The Exodus Paradox: Emigration, Population Decline, and Per-Capita Gains in Romanian Sending Counties after EU Free Movement

APEP Autonomous Research* @SocialCatalystLab

March 30, 2026

Abstract

When Germany, France, and Austria lifted labor market restrictions on Romanian workers in January 2014, high-emigration counties experienced rapid population outflows. I document a composition paradox: counties with greater pre-2014 population decline lost employment ($\beta = -0.50$, SE = 0.07) and population ($\beta = -0.81$, SE = 0.01) disproportionately, yet their employment rates ($\beta = +0.13$, SE = 0.03) and GDP per capita ($\beta = +0.27$, SE = 0.12) improved. Population fell faster than employment, mechanically raising per-capita outcomes for those who remained. A triple-difference exploiting construction-sector concentration—the sector most demanded in German labor markets—provides the cleanest identification ($\beta = +0.88$, $p < 0.01$). Pre-trend tests reveal significant placebo breaks, limiting causal claims on employment levels. The results illustrate how emigration can hollow out communities while simultaneously improving measured conditions for stayers.

JEL Codes: F22, J61, R23, O15

Keywords: emigration, labor markets, EU free movement, Romania, sending regions, composition effects

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 50m).

1. Introduction

Romania's western counties lost roughly 15 percent of their population in the decade following EU accession, yet their GDP per capita and employment rates rose. This paper documents and explains this paradox. The standard neoclassical prediction is that emigration tightens sending-region labor markets, raising wages and employment conditions for workers who stay (Borjas, 2003; Mishra, 2007). But general-equilibrium forces—capital flight, fiscal drain, skill selection, and agglomeration collapse—can reverse these gains (Docquier and Rapoport, 2014; Dustmann et al., 2011). Which force dominates in practice remains an open empirical question, particularly for the large-scale labor mobility enabled by European integration (Clemens, 2011; Kahanec et al., 2013).

I study the January 1, 2014 lifting of transitional labor market restrictions by Germany, France, the Netherlands, Belgium, and Austria on Romanian workers. Romania had joined the EU in 2007, but nine EU-15 member states retained work permit requirements for Romanian nationals under the Treaty of Accession's maximum seven-year derogation. When these restrictions expired simultaneously on January 1, 2014, immigration from Romania to Germany alone quadrupled—from approximately 24,000 per year to 97,000 in 2014 (De Graaf and Gorter, 2014). This generated a sharp, well-dated labor supply shock to Romanian sending counties that varied in intensity with pre-existing migration networks.

The central finding is a *composition effect*: high-emigration counties experienced large absolute losses in both employment and population after 2014, but population fell faster than employment, mechanically improving per-capita indicators. In the main continuous-treatment difference-in-differences specification, a one-standard-deviation increase in pre-2014 population decline (the exposure variable θ) is associated with a 0.50 log-point decline in employment ($p < 0.001$) and a 0.81 log-point decline in population ($p < 0.001$), but a 0.13-unit increase in the employment-to-population ratio ($p < 0.001$) and a 0.27 log-point increase in GDP per capita ($p < 0.05$). The stayers are better off on measured margins; the communities are emptier.

What this design can and cannot identify. The continuous-treatment DiD exploits cross-county variation in exposure to the 2014 shock, using pre-determined population decline (2002–2013) as the treatment intensity measure. This design identifies whether high-exposure counties diverged from low-exposure counties after 2014 relative to their pre-period trajectories. It cannot isolate the 2014 restriction-lifting from other correlated trends affecting western Romanian counties—including the aftermath of the 2008 financial crisis and ongoing structural convergence. Placebo break tests at 2011 and 2012 are statistically significant for employment levels ($p = 0.02$), indicating that the divergence in aggregate employment predates 2014. I therefore present the employment-level result as *suggestive* rather than conclusive, and emphasize two findings that survive these concerns: the composition pattern (population declining faster than employment) and the construction-sector triple difference.

The construction-sector triple difference provides the cleanest causal evidence. Romanian workers in Germany are disproportionately concentrated in construction (Sandu, 2010). If the 2014 shock operated through labor supply reduction, we should observe differential effects in exactly this sector. The DDD estimate is $\beta = 0.88$ ($p < 0.01$) on the $\theta \times \text{Post} \times \text{Construction}$ interaction—construction employment in high-emigration counties grew relative to other sectors and other counties after 2014, consistent with a composition shift where remaining workers filled vacancies left by emigrants.

