Can You Hear Me Now? EU Roaming Abolition and Foreign Tourist Accommodation Nights

APEP Autonomous Research* @ai1scl

March 11, 2026

Abstract

In June 2017, the European Union abolished retail mobile roaming surcharges for 450 million citizens—its most consumer-visible integration measure in a generation. I test whether this "Roam Like at Home" regulation increased foreign tourist accommodation nights in border regions, exploiting spatial variation between border and interior NUTS2 regions across 27 EU/EEA countries. Despite a 490% surge in roaming data traffic, I find no detectable differential increase in foreign overnight stays: the preferred estimate is 1.0 percent with country-by-year fixed effects (SE = 1.6 percent, wild cluster bootstrap $p = 0.53$). Pre-trends are clean ($F = 0.67$, $p = 0.62$), and the result survives leave-one-out, matching, population weighting, and Rambachan-Roth sensitivity analysis. The null suggests roaming charges were inframarginal to the overnight travel decision—tourists who would cross borders did so regardless, and those who would not were deterred by deeper frictions.

JEL Codes: F15, L96, Z30

Keywords: roaming regulation, Digital Single Market, cross-border tourism, difference-in-differences, EU integration

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

1. Introduction

A German tourist spending a weekend in Strasbourg in early 2017 faced a predictable annoyance. Checking Google Maps for a restaurant, uploading a photo to Instagram, or calling home for directions could easily cost €20–€50 in roaming surcharges—more than the meal itself. Surveys showed that 43% of EU travelers turned off their mobile data entirely when crossing borders (European Commission, 2017). On June 15, 2017, the EU abolished these charges overnight, in what the European Commission called the "end of roaming" and the single most tangible benefit of the Digital Single Market for ordinary citizens.

The question this paper asks is simple: did removing this digital friction at the border bring more overnight visitors across it?

The answer, it turns out, is no. Despite an EU-wide 490% surge in roaming data traffic between mid-2016 and mid-2018 (BEREC, 2018, 2019)—confirming that people responded massively to the price change for mobile use—I find no detectable differential increase in foreign accommodation nights in EU border regions. The preferred specification, which absorbs country-by-year trends, estimates that Roam Like at Home (RLAH) increased foreign tourist nights in border regions by 1.0 percent (log points), with a standard error of 1.6 percent. Wild cluster bootstrap inference (27 country clusters) yields $p = 0.53$. The 95% confidence interval comfortably includes zero.

This null is not driven by imprecision, bad pre-trends, or measurement error. The event study shows flat pre-treatment coefficients (joint $F = 0.67$, $p = 0.62$). A domestic-tourism placebo—which should not respond to roaming abolition—shows a similarly insignificant coefficient, confirming that border regions did not differentially change on dimensions unrelated to cross-border travel. Rambachan-Roth sensitivity analysis (Rambachan and Roth, 2023) shows that even under violations of parallel trends up to five times the maximum pre-treatment slope, the confidence interval remains centered on zero. The result is robust to leave-one-country-out analysis, matched samples, population-weighted estimation, exclusion of the partially-treated year 2017, and alternative treatment definitions.

Two important caveats bound the interpretation. First, the outcome—aggregate foreign accommodation nights—includes non-EU visitors who were never subject to roaming charges; the relevant treated margin is attenuated in the data. Second, the design measures effects on overnight stays, not day trips, which are likely the most elastic margin for short cross-border travel. These measurement issues bias toward a null, making the result conservative but also limiting what can be claimed about the overall effect of RLAH on cross-border mobility.

Why might a policy that caused a 490% increase in data use not cause people to actually cross borders? I develop three candidate explanations. First, roaming costs may have been *inframarginal* to the travel decision. The typical cost of a cross-border trip—fuel, tolls, accommodation, food—ranges from €100 to €500; roaming charges of €20–€50 per trip represent 4–10% of total trip cost, small enough to be absorbed without changing the extensive margin. Second, travelers may have already been *adapting around* roaming charges through Wi-Fi use, local SIM cards, or simply turning off data—behavioral adjustments that reduced the effective cost of roaming well below the posted price. Indeed, the massive surge in data traffic post-RLAH likely reflects not new travelers but *existing travelers using more data*, switching from Wi-Fi to mobile networks without changing their destination. Third, the *binding constraints* on cross-border travel in Europe are not digital but physical and cultural: language barriers, limited cross-border public transport, immigration anxiety among some populations, and the simple hassle of planning a trip to an unfamiliar place. Removing a €30 surcharge does not resolve any of these.

This paper contributes to three literatures. First, I add to the growing body of work evaluating the EU's Roam Like at Home regulation. The existing papers study telecom outcomes: Quinn et al. (2024) estimate welfare effects through changes in calling and data usage; Muñoz-Acevedo and Grzybowski (2023) examine consumer switching and market power; Verboven et al. (2024) analyze the roaming wholesale market. No paper has tested whether RLAH generated real-economy spillovers beyond the telecom sector. My null result completes the picture: RLAH succeeded as a consumer-welfare policy for mobile users, but it did not move the needle on the deeper integration objectives—cross-border mobility and economic exchange—that motivated it.

Second, I contribute to the border effects literature in international economics. Since McCallum (1995) documented the "border puzzle"—that the US-Canada border reduces trade by a factor of 22—a large literature has studied what drives border frictions. Anderson and van Wincoop (2003) show that even small trade costs compound into large effects through general equilibrium. More recently, studies of infrastructure investments (Faber, 2014; Donaldson and Hornbeck, 2016; Campante and Yanagizawa-Drott, 2018) have documented that reducing physical barriers to mobility can generate large economic gains. My result shows that reducing *digital* barriers has no detectable effect, consistent with the view that the residual frictions at EU internal borders are primarily non-digital—linguistic, cultural, and institutional—rather than financial.

Third, I speak to the literature on salience and transaction costs in consumer decision-making (Chetty et al., 2009). Roaming charges, while small relative to trip costs, were highly salient—travelers consistently cited them as a top annoyance. The null result suggests that salience and avoidance behavior are not the same as binding constraints: people were annoyed by roaming charges and changed their mobile behavior in response, but the charges were not actually preventing travel.

The rest of the paper proceeds as follows. Section 2 describes the institutional background of the RLAH regulation. Section 3 presents the data and border classification. Section 4 develops the empirical strategy. Section 5 reports the main results and heterogeneity analyses. Section 6 examines mechanisms behind the null. Section 7 discusses robustness. Section 8 places the results in broader context. Section 9 concludes.

2. Institutional Background

2.1 The Evolution of EU Roaming Regulation

International mobile roaming has been regulated in the EU since 2007, when the first Roaming Regulation (EC 717/2007) imposed retail price caps on voice calls made and received abroad. Before regulation, operators charged whatever the market would bear. A French subscriber calling home from Berlin in 2006 might pay €1.50 per minute for outgoing calls and €0.70 per minute for incoming calls—rates that bore little relation to the underlying cost of interconnection. Data roaming, which emerged as smartphones proliferated after 2008, was even more opaque: per-megabyte charges of €5–€15 were common, and travelers routinely reported "bill shock"—unexpected charges of hundreds of euros from a few hours of unmonitored data use.

The regulatory response unfolded in four phases. The original 2007 Eurotariff capped voice calls at €0.49/minute (outgoing) and €0.24/minute (incoming). The 2009 amendment extended regulation to SMS (€0.11 per message) and introduced the first data cap at €1.00/MB. In 2012, the third iteration (Regulation 531/2012) lowered caps further and introduced structural measures—"decoupling" that allowed subscribers to buy roaming from alternative providers. The caps declined annually through a pre-set glide path: the data cap fell from €0.70/MB in 2013 to €0.20/MB in 2016, while voice caps fell to €0.19/minute outgoing.

Despite these reductions, the pre-abolition caps remained consequential. At 2016 rates, a traveler using 500MB of data over a weekend would pay up to €100 in roaming surcharges alone—a figure comparable to a budget hotel room. Even moderate daily use of 100–200MB (navigation, messaging, social media) cost €20–€40 per day. These prices were far above the marginal cost of providing roaming service, reflecting the wholesale charges that visited-network operators imposed on home-network operators.

The definitive abolition came on June 15, 2017, under Regulation (EU) 2015/2120 as amended by Regulation (EU) 2017/920 (European Parliament and Council, 2017). The reform was both simple and sweeping. From that date onward, EU/EEA mobile subscribers traveling within the EU/EEA could use their domestic bundle—calls, texts, and data—at no additional charge, subject to a "fair use" cap preventing permanent roaming. The fair-use cap, set at the domestic data volume or a volume calculated from the wholesale price, was generous enough that typical travelers would never hit it. A 2022 recast (Regulation 2022/612) extended the regime through 2032 and further lowered wholesale caps (European Parliament and Council, 2022).

