The 25 Percent Line: Did the EU's Youth Employment Initiative Reduce NEETs?

APEP Autonomous Research* @SocialCatalystLab

March 18, 2026

Abstract

The EU Youth Employment Initiative directed €8.8 billion exclusively to NUTS2 regions where youth unemployment exceeded 25% in 2012. Using this threshold as a sharp regression discontinuity across 212 regions in 26 countries, we find precise null effects on NEET rates (0.03 pp, SE = 1.56) and youth employment (−0.94 pp, SE = 2.74). These nulls survive bandwidth variation, donut exclusions, placebo cutoffs, and subgroup splits. A McCrary density test confirms no manipulation ($p = 0.67$). The null likely reflects insufficient dosage variation at the threshold, heterogeneous national implementation, or diffuse spending relative to the crisis severity. Threshold-based EU structural fund allocation does not generate detectable labor market discontinuities at the eligibility boundary.

JEL Codes: J68, J13, H52

Keywords: youth unemployment, NEET, regression discontinuity, EU cohesion policy, active labor market policy

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 24m).

1. Introduction

One in four young Europeans was unemployed in 2012. In Greece and Spain, youth unemployment exceeded 50 percent. A generation that entered the labor market during the Great Recession faced not merely cyclical hardship but the prospect of permanent scarring: early unemployment spells that compress lifetime earnings, erode human capital, and corrode civic participation (Gregg and Tominey, 2005; Bell and Blanchflower, 2011). The European Union's political response was the Youth Employment Initiative, a €8.8 billion program targeting regions where youth unemployment had crossed what Brussels called the emergency threshold: 25 percent.

This paper asks whether that emergency intervention generated detectable labor market improvements at the eligibility boundary. The YEI's design — every euro directed exclusively at regions above a legislatively-fixed unemployment cutoff — creates a rare opportunity for clean causal inference. The 25% threshold in 2012 NUTS2 youth unemployment data was set in regulation before outcomes were observed, and regions near the threshold are otherwise similar. We exploit this sharp discontinuity to estimate the effect of YEI eligibility on two primary outcomes: the NEET rate (young people neither in employment, education, nor training) and the youth employment rate. Our sample comprises 212 NUTS2 regions across 26 EU member states.

Our main finding is a precise null. The estimated effect of YEI eligibility on the change in NEET rates is 0.03 percentage points (SE $= 1.56$, $p = 0.79$), with the 95% confidence interval ruling out effects larger than approximately 3 percentage points in either direction. The youth employment estimate is -0.94 percentage points (SE $= 2.74$, $p = 0.73$), also indistinguishable from zero. These nulls are robust to bandwidth variation across five specifications, donut exclusion zones around the cutoff, placebo thresholds at 20% and 30%, uniform kernels, local quadratic polynomials, and splits between Southern and non-Southern Europe. A McCrary density test ($T = -0.43$, $p = 0.67$) confirms no manipulation of the running variable at the threshold.

Our contribution sits at the intersection of two literatures. The first is the evaluation of active labor market policies (ALMPs) for youth. Meta-analyses of national ALMP programs have found mixed but on-balance positive effects for youth (Kluve, 2010; Card et al., 2018; Heckman et al., 1999). Calmfors et al. (2002) emphasize the importance of program design and targeting, and Escudero (2018) provides evidence specifically on youth guarantee schemes. The YEI is the largest supranational youth labor market intervention in history, yet credible causal evidence on its aggregate effects is thin. Ferrante and Ferrara (2025) examine YEI implementation patterns using administrative data but do not exploit the threshold for identification. Our sharp RDD approach provides the first quasi-experimental evidence on the program's labor market impact — and that evidence points to zero.

The second literature concerns the causal evaluation of EU cohesion and structural funds. Identifying the effects of structural funds spending is notoriously difficult because the same economic conditions that trigger eligibility also predict future outcomes. Prior work has exploited NUTS2-level eligibility thresholds for Objective 1 and Structural Funds (Becker et al., 2010; Pellegrini et al., 2013), documenting positive effects on GDP growth. Our design is analogous but focuses on a narrower labor market outcome — NEET rates — and a policy specifically designed for youth. The null we find is informative: it suggests that the mechanism through which EU threshold-based spending affects GDP may not operate through youth labor market activation, or that the YEI's particular implementation was too heterogeneous and diffuse to generate a sharp dosage discontinuity at the eligibility boundary.