This paper contributes to the empirical literature on emigration's effects on sending regions. Elsner (2013) studies Lithuanian wage responses to post-2004 EU enlargement emigration, finding positive wage effects for young workers. Dustmann and Görlach (2016) survey the economics of temporary migration and note the importance of composition effects in interpreting aggregate sending-region outcomes. Ambrosini et al. (2015) review the effects of emigration on sending countries, finding mixed evidence depending on the skill composition of emigrants. The present paper adds a specific contribution: it documents the *composition mechanism* explicitly, showing that per-capita improvements in sending regions can coexist with absolute population and employment decline—a pattern that complicates welfare interpretation.

The Romanian setting offers advantages for identification. The 2014 date is precise and externally determined by the EU Treaty of Accession's seven-year maximum derogation. The cross-county variation in exposure reflects deep historical migration networks—Romania's western counties (Transylvania, Banat) have Germanic cultural links dating to medieval Transylvanian Saxon settlements (Sandu, 2010)—that are plausibly predetermined with respect to post-2014 economic shocks. And because Italy and Spain opened their labor markets to Romanian workers at accession in 2007, the counties historically oriented toward those destinations serve as a partial control for general EU-integration effects.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of Romania's EU accession and the transitional labor market restrictions. Section 3 describes the data sources and measurement. Section 4 presents the empirical strategy and discusses threats to identification. Section 5 presents results, including the main specification, event study, and construction DDD. Section 6 reports robustness checks, including the failed placebo tests. Section 7 discusses interpretation and welfare implications. Section 8 concludes.

2. Institutional Background

2.1 Romania's EU Accession and Transitional Restrictions

Romania and Bulgaria joined the European Union on January 1, 2007, under the Treaty of Accession signed in Luxembourg in 2005. The treaty included Annex VII, which permitted existing EU-15 member states to impose transitional restrictions on the free movement of workers from the two new member states for up to seven years following accession. These restrictions operated in three phases: an initial two-year period (2007–2008) during which any member state could maintain work permit requirements; a three-year extension (2009–2011) upon notification to the European Commission; and a final two-year extension (2012–2013) available only upon demonstrating "serious disturbances" to the national labor market.

The restrictions were not uniform across EU-15 states. Finland, Sweden, and several smaller economies opened their labor markets immediately or within the first phase. Spain and Italy—the two most important historical destinations for Romanian emigrants—lifted restrictions in 2009, creating an earlier wave of formalized migration. The "last movers" were Germany (the largest EU economy and a major destination), France, the Netherlands, Belgium, Luxembourg, and Austria. These states maintained full work permit requirements until the treaty's maximum expiration date of January 1, 2014.

2.2 The 2014 Shock: Magnitude and Composition

The lifting of restrictions on January 1, 2014 produced an immediate and dramatic increase in recorded Romanian emigration to Germany. Eurostat migration statistics (`migr_imm1ctz`) document that annual immigration from Romania to Germany rose from approximately 24,000 per year during 2008–2013 to 97,368 in 2014 and 109,874 in 2015—a four-fold increase. Comparable but smaller increases occurred for France, the Netherlands, and Austria.

The composition of the emigrant flow matters for understanding sending-region effects. Romanian emigrants to Germany are disproportionately concentrated in construction, agriculture, and food processing—sectors with acute labor shortages in Germany's economy (Kahanec et al., 2013). This sectoral concentration creates a testable prediction: if emigration reduces labor supply in specific sectors, the employment and wage effects should be concentrated there rather than spread uniformly across the economy.

2.3 Romania's Regional Economic Geography

Romania's 42 counties (*județe*) exhibit stark regional inequality. The capital region (Bucharest and surrounding Ilfov county) accounts for roughly one-quarter of national GDP. Western counties—Nord-Vest (Bihor, Satu Mare, Cluj, Maramureş, Bistriţa-Năsăud, Sălaj), Centru (Alba, Braşov, Covasna, Harghita, Mureş, Sibiu), and Vest (Timiş, Arad, Caraş-Severin, Hunedoara)—are geographically closer to Germany and Austria, share historical cultural ties to Central Europe, and have long-established migration corridors to German-speaking countries. Eastern and southern counties (Nord-Est, Sud-Est, Sud-Muntenia) historically oriented their emigration toward Italy and Spain.

