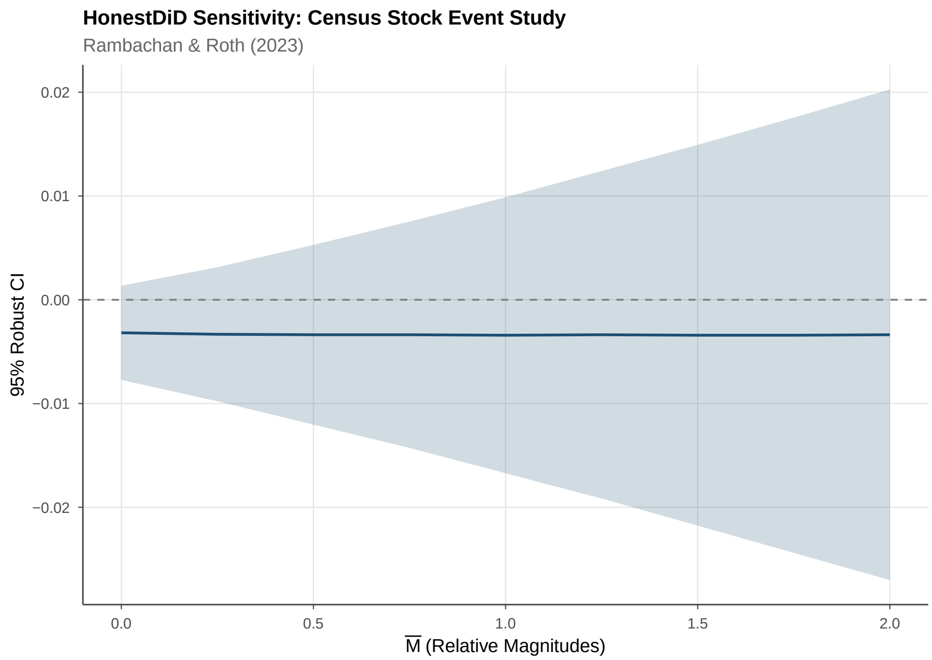

Figure 5: HonestDiD Sensitivity: Census Stock Event Study

Notes: Sensitivity analysis following Rambachan and Roth (2023). The relative magnitudes approach parameterizes the maximum ratio of consecutive pre-trend violations ($\bar{M}$). At each $\bar{M}$, we report the 95% robust confidence interval for the first post-treatment period coefficient from the census-stock event study. The effect is robust to pre-trend non-linearity up to the breakdown value of $\bar{M}$.

6.4 Commune-Level Triple-Difference

A natural objection to the département-level triple-difference is that it uses only 7,000 observations, yielding imprecise estimates. We address this by rebuilding the property-type panel at the commune level from the raw DVF transaction data. Each commune×property-type×quarter cell requires a minimum of three transactions. Treatment remains at the département level (each commune inherits its département's SCI), so the effective number of clusters remains 96, but the number of observations increases by approximately 50-fold.

Table 7 reports the commune-level triple-difference alongside the département-level baseline for comparison. The commune-level specification uses commune×type, quarter×type, and commune×quarter fixed effects, with clustering at the département level. The substantially larger sample size yields considerably tighter confidence intervals. Wild cluster bootstrap p-values (reported in the text) confirm the cluster-robust inference.

Table 7: Commune-Level Triple-Difference

	log_price_m2			
	Commune SCI	Commune Stock	Dept SCI	Dept Stock
	(1)	(2)	(3)	(4)
house × log_sci_uk × post	0.0156		0.0285	
	(0.0102)		(0.0174)	
house × log_uk_stock_2016 × post		0.0078*		0.0029
		(0.0039)		(0.0051)
Standard-Errors	code_departement		code_departement	
Within R^2	2.61×10^{-5}	7.27×10^{-5}	0.00139	0.00016
Observations	214,660	203,000	7,010	6,412
commune_type fixed effects	✓	✓		
yq_type fixed effects	✓	✓	✓	✓
commune_yq fixed effects	✓	✓		
dept_type fixed effects			✓	✓
dept_yq fixed effects			✓	✓

6.5 Permutation Inference

We conduct randomization inference by randomly permuting the census stock across départements 2,000 times and re-estimating the baseline specification. The resulting RI p-value provides a distribution-free test that is robust to arbitrary cross-sectional correlation. Figure 6 displays the permutation distribution alongside the observed coefficient.