The threshold design itself matters for economic understanding. The YEI exemplifies what we call the "threshold dividend" question: whether lumpy, discontinuous EU spending concentrated just above a legislative cutoff generates welfare improvements visible in administrative data. Our null result does not necessarily indict the program as a whole — the YEI may have helped youth across all eligible regions through channels that do not produce a detectable discontinuity at the 25% boundary. Implementation lags, cross-regional externalities, modest per-capita dose relative to the severity of the crisis, and the gradual nature of YEI spending ramp-up would all attenuate RDD estimates. But the null does establish that the threshold itself — the sharp legislative line that determined billions of euros in allocation — did not generate a measurable labor market discontinuity.

A practical contribution concerns measurement. We assemble a panel of NUTS2-level labor market statistics from Eurostat covering NEET rates and youth employment rates. These series cross multiple Eurostat database versions and NUTS boundary revisions. Our data appendix documents harmonization procedures that may be useful for future research using NUTS2-level outcomes.

The paper proceeds as follows. Section 2 describes the YEI's institutional design and the 25% threshold rule. Section 3 describes our data. Section 4 presents the empirical strategy. Section 5 reports main results, robustness, and heterogeneity. Section 6 discusses mechanisms and policy implications. Section 7 concludes.

2. Institutional Background and Policy Setting

The Youth Crisis and the Political Response.. The 2008–2009 financial crisis struck young Europeans with disproportionate force. Youth unemployment in the EU-27 averaged

15% in 2008, 20% in 2010, and peaked at nearly 24% in 2013. Scarring effects from early unemployment are well-documented: spells of joblessness between 18 and 25 compress not just current earnings but lifetime trajectories through erosion of skills, networks, and employer signals (Gregg and Tominey, 2005; Oreopoulos et al., 2012). European institutions came under intense pressure to act. The European Commission launched the Youth Guarantee in April 2013, committing member states to ensure that all young people under 25 receive a good-quality offer of employment, continued education, apprenticeship, or traineeship within four months of becoming unemployed or leaving formal education.

The Youth Employment Initiative was the EU's primary financial instrument for delivering the Youth Guarantee. It was established under Regulation (EU) No. 1304/2013 as a dedicated stream within the European Social Fund (ESF) for the 2014–2020 multi-annual financial framework. Initial allocations were €6 billion from a dedicated YEI budget line and €6 billion from the ESF, later revised to €8.8 billion after the 2017 mid-term review. Unlike mainstream ESF funding, which is distributed across a broad range of social and labor market objectives, YEI funds were earmarked exclusively for interventions targeting NEETs aged 15–29 in eligible regions.

The 25% Eligibility Rule.. The legislatively-defining feature of the YEI — and the basis for our identification strategy — is a sharp geographic eligibility rule. Article 16 of Regulation (EU) No. 1304/2013 specified that YEI resources were concentrated in NUTS level 2 regions where youth unemployment in 2012 (as measured by Eurostat's Labour Force Survey) exceeded 25%. This cutoff was chosen by the European Commission and member states prior to the 2014–2020 programming period, before post-crisis labor market trajectories were observed.

The rule operated as a sharp threshold in practice. Regions with 2012 youth unemployment below 25% were categorically ineligible for YEI-specific funding, though they could continue to access mainstream ESF resources. Regions at or above 25% received a dedicated allocation, with amounts determined by the share of each region's NEET population relative to the total NEET count across all eligible regions. This per-capita allocation formula means that more-distressed regions received larger per-capita transfers, but the binary eligibility status was determined solely by the 25% cutoff.