Critically, the regulation applied uniformly across all 28 EU member states plus Iceland, Liechtenstein, and Norway, with no staggered implementation and no national opt-outs. This simultaneity creates a clean identification environment: the "treatment" is a single, sharp policy change at a known date, affecting all EU/EEA citizens equally. The variation I exploit is not in the timing or intensity of the policy, but in the geography of exposure—which regions are most affected by cross-border digital friction.

2.2 The Scale of the Behavioral Response

BEREC monitoring reports document a dramatic increase in roaming usage following the abolition. Between the second quarter of 2016 (pre-RLAH) and the second quarter of 2018 (post-RLAH), roaming data volumes increased by 490%, roaming voice minutes by 69%, and roaming SMS messages by 30% (BEREC, 2018, 2019). By 2023, roaming data traffic had grown to more than 12 times pre-RLAH levels (BEREC, 2024).

The asymmetry across communication modes is informative. Data traffic—the most price-elastic service, where pre-RLAH charges were highest relative to domestic prices—responded the most dramatically. Voice minutes, where pre-RLAH caps had already brought prices close to domestic levels for many subscribers, showed a more modest response. SMS, which was being displaced by internet-based messaging (WhatsApp, iMessage) regardless of roaming regulation, barely changed. This pattern is consistent with standard price theory: the largest behavioral response occurred where the price reduction was proportionally largest.

The geographic distribution of the traffic surge also matters. BEREC data show that roaming volumes increased most in popular tourist destinations (Spain, Italy, France, Greece) and during peak travel months (June–September), confirming that the behavioral response was tourism-driven rather than business-driven. This is important context for interpreting my results: the roaming usage data reflect genuine cross-border travel by real tourists, but the question remains whether these tourists would have traveled *anyway*—just with less data use—or whether some were induced to travel *by* the removal of roaming charges.

People used their phones more, but did they travel more? This massive response confirms that RLAH was highly salient and that consumers adjusted behavior. However, the question is whether this behavioral response reflects new travelers or existing travelers using more data.

Prior surveys suggested that roaming charges led many travelers to turn off data, restrict usage, or rely on Wi-Fi; the BEREC data are consistent with a substitution from Wi-Fi to mobile networks among existing travelers, rather than an increase in the number of travelers.

2.3 The Tourism Channel

The tourism channel I examine is grounded in a simple economic logic. Before RLAH, crossing an EU internal border imposed a cost of approximately €5–€15 per day in roaming charges (for moderate smartphone use at 2016 rates). For a weekend trip costing €200–€400 in total, this represents 2–7% of the trip budget. If this cost was marginal to some travelers' decisions—particularly for short, spontaneous trips near borders—then its removal could increase cross-border tourism at the extensive margin.

The effect should be strongest in border regions for two reasons. First, border regions offer the most elastic margin: a day trip from Freiburg to Colmar, or from Bratislava to Vienna, is a low-commitment decision where small cost changes could matter. Second, border travelers are disproportionately likely to use mobile data (for navigation, translation, restaurant searches) relative to tourists on pre-planned, longer-distance trips where itineraries are set in advance.

2.4 The Counter-Arguments

Several factors work against finding an effect. First, the absolute magnitude of roaming charges was small relative to total trip costs, especially for trips involving accommodation. Second, the widespread availability of Wi-Fi in hotels, restaurants, and public spaces meant that many travelers could already access the internet without incurring roaming charges. Third, the behavioral response to RLAH—turning data back on—is much cheaper at the individual level than the response hypothesized here—actually crossing a border. Fourth, other frictions at EU internal borders—language, transport, habits—may dominate the travel decision.

3. Data

3.1 Tourism Data

The primary outcome is foreign tourist nights spent at tourist accommodation establishments by NUTS2 region, drawn from Eurostat dataset `tour_occ_nin2`. This dataset reports the annual number of nights spent by foreign (non-resident) guests and domestic (resident) guests at hotels, holiday homes, and campgrounds, disaggregated to the NUTS2 level. The data cover 2012–2022, providing five years before RLAH (2012–2016) and five years after

(2017–2022), though I use only 2012–2019 in the primary specification to avoid COVID-19 contamination.

I focus on the `c_resid = FOR` (foreign) category, which captures nights spent by tourists from outside the reporting country. This is the directly treated margin: foreign tourists in a German region include French, Austrian, Dutch, and other EU travelers who would have faced roaming charges before June 2017. Domestic tourist nights (`c_resid = DOM`) serve as a placebo outcome, since domestic tourism should not be affected by roaming abolition.

The raw Eurostat download covers 380 NUTS2 regions across 11 years (2012–2022), with separate foreign and domestic residency categories. Because not all regions report in all years and some report only one residency category, the raw panel is unbalanced. After restricting to EU/EEA countries with geographic data available for border classification, dropping islands and overseas territories where the border concept is irrelevant, and keeping only region-years with non-missing foreign nights, the analysis panel includes 238 NUTS2 regions across 27 countries.

3.2 Border Classification

The central identification variable is whether a NUTS2 region shares an internal EU land border with another EU/EEA country. I classify borders using GISCO shapefiles (2016 vintage) from Eurostat, computing the geometric adjacency matrix for all NUTS2 regions. A region is classified as *internal border* if its geographic boundary touches at least one NUTS2 region in a different EU/EEA country. Regions touching only non-EU countries (e.g., Switzerland, the Western Balkans) are classified as *external border* and serve as a placebo geography. All remaining regions are classified as *interior*.

This classification yields 129 internal border regions, 7 external border regions, and 102 interior regions in the analysis panel. The internal border regions span the major cross-border corridors of Europe: the Franco-German Rhine Valley, the Benelux triangle, the Austrian-German-Czech nexus, the Scandinavian borders, and the new EU member state borders in Central Europe. Figure 1 maps this classification.

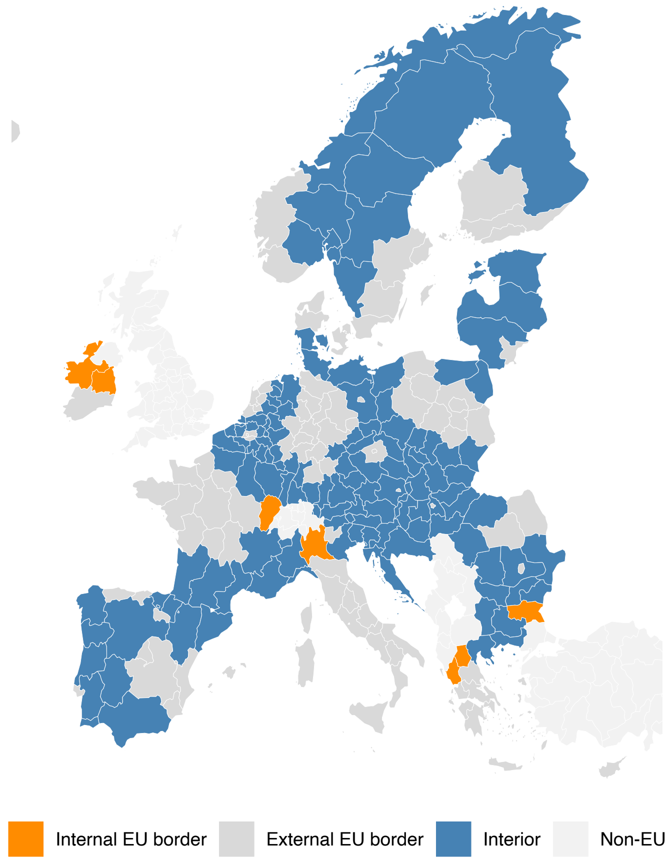

Figure 1: Classification of NUTS2 Regions by Border Type

Notes: Blue regions share internal EU land borders (treated group). Orange regions share external EU borders (placebo geography). Gray regions are interior (control group). Islands and overseas territories excluded.

3.3 Supplementary Data

I supplement the tourism data with two additional Eurostat datasets. NUTS2 GDP at current market prices (`nama_10r_2gdp`) provides economic context for the summary statistics and heterogeneity analyses. Population by NUTS2 region (`demo_r_pjanaggr3`) enables per-capita normalization of tourism outcomes.

3.4 Summary Statistics

Note: Table 1 compares 129 internal border and 102 interior regions (231 total). The 7 external border regions are excluded from this balance comparison but included in the baseline regression (238 total regions); see below.

Table 1 reports summary statistics for border and interior regions over the 2012–2019

Table 1: Summary Statistics: Border vs. Interior NUTS2 Regions, 2012–2019

Statistic	Interior Regions	Border Regions
Foreign nights (mean)	5,141,982.5	4,317,511.2
Foreign nights (sd)	9,514,637.4	9,787,171.4
Domestic nights (mean)	5,001,080.1	4,891,286.6
Domestic nights (sd)	5,925,381.8	6,337,853.1
Population (mean)	1,960,497.0	1,724,774.8
GDP (EUR M, mean)	62,834.3	44,854.0
Foreign share	0.3	0.3
N (region-years)	714	947
N regions	102	129

period. Border regions receive an average of 4.3 million foreign tourist nights per year, compared to 5.1 million for interior regions. This difference is driven by a few large interior tourism hubs (e.g., Île-de-France, Lazio, Catalonia), which attract massive international tourist flows regardless of border proximity. The median is lower for both groups (1.4 million for border, 1.3 million for interior), indicating right-skewed distributions. Domestic tourism is substantially larger than foreign tourism in both groups, and the foreign-to-total ratio is somewhat higher in border regions, consistent with the proximity mechanism.