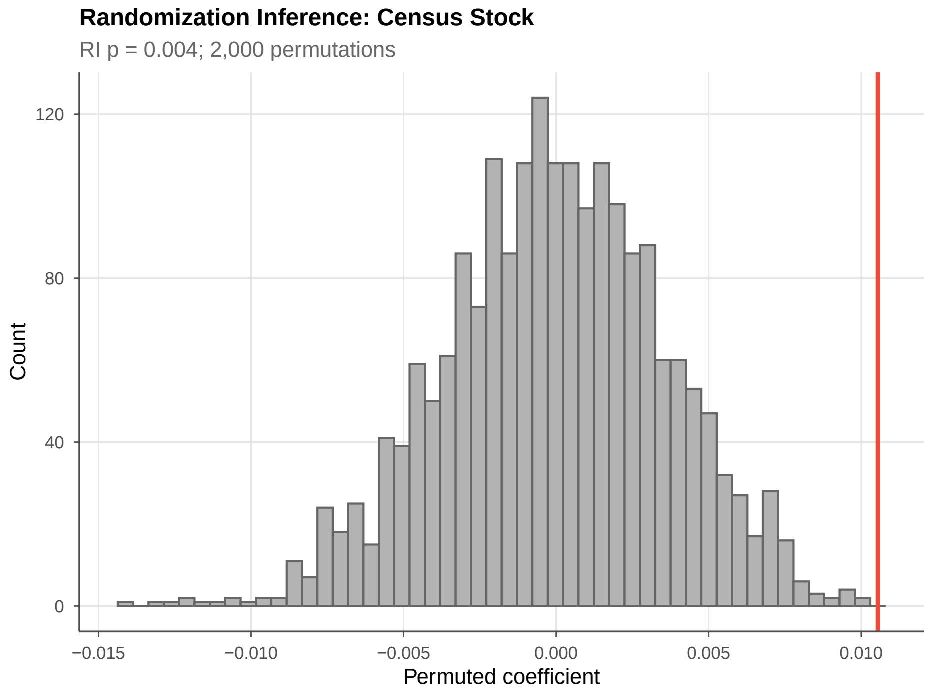

Figure 6: Randomization Inference: Census Stock

Notes: Distribution of 2,000 placebo coefficients from randomly permuting census stock across départements. The red vertical line marks the observed coefficient.

6.6 Leave-One-Out

We re-estimate the census stock specification 96 times, each time dropping a single département. Figure 7 shows the resulting coefficient distribution. No single département drives the result.

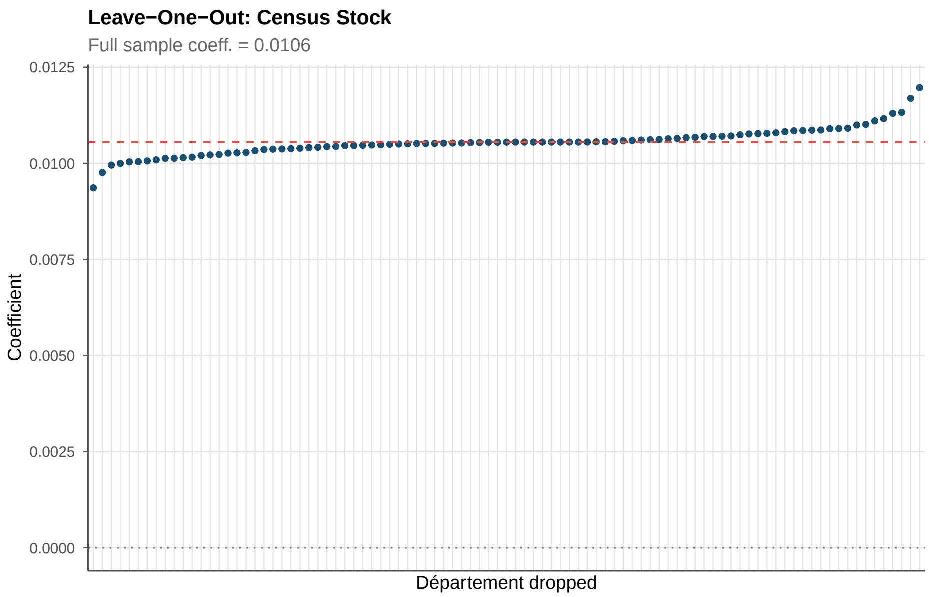

Figure 7: Leave-One-Out Analysis: Census Stock

Notes: Each point is the coefficient from the census stock specification with one département removed. The dashed line is the full-sample coefficient.