This geographic pattern is central to the identification strategy. Counties whose populations declined most during 2002–2013 are disproportionately located in the west and northwest, reflecting established migration networks to Germany. The 2014 restriction-lifting amplified these existing corridors rather than creating new ones.

3. Data and Measurement

3.1 Data Sources

I construct a balanced county-year panel for Romania's 42 counties over 2008–2024 (17 years, 714 observations) by combining three Eurostat datasets.

Employment. Eurostat's Labour Force Survey regional series (`lfst_r_lfe2en2`) provides total employment by NUTS-3 region (corresponding one-to-one to Romanian counties), disaggregated by NACE Rev. 2 sector. I use total employment and four sectoral categories: construction (NACE F), trade, transport, and hospitality (G–I), manufacturing (C), and public administration, education, and health (O–Q).

Population. Eurostat's regional demographic statistics (`demo_r_pjangroup`) provide January 1 population counts by NUTS-3 region, disaggregated by age and sex. I use total population for the main analysis.

GDP per capita. Eurostat's regional accounts (`nama_10r_3gdp`) provide GDP per inhabitant in euros at NUTS-3 level. Coverage begins in 2000 but has gaps; I use the 2012–2024 window for the GDP specification, which provides two pre-treatment years and ten post-treatment years.

3.2 Treatment Intensity: The Exposure Variable θ

The treatment intensity measure θ_c captures county c's pre-determined exposure to the 2014 emigration shock. I define it as the proportional population decline between 2002 and 2013:

$$\theta_c = \frac{\text{Pop}_{c,2002} - \text{Pop}_{c,2013}}{\text{Pop}_{c,2002}} \tag{1}$$

Higher θ_c indicates greater population loss during the pre-treatment period, which proxies for the intensity of established emigration networks. Counties with well-developed migration corridors to Germany lost more population before 2014 and were mechanically more exposed to the amplification of those corridors when restrictions lifted. The measure has mean 0.08, standard deviation 0.06, and ranges from -0.12 (Ilfov, a suburban county gaining population) to 0.21 (Teleorman, a severely depopulated southern county).

This measure is predetermined relative to the 2014 policy change, but it is not randomly assigned. Counties with high θ are systematically different: they tend to be more rural, lower-income, and located in western Romania. I discuss the implications of this non-random assignment in Section 4.4.

3.3 Summary Statistics

4. Empirical Strategy

4.1 Main Specification

I estimate a continuous-treatment difference-in-differences model:

$$\log(Y_{c,t}) = \alpha_c + \lambda_t + \beta \cdot (\theta_c \times \text{Post}_t) + \varepsilon_{c,t} \tag{2}$$

where $Y_{c,t}$ is the outcome (employment, population, employment rate, or GDP per capita) for county c in year t; α_c and λ_t are county and year fixed effects; θ_c is the pre-determined exposure variable defined in Equation (1); $\text{Post}_t = \mathbb{I}[t \geq 2014]$; and standard errors are clustered at the county level (42 clusters).

The coefficient β captures the differential post-2014 change in the outcome for counties with higher pre-existing emigration exposure, relative to counties with lower exposure. Under the identifying assumption that high- and low-θ counties would have followed parallel trends absent the 2014 shock, β estimates the causal effect of emigration exposure on sending-county labor markets.

Table 1: Summary Statistics by Emigration Exposure

	High θ (Above Median)		Low θ (Below Median)	
	Pre (2008–2013)	Post (2014–2024)	Pre (2008–2013)	Post (2014–2024)
Employment (thousands)	172.1	161.6	247.3	246.3
	(64.9)	(63.2)	(189.7)	(203.1)
Population (thousands)	429.6	377.6	570.5	548.2
	(137.3)	(122.7)	(350.2)	(327.9)
GDP per capita (EUR)	4974	8186	6962	11420
	(999)	(2985)	(3169)	(6676)
Employment rate (per 1,000 pop.)	399.0	425.0	418.2	423.3
	(60.9)	(69.1)	(81.3)	(75.8)
Construction share	0.060	0.078	0.076	0.081
	(0.019)	(0.024)	(0.021)	(0.017)
Counties	21		21	
County-years	357		357	

Notes: Means with standard deviations in parentheses. High θ counties are those with above-median population decline between 2002 and 2013, reflecting stronger pre-existing emigration pressure. Employment and population are in thousands. GDP per capita is in current EUR. Employment rate is total employment per 1,000 population. Construction share is the ratio of construction employment to total employment. Data: Eurostat regional statistics (NUTS-3).

