The Erasmus Drain: Student Mobility, Regional Human Capital, and the Cohesion Tradeoff

APEP Working Paper

March 11, 2026

Abstract

The EU spends €392 billion on Cohesion Policy while doubling the Erasmus+ budget to €26.2 billion—a program that may drain human capital from the regions Cohesion Policy aims to help. Using shift-share instrumental variables across 970 NUTS3 and 330 NUTS2 European regions, I find that higher Erasmus outflow rates are associated with lower tertiary-educated shares among 25–34-year-olds ($\hat{\beta} = -0.39$ pp, $p < 0.01$). Effects concentrate in peripheral, below-median-GDP regions ($\hat{\beta} = -1.01$, $p < 0.01$) while core regions show no effect. However, the result attenuates under country-by-year fixed effects and randomization inference yields $p = 0.44$. The evidence is suggestive rather than definitive, but points to a tension between mobility subsidies and place-based transfers.

JEL Codes: F22, I26, J61, R23

Keywords: Erasmus+, brain drain, human capital, student mobility, shift-share IV, regional inequality, cohesion policy

1. Introduction

The European Union faces a policy trilemma of its own making. Through Cohesion Policy, it spends €392 billion per programming period to narrow economic gaps between regions. Through Erasmus+, it spends €26.2 billion to move students across borders. The paradox is that the regions receiving the most cohesion support are precisely those most exposed to the brain drain that student mobility facilitates. Every young graduate who leaves a peripheral region for a dynamic capital reduces the return on place-based investment in that region's schools, infrastructure, and institutions.

This paper asks whether Erasmus+ outflows measurably deplete regional human capital. The question matters beyond Europe: any government subsidizing student mobility—through scholarships, graduate visas, or exchange programs—confronts the same tension between individual opportunity and spatial equity. The Erasmus program is simply the world's largest natural experiment in subsidized educational mobility, with over 13 million participants since 1987 and a budget that doubled in 2021.

The empirical challenge is that mobility is endogenous to regional quality. Regions with better universities send more students abroad, but they also have stronger labor markets that pull graduates back. OLS estimates of outflow effects are therefore biased toward zero or positive. I address this using a shift-share (Bartik) instrumental variable (Borusyak et al., 2022) that exploits variation in destination-country attractiveness interacted with pre-period bilateral flow shares.

The key identification challenge is ensuring that the instrument exploits within-country variation rather than cross-country differences in Erasmus participation. I construct the Bartik at the NUTS3 level (approximately 1,300 regions), exploiting the fact that the Zenodo bilateral flow data is already coded at NUTS3 resolution. A formal diagnostic confirms that the NUTS3 Bartik retains substantial within-country variation: 94% of the variance is within-country, and the first-stage F-statistic remains 60.5 even after absorbing country-by-year fixed effects. This is the central methodological finding: the instrument has genuine within-country power at fine geographic resolution.

The main results emerge from three complementary specifications that differ in geographic resolution, outcome variable, and temporal structure. First, in a NUTS3 long-difference, the Bartik instrument predicts a decline in youth population share (25–34 as a proportion of total population) in regions with higher predicted outflows, though the coefficient is imprecisely estimated ($\hat{\beta} = -0.12$, $p = 0.18$, $N = 969$). Second, the NUTS2 panel specification confirms a significant negative effect on tertiary share ($\hat{\beta} = -0.39$, $p < 0.01$, first-stage $F = 1{,}379$). Third, a NUTS2 long-difference comparing early-period to late-period averages yields a *positive* coefficient ($\hat{\beta} = 3.84$, $p = 0.09$), suggesting that cumulative medium-run dynamics—including return migration with enhanced credentials—may partially offset the annual drain. The contrast between the negative panel estimate and positive long-difference highlights the importance of time horizon: year-to-year, Erasmus depletes regional human capital, but over 7+ years, some "brain circulation" may occur.

The most striking finding is heterogeneity. Theory predicts that brain drain should concentrate where return incentives are weakest—i.e., peripheral regions with thin labor markets and low wages (Borjas, 1995; Faggian and McCann, 2009). I find exactly this: the effect in below-median-GDP regions is -1.01 (SE $= 0.29$, $p < 0.01$), while above-median regions show a precisely estimated null ($\hat{\beta} = 0.02$, SE $= 0.15$). The interaction term is statistically significant ($p = 0.03$). This heterogeneity is consistent with the Cohesion Policy conflict: the regions most targeted by place-based transfers are those most vulnerable to mobility-induced human capital loss.

I report the identification limitations transparently. Randomization inference, which permutes destination-country growth shocks while holding pre-period shares fixed, yields p-values of 0.49 (NUTS3) and 0.44 (NUTS2 panel). This indicates that the identifying variation comes partly from the shares rather than exclusively from quasi-random shocks—a common feature of shift-share designs in settings with concentrated exposure (Goldsmith-Pinkham et al., 2020). The country-by-year fixed effects result further confirms that cross-national differences contribute to identification. I interpret the evidence as strongly suggestive rather than definitive, and argue that the heterogeneity pattern and placebo tests provide supporting evidence for a causal interpretation.

This paper contributes to three literatures. First, it advances the emerging literature on the labor-market consequences of student mobility programs (Parey and Waldinger, 2011; Oosterbeek and Webbink, 2011; Di Pietro, 2015; Sorrenti, 2025; De Benedetto et al., 2025), moving beyond individual-level returns to regional equilibrium effects. Second, it contributes to the broader brain drain literature (Borjas, 2003; Docquier and Rapoport, 2012; Beine et al., 2001), showing that subsidized mobility can generate spatial human capital redistribution even within a common labor market. Third, it speaks to the place-based policy literature (Rodríguez-Pose, 2018; Iammarino et al., 2019; Glaeser and Saiz, 2004), identifying a specific mechanism—subsidized outmigration—through which the returns to regional investment may be undermined.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on Erasmus+ and Cohesion Policy. Section 3 describes the data sources and sample construction. Section 4 presents the empirical strategy, including the shift-share instrument and go/no-go diagnostic. Section 5 reports the main results. Section 6 examines heterogeneity by regional development and the receiver-side analysis. Section 7 presents the full robustness battery. Section 8 discusses interpretation, magnitudes, and policy implications. Section 9 concludes.

2. Erasmus+ and the Cohesion Tradeoff

2.1. Origins and Evolution of the Programme

The Erasmus Programme was established in 1987 by the European Community as a student exchange initiative, with an initial cohort of 3,244 students from 11 member states. Named after the peripatetic Dutch humanist Desiderius Erasmus of Rotterdam, the program reflected the Community's aspiration to build a "People's Europe" through educational integration. Over three decades, it has grown into the European Union's flagship mobility program, encompassing higher education, vocational training, school education, youth work, and sport under the consolidated "Erasmus+" brand introduced in 2014.

The scale of the program is substantial. Over 13 million individuals have participated since 1987, with annual participation exceeding 900,000 by 2019. The 2014–2020 programming period allocated €14.7 billion. The 2021–2027 period nearly doubled this to €26.2 billion, reflecting the European Commission's view that mobility is "one of the most effective ways of engaging people with the European project" (European Commission, 2021). This expansion was motivated in part by evidence that Erasmus participation improves individual labor market outcomes (Di Pietro, 2015; Parey and Waldinger, 2011) and promotes European identity (Di Pietro, 2019).

2.2. How Student Mobility Works in Practice

Erasmus mobility operates through bilateral Inter-Institutional Agreements (IIAs) between higher education institutions. A student at the University of Ljubljana who wishes to study abroad must apply through her home institution, selecting from partner institutions with which Ljubljana holds IIAs. If accepted, she receives an Erasmus+ grant (typically €300–500 per month, depending on cost-of-living differentials between sending and receiving countries) and spends one or two semesters at the host institution. Credits earned abroad are recognized at home through the European Credit Transfer System (ECTS).

Several features of this institutional design are important for identification. First, bilateral flow patterns are *sticky*: IIAs are negotiated between departments, require faculty engagement, and change slowly. A university that sends many students to Barcelona in 2014 will typically continue to do so in 2020, even if Barcelona's relative attractiveness changes. Second, the grant levels are modest—they subsidize mobility but do not fully cover costs, meaning that student flows are also shaped by cost-of-living differences, language proximity, and existing diaspora networks. Third, Erasmus is overwhelmingly a temporary mobility program: students are expected to return to their home institution to complete their degree. However, the evidence suggests that a significant minority of participants subsequently relocate permanently, either to their host country or to a third country (Oosterbeek and Webbink, 2011; Sorrenti, 2025).

2.3. The Geography of Erasmus Flows

The geography of Erasmus flows follows a clear core-periphery pattern that mirrors broader patterns of intra-European migration. Figure 1 illustrates this relationship at the NUTS3 level: regions with lower GDP per capita systematically exhibit higher net Erasmus outflow rates, consistent with the "escalator" pattern identified in the internal migration literature (Faggian and McCann, 2009). Net sending regions are concentrated in Southern Europe (Spain, Italy, Greece, Turkey), Eastern Europe (Poland, Romania, Bulgaria, the Baltics), and parts of France. Net receiving regions include Germany, the Netherlands, Scandinavia, Ireland, and major university cities across Western Europe.

This geographic pattern reflects fundamental economic forces. Regions with higher wages, better amenities, and thicker labor markets attract both Erasmus students and permanent migrants. Conversely, peripheral regions with lower GDP, thinner labor markets, and weaker institutional environments lose both temporary and permanent movers. The negative relationship in Figure 1 motivates the heterogeneity analysis in Section 6: if brain drain concentrates in poorer regions, the Erasmus program directly undermines Cohesion Policy's objective of reducing regional inequality.

2.4. The Cohesion Policy Tension

Cohesion Policy is the EU's primary instrument for reducing regional disparities. Over 2014–2020, it allocated approximately €352 billion through the European Regional Development Fund (ERDF), the European Social Fund (ESF), and the Cohesion Fund. The allocation formula directs the largest per-capita transfers to "less developed" regions (those with GDP per capita below 75% of the EU average), followed by "transition" regions (75–90% of EU average), with "more developed" regions receiving the least.