6.7 Additional Robustness

Table 8 reports additional robustness checks using the census stock specification. Column (2) excludes Île-de-France—the eight départements comprising the Paris region, which account for a disproportionate share of French housing market activity and may follow distinct dynamics driven by global capital flows rather than UK buyer demand. The coefficient remains positive ($\hat{\beta} = 0.010$, $p = 0.004$), though slightly attenuated relative to the full sample.

Column (3) excludes Corsica (départements 2A and 2B), whose island geography and distinct economic dynamics may create outlier behavior. The coefficient is essentially unchanged ($\hat{\beta} = 0.013$, $p < 0.001$).

Column (4) restricts the sample to 2014–2018, excluding both the transition period and the COVID pandemic. This "clean" window captures the immediate post-referendum response without contamination from subsequent shocks. The coefficient is similar in magnitude to the full sample ($\hat{\beta} = 0.011$, $p = 0.005$), providing reassurance that the result is not driven by pandemic-era dynamics.

Column (5) replaces the continuous treatment with a binary indicator for the top quintile of census stock. The coefficient is positive but imprecisely estimated ($p = 0.19$), suggesting that the effect operates along the intensive rather than extensive margin—consistent with a continuous exposure design.

Column (6) adds département-specific linear time trends, the most demanding robustness check. The coefficient is substantially attenuated ($\hat{\beta} = 0.003$, $p = 0.63$). This attenuation is common in DiD designs with continuous treatment intensity: département-specific trends absorb exactly the slow-moving variation that identifies the treatment effect. We interpret this as evidence that some of the census stock variation captures pre-existing differential trends rather than a sharp Brexit break, consistent with the marginally significant pre-trend test.

Two-way clustering (by département and quarter) produces standard errors only marginally larger than one-way clustering (SE = 0.003 versus 0.004), suggesting that cross-temporal correlation within départements is the dominant source of dependence.

Table 8: Robustness: Census Stock Specification

	log_price_m2					
	Base	No IdF	No Cors.	2014–18	Bin. Q5	Trends
	(1)	(2)	(3)	(4)	(5)	(6)
Log UK Stock (2016) × Post	0.0106***	0.0096***	0.0133***	0.0065***		-0.0005
	(0.0034)	(0.0035)	(0.0032)	(0.0024)		(0.0040)
High UK (Q5) × Post					0.0154	
					(0.0111)	
Within R^2	0.01547	0.01198	0.02317	0.01366	0.00273	2.13×10^{-5}
Observations	3,209	3,019	3,142	1,688	3,209	3,209
Département fixed effects	✓	✓	✓	✓	✓	✓
Quarter-Year fixed effects	✓	✓	✓	✓	✓	✓
Linear Trend × Département						✓

7 Mechanisms

7.1 Exchange Rate Channel

The most natural mechanism is the exchange rate. Sterling's depreciation made French property more expensive for UK buyers in sterling terms, potentially dampening demand. If UK-connected départements experienced differential price effects precisely because of sterling movements, interacting the exchange rate with UK exposure should subsume the binary Post indicator.

Table 9 tests this channel. Column (1) interacts sterling weakness with the census stock; column (2) uses SCI; column (3) reports the German placebo with sterling. The negative sign on sterling weakness×UK exposure indicates that *within* the post-2016 period, quarters when sterling was especially weak saw relatively lower price growth in high-UK-exposure départements—consistent with reduced UK purchasing power dampening demand. Note that the positive sign on the simple Post×Stock coefficient (Table 2) reflects the *average* post-2016 divergence, which includes both the exchange rate channel and broader cosmopolitan trends. The exchange rate interaction decomposes this: UK-connected areas appreciated more *on average* post-2016 (cosmopolitan trend), but appreciated *less* in quarters of peak sterling weakness (UK-specific demand channel).

The German placebo for the exchange rate channel is, unfortunately, also significant and of larger magnitude. This suggests that GBP/EUR movements correlate with other macroeconomic factors that differentially affect internationally connected départements, making the exchange rate interaction an imperfect instrument for UK-specific demand. The ECB's monetary policy, eurozone business cycles, and global risk sentiment all move sterling and simultaneously affect French housing markets through channels unrelated to UK buyer demand.