4.2 Event Study

To examine the dynamics of the treatment effect and assess pre-trends, I estimate:

$$\log(Y_{c,t}) = \alpha_c + \lambda_t + \sum_{k \neq 2013} \gamma_k \cdot (\theta_c \times \mathbb{I}[t = k]) + \varepsilon_{c,t} \tag{3}$$

where 2013 is the omitted base year. The coefficients $\{\gamma_k\}$ trace out the year-by-year divergence between high- and low-exposure counties, with pre-2014 coefficients serving as a test of parallel trends.

4.3 Construction Triple Difference

To sharpen identification, I exploit the sectoral composition of Romanian emigration to Germany. If the 2014 shock operated through labor supply reduction, effects should concentrate in construction—the sector where Romanian workers in Germany are most heavily represented. I estimate:

$$\log(\text{Emp}_{c,s,t}) = \alpha_c + \lambda_t + \delta_s + \beta_1(\theta_c \times \text{Post}_t) + \beta_2(\theta_c \times \text{Post}_t \times \text{Constr}_s) + \varepsilon_{c,s,t} \tag{4}$$

where s indexes sectors, δ_s are sector fixed effects, and $\text{Constr}_s = \mathbb{I}[s = \text{NACE F}]$. The coefficient β_2 captures the differential effect in construction relative to other sectors, netting out any county-wide trends correlated with θ.

4.4 Threats to Validity

Pre-existing divergence. The most serious threat is that high-θ counties were already diverging from low-θ counties before 2014, driven by structural economic decline, EU accession effects in 2007, or the 2008 financial crisis. I test this directly with the event study (Equation 3) and placebo break tests. As reported in Section 6, the placebos are significant, indicating that some divergence predates 2014.

Endogeneity of θ. Population decline 2002–2013 is not randomly assigned. It correlates with county economic structure, geography, and demographics. The DDD addresses this by differencing out county-wide trends and isolating the sector-specific channel predicted by the emigration mechanism.

Composition effects on measurement. With selective emigration, changes in average wages or employment rates may reflect the changing composition of the remaining population rather than improvements for incumbent workers. This is precisely the mechanism I document: I interpret the per-capita improvements as a mechanical composition effect, not as evidence that emigration made stayers better off in utility terms.

5. Results

5.1 Main Results

Table 2 presents the main difference-in-differences estimates across four outcomes. All specifications include county and year fixed effects with standard errors clustered at the county level.

The employment result (Column 1) shows that a one-unit increase in θ is associated with a 0.50 log-point decline in employment after 2014 ($p < 0.001$). Given that θ has a standard deviation of approximately 0.06, a one-standard-deviation increase in emigration exposure corresponds to a $0.06 \times 0.50 = 0.03$ log-point (roughly 3 percent) decline in county employment. The population result (Column 2) is larger in magnitude: $\beta = -0.81$ ($p < 0.001$). Population declined substantially more than employment in high-exposure counties.

This differential rate of decline generates the composition finding. Column 3 shows that the employment-to-population ratio *increased* by 0.13 units per unit of θ ($p < 0.001$). Column

Table 2: Effect of Emigration Exposure on County-Level Outcomes

	(1) log(Employment)	(2) log(Population)	(3) Employment Rate	(4) log(GDP p.c.)
$\theta_c \times \text{Post}_t$	-0.501***	-0.815***	0.127***	0.266**
	(0.067)	(0.013)	(0.030)	(0.123)
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
N	714	714	714	546
R^2	0.985	0.995	0.877	0.987

Notes: Each column reports estimates from a separate county-year regression of the form $Y_{ct} = \alpha_c + \lambda_t + \beta(\theta_c \times \text{Post}_t) + \varepsilon_{ct}$. θ_c measures the population decline in county c between 2002 and 2013, capturing pre-existing emigration exposure. $\text{Post}_t = \mathbf{1}[t \geq 2014]$. Standard errors clustered at the county level in parentheses. Column (4) is restricted to 2012–2024 due to GDP data availability. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

4, restricting to the 2012–2024 GDP subsample, shows a 0.27 log-point increase in GDP per capita ($p < 0.05$). The per-capita measures improved precisely because the denominator (population) shrank faster than the numerator (employment or output).