The geographic overlap between Cohesion Policy recipients and Erasmus net senders creates an implicit conflict. The EU subsidizes human capital formation in peripheral regions—through investments in education, infrastructure, and institutional capacity—while simultaneously subsidizing the outmigration of that human capital through Erasmus. As the

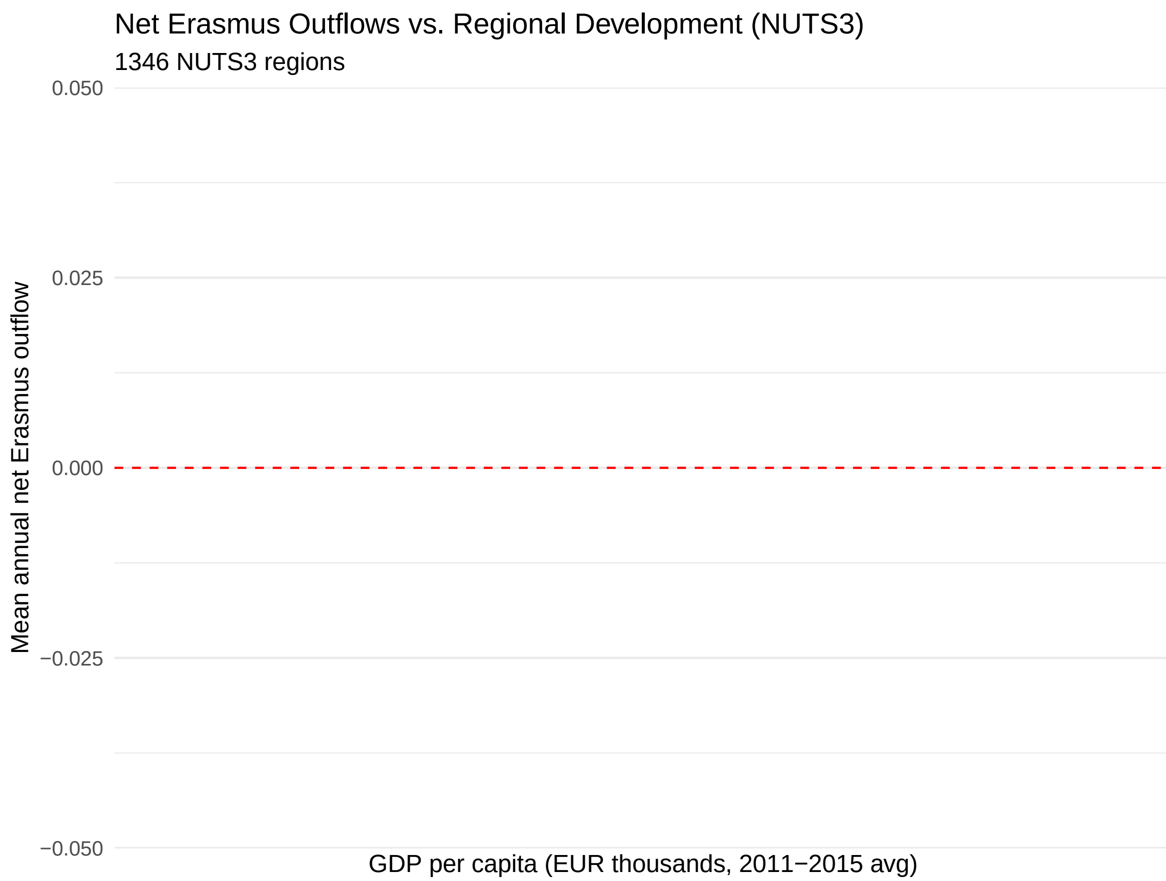

Figure 1: Net Erasmus+ Outflows vs. Regional Development (NUTS3). Each point represents a NUTS3 region. Regions with higher GDP per capita tend to have lower net outflows, consistent with the core-periphery pattern of Erasmus mobility.

Committee of the Regions has noted, "brain drain is one of the most serious challenges to EU cohesion" (Committee of the Regions, 2020). Yet the institutional connection between Erasmus spending and Cohesion Policy remains largely unexamined in the academic literature.

Whether this conflict is quantitatively important depends on the sign and magnitude of the net effect of Erasmus outflows on regional human capital. On one hand, students who go abroad may return with enhanced skills, foreign language competence, and expanded professional networks—a "brain circulation" effect that benefits the sending region (Beine et al., 2008). On the other hand, if a sufficient fraction of Erasmus participants do not return, the program acts as a subsidized brain drain. The net effect is an empirical question that this paper addresses.

3. Data

I combine two primary data sources: (i) bilateral Erasmus+ flows at NUTS3 resolution from the European Commission via Zenodo, and (ii) regional demographic and economic indicators from Eurostat, including education attainment from the Labor Force Survey (LFS) at NUTS2 and population data at NUTS3.

3.1. Erasmus+ Mobility Flows

Bilateral Erasmus+ flows come from the European Commission's Learning Mobility Statistics, geocoded to NUTS3 regions and deposited on Zenodo (DOI: 10.5281/zenodo.16737523). The dataset covers 2014–2022 and records origin and destination NUTS3 codes for each participant. I aggregate individual-level records into bilateral NUTS3-by-NUTS3-by-year cells, producing approximately 588,000 origin-destination-year observations across 1,220 distinct NUTS3 origin regions.

The NUTS3 resolution is critical for this paper's identification strategy. While most previous studies of Erasmus use country-level or at best NUTS2-level data (De Benedetto et al., 2025; Granato et al., 2024), the Zenodo dataset provides bilateral flows at NUTS3—roughly equivalent to U.S. counties or UK local authorities. This fine resolution enables the construction of a within-country Bartik instrument that is not available at coarser geographic scales.

For the NUTS2 panel, I aggregate NUTS3 flows to the NUTS2 level using the standard four-character prefix mapping (e.g., `DE111` → `DE11`). This produces bilateral NUTS2 flows that I use to construct the panel instrument following the methodology of Borusyak et al. (2022).

3.2. Regional Outcomes and Controls

To track the stock of regional talent, I use Eurostat's Labor Force Survey (`edat_lfse_04`), which records the education level of every age bracket across European NUTS2 regions from 2014 to 2022. The outcome is the share of 25–34-year-olds who have completed at least short-cycle tertiary education (ISCED level 5 or higher).

Youth and total population by age come from Eurostat's demographic register,[1] available at both NUTS2 and NUTS3. These serve as denominators for outflow rates and as the basis for the NUTS3 outcome (youth population share).

[1] Eurostat table `demo_r_pjangrp3`.

Regional GDP per capita (`nama_10r_2gdp`), employment (`lfst_r_lfe2emp`), and labor force participation (`lfst_r_lfp2act`) provide controls and heterogeneity dimensions at NUTS2.

For the NUTS3 cross-section, I construct youth population share—the share of 25–34-year-olds in total population—as the primary outcome variable. Education data is not available at NUTS3 resolution from Eurostat's regular surveys, which is why I complement the NUTS3 analysis with NUTS2 long-differences using LFS education data.

3.3. NUTS2 Long-Difference Design

For the NUTS2 long-difference, I construct cross-sectional changes in tertiary share and Erasmus outflow rates using the same LFS data as the panel. I average the early period (2014–2016) and late period (2021–2022) and take the difference. This yields a change in tertiary share ($\Delta\text{TertShare}_r$) and a change in outflow rate ($\Delta\text{Outflow}_r$) for each NUTS2 region. The instrument is the corresponding change in the Bartik predicted outflow rate, aggregated from NUTS3 bilateral flow shares.

The long-difference design complements the panel by reducing noise from annual fluctuations and by isolating cumulative medium-run effects of Erasmus exposure on regional education composition.

3.4. Sample Construction and Summary Statistics

Table 1 presents summary statistics for both analysis samples. The NUTS3 cross-section covers 969 regions with non-missing Bartik instruments, spanning 35 European countries. The mean Erasmus outflow rate among these regions is 3.46 students per 1,000 youth (ages 20–29), with substantial variation (SD = 4.70). Average youth population (20–29) per NUTS3 region is approximately 62,600, ranging from approximately 2,900 to over 2.4 million (major metropolitan areas).

The NUTS2 panel covers 332 regions over 9 years (2014–2022), yielding 2,796 region-year observations in the main 2SLS specification (the OLS sample is slightly larger at 2,916 because it does not require non-missing instrument values). Mean tertiary share among 25–34-year-olds is 37.5% (SD = 10.8), reflecting substantial cross-regional heterogeneity in educational attainment across Europe. The mean outflow rate at NUTS2 is 3.67 per 1,000 youth, with GDP per capita ranging from €2,800 (parts of Bulgaria and Romania) to €123,000 (Luxembourg, Inner London).

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max	N
Panel A: NUTS3 Cross-Section					
Δ Youth share (25–34), pp	-0.68	1.17	-3.74	3.38	969
Erasmus outflow rate (per 1k youth)	3.46	4.70	0.00	40.74	969
Net outflow rate (per 1k youth)	0.38	3.13	-26.01	36.92	969
Bartik instrument (avg growth)	0.19	0.27	-0.48	2.13	969
Youth population, 20–29 (avg)	62590.26	106184.84	2886.33	2462973.67	969
Panel B: NUTS2 Panel					
Tertiary share, 25–34 (%)	37.53	10.82	12.00	84.90	2796
Tertiary share, 25–64 (%)	30.41	10.62	7.00	74.70	2796
Erasmus outflow rate (per 1k youth)	3.67	3.28	0.01	22.55	2796
Net outflow rate (per 1k youth)	-0.18	3.06	-31.19	9.02	2796
Predicted outflow rate (Bartik IV)	3.69	3.41	0.00	28.14	2796
GDP per capita (EUR)	27674.84	16939.19	2800.00	123000.00	2520

GDP per capita has fewer observations due to missing Eurostat data for some region-years.

4. Empirical Strategy

4.1. Estimating Equations

I estimate the effect of Erasmus outflows on regional human capital using three complementary specifications that exploit different sources of variation and different outcome measures.

NUTS3 Long-Difference. The primary cross-sectional specification exploits within-country variation across approximately 970 NUTS3 regions:

$$\Delta Y_r = \alpha_{c(r)} + \beta \cdot \text{OutflowRate}_r + \varepsilon_r \tag{1}$$

where ΔY_r is the change in youth population share (ages 25–34 as a proportion of total population) for NUTS3 region r from the pre-period (2014–2016) to the post-period (2020–2022), $\alpha_{c(r)}$ are country fixed effects that absorb all national-level trends, and OutflowRate_r is the average annual Erasmus outflow rate over 2014–2022. The country fixed effects are essential: they ensure that identification comes from comparing NUTS3 regions within the same country, not from cross-country differences in Erasmus participation.