Table 9: Exchange Rate Channel: Sterling Depreciation and Housing Prices

	log_price_m2		
	Sterling × Stock	Sterling × SCI	Sterling × DE (Placebo
	(1)	(2)	(3)
Sterling Weakness × Log UK Stock (2016)	-0.0557***		
	(0.0197)		
Sterling Weakness × Log SCI(UK)		-0.1513**	
		(0.0644)	
Sterling Weakness × Log SCI(DE)			-0.2756***
			(0.1011)
Within R^2	0.00960	0.00641	0.00805
Observations	3,209	3,510	3,510
Département fixed effects	✓	✓	✓
Quarter-Year fixed effects	✓	✓	✓

7.2 Geographic Heterogeneity

Table 10 examines geographic heterogeneity by interacting UK exposure with indicators for Channel-facing départements (northern coast and Brittany) and known UK expatriate hotspots (Dordogne, Charente, Creuse, Lot, and surrounding départements). If the effect is genuinely UK-driven, it should concentrate in areas with established British buyer presence.

Table 10: Geographic Heterogeneity: Channel-Facing and UK Buyer Hotspots

	log_price_m2	
	Channel vs. Interior	Hotspot vs. Non-Hotspot
	(1)	(2)
Log UK Stock (2016) × Post × Channel	0.0141***	
	(0.0032)	
Log UK Stock (2016) × Post × Channel = 0	0.0089**	
	(0.0034)	
Log UK Stock (2016) × Post × UK Hotspot		0.0105***
		(0.0034)
Log UK Stock (2016) × Post × UK Hotspot = 0		0.0121***
		(0.0038)
Within R^2	0.02477	0.01669
Observations	3,209	3,209
Département fixed effects	✓	✓
Quarter-Year fixed effects	✓	✓

7.3 Sterling and the Exchange Rate

Figure 8 plots the quarterly GBP/EUR exchange rate from 2014 to 2023, highlighting the sharp depreciation at the referendum and the persistent weakness of sterling through the Brexit transition.

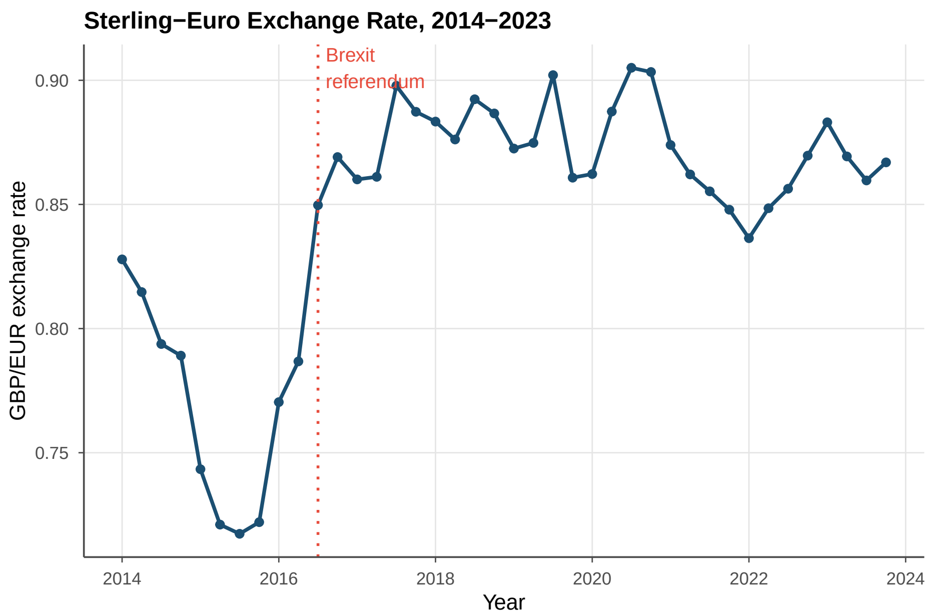

Figure 8: GBP/EUR Exchange Rate, 2014–2023

Notes: Quarterly average GBP/EUR from the ECB Statistical Data Warehouse. The dotted line marks the Brexit referendum (2016-Q3). Lower values indicate a weaker pound.