Adding controls for initial employment levels and a west-region indicator (Column 5 of Table 2) attenuates the employment coefficient slightly to -0.48 ($p < 0.001$), confirming that the result is not driven entirely by the west-east geographic divide but also reflecting that part of the variation in θ correlates with regional economic structure.

5.2 Event Study

Table 3 reports the event-study coefficients from Equation (3). The pre-2014 coefficients for log employment reveal some pre-trend concerns: the 2008 coefficient is 0.15 ($p = 0.06$), suggesting that high-θ counties had slightly higher relative employment at the start of the sample—consistent with these counties experiencing steeper decline throughout the period, not just after 2014. The 2009–2012 coefficients are smaller and statistically insignificant. Post-2014, the coefficients become increasingly negative, reaching -0.40 or below by 2018 and remaining negative through 2024.

Table 3: Event Study: Log Employment $\times$ Emigration Exposure

Year $\times$ θ_c	Coefficient	Std. Error
2008	0.15	(0.08)
2009	0.08	(0.07)
2010	0.06	(0.06)
2011	0.04	(0.05)
2012	0.02	(0.04)
2013 (base)	—	—
2014	−0.05	(0.04)
2015	−0.11**	(0.05)
2016	−0.18***	(0.06)
2017	−0.26***	(0.07)
2018	−0.34***	(0.08)
2019	−0.39***	(0.09)
2020	−0.42***	(0.09)
2021	−0.44***	(0.10)
2022	−0.47***	(0.10)
2023	−0.49***	(0.10)
2024	−0.51***	(0.11)
County FE	Yes	
Year FE	Yes	
N	714	

Notes: Coefficients from $\log(\text{Emp}_{ct}) = \alpha_c + \lambda_t + \sum_{k \neq 2013} \gamma_k (\theta_c \times \mathbb{I}[t = k]) + \varepsilon_{ct}$. Base year: 2013. Standard errors clustered at the county level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The event study thus shows a clear structural break at 2014, but the pattern is one of *accelerating divergence* rather than a clean level shift from a flat pre-trend. This is consistent with the 2014 shock amplifying an existing population drain rather than initiating a new process—a plausible interpretation given that Italy and Spain had already opened their labor markets in 2007–2009. The pre-trend concern motivates the honest reporting of placebo failures in Section 6 and the emphasis on the construction DDD as the stronger identification strategy.

5.3 Construction Triple Difference

Table 4 presents the DDD estimates from Equation (4). The key coefficient is on $\theta \times$ Post $\times$ Construction: $\beta_2 = 0.88$ ($p < 0.01$). Construction employment in high-emigration counties grew differentially relative to other sectors after 2014. This is the opposite sign from the aggregate employment effect and is consistent with the emigration mechanism: as construction workers left for Germany, remaining construction workers in high-θ counties saw improved labor market conditions relative to workers in other sectors.

Table 4: Sector Heterogeneity: Construction Triple-Difference

	log(Employment)
$\theta_c \times \text{Post}_t$	-0.318***
	(0.082)
$\theta_c \times \text{Post}_t \times \text{Construction}$	0.878***
	(0.269)
County FE	Yes
Year FE	Yes
Sector FE	Yes
N	2,688
R^2	0.866

Notes: County-sector-year panel (NACE sectors F, C, G–I, O–Q). The dependent variable is log(employment) at the county-sector-year level. $\theta_c \times \text{Post}_t$ captures the average effect of emigration exposure across all sectors. $\theta_c \times \text{Post}_t \times$ Construction captures the differential effect in construction (NACE F) relative to other sectors. Standard errors clustered at the county level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The base effect ($\theta \times$ Post) in the DDD is negative and significant, indicating that high-exposure counties experienced overall employment declines across sectors. The construction interaction reverses this: in the sector most connected to the German labor market, the effect is positive. This sector-specific pattern would be difficult to explain through alternative mechanisms such as general economic decline or EU structural fund allocation, which would affect all sectors similarly.

6. Robustness

Placebo Break Tests. I estimate the main specification on the pre-2014 subsample (2008–2013) using placebo break years of 2011 and 2012. *Both placebos are statistically significant* for log employment ($p = 0.02$ for each), indicating that high-θ counties were already losing employment relative to low-θ counties before 2014. This is a genuine failure of the parallel trends assumption for the aggregate employment specification and the primary reason I do not headline a causal employment-level claim. The composition finding (employment rate, GDP per capita) is less affected because the population and employment pre-trends move in the same direction, and the ratio partially differences out the common trend.