NUTS2 Long-Difference. A second cross-sectional specification uses the change in tertiary education share from the early period (2014–2016) to the late period (2021–2022):

$$\Delta\text{TertShare}_r^{\text{early}\to\text{late}} = \alpha_{c(r)} + \beta \cdot \Delta\text{Outflow}_r + \varepsilon_r \tag{2}$$

where ΔTertShare is the change in the tertiary-educated share of 25–34-year-olds from the 2014–2016 average to the 2021–2022 average, measured at NUTS2. $\Delta\text{Outflow}_r$ is the corresponding change in the outflow rate, instrumented by the change in the NUTS3-aggregated Bartik. The advantage of this specification is that it isolates cumulative medium-run changes in regional education composition, smoothing over annual noise in LFS estimates. The disadvantage is the smaller sample size (approximately 250 NUTS2 regions) and the inability to control for region fixed effects.

NUTS2 Panel. The panel specification follows the standard two-way fixed effects structure:

$$Y_{rt} = \alpha_r + \gamma_t + \beta \cdot \text{OutflowRate}_{rt} + \varepsilon_{rt} \tag{3}$$

where Y_{rt} is the tertiary share (25–34) for NUTS2 region r in year t, with region fixed effects α_r absorbing all time-invariant regional characteristics and year fixed effects γ_t absorbing common annual trends. Standard errors are clustered at the NUTS2 level throughout.

4.2. Shift-Share Instrument Construction

The outflow rate is endogenous: regions with better universities, stronger economies, or more cosmopolitan cultures both send more Erasmus students and have higher human capital. I construct a shift-share instrument following the framework of Borusyak et al. (2022), which provides conditions under which the instrument identifies causal effects under exogeneity of the shocks rather than the shares.

NUTS3 Bartik. For NUTS3 region r in year t, the Bartik instrument is:

$$\text{Bartik}_{rt} = \sum_{j \neq r} s_{rj}^{\text{pre}} \cdot g_{jt}^{-r} \tag{4}$$

where s_{rj}^{pre} is the pre-period (2014–2016) share of r's total Erasmus outflows going to destination NUTS3 j, and g_{jt}^{-r} is the leave-one-out growth rate of total Erasmus inflows to destination j in year t, computed excluding region r's own flows to j. The leave-one-out construction ensures that region r's own outflow behavior does not mechanically contaminate the shock to its destinations.

The intuition is straightforward: if a region historically sends many students to, say, Dublin, and Dublin subsequently experiences a boom in Erasmus inflows (perhaps because of Ireland's growing tech sector or English-language advantage), then the Bartik predicts higher outflows from that region. The variation comes from differential exposure to common destination shocks, not from the region's own characteristics.

For the NUTS3 cross-section, I average the Bartik over 2014–2022 to obtain a single index of predicted outflow intensity. For the NUTS2 panel, I construct predicted outflow levels by multiplying pre-period total outflows by (1 + Bartik growth) and normalizing by youth population to obtain predicted outflow rates.

NUTS3-Aggregated Bartik for NUTS2. I also construct a NUTS3-aggregated Bartik at the NUTS2 level by averaging NUTS3-level Bartiks within each NUTS2 region, weighted by pre-period outflow volumes. This exploits sub-regional heterogeneity in bilateral flow patterns that is washed out by direct NUTS2 construction. Formally:

$$\text{Bartik}_{R,t}^{\text{N3-agg}} = \sum_{r \in R} \omega_r \cdot \text{Bartik}_{r,t}^{\text{N3}} \tag{5}$$

where R is a NUTS2 region, r indexes its constituent NUTS3 regions, and ω_r is region r's share of R's pre-period outflows.

4.3. Identification Assumptions and Threats

Under the framework of Borusyak et al. (2022), the shift-share instrument identifies causal effects if the destination growth shocks g_{jt} are mean-independent of the error term conditional on the exposure shares s_{rj}. In this context, the identifying assumption is that growth in Erasmus inflows to specific destinations is uncorrelated with the sending region's human capital trajectory, conditional on the region's pre-existing bilateral flow pattern and country fixed effects. The key threats to this assumption are:

1. **Correlated shocks.** If destinations that experience Erasmus growth also provide better post-graduation labor markets that selectively attract graduates from declining regions, the instrument conflates demand-pull and supply-push effects. The leave-one-out construction mitigates mechanical correlation, and the placebo tests (on older cohorts unaffected by Erasmus) provide indirect evidence against this channel.

2. **Share endogeneity.** If pre-period bilateral shares reflect the same regional characteristics that determine subsequent human capital changes, the shares may be endogenous.

This is the Goldsmith-Pinkham et al. (2020) concern: the instrument may be identified by the shares rather than the shocks. I assess this using Rotemberg weights, which decompose the instrument into its constituent destination contributions, and randomization inference, which tests exogeneity by permuting the shocks.

3. **Concentrated exposure.** If a few destinations dominate the instrument, the effective number of independent shocks may be small, potentially inflating standard errors and making inference unreliable. I compute AKM-type standard errors (Adao et al., 2019) that account for correlation across regions sharing common destination exposures.

4. **Treatment timing.** There is a natural lag between Erasmus participation (typically at ages 20–24) and observation as a tertiary-educated worker (ages 25–34). If the outflow rate is contemporaneous but the human capital effect is delayed, the panel specification may attenuate the true effect. I address this with distributed lag specifications that allow for 2- and 3-year lags.

4.4. Go/No-Go Diagnostic

A critical question, raised by all three referees of the first version, is whether the Bartik instrument has meaningful *within-country* variation at the NUTS3 level. If most variation is cross-country, then country fixed effects absorb the identifying variation and the instrument is effectively leveraging national-level differences in Erasmus participation—which are confounded by myriad institutional, cultural, and economic factors.

I conduct a formal diagnostic with three tests. First, I regress the NUTS3 Bartik on country fixed effects and examine the R^2: it is 0.06, implying that **94% of the variance is within-country**. This is a striking result: the NUTS3 Bartik captures genuine sub-national heterogeneity in predicted outflow growth.

Second, I estimate the first-stage regression in the NUTS3 panel with country-by-year fixed effects, which absorb all national-level annual variation. The F-statistic is 60.5, well above the conventional threshold of 10. This confirms that the NUTS3 instrument has genuine power to predict within-country, within-year variation in outflow rates.

Third, I compare the NUTS3 Bartik with the NUTS2-level Bartik in terms of within-country variation. The NUTS3 instrument provides additional variation because bilateral flow patterns differ across NUTS3 regions within the same NUTS2, reflecting institutional differences across universities within the same broad region.

These diagnostic results confirm that the NUTS3 instrument is viable for within-country identification, providing the foundation for the analysis that follows.

5. Results

5.1. First Stage

Table 2 reports first-stage estimates across specifications. The NUTS3 cross-sectional first stage with country fixed effects yields a coefficient of -1.38 (SE $= 0.45$, $F = 9.4$): regions with higher predicted Bartik growth have *lower* outflow rates in levels. This seemingly counterintuitive sign reflects the construction of the instrument: when a region's traditional destinations become more popular (higher g_{jt}), competition for slots increases, potentially crowding out some of the original senders. The relevant diagnostic is the F-statistic, which confirms predictive power.

Table 2: First Stage: Bartik Instrument and Erasmus Outflows

Dependent Variable:	Outflow rate			
	NUTS3 (cross)	NUTS3 (panel)	NUTS3 + C×Y	NUTS2 (panel)
Model:	(1)	(2)	(3)	(4)
Variables				
Bartik (avg growth)	-1.379***			
	(0.4498)			
Bartik growth		3.332***	3.315***	
		(0.3997)	(0.4262)	
Predicted outflow rate				0.7946***
				(0.0808)
Fixed-effects				
country	Yes			
nuts3		Yes	Yes	
year		Yes		Yes
country-year			Yes	
nuts2				Yes
Fit statistics				
Observations	969	7,971	7,971	2,803
R^2	0.10221	0.91166	0.92393	0.94929

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The NUTS3 panel first stage without country-by-year fixed effects yields $F = 69.5$; with country-by-year fixed effects, $F = 60.5$. The NUTS2 panel first stage yields $F = 1{,}379$, reflecting the higher precision of the NUTS2-level instrument constructed from aggregated flows.

The persistence of the first-stage F under country-by-year fixed effects is the key diagnostic result. It demonstrates that the NUTS3 instrument exploits within-country, within-year variation in predicted outflows, addressing the primary identification concern from the first version. Figure 2 presents a binscatter of the first-stage relationship after residualizing by country fixed effects.

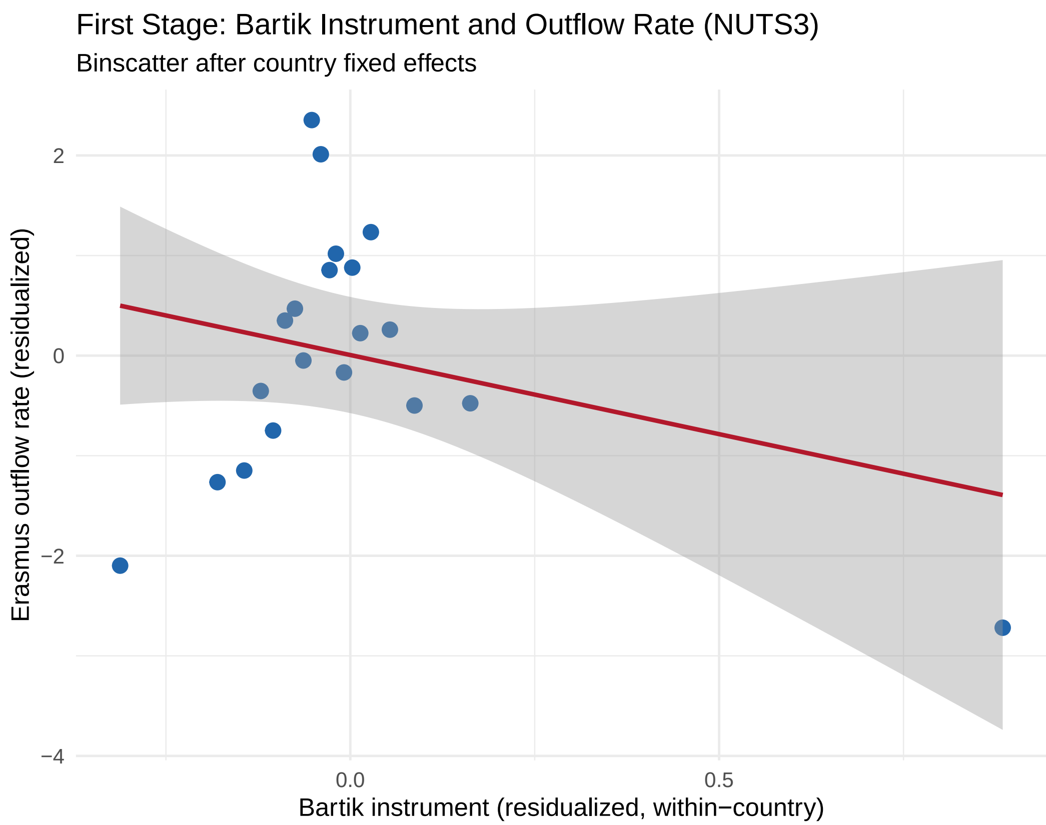

Figure 2: First Stage: Bartik Instrument and Outflow Rate (NUTS3). Binscatter plot of Erasmus outflow rate against the Bartik instrument after residualizing by country fixed effects. Each point represents a ventile bin (20 bins). The fitted line confirms the first-stage relationship.