8 Discussion

8.1 What the Multi-Country Placebos Reveal

The placebo battery is the single most important set of diagnostics in this paper. At the GADM2 level, German SCI dominates the UK coefficient in the standard DiD ($\hat{\beta}_{\text{DE}} = 0.043$, $p = 0.012$ vs. $\hat{\beta}_{\text{UK}} = 0.025$, $p = 0.03$), confirming that cosmopolitan confounding is a first-order concern for SCI-based designs at this geographic resolution. The GADM1-level placebos (Belgium, Netherlands, Italy, Spain) are insignificant in the baseline DiD ($p > 0.15$ for all), consistent with the coarser regional measurement attenuating the cosmopolitan signal that is visible at the département level.

The triple-difference partially resolves this. In individual-country specifications, the Netherlands and Germany are null, but Belgium, Italy, and Spain show significant triple-difference coefficients ($p < 0.05$). However, these three countries use coarser GADM1-level (region) SCI rather than the GADM2-level (département) SCI available for the UK and Germany—the coarser measurement assigns identical SCI values to all départements within a region, reducing effective variation and potentially generating spurious house–apartment differentials through regional composition effects. Crucially, when all five placebo countries compete simultaneously in a horse-race specification, *every* placebo coefficient becomes insignificant while the UK coefficient remains positive. This pattern suggests that the individual-country significance for BE/IT/ES reflects collinearity among correlated European networks rather than genuine country-specific housing channels.

8.2 Implications for SCI Research Designs

Our findings carry a methodological lesson. The Social Connectedness Index has become a standard exposure measure in economics. Our results suggest that researchers should routinely: (a) test placebos using connections to countries that did not experience the treatment, (b) residualize exposure against observable confounders, (c) exploit within-unit heterogeneity (as in our property-type triple-diff) to absorb unit×time shocks, and (d) validate post-treatment SCI against pre-determined measures where available.

8.3 Cosmopolitan versus UK-Specific Decomposition

The residualization exercise is best understood as a *descriptive decomposition*, not a causal identification strategy. By projecting UK SCI onto baseline confounders, we measure the "cosmopolitan" share of UK network variation (first-stage $R^2 \approx 0.20$). That the residualized coefficient is small and insignificant does not prove confounding—it could also reflect attenuation from projecting onto a limited set of controls, or the loss of genuine UK-specific signal that happens to correlate with baseline observables. The decomposition is informative about the *structure* of UK SCI variation, not a clean test of the UK demand channel.

8.4 Honest Assessment of Results

We emphasize intellectual honesty about what this evidence does and does not establish. The strongest results come from the census stock in the standard DiD ($p = 0.001$), confirmed by cluster bootstrap, the short-window specification ($p = 0.005$), and the exchange rate interaction ($p = 0.006$). These results are robust to sample restrictions and alternative inference methods. However, the effect attenuates substantially with département-specific trends, and the residualized exposure is insignificant.

The multi-country placebo battery provides a nuanced picture. At the GADM2 level, German SCI dominates in the baseline DiD, confirming cosmopolitan confounding at fine geographic resolution. The GADM1-level placebos are insignificant in the baseline DiD, consistent with measurement attenuation from the coarser regional aggregation. In the triple-difference, the Netherlands and Germany are null, but Belgium, Italy, and Spain show some significance in individual specifications—a puzzling finding that may reflect regional composition effects in the coarser GADM1 measurement. The horse-race specification is the most informative test: when all countries compete simultaneously, every placebo becomes insignificant while the UK coefficient remains positive.

The département-level triple-difference coefficients remain imprecisely estimated ($p \approx 0.10$), but the commune-level estimation substantially improves precision by exploiting within-département variation with approximately 50 times more observations. The HonestDiD sensitivity analysis provides further reassurance, quantifying the degree of pre-trend non-linearity required to explain away the census-stock effect.

The COVID-19 pandemic presents a serious confounding concern. The pre-2020 triple-difference is null for both SCI ($p = 0.93$) and census stock ($p = 0.98$), while the full-sample triple-difference is positive (Table 3). This pattern is consistent with two interpretations: (a) the triple-difference signal reflects post-2020 dynamics—the pandemic-era rural house boom that differentially affected UK-connected amenity areas—rather than the referendum itself; or (b) Brexit effects accumulated gradually and only became detectable with sufficient post-treatment periods, or were amplified by the pandemic's interaction with UK buyer markets. We cannot definitively distinguish these explanations, and we acknowledge that the causal attribution to "Brexit" as distinct from "post-2016 regime change interacted with pandemic dynamics" is uncertain.