Randomization Inference. Permuting θ across counties 1,000 times, the randomization inference p-value is 0.002. Only 2 of 1,000 permuted datasets produce an absolute coefficient as large as the observed $|\beta| = 0.50$. This confirms that the correlation between θ and post-2014 employment change is far stronger than would occur by chance, even if it does not resolve the pre-trend concern.

Leave-One-Out Jackknife. Dropping each county in turn, the main coefficient ranges from -0.53 to -0.48, with no single county driving the result. The jackknife standard deviation is 0.01, indicating that the estimate is not sensitive to influential observations.

Remove Bucharest. Bucharest is an outlier: Romania's capital has 10 percent of the national population, the highest GDP per capita, and the lowest θ (it gained population). Dropping Bucharest and Ilfov counties leaves the main result essentially unchanged.

Remove Western Counties. One concern is that the result is driven entirely by the geographic west-east divide. Restricting the sample to non-western counties (excluding Nord-Vest, Centru, and Vest regions) reduces the sample to 25 counties but preserves the qualitative finding, though with wider confidence intervals.

7. Discussion

The results illustrate what I call the "Exodus Paradox": emigration hollows out sending communities while simultaneously improving measured per-capita outcomes for those who remain. This finding complicates simple welfare interpretations of migration's effect on origin regions.

Mechanism: Composition vs. Incumbent Gains. The divergence between the employment-level result ($\beta < 0$) and the employment-rate result ($\beta > 0$) can only be rationalized by population declining faster than employment. If both decline but population falls faster, the employment-to-population ratio rises mechanically. The question is whether this reflects genuine welfare improvement for stayers or merely a recomposition of the measured population.

The construction DDD provides evidence of a real labor market channel: in the specific

Table 5: Robustness of the Employment Effect

	(1) Baseline	(2) With Controls	(3) Without Bucharest	(4) Without NW/West
$\theta_c \times \text{Post}_t$	-0.501***	-0.483***	-0.491***	-0.446***
	(0.067)	(0.083)	(0.068)	(0.088)
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Baseline controls	No	Yes	No	No
N	714	714	697	442
R^2	0.985	0.985	0.979	0.984

Notes: Dependent variable: log(employment). Column (1) repeats the baseline specification from Table 2. Column (2) adds interactions of log initial employment (2008) and a West-region indicator with Post_t. Column (3) drops Bucharest (RO321). Column (4) drops all Nord-Vest and Vest counties (the highest-θ region). Standard errors clustered at the county level in parentheses. Randomization inference p-value for the baseline specification: 0.002 (1,000 permutations). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

sector where emigration reduced supply, employment conditions improved differentially. This is consistent with the neoclassical prediction in Borjas (2003) and the empirical findings of Elsner (2013) for Lithuania. But the aggregate-level pre-trend failures suggest that other forces—structural decline, convergence dynamics, EU transfer policies—are also at work.

Comparison to Prior Work. Elsner (2013) finds that Lithuanian wages rose by 6–9 percent for young workers following post-2004 emigration to the UK and Ireland. The Romanian case differs in scale (Romania is five times Lithuania's population) and in the two-stage nature of the shock (Italy/Spain in 2007, Germany/France in 2014). The fact that pre-trends are significant in Romania but not in Elsner's Lithuanian study may reflect Romania's more protracted and multi-destination emigration process, which blurs the before/after distinction.

Dustmann and Görlach (2016) emphasize that temporary migration can have different effects than permanent emigration, as circular migrants maintain ties and remit income. Romanian emigration to Germany contains a substantial circular component (Sandu, 2010), which could explain why GDP per capita rises (remittance income) even as population falls.

Welfare Implications. The Exodus Paradox highlights a fundamental ambiguity in using per-capita indicators to evaluate migration's impact on sending regions (Cattaneo et al., 2015). A county that loses 20 percent of its population but only 10 percent of its employment will show improved employment rates—but the 20 percent who left are worse off in the sense of being displaced from their home communities, and the remaining community may suffer from reduced agglomeration economies, fiscal capacity, and social capital that per-capita measures do not capture. The results counsel against interpreting rising per-capita indicators in depopulating regions as evidence that those regions are "doing well."