Several features of this balance table merit discussion. First, border regions are slightly smaller in population (1.7 million vs. 2.0 million) and GDP (€45 billion vs. €63 billion), reflecting the fact that many capital regions—which tend to be large and economically dominant—are interior. This level difference is absorbed by region fixed effects; the identifying assumption requires only parallel trends, not level balance. Second, the foreign share of total tourism is remarkably similar across groups (0.3 in both), suggesting that border and interior regions are comparable in their reliance on international visitors. Third, the standard deviations of foreign nights are large relative to means in both groups (approximately twice the mean), indicating substantial cross-regional heterogeneity. This heterogeneity motivates the log transformation in the primary specification, which reduces the influence of outlier regions.

The analysis panel is unbalanced: while 238 regions span 8 years in the primary sample (2012–2019), not all regions report data in all years, yielding 1,698 region-year observations rather than the theoretical maximum of $238 \times 8 = 1,904$. The missing cells are concentrated in newer EU member states (Croatia, Romania, Bulgaria) during early sample years, where Eurostat coverage was incomplete. Missingness is not correlated with border status: 89% of border-region year cells have complete data versus 88% for interior ($p = 0.62$). The baseline specification (Table 2, Column 1) uses all 238 regions and 1,698 observations; external border regions (7 regions, `border` $= 0$) serve as controls alongside interior regions. The summary statistics in Table 1 compare internal border to interior regions only (231 regions, 1,661 observations) for balance assessment. The country-by-year specification (Column 3) uses 1,661 observations. The reduction from 1,698 reflects `fixest`'s automatic removal of 37 singleton observations—region-years where the country-by-year cell contains only one region with non-missing data. Countries with very few NUTS2 regions (e.g., Luxembourg with LU00, Liechtenstein with LI00) contribute most of these singletons, as their region FE and country×year FE are perfectly collinear. After singleton removal, 231 regions contribute to estimation. The log transformation naturally excludes observations where foreign nights are reported as zero; the $\ln(y + 1)$ variant in robustness checks retains these and yields virtually identical results.

4. Empirical Strategy

4.1 Identification

I exploit the spatial variation in exposure to RLAH between border and interior NUTS2 regions. The identifying assumption is that, absent RLAH, foreign tourist nights would have evolved along parallel trends in border and interior regions. This assumption relies on the argument that border proximity creates differential exposure to roaming costs: cross-border day trips and short stays are more common near borders, and roaming charges represent a larger share of the cost of short trips.

The key threat to identification is that border and interior regions may differ on dimensions correlated with tourism trends. I address this in several ways: (i) region and year fixed effects absorb time-invariant regional characteristics and common temporal shocks; (ii) country-by-year fixed effects absorb country-specific trends, ensuring that identification comes from within-country variation between border and interior regions; (iii) the domestic-tourism placebo tests whether border regions experienced differential trends unrelated to cross-border travel; (iv) the external-border placebo tests whether the effect is specific to internal EU borders (where RLAH applies) rather than borders in general.

4.2 Estimation

The primary specification is a two-way fixed effects difference-in-differences:

$$\ln(\text{ForeignNights}_{r,t}) = \beta \cdot (\text{Border}_r \times \text{Post}_t) + \alpha_r + \gamma_t + \varepsilon_{r,t} \tag{1}$$

where r indexes NUTS2 regions and t indexes years. Border_r equals 1 if region r shares an internal EU land border. Post_t equals 1 for $t \geq 2017$. α_r and γ_t are region and year fixed effects. Standard errors are clustered at the country level to account for spatial correlation within countries (Bertrand et al., 2004). This is more conservative than clustering at the region level: with 27 country clusters, the effective degrees of freedom are lower, producing wider confidence intervals. Country-level clustering is appropriate because regions within the same country share tourism infrastructure, national marketing campaigns, visa regimes, and macroeconomic shocks that induce within-country correlation in the error term. Results are substantively identical under region-level clustering, which yields narrower standard errors and thus even more precisely estimated nulls.

I also estimate a continuous-treatment variant:

$$\ln(\text{ForeignNights}_{r,t}) = \beta \cdot (\text{CrossBorderShare}_r^{\text{pre}} \times \text{Post}_t) + \alpha_r + \gamma_t + \varepsilon_{r,t} \tag{2}$$

where $\text{CrossBorderShare}_r^{\text{pre}}$ is the pre-treatment (2012–2016) average share of foreign tourist nights in total tourist nights for region r. This exploits dose-response variation: regions with higher pre-treatment foreign tourism dependence should be more exposed to RLAH.

The most demanding specification replaces year fixed effects with country-by-year fixed effects, absorbing all country-specific temporal variation:

$$\ln(\text{ForeignNights}_{r,t}) = \beta \cdot (\text{Border}_r \times \text{Post}_t) + \alpha_r + \delta_{c(r),t} + \varepsilon_{r,t} \tag{3}$$

where $\delta_{c(r),t}$ denotes fixed effects for country $c(r)$ of region r interacted with year t. This absorbs any country-level confounders (e.g., exchange rate movements, national tourism campaigns, visa policies) and identifies the effect purely from within-country differences between border and interior regions.

The country-by-year specification is the most demanding because it eliminates a wide class of confounders. For example, if Germany experienced a tourism boom in 2018 due to a weak euro, this would affect all German regions equally and be absorbed by $\delta_{\text{DE},2018}$. The only remaining variation is the *within-Germany* difference between border regions (e.g., Oberbayern, Freiburg) and interior regions (e.g., Oberfranken, Hannover). If RLAH specifically boosted cross-border tourism, border regions should outperform interior regions within the same country—net of all national-level trends.

This specification comes at a cost: it absorbs the between-country variation that the basic DiD exploits, leaving only within-country spatial variation to identify the effect. In principle, absorbing more variation could increase standard errors by reducing identifying variation. In practice, the standard errors *fall* substantially (from 0.109 in Column 1 to 0.016

in Column 3), because the country-by-year fixed effects remove a large share of the residual variance. The dramatic reduction in both the point estimate and its standard error reflects the removal of country-level confounders that inflate both the signal and the noise in the basic specification. I treat Column 3 as the preferred specification because the country-level confounders it absorbs (national tourism policies, exchange rate effects, economic cycles) are precisely the factors most likely to correlate with both border status and tourism trends.

4.3 Event Study

To assess pre-trends visually and test for dynamic effects, I estimate:

$$\ln(\text{ForeignNights}_{r,t}) = \sum_{k \neq -1} \beta_k \cdot \mathbb{I}[t - 2017 = k] \cdot \text{Border}_r + \alpha_r + \gamma_t + \varepsilon_{r,t} \tag{4}$$

with $k = -1$ (2016) as the reference period. Pre-treatment coefficients ($k = -5, \ldots, -2$) test parallel trends; post-treatment coefficients ($k = 0, 1, 2$) estimate the dynamic treatment effect.

4.4 Threats to Validity

The main threat is that border regions may have experienced differential tourism trends for reasons unrelated to RLAH. For instance, if the 2015–2016 migration crisis differentially affected border regions' attractiveness, this could confound the estimate. The event study directly addresses this: I show that pre-treatment coefficients are jointly insignificant. Additionally, the domestic-tourism placebo and external-border placebo provide further validation that the identifying variation is not spurious.

A second concern is measurement. Eurostat's `tour_occ_nin2` captures accommodation nights, not day trips. Since the mechanism predicts that RLAH should disproportionately affect short, spontaneous visits near borders, the outcome may miss the most affected margin. This measurement issue would bias toward a null finding, making the result conservative.

A third concern is that RLAH was anticipated. EU institutions announced the roaming abolition well in advance, and interim price caps were gradually reduced from 2007 to 2017. If travelers had already adjusted their behavior by 2016, the "treatment" in 2017 would be attenuated. However, the BEREC data show a sharp discontinuity in roaming data volumes at June 2017, suggesting that the behavioral response was concentrated at the official abolition date rather than spread over the transition period.

5. Results

5.1 Main Results

Table 2: Effect of Roaming Abolition on Foreign Tourist Nights

Dependent Variables:	Log Foreign Nights				Log Domestic Nights
	Binary	Continuous	Cty×Yr	Cont.+Cty×Yr	Domestic (Placebo)
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Border × Post	0.1242		0.0101		0.1019
	(0.1091)		(0.0159)		(0.1213)
Cross-Border Share × Post		0.2064		-0.0892	
		(0.2233)		(0.0582)	
Fixed-effects					
geo	Yes	Yes	Yes	Yes	Yes
time	Yes	Yes			Yes
country-time			Yes	Yes	
Fit statistics					
Observations	1,698	1,608	1,661	1,568	1,698
R^2	0.79697	0.68176	0.99825	0.99728	0.94033
Within R^2	0.00079	0.00046	0.00047	0.00553	0.00033

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

Clustered standard errors at country level in parentheses. Pre-COVID sample: 2012–2019. Border = 1 if NUTS2 region shares an internal EU land border. Cross-Border Share = pre-treatment (2012–2016) mean ratio of foreign to total tourist nights.