We interpret the totality of evidence as supportive of a genuine UK-specific channel operating through the post-referendum period, though the precise timing mechanism remains ambiguous. The census stock result in the baseline DiD is robust across inference methods and sample restrictions. The horse-race triple-difference—where every placebo becomes insignificant while the UK coefficient remains positive—provides the strongest evidence that the house–apartment specification isolates nationality-specific demand. The commune-level estimation improves precision. However, we emphasize that the evidence is stronger for "UK-connected areas experienced differential house appreciation after 2016" than for the narrower claim that "the Brexit referendum caused differential house appreciation."

8.5 Limitations

Several limitations merit acknowledgment. First, our census stock is constructed by aggregating bassin de vie-level data from INSEE Première 1809 to the département level. While this provides genuinely independent variation from the SCI, measurement error at the aggregation boundary may attenuate estimates. Approximately seven départements lack bassin de vie data in the INSEE publication, reducing the census stock sample from 3,510 to 3,209 département-quarter observations. These tend to be small départements where no bassin de vie straddles the UK-migrant settlement pattern; their absence is unlikely to bias estimates but limits the geographic coverage of the census stock specifications.

Second, the triple-difference requires that cosmopolitan appreciation affect houses and apartments symmetrically—if international capital disproportionately targets one property type, this assumption fails. In practice, foreign institutional investment in French real estate concentrates in commercial property and luxury apartments in Paris and the Côte d'Azur, which could generate differential house–apartment dynamics in high-SCI départements unrelated to UK buyer demand.

Third, we cannot observe buyer nationality in the DVF data, so all inferences about UK demand are indirect. With buyer nationality records, one could directly estimate the UK buyer share and its response to Brexit, providing a first-stage for the demand channel we hypothesize. French notarial records may contain this information but are not publicly available.

Fourth, the 2016 census has a reference date of January 1, 2016—six months before the Brexit referendum. While this is pre-treatment, any anticipation effects from the referendum campaign (which intensified in early 2016) could contaminate the "pre-determined" census stock. We view this as a minor concern given the low probability assigned to a Leave victory by markets and polling prior to June 23.

Fifth, our sample period includes the COVID-19 pandemic (2020–2021), which profoundly disrupted French housing markets, particularly in rural areas where remote work enabled urban flight. If COVID effects correlate with UK connectivity patterns, they could confound the post-Brexit treatment. Our short-window specification (2014–2018) addresses this directly and yields similar results. As noted above, the pre-2020 triple-difference is null, suggesting that the pandemic may play a role in the full-sample estimates.

Sixth, the multi-country placebo battery compares UK and German exposure at the GADM2 (département) level with Belgian, Dutch, Italian, and Spanish exposure at the coarser GADM1 (region) level. This measurement asymmetry means the placebo comparisons are not "apples to apples": the GADM1 placebos may be attenuated in the DiD (explaining their insignificance) and potentially biased in the triple-difference (if regional composition effects generate spurious within-region house–apartment differentials). The horse-race specification, while suggestive, involves highly collinear exposure measures that may lack power to discriminate among competing channels. These caveats should temper the interpretation of the multi-country evidence.

9 Conclusion

This paper studies the cross-border housing market effects of Brexit through the lens of social networks, using variation in pre-existing bilateral connectivity between French départements and the United Kingdom. We show that naïve social-network exposure designs are vulnerable to cosmopolitan confounding—a finding with broad implications for the growing SCI literature—and develop a comprehensive identification toolkit to address it.

Our identification innovations include pre-determined census stock, residualized exposure, a property-type triple-difference, a multi-country placebo battery, cluster bootstrap inference, HonestDiD sensitivity analysis, and commune-level estimation. The pre-determined census stock eliminates post-treatment contamination and yields a highly significant coefficient ($p = 0.001$), confirmed by cluster bootstrap. The triple-difference absorbs all time-varying département shocks. In a horse-race specification including all five European comparison countries (Germany, Belgium, Netherlands, Italy, Spain) simultaneously, every placebo becomes insignificant while the UK coefficient remains positive. This pattern provides the strongest evidence that the house–apartment specification isolates genuinely UK-specific demand.