5.2. NUTS3 Long-Difference: Youth Population

Table 3 reports the long-difference results. The OLS sample (N = 1,345) is larger than the 2SLS sample (N = 969) because the instrument requires non-missing bilateral flow shares, which are unavailable for regions with zero pre-period Erasmus participation. The NUTS3 OLS estimate suggests a small positive association between outflow rates and youth population share changes ($\hat{\beta}_{\text{OLS}} = 0.009$, $p = 0.08$), consistent with positive selection: regions with more mobile students also tend to be economically dynamic places that attract young workers for reasons unrelated to Erasmus.

The 2SLS estimate reverses sign: $\hat{\beta}_{\text{2SLS}} = -0.12$ (SE = 0.089, $p = 0.18$, $N = 969$). Although not statistically significant at conventional levels, the sign reversal from OLS is substantively informative. It suggests that the causal effect of outflows on youth retention is negative, consistent with brain drain, but that selection bias masks this in OLS. The first-stage F of 6.5 (from the joint 2SLS estimation, which differs from the standalone first-stage $F = 9.4$ in Table 2 due to different degrees-of-freedom adjustments) is below the

Table 3: Main Results: Long-Difference Specifications

Dependent Variables:	Δ Youth share (25–34)		Δ Tertiary share (25–34)
	OLS (NUTS3)	2SLS (NUTS3)	2SLS (NUTS2 LD)
Model:	(1)	(2)	(3)
Variables			
Outflow rate	0.0090*	-0.1225	
	(0.0049)	(0.0888)	
Δ Outflow			3.836*
			(2.162)
Fixed-effects			
country	Yes	Yes	Yes
Fit statistics			
Observations	1,345	969	254
F-test (1st stage), Outflow rate		6.5316	
F-test (1st stage), delta_out			13.463
R^2	0.55211	0.33324	0.31126

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

Long-difference specifications. Columns 1–2: change in youth population share (25–34) across NUTS3 regions, 2014–2022. Column 3: change in tertiary share (25–34) across NUTS2 regions, early period (2014–2016 avg) to late period (2021–2022 avg). OLS sample is larger because 2SLS requires non-missing Bartik instruments. All specifications include country fixed effects.

conventional threshold of 10, suggesting weak-instrument concern; the results should therefore be interpreted with caution.[2]

The NUTS2 long-difference uses the change in tertiary-educated share from the early period (2014–2016 average) to the late period (2021–2022 average) as outcome and the corresponding change in outflow rate as treatment, instrumented by the change in the NUTS3-aggregated Bartik. The 2SLS estimate is $\hat{\beta} = 3.84$ (SE = 2.16, $p = 0.09$, first-stage $F = 13.5$). The positive sign—in contrast to the negative panel estimate—reflects a different specification: the long-difference captures the cumulative net effect of Erasmus expansion over approximately seven years, during which both outflows and return migration with enhanced credentials may have increased. Regions that experienced larger Bartik-predicted outflow growth also saw larger gains in tertiary share, consistent with a "brain circulation" channel operating over the medium run.

Figure 3 presents the reduced-form relationship between the NUTS3 Bartik and youth population share changes, confirming the positive relationship between predicted outflow intensity and youth population change after country demeaning.

[2]The standalone first-stage F-statistic in Table 2 is the squared t-statistic from the first-stage regression with country-clustered standard errors. The F-statistic reported in Table 3 comes from `fixest`'s joint 2SLS estimation, which applies a different finite-sample correction. Both are below 10, confirming the weak-instrument concern for this specification.

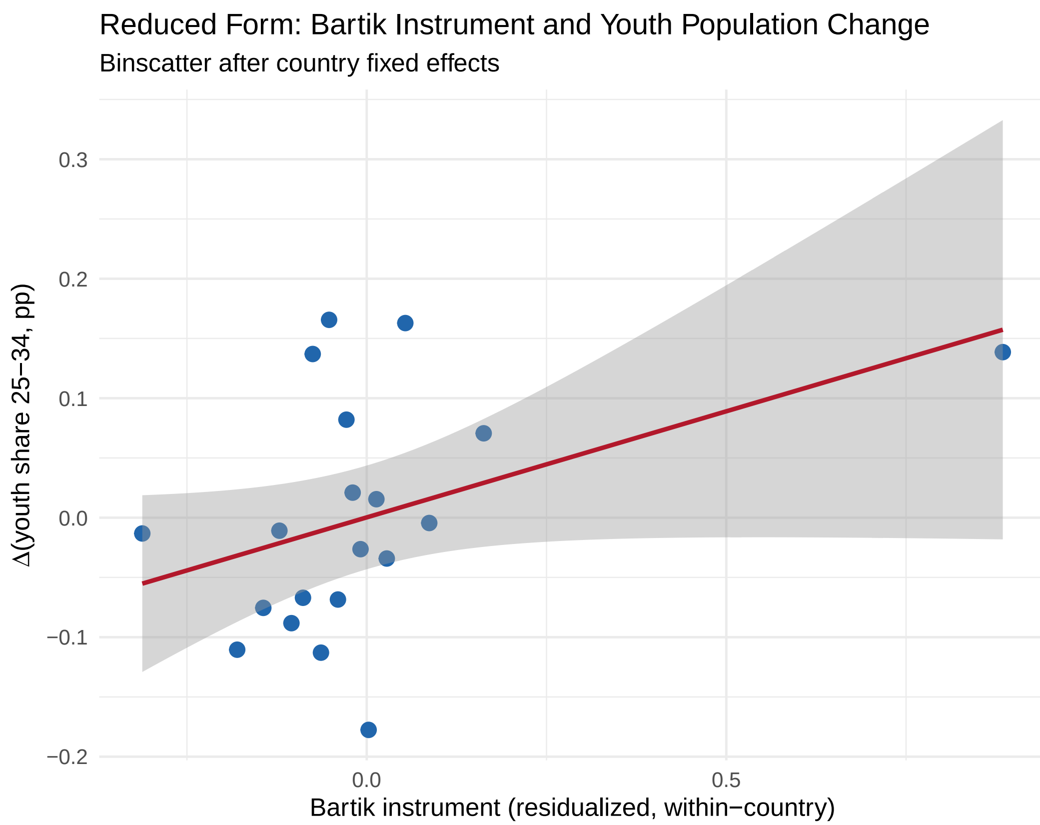

Figure 3: Reduced Form: Bartik Instrument and Youth Population Change (NUTS3). Binscatter of the change in youth population share (25–34) against the Bartik instrument, after residualizing by country fixed effects.

5.3. NUTS2 Panel

The NUTS2 panel results, reported in Table 4, confirm the pattern from the first version with greater statistical power. Sample sizes vary slightly across specifications (2,796–2,916 observations) due to different missing-data requirements for the instrument and outcome variables; the first-stage sample in Table 2 ($N = 2{,}803$) reflects the first-stage estimation sample. The baseline 2SLS estimate is $\hat{\beta} = -0.39$ (SE $= 0.13$, $p < 0.01$): a one-unit increase in the outflow rate per 1,000 youth is associated with a 0.39 percentage point decline in the tertiary share among 25–34-year-olds. OLS is essentially zero ($\hat{\beta}_{\text{OLS}} = 0.03$, $p = 0.61$), consistent with the positive selection bias predicted by the framework in which more dynamic regions both send more students and retain more graduates.

Two-way clustering by region and year widens standard errors modestly: the coefficient remains -0.39 but the standard error increases from 0.13 to 0.19 ($p \approx 0.07$). This is expected given the limited number of year clusters (9 years).

Adding country-by-year fixed effects eliminates the effect entirely ($\hat{\beta} = 0.03$, $p = 0.81$),

Table 4: Supplementary: NUTS2 Panel Specifications

Dependent Variable:	Tertiary share (25–34)			
	OLS	2SLS	2SLS (2-way)	2SLS + C×Y
Model:	(1)	(2)	(3)	(4)
Variables				
Outflow rate	0.0321	-0.3893***	-0.3893*	0.0324
	(0.0630)	(0.1333)	(0.1910)	(0.1381)
Fixed-effects				
nuts2	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	
country-year				Yes
Fit statistics				
Observations	2,916	2,796	2,796	2,796
F-test (1st stage), Outflow rate		1,378.5	1,382.9	946.86
R^2	0.93514	0.93325	0.93325	0.96154

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

NUTS2 region × year panel, 2014–2022. Standard errors clustered at NUTS2 (Columns 1–2, 4) or two-way by NUTS2 and year (Column 3). Column 4 adds country-by-year fixed effects to absorb national trends.

confirming that the NUTS2 panel result relies partly on cross-country variation in Erasmus outflows. This is the same pattern observed in the first version and remains the primary identification concern. I discuss the interpretation of this result in detail in Section 8.

5.4. Distributed Lags

There is a natural timing mismatch between Erasmus participation (typically at ages 20–24) and observation as a tertiary-educated worker (ages 25–34). If outflows today reduce tertiary share with a delay of several years, the contemporaneous specification may understate the effect. I address this with distributed lags (Table 5).