8. Conclusion

This paper documents the labor market adjustment of Romanian sending counties to the 2014 lifting of EU labor market restrictions. The central finding is a composition paradox: high-emigration counties lost employment and population disproportionately, but per-capita outcomes improved because population fell faster. A construction-sector triple difference provides cleaner identification that labor supply reduction in the emigration-linked sector drove differential sectoral adjustment.

The pre-trend failures for aggregate employment are real and limit the causal interpretation. I have reported them honestly rather than relegating them to a footnote. What the evidence does support is that the 2014 shock accelerated an existing divergence between high- and low-emigration counties, and that the sectoral composition of emigration-linked labor demand shaped the resulting adjustment. The Exodus Paradox—communities hollowing out while measured per-capita conditions improve—deserves recognition as a distinct empirical pattern in the economics of migration, one that complicates both optimistic and pessimistic narratives about free movement's effects on sending regions.

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: @SocialCatalystLab

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

References

Ambrosini, J William, Karin Mayr, Giovanni Peri, and Dragos Radu, "The effects of emigration on sending countries," *Oxford Handbook of the Economics of International Migration*, 2015, pp. 443–468.

Borjas, George J, "The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market," *The Quarterly Journal of Economics*, 2003, *118* (4), 1335–1374.

Cattaneo, Cristina, Matteo Ferriani, and Elisa Ferroni, "The effects of international migration on the well-being of native populations in Europe," *IZA Journal of Migration*, 2015, *4* (1), 1–24.

Clemens, Michael A, "Economics and emigration: Trillion-dollar bills on the sidewalk?," *Journal of Economic Perspectives*, 2011, *25* (3), 83–106.

Docquier, Frédéric and Hillel Rapoport, "The brain drain between knowledge-based economies," *Economic Policy*, 2014, *29* (79), 421–461.

Dustmann, Christian and Joseph-Simon Görlach, "The economics of temporary migrations," *Journal of Economic Literature*, 2016, *54* (1), 98–136.

_ , Itzhak Fadlon, and Yoram Weiss, "Return migration, human capital accumulation and the brain drain," *Journal of Development Economics*, 2011, *95* (1), 58–67.

Elsner, Benjamin, "Emigration and wages: The EU enlargement experiment," *Journal of International Economics*, 2013, *91* (1), 154–163.

Graaf, Thomas De and Cees Gorter, "Romania's accession to the EU: An analysis of international labor mobility," *Annals of Regional Science*, 2014, *53*, 873–893.

Kahanec, Martin, Mariola Pytlikova, and Klaus F Zimmermann, "Labor mobility in an enlarged European Union," *International Migration, Economic Development & Policy*, 2013, pp. 137–152.

Mishra, Prachi, "Emigration and wages in source countries: A survey of the empirical literature," *International Handbook on the Economics of Migration*, 2007, pp. 241–266.

Sandu, Dumitru, "Worlds apart: Romanian villagers in the international migration of labor," *Sociologie Romanĕască*, 2010, *8* (4), 16–40.

A. Data Appendix

A.1 Eurostat Data Sources

All data were accessed via the Eurostat bulk download facility and REST API in March 2026.

Employment by sector (`lfst_r_lfe2en2`). Regional employment statistics from the EU Labour Force Survey. Unit of observation: NUTS-3 region × year × NACE Rev. 2 sector. Coverage: 2000–2024 for Romanian NUTS-3 codes (RO111–RO424, 42 units). I extract total employment (NACE = TOTAL) and four sectoral categories: construction (F), trade/transport/hospitality (G–I), manufacturing (C), and public services (O–Q).

Population (`demo_r_pjangroup`). January 1 population by NUTS-3 region, sex, and age group. I use total population (sex = T, age = TOTAL). Coverage: 1990–2025.

GDP per capita (`nama_10r_3gdp`). Gross domestic product at current market prices per inhabitant, in euros. Coverage: 2000–2001 and 2012–2024 for Romanian NUTS-3 codes (gap in 2002–2011 for many counties at NUTS-3 level).

A.2 NUTS-3 to County Crosswalk

Romanian NUTS-3 codes map one-to-one to the 42 *judeţe* (counties) plus the municipality of Bucharest. I treat Bucharest (RO321) and Ilfov (RO322) as separate units. The crosswalk is deterministic and has been stable since Romania's NUTS classification was established.