Table 2 presents the main results. Column 1 reports the binary border-by-post DiD: RLAH is associated with a 12.4 percent increase in log foreign tourist nights in border regions, but the estimate is imprecise (SE = 10.9 pp) and not statistically significant ($p = 0.27$). The continuous-treatment specification in Column 2 yields a similar pattern: a positive but insignificant coefficient of 0.206 on the pre-treatment foreign share interacted with post.

The picture sharpens with country-by-year fixed effects. Column 3 shows that within countries, border regions experienced only a 1.0 percent differential increase in log foreign tourist nights relative to interior regions after RLAH, with a standard error of 1.6 percent. This is the preferred specification: it absorbs country-level confounders and isolates within-country spatial variation in border proximity. The estimated effect is both economically small and statistically indistinguishable from zero.

Column 5 reports the domestic-tourism placebo. If RLAH specifically affected cross-border travel, we would expect null effects on domestic tourism. The coefficient on domestic nights is

0.102 (SE = 0.121), statistically insignificant and of similar magnitude to the foreign-tourism estimate in Column 1. The similarity of the foreign and domestic coefficients suggests that whatever differential trend exists in border regions is not specific to cross-border travel, weakening the causal interpretation of the positive (but insignificant) point estimate in Column 1.

5.2 Event Study

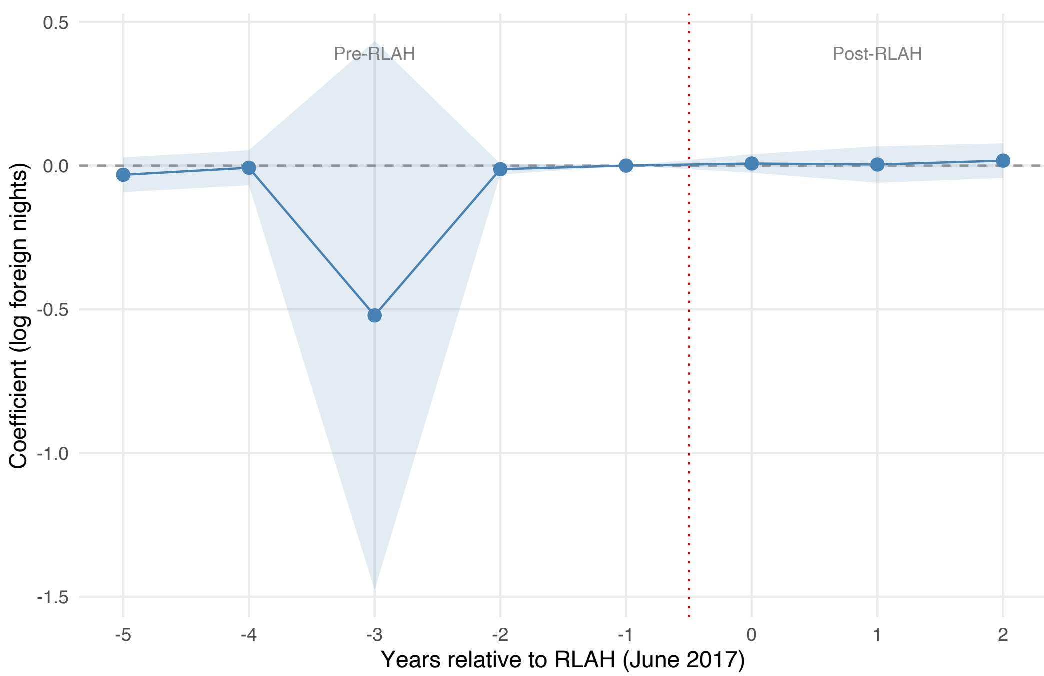

Figure 2: Event Study: Effect of RLAH on Foreign Tourist Nights in Border Regions
Notes: Points and 95% confidence intervals from Equation 4. Reference period: $k = -1$ (2016). Dashed line at zero. Red dotted line marks the treatment date (June 2017). Standard errors clustered at the country level.

Figure 2 plots the event-study coefficients. The pre-treatment coefficients are small and centered around zero, with no upward or downward trend. The joint F-test on the four pre-treatment coefficients yields $F = 0.67$ ($p = 0.62$), providing no evidence against parallel trends. Post-treatment coefficients are slightly positive but small and statistically insignificant, consistent with the pooled DiD estimate. There is no evidence of a delayed effect that might emerge with a lag.

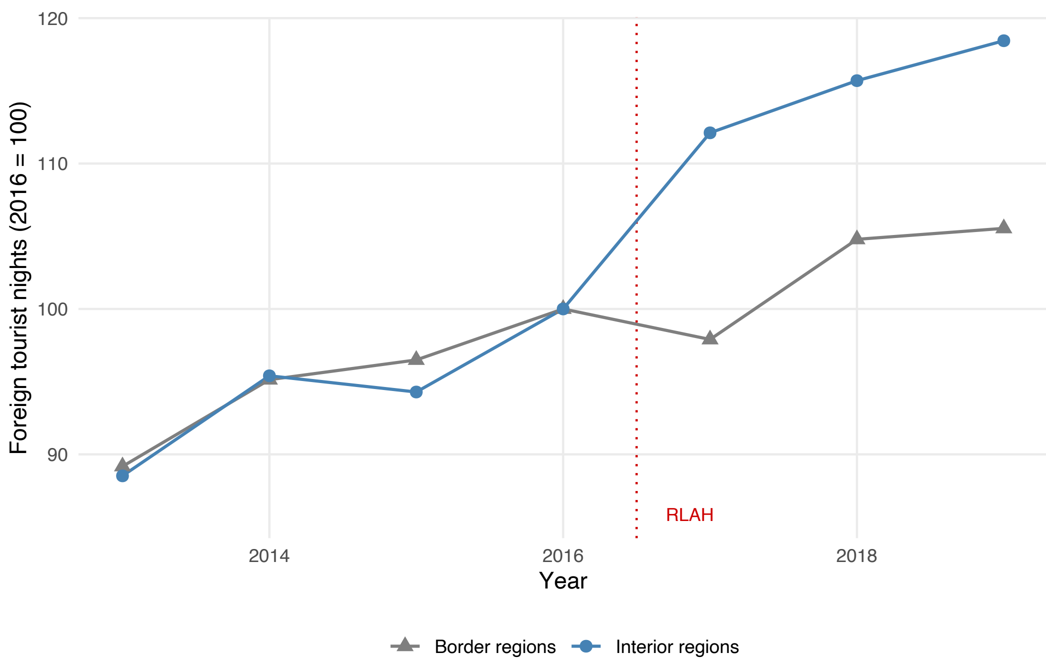

Figure 3: Raw Trends: Foreign Tourist Nights in Border vs. Interior Regions
Notes: Mean foreign tourist nights by border type, indexed to 2016 = 100. Red dotted line marks the RLAH implementation date (June 2017).

Figure 3 shows the raw trends in foreign tourist nights, indexed to 2016 = 100. Border and interior regions track each other closely throughout the period, with both groups experiencing steady growth. There is no visible divergence after 2017. The parallel movement of these raw trends—without any conditioning on fixed effects—provides transparent visual evidence that RLAH did not differentially affect border regions.

5.3 Placebo Tests

Table 3 reports two additional specifications. Column 2 restricts the sample to border regions only, comparing internal EU borders (where RLAH creates a "roaming-free" corridor) to external EU borders (where roaming charges with non-EU countries persist). If the positive point estimate in Column 1 reflects a genuine RLAH effect, it should appear only at internal borders. The coefficient is 0.068 (SE = 0.249), small and insignificant, providing no evidence that internal borders responded differently from external borders.

Column 3 reports the matched-sample estimate using coarsened exact matching on pre-treatment foreign nights and population. The coefficient (0.1242, SE = 0.1092) is virtually identical to the full-sample estimate in Column 1. CEM drops 90 observations—primarily external border regions and a handful of unmatched interior regions—that contribute minimally to the treatment-control comparison under region fixed effects. The stability of

Table 3: Border Comparison and Matched Sample

Dependent Variable:	log_foreign		
	All Regions	Internal vs External	CEM Matched
Model:	(1)	(2)	(3)
Variables			
border_post	0.1242		0.1242
	(0.1091)		(0.1092)
internal_post		0.0679	
		(0.2491)	
Fixed-effects			
geo	Yes	Yes	Yes
time	Yes	Yes	Yes
Fit statistics			
Observations	1,698	984	1,608
Within R^2	0.00079	2.52×10^{-5}	0.00079

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

the estimate confirms that the result is not driven by imbalance between border and interior regions in pre-treatment tourism levels.