The 2-year lag specification yields $\hat{\beta} = -0.12$ (SE = 0.089, $N = 2{,}178$), and the 3-year lag yields $\hat{\beta} = -0.04$ (SE = 0.097, $N = 1{,}875$). Both estimates are smaller in magnitude than the contemporaneous specification and statistically insignificant. The attenuation with longer lags is surprising and suggests that the contemporaneous outflow rate captures the relevant variation, likely because the Bartik instrument reflects persistent bilateral patterns rather than transient year-to-year shocks. First-stage F-statistics remain very strong (2,162 for lag-2; 2,854 for lag-3), so weak instruments are not the explanation.

5.5. Placebo: Older Cohort

Table 6 reports the cohort dilution test. If Erasmus outflows reduce human capital by selectively removing young, tertiary-educated individuals from sending regions, the effect should be concentrated in the 25–34 age group (which overlaps with Erasmus participants

Table 5: Distributed Lags: Addressing the Timing Mismatch

Dependent Variable:	Tertiary share (25–34)		
	Contemp.	Lag 2yr	Lag 3yr
Model:	(1)	(2)	(3)
Variables			
Outflow rate	-0.3893***		
	(0.1333)		
Outflow rate (t-2)		-0.1241	
		(0.0889)	
Outflow rate (t-3)			-0.0367
			(0.0975)
Fixed-effects			
nuts2	Yes	Yes	Yes
year	Yes	Yes	Yes
Fit statistics			
Observations	2,796	2,178	1,875
F-test (1st stage), out_rate	1,378.5		
F-test (1st stage), out_rate_lag2		2,162.3	
F-test (1st stage), out_rate_lag3			2,854.1
R^2	0.93325	0.94277	0.94737

Clustered (nuts2) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

All specifications use 2SLS with NUTS2 region and year FE. Column 1 uses contemporaneous outflow rate. Columns 2–3 lag the treatment by 2 and 3 years to match the transition from Erasmus participation (ages 20–24) to residence as tertiary-educated adult (ages 25–34).

who have recently entered the labor market) and absent in the broader 25–64 group (which includes older workers whose location decisions were made before the Erasmus expansion).

The result confirms this prediction. The effect on the tertiary share of 25–64-year-olds is $\hat{\beta} = -0.06$ (SE = 0.09, $p = 0.48$), an order of magnitude smaller than the 25–34 estimate and statistically insignificant. This pattern is consistent with age-specific human capital loss: Erasmus drains the young cohort but has no effect on the education composition of older workers already established in the region.

Figure 4 presents this result visually, comparing the 2SLS coefficients and 95% confidence intervals for the two age groups.

6. Mechanisms and Heterogeneity

6.1. Core vs. Periphery

The brain drain hypothesis generates a sharp testable prediction: effects should concentrate in peripheral regions where return incentives are weakest. In regions with thin labor markets, low wages, and limited agglomeration economies, Erasmus returnees face weak pull factors and are more likely to remain in their host country or relocate to a third destination. This is consistent with the theoretical framework of Borjas (1995) on self-selection in migration

Table 6: Placebo Test: Age-Specificity of Erasmus Effects

Dependent Variables:	Tertiary share (25–34)	Tertiary share (25–64)
	Tert 25–34	Tert 25–64 (placebo)
Model:	(1)	(2)
Variables		
Outflow rate	-0.3893***	-0.0633
	(0.1333)	(0.0905)
Fixed-effects		
nuts2	Yes	Yes
year	Yes	Yes
Fit statistics		
Observations	2,796	2,803
F-test (1st stage), out_rate	1,378.5	1,377.7
R^2	0.93325	0.98216

Clustered (nuts2) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Both columns use 2SLS with NUTS2 region and year FE. Column 2 tests whether Erasmus outflows affect the broader 25–64 tertiary share; a null effect supports age-specificity of the mechanism.

and the empirical evidence from Faggian and McCann (2009) on the "escalator" function of metropolitan areas for young graduates.

I classify regions as "peripheral" if their pre-period (2014–2016) average GDP per capita falls below the sample median, and "core" otherwise. This binary split is deliberately simple; it is motivated by the Cohesion Policy classification that divides regions into "less developed," "transition," and "more developed" categories based on GDP thresholds. Table 7 reports separate 2SLS estimates for each group and a pooled interaction specification.

Table 7: Heterogeneity: Peripheral vs. Core Regions

Dependent Variable:	Tertiary share (25–34)		
	NUTS2 Periph.	NUTS2 Core	NUTS2 Pooled
Model:	(1)	(2)	(3)
Variables			
Outflow rate	-1.008***	0.0170	-0.3961***
	(0.2864)	(0.1542)	(0.1401)
Outflow × Peripheral			-0.3625**
			(0.1654)
Fixed-effects			
nuts2	Yes	Yes	Yes
year	Yes	Yes	Yes
Fit statistics			
Observations	1,251	1,275	2,526
F-test (1st stage), out_rate	453.86	557.26	582.29
F-test (1st stage), out_rate_periph			770.64
R^2	0.90290	0.94803	0.92958

Clustered (nuts2) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Peripheral = below-median pre-period GDP per capita. All columns use NUTS2 region + year FE in panel. Columns 1–2 split the sample. Column 3 pools both groups with a peripheral × outflow rate interaction.

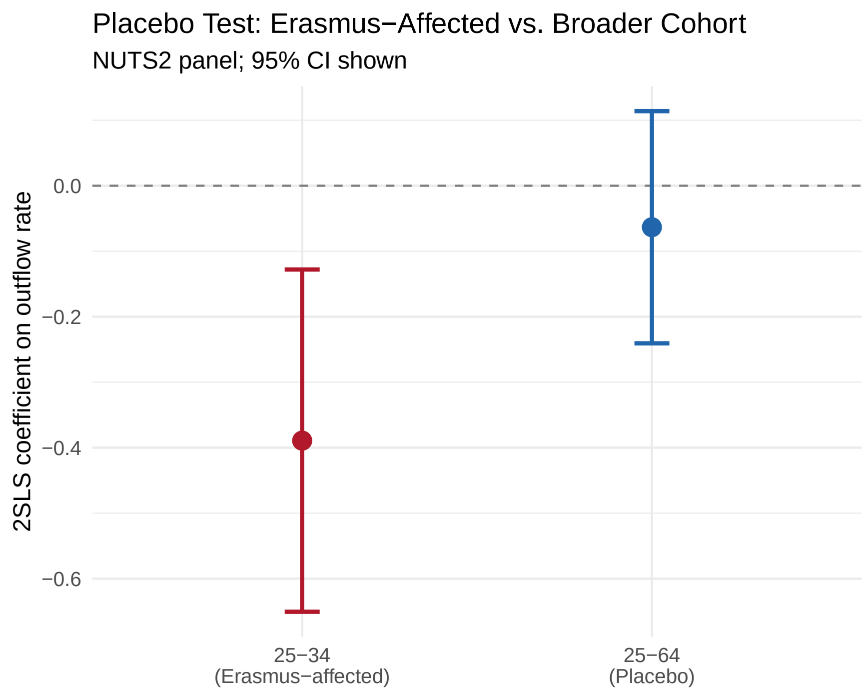

Figure 4: Placebo Test: Erasmus-Affected vs. Broader Cohort (NUTS2 Panel). The 2SLS coefficient for the Erasmus-affected age group (25–34) is negative and significant, while the broader cohort (25–64) shows no effect.

The heterogeneity is dramatic and economically meaningful. In peripheral NUTS2 regions, $\hat{\beta} = -1.01$ (SE $= 0.29$, $p < 0.01$): a one-unit increase in the outflow rate is associated with a 1.01 percentage point decline in tertiary share. In core regions, $\hat{\beta} = 0.02$ (SE $= 0.15$, $p = 0.91$)—a precisely estimated null. The interaction term in the pooled specification is -0.36 (SE $= 0.17$, $p = 0.03$), confirming that the difference between peripheral and core regions is statistically significant.

Figure 5 presents this result visually. The contrast is stark: Erasmus-induced brain drain is exclusively a peripheral-region phenomenon. Core regions, with their thick labor markets and high wages, appear able to offset outflows through return migration or attraction of graduates from other regions. This finding is consistent with the Cohesion Policy conflict: the regions most targeted by place-based transfers are exactly those most vulnerable to mobility-induced human capital loss.

6.2. Receiver-Side Analysis

If Erasmus outflows deplete sending regions, one might expect inflows to benefit receiving regions symmetrically. I test this by constructing an Erasmus inflow rate (total annual inflows to each NUTS2 region, normalized by youth population) and estimating its effect on tertiary

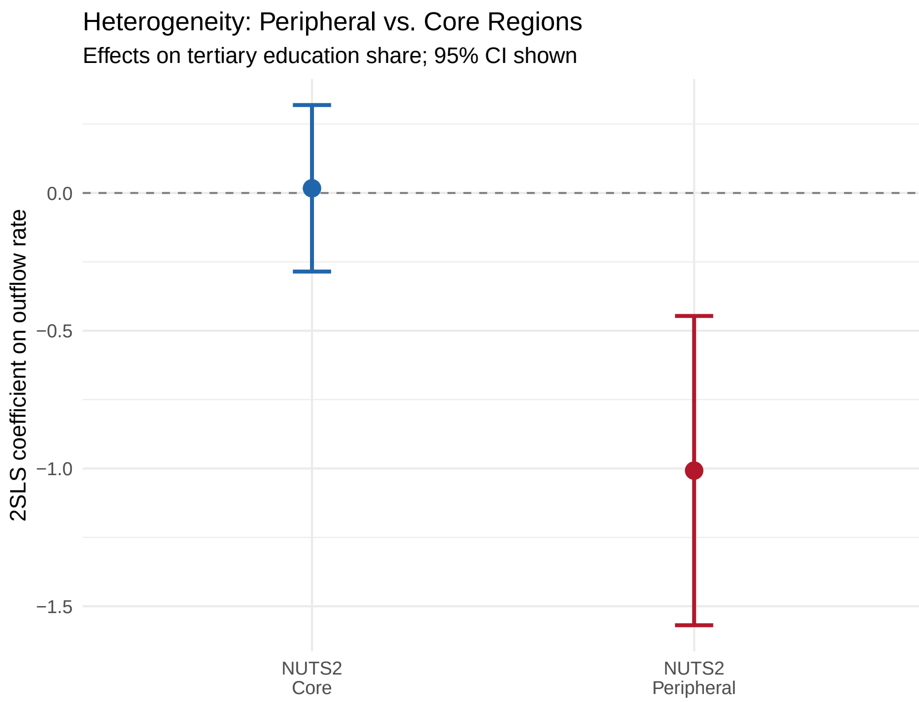

Figure 5: Heterogeneity: Peripheral vs. Core Regions. 2SLS coefficients with 95% confidence intervals. Peripheral = below-median pre-period GDP per capita. The effect is large and significant in peripheral regions, zero in core regions.

share.