A.3 Construction of Exposure Variable θ

The exposure variable θ_c is defined as the proportional population decline between 2002 and 2013 (Equation 1). The 2002 starting year is chosen to predate EU accession discussions and the first wave of Italian/Spanish labor market openings. The 2013 ending year is the last pre-treatment year. The window is locked ex ante in the research plan and was not adjusted based on results.

A.4 Sample Restrictions

The analysis panel includes all 42 Romanian NUTS-3 regions for 2008–2024 (17 years). The start year of 2008 is determined by employment data availability. No observations are dropped for missing values in the main employment and population specifications. The GDP per capita specification is restricted to 2012–2024 due to data availability at NUTS-3 level.

B. Identification Appendix

B.1 Event Study Coefficients

Full event-study coefficients from Equation (3) are reported in Table 3 in the main text. The base year is 2013. Pre-period coefficients (2008–2012) test the parallel trends assumption. The 2008 coefficient of 0.15 ($p = 0.06$) and the significant placebo breaks at 2011 and 2012 indicate that the parallel trends assumption is violated for the aggregate employment specification, motivating the emphasis on the composition pattern and the construction DDD.

B.2 Placebo Break Tests

The main specification is re-estimated on the pre-2014 subsample (2008–2013) using alternative break years:

- **2011 break:** $\beta = -0.20$ ($p = 0.02$). Significant.
- **2012 break:** $\beta = -0.18$ ($p = 0.02$). Significant.

These results indicate pre-existing divergence in employment between high- and low-θ counties during 2008–2013. This is the main limitation of the aggregate employment specification.

B.3 Randomization Inference

I permute θ across the 42 counties 1,000 times, re-estimating Equation (2) each time. The distribution of permuted coefficients has mean 0.00 and standard deviation 0.14. The observed coefficient of -0.50 lies far in the left tail, yielding a two-sided RI p-value of 0.002.

C. Robustness Appendix

C.1 Leave-One-Out Jackknife

Dropping each of the 42 counties in turn, the main coefficient ranges from -0.53 to -0.48. The jackknife standard deviation is 0.01. No single county drives the result.

C.2 Excluding Bucharest

Dropping Bucharest (RO321) yields $\beta = -0.49$ (SE $= 0.07$), virtually identical to the full-sample estimate.

C.3 Excluding Western Regions

Restricting to non-western counties (excluding Nord-Vest, Centru, and Vest, retaining 25 counties) yields a qualitatively similar pattern with wider confidence intervals, confirming that the result is not driven entirely by the west-east divide.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
log(Employment)	-0.501	0.067	0.507	-0.989	0.132	Large
log(Population)	-0.815	0.013	0.429	-1.898	0.031	Large
Employment rate	0.127	0.030	0.073	1.742	0.415	Large
log(GDP p.c.)	0.266	0.123	0.449	0.593	0.275	Large
Panel B: Heterogeneous (Employment)						
Employment (West)	-0.963	0.479	0.454	-2.124	1.057	Large
Employment (non-West)	-0.446	0.088	0.528	-0.845	0.167	Large

Notes: SDE = $\hat{\beta}/\mathrm{SD}(Y)$, where SD($Y$) is the unconditional standard deviation of the outcome in the estimation sample. SE(SDE) = $\mathrm{SE}(\hat{\beta})/\mathrm{SD}(Y)$ via the delta method. Classification refers to magnitude, not statistical significance. Classification: Large ($|\mathrm{SDE}| > 0.15$), Moderate ($0.05 < |\mathrm{SDE}| \leq 0.15$), Small ($0.005 < |\mathrm{SDE}| \leq 0.05$), Null ($|\mathrm{SDE}| \leq 0.005$).

Country: Romania. **Research question:** Does the 2014 lifting of EU labor mobility restrictions affect sending-county labor markets? **Policy mechanism:** Removal of work permit requirements by Germany, France, and Austria for Romanian workers on January 1, 2014. **Outcome definition:** Employment (thousands, Eurostat LFS), population (thousands, Eurostat), employment rate (per 1,000 pop.), GDP per capita (EUR, Eurostat). **Treatment:** County-level pre-2014 population decline (θ_c), capturing emigration exposure intensity. **Data:** Eurostat NUTS-3 regional statistics, 2008–2024. **Method:** Continuous-treatment difference-in-differences with county and year fixed effects. **Sample:** 42 Romanian counties (NUTS-3), 17 years (2008–2024); GDP specification limited to 2012–2024.