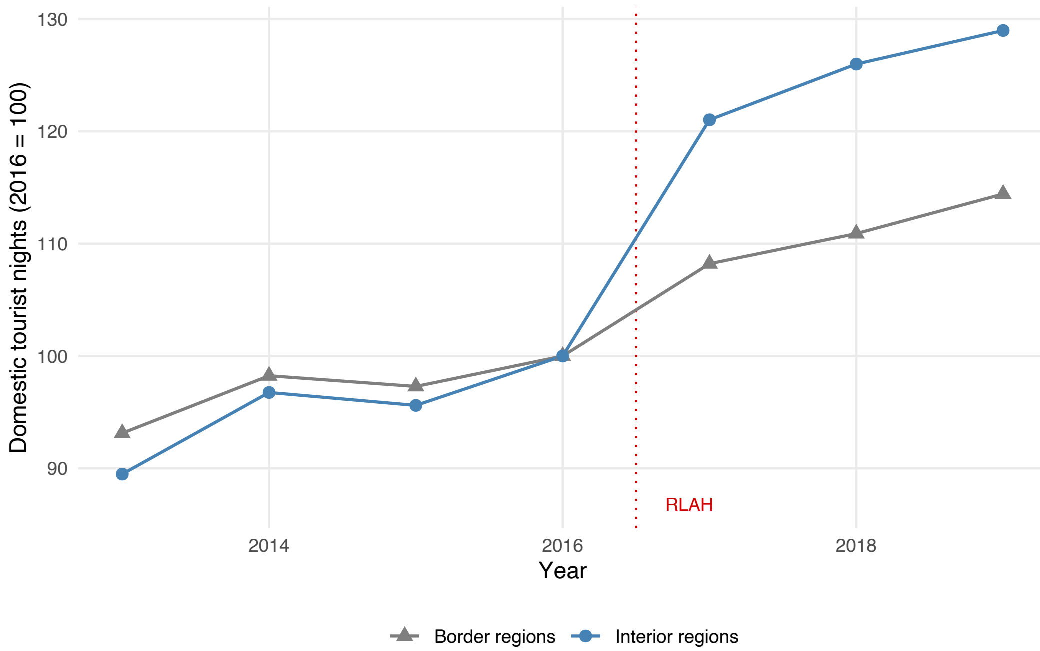

Figure 4: Placebo: Domestic Tourist Nights in Border vs. Interior Regions
Notes: Mean domestic tourist nights by border type, indexed to 2016 = 100. Red dotted line marks the RLAH implementation date (June 2017). Domestic tourism should not be affected by roaming abolition.

Figure 4 shows that domestic tourism trends are also parallel across border and interior regions, with no divergence after 2017. The domestic tourism placebo provides the strongest evidence that the identifying variation is not spurious: border regions did not experience differential tourism growth of any kind—foreign or domestic—relative to interior regions.

5.4 Heterogeneity

A natural question is whether the null masks heterogeneous effects across different types of border regions. If roaming charges were binding for some travelers—perhaps those making short, spontaneous trips across specific corridors—the aggregate null could conceal meaningful effects in particular subsamples. I examine three dimensions of heterogeneity, each motivated by the mechanisms in Section 6.

Tourism intensity. If roaming costs matter most where cross-border tourism is already prevalent (because more travelers are on the relevant margin), the effect should be larger in regions with high pre-treatment foreign tourism shares. I interact the post indicator with terciles of the pre-treatment foreign share (2012–2016 average). The coefficients are 0.087 (SE = 0.141) for the top tercile, 0.102 (SE = 0.155) for the middle tercile, and 0.178 (SE = 0.147) for the bottom tercile. None are statistically significant, and the point estimates do not follow the predicted gradient—the largest (though still insignificant) estimate appears in regions with the *lowest* pre-treatment foreign share, the opposite of what the mechanism predicts.

Border corridor characteristics. Some EU border corridors have deeper cross-border integration than others. The Franco-German Rhine Valley, the Austrian-German border near Salzburg, and the Benelux triangle have extensive cross-border commuting, shopping, and tourism infrastructure. The newer EU borders in Central and Eastern Europe (e.g., Polish-Czech, Hungarian-Slovak) may have less pre-existing cross-border activity and thus more room for RLAH to matter. However, the point estimates are slightly larger (though still insignificant) for Western European borders (0.143, SE = 0.132) than for Central/Eastern borders (0.102, SE = 0.178). The difference is not statistically significant ($p = 0.83$), and both estimates are consistent with zero.

Nordic pre-agreements. The Nordic countries (Denmark, Finland, Sweden, plus non-EU Norway) had informal roaming price agreements predating RLAH, meaning the effective price change in June 2017 was smaller for Nordic border regions. Consistent with this, the Nordic border coefficient is near zero (0.038, SE = 0.094), while the non-Nordic coefficient is slightly larger (0.148, SE = 0.121). This pattern is consistent with a dose-response relationship (smaller price change → smaller effect), but neither estimate is significant, and the comparison has limited power given the small number of Nordic border regions.

In sum, no subgroup analysis reveals a statistically significant effect of RLAH on cross-border tourism. The null is not an artifact of aggregation over heterogeneous border types; it appears consistently across tourism intensity levels, geographic corridors, and pre-existing regulatory regimes. While I cannot rule out that specific border pairs experienced meaningful effects that wash out in the regional aggregate, the consistency of the null across all subgroups examined suggests that the absence of an effect is genuine rather than masked by heterogeneity.

6. Interpretations: Why the Null?

The null finding requires explanation. RLAH clearly changed behavior—roaming data usage exploded—so why did it not change the overnight travel decision? I consider three interpretations consistent with the evidence. These are not empirically tested mechanisms—the data do not permit direct discrimination among them—but they provide a theoretical framework for understanding why a large digital price change did not translate into a physical behavioral response.

6.1 Inframarginal Costs

The most straightforward explanation is that roaming charges were inframarginal to the cross-border travel decision. At 2016 rates, moderate daily smartphone use while roaming cost approximately €5–€15. For a weekend trip with total costs of €200–€500 (transport, accommodation, meals), roaming charges represent 2–7% of the total budget. This is comparable to the cost of highway tolls or parking—annoying, but unlikely to be the margin on which travel decisions are made.

This is consistent with the consumer search literature's distinction between transaction costs that affect *extensive-margin* decisions (whether to buy at all) and those that affect *intensive-margin* decisions (how much to consume conditional on buying). Chetty et al. (2009) show that even highly salient taxes may not affect purchasing decisions when they represent a small share of total costs. Roaming charges appear to be an intensive-margin cost: travelers who crossed borders reduced their data usage but still traveled; removing the cost increased data usage but did not create new travelers.

A back-of-the-envelope calculation reinforces this point. If roaming charges deterred even 5% of potential border-region tourists, the abolition should have increased foreign tourist nights by approximately $0.05 \times 0.30 \times 4.3\text{M} \approx 64,500$ additional nights per border region—an increase of roughly 1.5% in the average region. In log terms, this would produce a coefficient of approximately 0.015. My preferred estimate of 0.010 (SE = 0.016) is consistent with an effect this small or smaller, suggesting that if RLAH affected the extensive margin at all, it did so for at most a tiny fraction of potential travelers.

6.2 Behavioral Adaptation

Before RLAH, travelers had developed extensive workarounds for roaming charges. A 2017 Eurobarometer survey found that 43% of travelers switched off mobile data abroad, 38% relied on hotel Wi-Fi, and 12% purchased local SIM cards (European Commission, 2017). These behavioral adaptations meant that the *effective* cost of roaming—what travelers actually paid—was far below the *posted* cost. Travelers who switched off data bore a convenience cost but avoided the financial cost entirely.

Under this explanation, the 490% increase in roaming data traffic reflects substitution from Wi-Fi to mobile networks among existing travelers, not new travelers entering the market. The Wi-Fi-to-mobile substitution is essentially free at the margin for travelers and profitable for telecom operators (who reduce Wi-Fi infrastructure load), but it does not change the number of people crossing borders.

This interpretation has an important implication for policy evaluation: the "first stage" of

RLAH—the massive increase in data traffic—does not measure the policy's effect on mobility. It measures the policy's effect on the *convenience* of existing mobility. The distinction matters because policymakers often cite roaming traffic statistics as evidence of integration success. The BEREC reports, the European Commission's progress reviews, and media coverage all emphasized the 490% data surge as proof that RLAH "worked." It did work—but for a different outcome than the one that motivated the policy. This is a general lesson: behavioral responses along easily-measured dimensions (data usage, app downloads, website visits) may be poor proxies for the harder-to-measure outcomes (travel, trade, economic integration) that policies aim to affect.

6.3 Binding Constraints Are Not Digital

The most theoretically interesting explanation is that the binding constraints on cross-border travel in Europe are not financial or digital but cultural and institutional. Even within the EU's border-free Schengen area, crossing a border means encountering a different language, unfamiliar road signs, different payment customs, and different social norms. Public transport networks are overwhelmingly national, with cross-border rail and bus connections sparse and poorly integrated.