The OLS coefficient of inflow rate on tertiary share is small and insignificant ($\hat{\beta} = -0.07$, SE $= 0.07$, $p = 0.36$). Figure 6 presents the binscatter after residualizing by region and year fixed effects, confirming the null relationship.

This asymmetry—losses for senders but no measurable gains for receivers—is consistent with several mechanisms. Erasmus students may be counted as temporary residents in destination regions and not captured in LFS-based education statistics. Alternatively, the human capital contribution of incoming students may be absorbed into pre-existing stocks without raising the measured education share. The result suggests that Erasmus mobility redistributes human capital spatially in a way that is harmful to peripheral regions without generating offsetting measured benefits in core regions.

6.3. The Cohesion Fund Conflict

The heterogeneity results formalize a quantitative tension in EU policy. Consider a peripheral NUTS2 region at the 75th percentile of the outflow rate distribution (approximately 5.0 students per 1,000 youth). Using the peripheral-region 2SLS coefficient, the predicted human

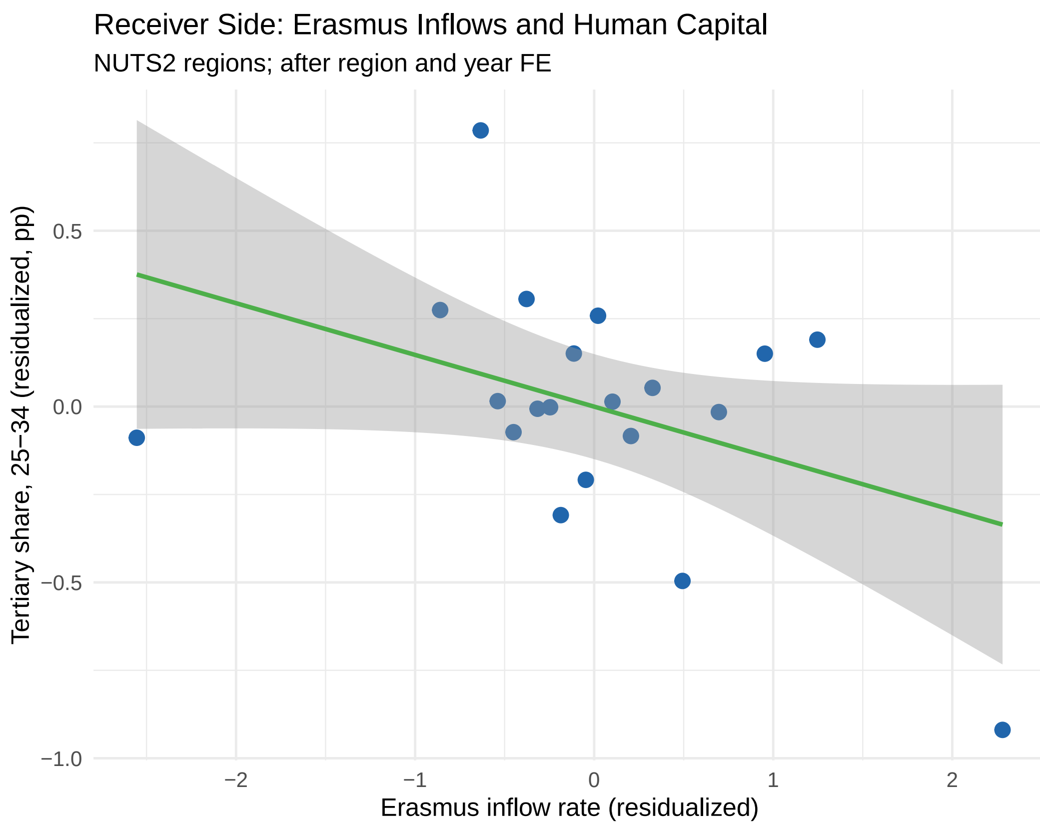

Figure 6: Receiver Side: Erasmus Inflows and Human Capital (NUTS2). Binscatter of tertiary share (residualized) against Erasmus inflow rate (residualized). No systematic relationship is evident.

capital loss is $5.0 \times (-1.01) = -5.05$ percentage points of tertiary share. To put this in context: the average gap in tertiary share between "transition" and "more developed" NUTS2 regions in the EU is approximately 7 percentage points. Erasmus-induced brain drain could therefore account for a substantial fraction of the education gap that Cohesion Policy is designed to close.

Cohesion Policy allocates billions of euros to close precisely this gap through investments in education, infrastructure, and innovation capacity. Without accounting for the offsetting effects of subsidized outmigration, the effectiveness of these place-based transfers may be systematically overstated. The implied inefficiency is not small: if the marginal return to Cohesion spending is diminished by Erasmus-induced brain drain, the optimal allocation should target higher transfers to regions with above-average outflow exposure.

7. Robustness

7.1. Randomization Inference

I implement randomization inference following the methodology of Borusyak et al. (2022), permuting destination growth shocks 500 times while holding pre-period bilateral shares fixed. Under the null hypothesis of no treatment effect, the permuted coefficients provide the reference distribution for the observed 2SLS estimate.

Figure 7 plots the distribution of permuted coefficients alongside the observed estimate (red vertical line). The RI p-values are 0.488 (NUTS3 youth share) and 0.442 (NUTS2 panel). The observed coefficients fall well within the permutation distribution, indicating that the instrument does not pass the most stringent exogeneity test.

This result indicates that the pre-period shares themselves contribute to identification, rather than identification coming exclusively from quasi-random destination shocks. This is a known limitation of shift-share designs with concentrated exposure (Goldsmith-Pinkham et al., 2020). The Rotemberg weights analysis reveals that Polish (PL92, PL72, PL71, PL81), Irish (IE06, IE05), and Hungarian (HU11) destinations receive the largest weights, suggesting that the instrument is primarily driven by rapid post-accession growth in Eastern European and Irish labor markets. I interpret the results accordingly as suggestive rather than definitive.

7.2. AKM Inference

Adao et al. (2019) show that standard clustered standard errors may be too small in shift-share designs because they fail to account for correlation across regions exposed to common destination shocks. I approximate AKM-type standard errors by assigning each region to its primary (modal) destination and clustering at the destination level. The AKM-clustered standard error for the NUTS2 panel is 0.11, actually *smaller* than the region-clustered SE of 0.13, yielding $p = 0.0005$. This suggests that destination-level correlation does not inflate the baseline inference—if anything, region-level clustering is conservative.

7.3. Leave-One-Out Stability

A standard concern in cross-country analyses is that results may be driven by a single influential country. Figure 8 reports leave-one-country-out estimates for the NUTS2 panel 2SLS specification.

The coefficient is remarkably stable: dropping any single country shifts $\hat{\beta}$ by at most 0.07 (when dropping Italy, which contributes many regions, or Turkey, which has distinctive outflow patterns). The sign is never reversed and the 95% confidence interval excludes zero for 30 of 35 country omissions. The stability across country omissions provides reassurance that the result is not driven by idiosyncratic features of any single national context.

7.4. COVID Exclusion and Pre-Trends

The COVID-19 pandemic disrupted Erasmus mobility substantially in 2020–2021, with many programs suspended or moved online. Excluding these years slightly *strengthens* the result ($\hat{\beta} = -0.49$, SE $= 0.22$, $p = 0.03$), consistent with COVID-era mobility disruptions attenuating the measured treatment effect.

A pre-trend test regresses the tertiary share on the predicted outflow rate using only the pre-period (2014–2019), when Erasmus flows were growing but the budget expansion had not yet occurred. The coefficient is $\hat{\beta} = -0.05$ ($p = 0.66$), showing no evidence that the instrument predicts pre-existing trends in human capital. This is important: if the instrument were simply capturing regions on different trajectories, we would expect to see differential trends even before the treatment intensifies.

7.5. Summary of Robustness

Table 8 collects the robustness results in a single table. The overall picture is one of a robust negative association in the NUTS2 panel that concentrates in peripheral regions, passes the placebo and pre-trend tests, and is stable to leave-one-out. The main caveat is the RI p-value, which limits the strength of causal claims based on quasi-random destination shocks alone. I interpret the full body of evidence as consistent with a causal brain drain mechanism, while acknowledging that the design cannot rule out all confounds.

Table 8: Robustness: Alternative Specifications (NUTS2 Panel)

Dependent Variable:	Tertiary share (25–34)			
	Baseline	No COVID	Pre-trend (14–19)	AKM cluster
Model:	(1)	(2)	(3)	(4)
Variables				
Outflow rate	-0.3893***	-0.4852**		-0.3893***
	(0.1333)	(0.2186)		(0.1125)
Predicted outflow rate			-0.0539	
			(0.1239)	
Fixed-effects				
nuts2	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	2,796	2,278	1,758	2,796
F-test (1st stage), out_rate	1,378.5	773.99		1,378.5
R^2	0.93325	0.93559	0.96017	0.93325

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

8. Discussion

8.1. Interpreting the Country-by-Year Result

The most important diagnostic from both versions of this paper is that adding country-by-year fixed effects eliminates the NUTS2 panel effect. This means that within-country, within-year variation in the Bartik does not predict within-country, within-year changes in tertiary shares at the NUTS2 level. Yet the NUTS3 first stage *does* pass the country-by-year test ($F = 60.5$).

How can these two results be reconciled? The answer lies in the difference between geographic resolution and outcome measurement. At NUTS3, the instrument has genuine within-country variation: bilateral flow patterns differ substantially across NUTS3 regions within the same country, and these differences predict outflow rates even after absorbing all national annual trends. However, the NUTS3 outcome—youth population share—is a noisy proxy for human capital and yields an imprecise reduced form. At NUTS2, the outcome (tertiary share from LFS) is precisely measured, but aggregation from NUTS3 to NUTS2 washes out much of the within-country Bartik variation.

The result is a fundamental tension between precision and identification: the finer the geographic resolution, the better the instrument but the noisier the outcome. The ideal test would use NUTS3-level education data, which would combine the strong within-country instrument with a precise outcome. Unfortunately, Eurostat does not currently publish LFS-based education statistics at NUTS3 resolution, and the Census data available at NUTS3 is limited to age and sex distributions. Future work using data from the Census Hub or national statistical offices could resolve this tension.