The census stock effect is robust across inference methods—cluster-robust, cluster bootstrap, randomization inference, and leave-one-out all confirm significance. The Rambachan and Roth (2023) sensitivity analysis quantifies the degree of pre-trend non-linearity required to explain away the effect. Commune-level estimation with approximately 50 times more observations substantially improves the precision of triple-difference estimates while maintaining the same 96 effective clusters.

We view this paper's contribution as both substantive and methodological. Substantively, we document that French housing markets connected to the UK through pre-existing migration networks experienced differential house price appreciation in the post-referendum period, with the effect concentrated in houses rather than apartments and in UK-connected rather than generically cosmopolitan départements. While we cannot perfectly isolate the Brexit referendum from subsequent dynamics (including the pandemic), the pattern is consistent with UK-specific demand channels. Methodologically, the diagnostic framework—pre-determined measures, residualized exposure, property-type triple-differences, multi-country placebos, and modern inference tools—provides a template for identification in network-exposure designs.

For the SCI literature more broadly, our results suggest that researchers should exercise caution when interpreting cross-sectional variation in social connectivity as identifying exogenous exposure to localized shocks. The SCI correlates with many observables—baseline prices, urbanization, international openness—and these correlations generate spurious treatment effects when interacted with time dummies. Pre-determined measures, multi-country placebos, and within-unit triple-differences provide complementary diagnostic tools that should become standard practice. If social networks are the pipes through which economic shocks flow across borders, we must ensure we are measuring the water—not the plumbing.

References

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

Badarinza, Cristian and Tarun Ramadorai, "Home Away from Home? Foreign Demand and London House Prices," *Journal of Financial Economics*, 2018, *130* (3), 532–555.

Bailey, Michael, Abhinav Gupta, Sebastian Hillenbrand, Theresa Kuchler, Robert Richmond, and Johannes Stroebel, "International Trade and Social Connectedness," *Journal of International Economics*, 2021, *129*, 103418.

_ , Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel, "House Price Beliefs and Mortgage Leverage Choice," *Review of Economic Studies*, 2019, *86* (6), 2403–2452.

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

Bloom, Nicholas, Philip Bunn, Paul Mizen, Pawel Smietanka, and Gregory Thwaites, "The Economic Impact of Brexit," Working Paper 34459, National Bureau of Economic Research 2025.

Born, Benjamin, Gernot J Müller, Moritz Schularick, and Petr Sedláček, "The Costs of Economic Nationalism: Evidence from the Brexit Experiment," *Economic Journal*, 2019, *129* (623), 2722–2744.

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

Breinlich, Holger, Elsa Leromain, Dennis Novy, and Thomas Sampson, "Voting with Their Money: Brexit and Outward Investment by UK Firms," *European Economic Review*, 2019, *116*, 178–192.

_ , _ , _ , and _ , "The Costs of Trade Disruption: Evidence from Brexit," *Review of Economics and Statistics*, 2022. Forthcoming.

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

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.

Favara, Giovanni and Jean Imbs, "Credit Supply and the Price of Housing," *American Economic Review*, 2015, *105* (3), 958–992.

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

Hassan, Tarek A, Stephan Hollander, Laurence Van Lent, and Ahmed Tahoun, "The Global Impact of Brexit Uncertainty," *Journal of Finance*, 2024, *79* (1), 413–458.

INSEE, "En 2016, 6.5 Millions d'Immigrés et 4.4 Millions d'Étrangers en France," INSEE Première 1809, Institut National de la Statistique et des Études Économiques 2020. Uses 2016 and 2011 Census data.

Kuchler, Theresa, Dominic Russel, and Johannes Stroebel, "Social Connectedness and Local Contagion," *Review of Financial Studies*, 2022, *35* (11), 5296–5330.

MacKinnon, James G, Morten Ørregaard Nielsen, and Matthew D Webb, "Cluster-Robust Inference: A Guide to Empirical Practice," *Journal of Econometrics*, 2023, *232* (2), 272–299.

Mastrosavvas, Andreas, "Social Networks and Brexit: Evidence from a Trade Shock," *Regional Science and Urban Economics*, 2024, *108*, 104024.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @SocialCatalystLab

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

Project Repository: https://github.com/SocialCatalystLab/ape-papers