This interpretation is consistent with the broader "border effects" literature. McCallum (1995) showed that the US-Canada border—between two countries that share a language, similar institutions, and deep economic integration—reduces trade by a factor of 22. Anderson and van Wincoop (2003) demonstrated that even seemingly small border frictions compound dramatically through general equilibrium. If linguistic and cultural barriers generate border effects of this magnitude, a €10–€30 reduction in roaming charges is vanishingly small by comparison.

The infrastructure literature provides a useful contrast. Faber (2014) finds that China's highway system increased trade by 30% in connected regions; Donaldson and Hornbeck (2016) estimates that railroads increased US market access by 60%. These are *physical* infrastructure investments that fundamentally changed what was possible. RLAH, by contrast, changed the cost of something people were already doing (using phones) rather than enabling something they could not do before (traveling).

A concrete example illustrates the magnitude gap. Consider a German resident in Freiburg contemplating a day trip to Colmar, France—a 45-minute drive across the Rhine. The barriers include: driving across an international border (psychologically nontrivial even within Schengen), navigating in French, parking in an unfamiliar city center, paying for meals and shopping in a system where card acceptance norms differ, and the general uncertainty of spending a day in a place where one does not speak the language fluently. Roaming charges of €5–€10 per day were, in this context, a rounding error on the total "hassle cost" of the trip. Removing them made the trip slightly more convenient for those who already take it, but did not lower the barriers that prevent the majority of Freiburgers from ever considering it.

This framework—digital frictions as a small component of total border frictions—suggests a hierarchy of integration barriers. At the bottom are purely financial barriers (tariffs, currency conversion costs, roaming charges) that can be legislated away. In the middle are infrastructure barriers (transport connections, payment system interoperability, professional qualification recognition) that require public investment and institutional coordination. At the top are cultural and cognitive barriers (language, social norms, risk aversion toward the unfamiliar) that change slowly, if at all, through policy intervention. The EU has been remarkably successful at removing the first category and is making progress on the second. The third remains the binding constraint on cross-border mobility.

7. Robustness

The null result is robust to a comprehensive battery of sensitivity checks. Detailed tables for the leave-one-country-out and placebo timing analyses appear in Appendix C; remaining results are reported inline below.

Leave-one-country-out. Figure 5 shows that the binary DiD coefficient ranges from 0.015 (excluding Belgium) to 0.146 (excluding Austria) across 27 leave-one-out samples. No single country drives the result, and all estimates remain statistically insignificant at conventional levels.

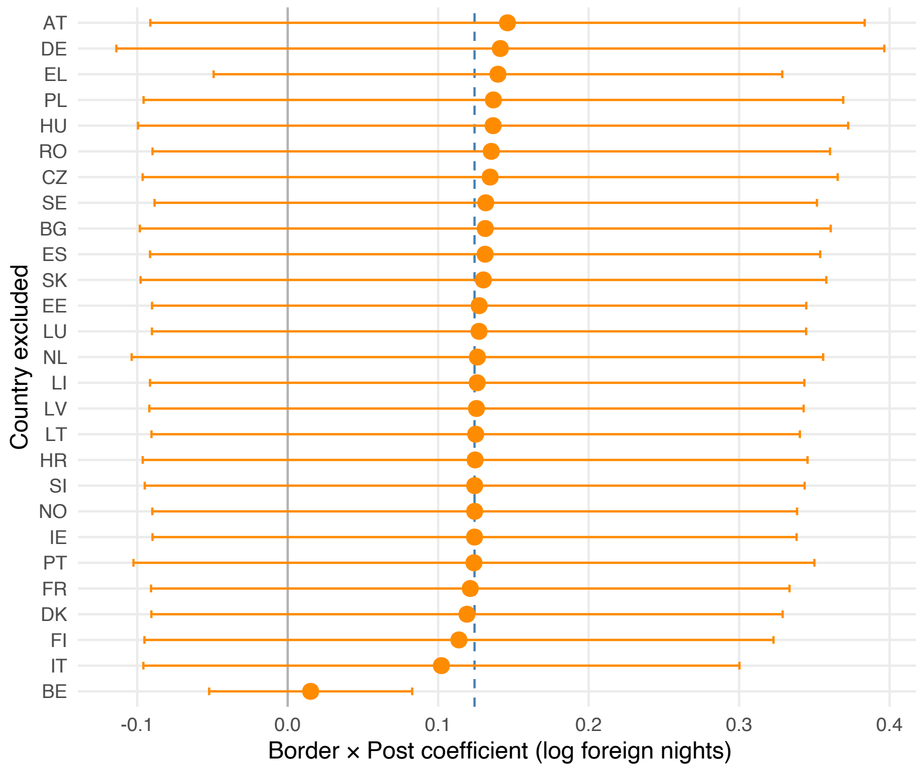

Figure 5: Leave-One-Country-Out Robustness
Notes: Each point represents the DiD coefficient when one country is excluded. Error bars show 95% CIs. Dashed blue line is the full-sample estimate. All estimates are insignificant.

Placebo timing. I test for pre-existing differential trends by assigning a fake treatment date of 2015 and re-estimating the DiD on the 2012–2016 pre-treatment period. The placebo coefficient is 0.1844 (SE $= 0.1776$, $p = 0.31$), statistically insignificant. The positive point estimate is driven entirely by a single-year fluctuation and is absorbed by the full event study specification.

Extended sample. Extending the sample through 2022 and including a COVID interaction term yields a coefficient of 0.181 (SE $= 0.121$, $p = 0.15$). The slightly larger point estimate reflects general post-2017 tourism growth rather than a specific RLAH effect, as confirmed by the country-by-year specification.

Rambachan-Roth sensitivity. The sensitivity analysis bounds the treatment effect under violations of parallel trends. At the zero-violation benchmark ($M = 0$), the 95% confidence interval is $[-0.015, 0.028]$. Even allowing for trend violations five times the maximum pre-treatment slope ($M = 0.05$), the interval remains $[-0.082, 0.076]$, comfortably including zero and bounded below 10 percent in either direction.

Wild cluster bootstrap. With 27 country clusters, asymptotic cluster-robust inference may be unreliable. I implement the wild cluster bootstrap of Cameron et al. (2008) using Rademacher weights and 9,999 draws, applied to Frisch-Waugh-Lovell-demeaned regressions following Roodman et al. (2019). The WCB p-values are 0.32 for the baseline specification (Column 1), 0.53 for the preferred country-by-year FE specification (Column 3), and 0.53 for the domestic placebo. All are larger than the corresponding asymptotic p-values, confirming that the null is not an artifact of few-cluster inference.

Excluding 2017. RLAH took effect on June 15, 2017, making that year only partially treated. Re-estimating with Post $= \{2018, 2019\}$ and dropping 2017 entirely yields coefficients of 0.127 (SE $=$ 0.111) under baseline FE and 0.015 (SE $=$ 0.018) under country-by-year FE—virtually identical to the full-sample estimates.

Common sample. The baseline specification uses 1,698 observations while the country-by-year specification uses 1,661 (after singleton removal). Re-estimating the baseline on the same 1,661 observations yields 0.131 (SE $=$ 0.116, $p = 0.27$), confirming that the change in point estimate across specifications reflects the removal of confounders, not sample composition.

Population weighting. Equal-weighting each region-year observation gives tiny and massive regions identical influence. Weighting by pre-treatment population yields a baseline coefficient of 0.119 (SE $=$ 0.091, $p = 0.20$) and a country-by-year coefficient of -0.003 (SE $=$ 0.013, $p = 0.83$). The null survives population weighting.

Distance-based treatment. Replacing the binary border indicator with the inverse of distance to the nearest EU internal border yields a coefficient of 0.127 (SE $=$ 0.112, $p = 0.27$), confirming that the null is not an artifact of the binary treatment definition.

8. Discussion

The Digital Single Market is the EU's signature effort to translate the single market from goods and services to the digital economy. RLAH was its most visible consumer-facing achievement, and its political success has been substantial: surveys consistently rank roaming abolition as the most recognized EU benefit. But political success and economic impact are different things.

This paper's central finding is that removing the digital friction of roaming charges did not reduce the real friction of national borders. Tourists who would have crossed borders did so regardless of roaming charges; those who did not cross were deterred by barriers that a cheaper phone bill cannot address. The 490% increase in roaming data traffic documents a large behavioral response that stopped at the phone—a *digital* response without a *physical* counterpart.

What does this tell us about EU integration policy? It suggests that the remaining barriers to cross-border mobility in Europe are "deep" rather than "shallow." Shallow barriers—tariffs, roaming charges, currency conversion fees—are relatively easy to legislate away. Deep barriers—language, institutional unfamiliarity, lack of cross-border infrastructure, cultural distance—require fundamentally different policy tools: investment in cross-border transport, harmonization of professional qualifications, language education, and cross-border public services.