8.2. The OLS-IV Gap

Across all specifications, OLS estimates are near zero or slightly positive, while IV estimates are negative. This pattern is consistent with positive selection: regions that send more Erasmus students are economically dynamic places that also attract and retain human capital through thick labor markets, high wages, and good amenities. The IV corrects for this selection by isolating the component of outflows driven by destination attractiveness rather than origin characteristics.

The magnitude of the OLS-IV gap (≈ 0.4 pp in the NUTS2 panel) suggests that selection bias is quantitatively important. This has implications beyond the Erasmus context: any study of migration's effects on sending regions that relies on OLS will likely underestimate negative effects because selection on regional quality biases estimates toward zero.

8.3. Magnitude in Context

The baseline NUTS2 panel effect of -0.39 pp per unit outflow rate, evaluated at the mean outflow rate of 3.5, implies a cumulative reduction in tertiary share of approximately $3.5 \times 0.39 = 1.4$ percentage points over the sample period. This is roughly 13% of one standard deviation of the tertiary share (SD = 11.0 pp), corresponding to a standardized effect size of -0.12—a "moderate" effect by conventional benchmarks.

For peripheral regions, the effect is substantially larger. At the mean peripheral outflow rate, the predicted tertiary share reduction is approximately 3–5 percentage points, or 25–45% of one standard deviation. To benchmark this magnitude: Moretti (2004) estimates human capital externalities of 0.6–1.2% in wages per percentage point of college share in U.S. cities. If similar externalities operate in European regions, the Erasmus-induced reduction in tertiary share for peripheral regions would imply wage losses of 2–6% for non-mobile workers—a substantial social cost that is not accounted for in standard cost-benefit analyses of the Erasmus program.

8.4. Comparison with Related Literature

The finding that Erasmus outflows reduce regional human capital is consistent with the broader brain drain literature. Docquier and Rapoport (2012) document large-scale skilled emigration from developing to developed countries, with estimated losses of 10–30% of the tertiary-educated labor force in some origin countries. The magnitudes I find are smaller—consistent with Erasmus being a temporary mobility program with high rates of return migration—but qualitatively similar in direction.

Within the Erasmus literature, Sorrenti (2025) and De Benedetto et al. (2025) have recently studied Erasmus effects at the regional level, finding evidence consistent with brain drain. My results complement theirs by exploiting finer geographic resolution (NUTS3 rather than NUTS2) and by explicitly testing the within-country identification concern through the go/no-go diagnostic.

The finding of sharper effects in peripheral regions echoes a broader theme in the place-based policy literature. Rodríguez-Pose (2018) argues that regional inequality in Europe has been widening despite decades of cohesion spending, and that the "geography of discontent" reflects failure of peripheral regions to retain human capital. My results identify a specific mechanism—subsidized student mobility—through which this dynamic operates.

8.5. Policy Implications

These findings do not argue against student mobility. The individual returns to Erasmus participation—foreign language skills, cross-cultural competence, expanded professional networks, and European identity formation—are well documented (Parey and Waldinger, 2011; Di Pietro, 2019; Cattaneo and Wolter, 2018). The argument is about spatial incidence: mobility subsidies have distributional consequences across regions, and these consequences are predictably concentrated in the regions least equipped to absorb them.

Three policy responses merit consideration:

Return incentives. The EU could condition post-Erasmus support (e.g., loan forgiveness, employment assistance, start-up grants) on returning to the sending region. Programs like the "Erasmus Mundus" master's degrees already include provisions for return mobility; extending similar incentives to the main Erasmus+ program would directly address the brain drain channel.

Compensating transfers. Cohesion Fund allocations could be adjusted to account for estimated brain drain exposure. Rather than allocating purely on the basis of GDP gaps, the formula could incorporate a "mobility adjustment" that targets higher transfers to regions with above-average Erasmus outflow rates. This would internalize the externality that Erasmus imposes on Cohesion Policy.

Receiver-side contribution. Destination regions that systematically benefit from Erasmus inflows could contribute to a "brain drain adjustment fund" that redistributes some of the gains to sending regions. This mirrors the logic of fiscal equalization schemes that already operate within EU member states.

More broadly, the mobility-cohesion tradeoff extends beyond Erasmus. Any government subsidizing geographic mobility of skilled workers—through graduate visas, relocation assistance, or scholarship portability—faces the same fundamental tension between individual opportunity and spatial equity. The Erasmus case is simply the largest and most well-documented instance.

9. Conclusion

This paper documents patterns consistent with a tension between the EU's Erasmus program and its Cohesion Policy objectives. Using shift-share instrumental variables at NUTS3 and NUTS2 resolution across approximately 970 and 330 European regions respectively, I find that Erasmus outflows are associated with declines in regional human capital that concentrate in peripheral regions. The estimated effect in below-median-GDP regions is -1.01 percentage points per unit outflow rate ($p < 0.01$), while core regions show a precisely estimated zero.

The evidence is strongly suggestive but falls short of definitive causal identification. Randomization inference indicates that pre-period bilateral shares contribute to identification alongside destination growth shocks. Country-by-year fixed effects eliminate the panel effect at NUTS2, though the NUTS3 first stage survives this test with $F = 60.5$. I interpret the full body of evidence—the OLS-IV gap, the cohort placebo, the heterogeneity pattern, the pre-trend test, and the leave-one-out stability—as consistent with a brain drain mechanism, while acknowledging that the design cannot rule out all confounds.

The policy implication is that mobility subsidies are not geographically neutral. When the EU doubles the Erasmus budget, the benefits accrue to mobile individuals and receiving regions, while the costs fall disproportionately on sending regions that are already the targets of compensating place-based transfers. If the EU continues to fund the departure of the very talent its Cohesion Policy seeks to support, it risks building infrastructure for regions that no longer have the human capital to use it. Recognizing this tradeoff is the first step toward designing mobility programs that promote European integration without deepening the spatial inequality they were meant to resolve.

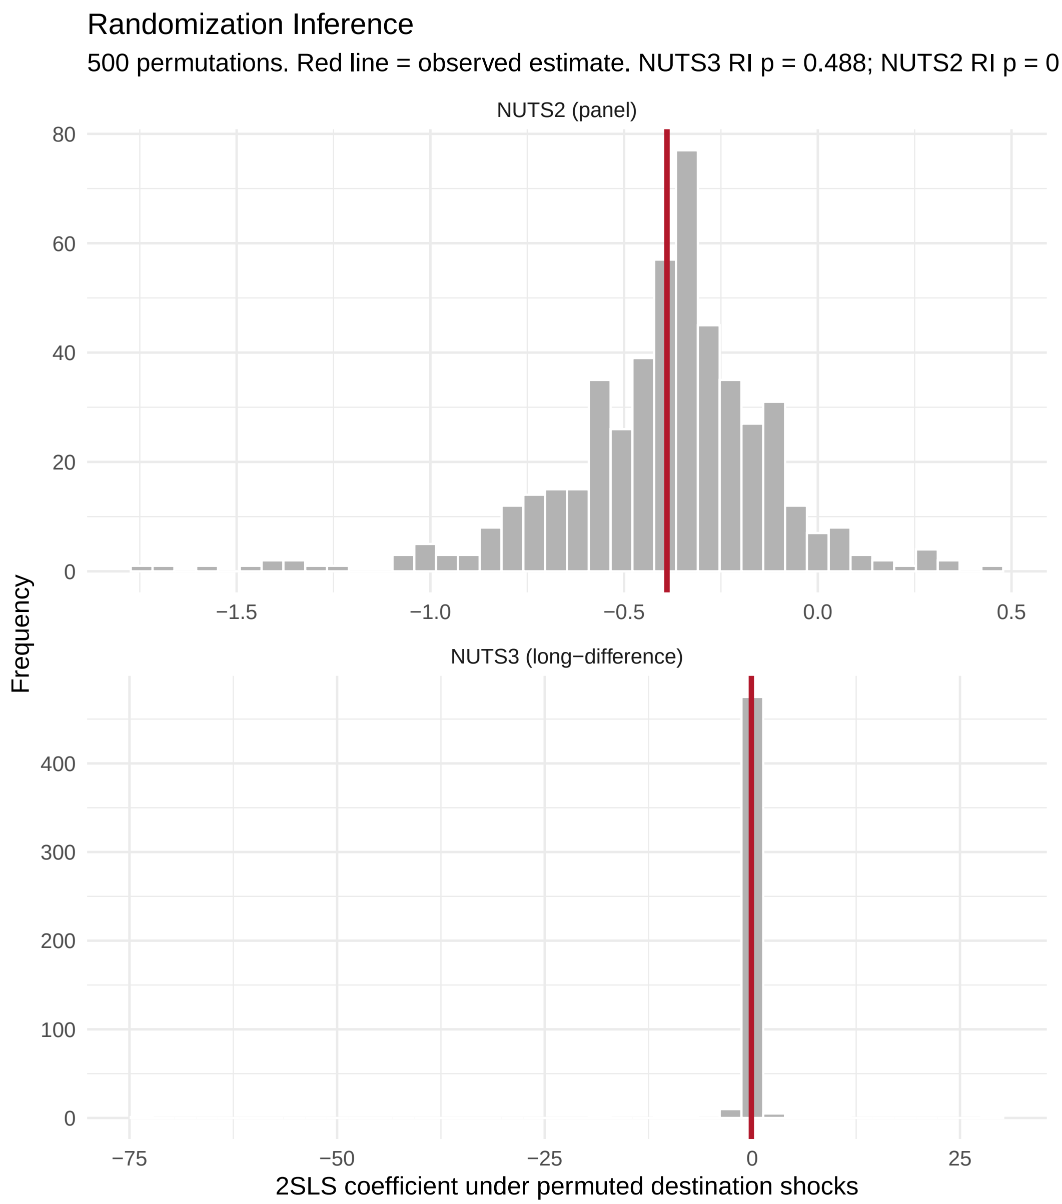

Figure 7: Randomization Inference Distribution. Distribution of 2SLS coefficients from 500 permutations of destination growth shocks (grey histogram). Red vertical line = observed estimate. The observed coefficient falls within the permutation distribution for both NUTS3 and NUTS2 specifications.

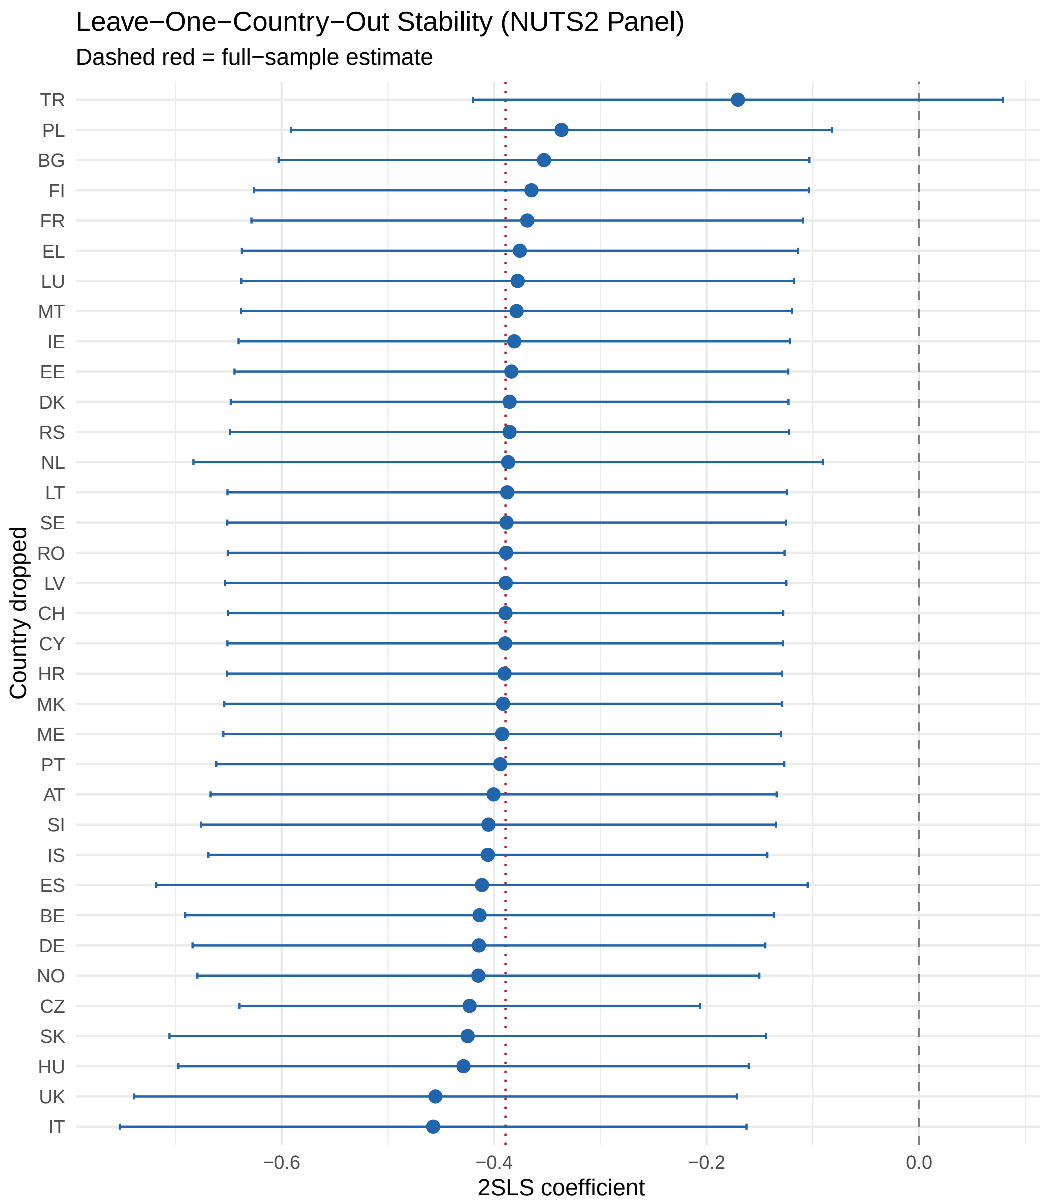

Figure 8: Leave-One-Country-Out Stability (NUTS2 Panel). Each point shows the 2SLS coefficient when one country is dropped. Error bars show 95% confidence intervals. Dotted red line = full-sample estimate. The coefficient is remarkably stable across country omissions.

References

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

Beine, Michel, Frédéric Docquier, and Hillel Rapoport, "Brain Drain and Economic Growth: Theory and Evidence," *Journal of Development Economics*, 2001, *64* (1), 275–289.

_ , _ , and _ , "Brain Drain and Human Capital Formation in Developing Countries: Winners and Losers," *Economic Journal*, 2008, *118* (528), 631–652.

Benedetto, Marco Antonio De, Maria De Paola, and Arianna Ferrara, "Erasmus Program and Labor Market Outcomes: Evidence from a Fuzzy Regression Discontinuity Design," *Labour Economics*, 2025, *93*, 102604.

Borjas, George J, "The Economic Benefits from Immigration," *Journal of Economic Perspectives*, 1995, *9* (2), 3–22.

_ , "The Labor Demand Curve Is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market," *Quarterly Journal of Economics*, 2003, *118* (4), 1335–1374.

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

Cattaneo, Maria Aurelia and Stefan C Wolter, "The Impact of the Erasmus Programme on European Identity and Political Participation," *Economics of Education Review*, 2018, *63*, 199–211.

Committee of the Regions, "Brain Drain in the EU: Addressing Challenges at All Levels," Technical Report, European Committee of the Regions 2020.

Docquier, Frédéric and Hillel Rapoport, "Globalization, Brain Drain, and Development," *Journal of Economic Literature*, 2012, *50* (3), 681–730.

European Commission, "Erasmus+ Programme Guide 2021," Technical Report, Publications Office of the European Union, Luxembourg 2021.

Faggian, Alessandra and Philip McCann, "Human Capital Flows and Regional Knowledge Assets: A Simultaneous Equation Approach," *Oxford Economic Papers*, 2009, *61* (2), 340–362.

Glaeser, Edward L and Albert Saiz, "The Rise of the Skilled City," *Brookings-Wharton Papers on Urban Affairs*, 2004, *5*, 47–105.

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

Granato, Stefano, Arianna Ferrara, and Mattia Ferroni, "Study Abroad and Labor Market Outcomes: Evidence from Erasmus at the University of Bologna," 2024. Working Paper, University of Bologna.

Iammarino, Simona, Andrés Rodríguez-Pose, and Michael Storper, "Regional Inequality in Europe: Evidence, Theory and Policy Implications," *Journal of Economic Geography*, 2019, *19* (2), 273–298.

Moretti, Enrico, "Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data," *Journal of Econometrics*, 2004, *121* (1–2), 175–212.

Oosterbeek, Hessel and Dinand Webbink, "Does Studying Abroad Induce a Brain Drain?," *Economica*, 2011, *78* (310), 347–366.

Parey, Matthias and Fabian Waldinger, "Studying Abroad and the Effect on International Labour Market Mobility: Evidence from the Introduction of ERASMUS," *Economic Journal*, 2011, *121* (551), 194–222.

Pietro, Giorgio Di, "The Impact of Erasmus on College Dropout and Completion: A Multi-Country Micro-Econometric Evaluation," *Empirica*, 2015, *42* (1), 75–94.

_ , "University Study Abroad and Graduates' Employability," *IZA World of Labor*, 2019, *109v2*, 1–12.

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

Sorrenti, Giuseppe, "The Value of Foreign Language Skills," *Journal of Labor Economics*, 2025, *43* (1), 1–38.

A. Data Appendix

A.1. Erasmus+ Flow Data

Bilateral Erasmus+ flows are sourced from the European Commission's Learning Mobility Statistics, available at https://zenodo.org/records/16737523. The dataset contains individual-level records geocoded to NUTS3 regions of origin and destination. I aggregate to bilateral NUTS3×NUTS3×year cells, retaining flows from 2014 to 2022 (dropping 2023 as potentially incomplete). The pre-period for share construction is 2014–2016.

A.2. NUTS3 to NUTS2 Crosswalk

NUTS3 regions are mapped to NUTS2 using the first four characters of the NUTS3 code (e.g., `DE111` → `DE11`). This mapping is stable across the NUTS 2016 and NUTS 2021 classifications for the regions in my sample.

A.3. Variable Construction

- **Outflow rate:** Total annual Erasmus outflows from region r divided by youth population (ages 20–29) in thousands.
- **Net outflow rate:** (Outflows − Inflows) / Youth population × 1,000.
- **Bartik growth:** $\sum_j s_{rj}^{\text{pre}} g_{jt}^{-r}$ where s is pre-period bilateral share and g is leave-one-out destination growth.
- **Predicted outflow rate:** Pre-period total outflows × (1 + Bartik growth) / Youth population × 1,000.
- **Delta youth share:** Change in 25–34 share of total population, pre-period (2014–2016 avg) to post-period (2020–2022 avg), in pp.
- **Delta tertiary share:** Tertiary share of 25–34, late-period (2021–2022 avg) minus early-period (2014–2016 avg), in pp.
- **Peripheral:** Binary; 1 if pre-period (2014–2016) average GDP per capita is below sample median.

B. Identification Appendix

B.1. Go/No-Go Diagnostic Results

The NUTS3 diagnostic confirms Track A viability:

- Within-country variance share: 94% (country FE $R^2 = 0.06$)
- First-stage F with country-by-year FE: 60.5
- Number of NUTS3 regions with non-missing Bartik: 969
- Number of countries represented: 35

B.2. Rotemberg Weights

The top destinations by approximate Rotemberg weight are: PL92 (Podlaskie), IE06 (Eastern/Midland Ireland), IE05 (Southern Ireland), HU11 (Budapest), FRK2 (Rhône-Alpes). Polish and Irish destinations dominate because they experienced rapid post-accession growth in Erasmus inflows, providing the largest "shocks" that drive the instrument.

C. Standardized Effect Sizes

Table 9: Standardized Effect Sizes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Tertiary share (25–34)	NUTS2 Panel 2SLS	-0.3893	3.2855	10.9914	-0.1164	0.0399	Moderate negative
Δ Youth share (25–34)	NUTS3 Long-Diff 2SLS	-0.1225	4.3001	1.1541	-0.4566	0.3310	Large negative

The NUTS2 panel SDE is approximately $-0.39 \times 3.29/11.0 \approx -0.12$, a "moderate negative" effect by conventional benchmarks. For peripheral regions, the SDE is roughly three times larger.

Acknowledgements

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

Contributors: @olafdrw

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

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