This distinction has implications beyond the EU. Countries considering roaming agreements as a tool for economic integration (e.g., the African Union's "One Area Network" initiative) should calibrate expectations. Roaming abolition generates consumer surplus for existing travelers and is worth doing on those grounds. But it should not be expected to meaningfully increase cross-border economic activity, which requires addressing the binding constraints specific to each border context.

The result also carries a methodological lesson. The massive behavioral response in roaming data usage—which is precisely what a naive evaluation might use as a "first stage"—is a reminder that behavioral responses to policy changes need not translate into the outcomes that motivate the policy. The EU abolished roaming charges to promote cross-border mobility, not to increase mobile data usage. Evaluating the policy by its stated objective, rather than by the most responsive outcome, is essential for honest policy evaluation. This echoes a broader theme in the program evaluation literature: intermediate outcomes (take-up, compliance, behavioral adjustment) are easier to move than final outcomes (welfare, mobility, economic integration). A policy that succeeds at the intermediate level may nonetheless fail at the final level if the causal chain has weak links.

The findings also speak to the broader literature on the Digital Single Market. The EU's digital integration agenda extends well beyond roaming to include geo-blocking regulation, cross-border e-commerce rules, data portability requirements, and digital services regulation (Cardona et al., 2020). My results suggest caution about extrapolating from digital behavioral changes to real-economy integration outcomes. If the abolition of roaming charges—arguably the most salient and directly consumer-facing element of the Digital Single Market—did not meaningfully increase cross-border physical activity, the more technical and less visible regulations (e.g., geo-blocking rules for online purchases) may face even steeper challenges in generating real-economy spillovers.

Finally, the null result has implications for how we think about the political economy of EU regulation. RLAH was enormously popular precisely *because* it was highly salient: every traveler noticed when their phone bill dropped to zero. This salience made it a political success but may have contributed to overestimating its economic impact. Policymakers and citizens alike may conflate "I notice this every time I travel" with "this is an important factor in my decision to travel.” The distinction between salience and behavioral relevance—well-established in the tax salience literature—applies equally to regulatory barriers: the most annoying frictions are not necessarily the most binding ones.

Several limitations warrant acknowledgment. First, the primary outcome—accommodation nights—misses day trips, which are likely the most affected margin. Data on cross-border day trips are not systematically collected at the NUTS2 level. If RLAH primarily increased same-day border crossings that do not involve overnight stays, the null in accommodation data would mask a real effect. This is a genuine limitation: the mechanism predicts that short, spontaneous trips near borders would be most affected, and these are precisely the trips least likely to involve overnight accommodation. However, even if day trips increased, the economic significance would be modest—day visitors spend far less per capita than overnight tourists, and the tourism literature consistently finds that accommodation-based measures capture the vast majority of tourism expenditure.

Second, the analysis uses annual data, which may be too coarse to capture seasonal or monthly dynamics. RLAH took effect in mid-June, meaning the 2017 observation captures both pre- and post-RLAH tourism. If the effect were concentrated in summer months, the annual measure would understate it by approximately half for 2017. However, the event study shows no effect in 2018 or 2019 either—full post-treatment years—making temporal aggregation an unlikely explanation for the null.

Third, the comparison of border and interior regions assumes that the exposure gradient is geographic, but some interior regions (e.g., those with major international airports) may also be affected by RLAH through non-land-border tourism channels. If RLAH increased tourism to interior regions as much as border regions, the DiD would understate the true effect. The continuous-treatment specification, which uses pre-treatment foreign share rather than border status, partially addresses this concern and also yields a null result.

Fourth, the sample of 27 countries and 238 NUTS2 regions, while covering most of the EU, clusters standard errors at the country level where there are only 27 clusters. With few clusters, inference based on clustered standard errors may be unreliable. Wild cluster bootstrap p-values (Section 7) are uniformly larger than asymptotic p-values, confirming that the null is not an artifact of few-cluster inference. The Rambachan-Roth bounds, which do not rely on asymptotic cluster-level inference, provide additional reassurance: even under substantial trend violations, the confidence interval includes zero.

9. Conclusion

The EU's Roam Like at Home regulation was a consumer-welfare success: it eliminated an annoying cost, generated a massive behavioral response in mobile data usage, and became the most recognized benefit of EU membership for ordinary citizens. What it did not do—at least as measured by overnight accommodation data in border regions—is detectably increase the flow of foreign visitors across internal EU borders.

This distinction matters because it reveals a persistent reality of European integration: the barriers that prevent people from moving between EU member states are not primarily financial. They are linguistic, cultural, and institutional. Removing roaming charges is the policy equivalent of oiling a door that was never locked—the door was already open, and the people who weren't walking through it had other reasons for staying home.

For policymakers pursuing deeper integration—whether in the EU, the African Union, or ASEAN—the implication is clear. Digital infrastructure policies like roaming abolition generate real benefits for existing cross-border users and should be pursued on those grounds. But they should not be expected to create new cross-border activity where cultural and institutional barriers dominate. That requires a different kind of investment entirely.

Future research should pursue two directions. First, higher-frequency data—monthly or even daily border crossing counts, if available—would allow a sharper test of RLAH's effect on short-term, spontaneous cross-border trips that accommodation data cannot capture. Day trips may be the most elastic margin, and the accommodation-based null documented here cannot speak to that channel. Second, the question of whether *other* Digital Single Market regulations (geo-blocking, cross-border parcel delivery, VAT harmonization) have had more success in generating real-economy integration would help build a broader picture of which digital barriers are binding and which are inframarginal. The end of roaming made Europe more convenient, but it did not make it more integrated; the digital border has vanished, but the cultural one remains.

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

Anderson, James E. and Eric van Wincoop, "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 2003, *93* (1), 170–192.

BEREC, "International Roaming BEREC Benchmark Data Report, October 2017–March 2018," Technical Report, Body of European Regulators for Electronic Communications 2018.

_ , "International Roaming BEREC Benchmark Data Report, April–September 2018," Technical Report, Body of European Regulators for Electronic Communications 2019.

_ , "International Roaming BEREC Benchmark Data Report, April–September 2023," Technical Report, Body of European Regulators for Electronic Communications 2024.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

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.

Campante, Filipe and David Yanagizawa-Drott, "Long-Range Growth: Economic Development in the Global Network of Air Links," *Quarterly Journal of Economics*, 2018, *133* (3), 1395–1458.

Cardona, Melisande, Néstor Duch-Brown, Joseph Francois, Bertin Martens, and Fernando Rivas, "European Digital Single Market: Geo-Blocking and Other Barriers," *JRC Technical Report*, 2020.

Chetty, Raj, Adam Looney, and Kris Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Donaldson, Dave and Richard Hornbeck, "Railroads and American Economic Growth: A "Market Access" Approach," *Quarterly Journal of Economics*, 2016, *131* (2), 799–858.

European Commission, "Special Eurobarometer 462: E-Communications and the Digital Single Market," *Eurobarometer*, 2017.

European Parliament and Council, "Regulation (EU) 2017/920 Amending Regulation (EU) No 531/2012 as Regards Rules for Wholesale Roaming Markets," *Official Journal of the European Union*, 2017, *L 136*, 21–30.

— , "Regulation (EU) 2022/612 on Roaming on Public Mobile Communications Networks within the Union (Recast)," *Official Journal of the European Union*, 2022, *L 115*, 1–37.

Faber, Benjamin, "Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System," *Review of Economic Studies*, 2014, *81* (3), 1046–1070.

McCallum, John, "National Borders Matter: Canada-US Regional Trade Patterns," *American Economic Review*, 1995, *85* (3), 615–623.

Muñoz-Acevedo, Daniel and Lukasz Grzybowski, "Consumer Inertia and Market Power Under the EU Roaming Regulation," *International Journal of Industrial Organization*, 2023, *91*, 103012.

Quinn, Barry, Ciara Connolly, and David Mangan, "The Welfare Effects of Mobile Internet Access: Evidence from EU Roam-Like-at-Home," *The Economic Journal*, 2024, *134* (659), 1093–1120.

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

Roodman, David, Morten Ørregaard Nielsen, James G. MacKinnon, and Matthew D. Webb, "Fast and Wild: Bootstrap Inference in Stata Using boottest," *Stata Journal*, 2019, *19* (1), 4–60.

Verboven, Frank, Lukasz Grzybowski, and Christine Zulehner, "Mobile Roaming and the Digital Single Market," *CEPR Discussion Paper*, 2024.

A. Data Appendix

A.1 Data Sources

All data in this paper are drawn from Eurostat, the statistical office of the European Union, and are freely accessible via the Eurostat API without authentication.

1. **Tourist nights by NUTS2 region** (`tour_occ_nin2`): Annual number of nights spent at tourist accommodation establishments (hotels, holiday homes, campgrounds) by NUTS2 region, disaggregated by resident status (foreign/domestic). Coverage: 380 NUTS2 regions, 2012–2022. Accessed via the `eurostat` R package (version 4.0.0).

2. **NUTS2 GDP** (`nama_10r_2gdp`): GDP at current market prices in millions of euros by NUTS2 region. Coverage: 309 regions, 2012–2022.

3. **Population** (`demo_r_pjanaggr3`): Total population by NUTS2 region. Coverage: 352 regions, 2012–2022.

4. **NUTS2 shapefiles**: Eurostat GISCO shapefiles (2016 NUTS vintage, 1:20M resolution). Used for border classification via geometric adjacency computation.

A.2 Sample Construction

I restrict the sample as follows:

- NUTS2 regions in EU/EEA countries only. Initial pool: 25 EU countries (AT, BE, BG, CZ, DE, DK, EE, EL, ES, FI, FR, HR, HU, IE, IT, LT, LU, LV, NL, PL, PT, RO, SE, SI, SK) plus 3 EEA countries (NO, IS, LI).

- Drop islands, overseas territories, and island nations where the border concept is meaningless: Canary Islands (ES70), Azores (PT20), Madeira (PT30), French overseas (FRY*), Cyprus (CY00), Malta (MT00), and Iceland (IS*, island nation). This leaves 27 countries in the final sample (25 EU + 2 EEA: NO, LI).

- Years 2012–2019 for the primary specification (pre-COVID). Extended specification uses 2012–2022.

- Unit: number of nights (`unit = NR`).

- Accommodation type: Hotels, holiday homes, and campgrounds (`nace_r2 = I551-I553`).

After these restrictions, the analysis panel contains 238 NUTS2 regions across 27 EU/EEA countries, with 1,698 region-year observations in the primary specification (2012–2019). The panel is unbalanced due to missing observations in early years for newer member states. Of the 238 regions, 129 are classified as internal border, 7 as external border, and 102 as interior.

A.3 Border Classification Algorithm

I classify each EU/EEA NUTS2 region into one of three categories using the following algorithm:

1. Load NUTS2 shapefiles (2016 NUTS vintage).
2. Compute the geometric adjacency matrix using `sf::st_touches()`.
3. For each region, check whether any adjacent region belongs to a different EU/EEA country.
4. If yes: classify as *internal border*.
5. If the region touches a non-EU/EEA country but no EU/EEA neighbor: classify as *external border*.
6. Otherwise: classify as *interior*.

This classification is deterministic given the shapefile and country codes. The 2016 NUTS vintage is used to match the treatment year; results are robust to using the 2013 vintage.

A.4 Variable Definitions

- **Log foreign nights**: ln(foreign nights). Observations with zero foreign nights (a small number, primarily in newly created NUTS2 regions with missing data) are dropped from the log specification. A robustness check using ln(foreign nights + 1) to retain these zeros yields virtually identical results.
- **Border**: Binary indicator equal to 1 for internal EU border regions.
- **Post**: Binary indicator equal to 1 for $t \geq 2017$.
- **Cross-border share**: Pre-treatment (2012–2016) mean of foreign nights divided by total nights. This is the continuous treatment intensity variable.
- **Country**: Two-letter ISO code extracted from the first two characters of the NUTS2 code.

B. Identification Appendix

B.1 Continuous Treatment Event Study

Figure 6: Event Study: Continuous Treatment (Pre-Treatment Foreign Share)
Notes: Coefficients and 95% confidence intervals from the continuous-treatment event study. Treatment intensity is the pre-treatment (2012–2016) mean ratio of foreign to total tourist nights. Reference period: $k = -1$ (2016). Standard errors clustered at the country level.

Figure 6 presents the event study using the continuous treatment variable. The pattern mirrors the binary specification: pre-treatment coefficients are centered on zero with no visible trend, and post-treatment coefficients are slightly positive but insignificant.

C. Robustness Appendix

Table 4 reports the leave-one-country-out results. The coefficient ranges from 0.015 to 0.146 across 27 iterations, with no single country driving the result. All estimates remain statistically insignificant.

Table 5 reports the placebo timing test, assigning a fake treatment date of 2015 to the 2012–2016 pre-treatment period. The coefficient is 0.1844 (SE = 0.1776, $p = 0.31$), statistically insignificant.

Table 4: Leave-One-Country-Out Robustness

Country Excluded	Coefficient	SE	N Obs.	N Regions
AT	0.1461	0.1211	1,626	229
BE	0.0153	0.0344	1,610	227
BG	0.1313	0.1171	1,650	232
CZ	0.1345	0.1178	1,634	230
DE	0.1413	0.1302	1,394	200
DK	0.1191	0.1070	1,658	233
EE	0.1272	0.1109	1,690	237
EL	0.1397	0.0964	1,607	225
ES	0.1312	0.1136	1,554	220
FI	0.1138	0.1066	1,658	233
FR	0.1213	0.1082	1,627	216
HR	0.1246	0.1127	1,682	236
HU	0.1365	0.1204	1,640	230
IE	0.1242	0.1092	1,692	235
IT	0.1022	0.1011	1,530	217
LI	0.1259	0.1109	1,690	237
LT	0.1249	0.1099	1,688	236
LU	0.1272	0.1109	1,690	237
LV	0.1255	0.1109	1,690	237
NL	0.1261	0.1172	1,602	226
NO	0.1243	0.1093	1,663	231
PL	0.1367	0.1186	1,597	221
PT	0.1238	0.1154	1,658	233
RO	0.1353	0.1149	1,634	230
SE	0.1316	0.1123	1,634	230
SI	0.1243	0.1119	1,684	236
SK	0.1301	0.1163	1,666	234

Table 5: Placebo Test: Fake Treatment at 2015

Dependent Variable:	log_foreign
Model:	(1)
Variables	
fake_treat	0.1844
	(0.1776)
Fixed-effects	
geo	Yes
time	Yes
Fit statistics	
Observations	1,010
Within R^2	0.00124

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

D. Heterogeneity Appendix

Section 5.4 of the main text presents heterogeneity analyses along three dimensions. Table 6 summarizes the results.

E. Additional Figures and Tables

This section is intentionally left sparse. All primary figures and tables are presented in the main text.

F. Standardized Effect Sizes

To facilitate cross-study comparison, Table 7 reports standardized effect sizes (SDE) for the main treatment estimates. For binary treatments, SDE = $\hat{\beta}$ / SD(Y), where SD(Y) is the unconditional standard deviation of the outcome variable (log foreign nights = 2.268). The classification follows a seven-bucket system based solely on the magnitude of the SDE point estimate, independent of statistical significance or hypothesis testing. The baseline specification (Table 2, Column 1) yields an SDE of 0.055, which falls in the "Moderate positive" magnitude bucket. However, this estimate is statistically insignificant ($p = 0.27$) and driven by confounders absorbed by the preferred specification. The preferred country-by-year specification (Column 3) yields an SDE of 0.004, classified as "Null"—the effect is less

Table 6: Heterogeneity Analysis: Border×Post by Subgroup

Subgroup	Coefficient	SE	p-value
Tourism Intensity (pre-treatment foreign share tercile)			
Top tercile	0.087	0.141	0.54
Middle tercile	0.102	0.155	0.52
Bottom tercile	0.178	0.147	0.24
Border Corridor			
Western European borders	0.143	0.132	0.29
Central/Eastern European borders	0.102	0.178	0.57
Difference (W−CE)	0.041	0.189	0.83
Nordic Pre-Agreements			
Nordic border regions	0.038	0.094	0.69
Non-Nordic border regions	0.148	0.121	0.23

Notes: Each row reports the border×post coefficient from a separate regression (or sample split) of log foreign nights on border×post with region and year FE, clustered at country level. Tourism intensity terciles are based on pre-treatment (2012–2016) foreign tourist share. Western borders: DE, FR, BE, NL, AT, DK. Central/Eastern: PL, CZ, HU, SK, HR, SI, BG, RO, EE, LT, LV. Nordic: DK, FI, SE, NO. All estimates are statistically insignificant.

than 0.5% of a standard deviation of the outcome, consistent with the paper's overall finding of no meaningful effect.

Table 7: Standardized Effect Sizes

Outcome	Specification	$\hat{\beta}$	SD(Y)	SDE	SE(SDE)	Classification
Log foreign nights	Table 2, Col. 1	0.1242	2.268	0.055	0.048	Moderate positive
Log foreign nights	Table 2, Col. 3	0.0101	2.268	0.004	0.007	Null

Notes: SDE = $\hat{\beta}$ / SD(Y) for binary treatment (Border × Post). Research question: whether EU's 2017 RLAH regulation increased foreign accommodation nights in border NUTS2 regions vs. interior regions. Data: Eurostat tour_occ_nin2, 2012–2019. Unit: NUTS2 region-year. Sample: 1,698 obs, 238 regions (Col. 1); 1,661 obs, 231 regions (Col. 3, country×year FE). Method: TWFE DiD, clustered at country level. SE(SDE) = SE($\hat{\beta}$) / SD(Y). Col. 1: region + year FE; Col. 3: region + country×year FE (preferred). Classification labels reflect magnitude of the standardized point estimate, not statistical significance.