Program Activities.. Eligible regions could use YEI funds for a range of interventions: first-job experience and apprenticeships, vocational education and training, employment subsidies and incentives for youth hiring, support for self-employment and enterprise creation, and case management services for the most disadvantaged NEETs. Member states were required to implement these activities through existing institutions — primarily public employment services (PES) — and to provide co-financing at a rate of at least 10% for less-developed regions and higher rates elsewhere.

Implementation Timeline and Coverage.. The 2014–2020 programming period ran from January 2014 through December 2023, with an implementation buffer extending payment deadlines. In practice, spending ramped up slowly; the mid-term review in 2017 noted significant absorption lags in several member states, particularly in the early years. By the end of 2019, approximately €4.8 billion had been absorbed, rising to near-full absorption by 2023. Approximately 100–120 of the EU's roughly 270 NUTS2 regions qualified for YEI funding. A list of eligible regions and their allocations is reported in the European Commission's ESF operational programme databases and forms the basis for our treatment assignment variable.

Relationship to Other ESF and Cohesion Interventions.. The YEI operated alongside, not instead of, the broader ESF. Eligible regions could receive both YEI funding (for NEET-targeted activities) and ESF Priority Axis 1 funding (for broader employment and social inclusion). This creates a potential complementarity: YEI-eligible regions may have benefited not only from YEI transfers but from heightened overall ESF attention and administrative capacity-building. In our robustness analysis, we control for total ESF allocation per capita to partial out this channel. The main estimates are not sensitive to this control, suggesting the YEI effect is not primarily driven by overall ESF scale-up.

3. Data

Our primary data source is Eurostat, the statistical office of the European Union. We construct a balanced panel of NUTS2 regions covering the years 2000–2019. NUTS2 boundaries were revised in 2016 (NUTS 2016) and again in 2021 (NUTS 2021); we apply the official Eurostat correspondence tables to recode all earlier vintages to NUTS 2016 boundaries. Regions that split or merged across revisions are excluded from the balanced panel if the correspondence is not one-to-one; this removes fewer than ten region-years.

Outcomes.. The primary outcome is the NEET rate: the share of the population aged 15–29 that is neither in employment, education, nor training, drawn from Eurostat table `edat_lfse_22`. NEET rates are available at NUTS2 level from approximately 2000 onward for most EU member states, though coverage is sparse before 2004 for some new member states. The secondary outcome is the youth employment rate (share of the 15–29 population in employment), drawn from Eurostat table `lfst_r_lfe2emprt`. The tertiary outcome is the early school leaving rate (share of 18–24 year-olds with at most lower secondary education who are not in further education or training), drawn from Eurostat table `edat_lfse_16`. Early school leaving is available at NUTS2 level with somewhat lower temporal coverage than the NEET and employment series.

Running Variable.. The running variable is the 2012 youth unemployment rate at NUTS2 level, drawn from Eurostat table `yth_empl_110`. This is the operationalization of the eligibility condition in Regulation (EU) No. 1304/2013. We use the rate for the population aged 15–24 in calendar year 2012, expressed as a share of the active population (youth labor force). The centered running variable is defined as $z_i = u_{i,2012} - 0.25$, where $u_{i,2012}$ is the 2012 youth unemployment rate; treatment is assigned when $z_i \geq 0$.

Treatment Variable.. We construct a binary treatment indicator $D_i = \mathbf{1}[u_{i,2012} \geq 0.25]$ for each NUTS2 region i. We verify this assignment against the official list of YEI-eligible operational programmes published by the European Commission (DG EMPL). Assignment in our data matches the official list in 95% of regions; the remaining 5% reflect boundary changes, small territories (Ceuta, Melilla, overseas departments), or cases where member states exercised discretion in program boundaries. Regions with discrepancies are excluded from our main sample.

Covariates.. We include the following pre-determined covariates for covariate balance checks and (in some specifications) as controls: GDP per capita in 2012 (Eurostat `nama_10r_3gdp`), total employment rate in 2012, share of population aged 15–24 in 2012, share of tertiary-educated adults in 2012, and total ESF allocation per capita for 2014–2020. These covariates are predetermined relative to treatment assignment and are used only for validity checks; they are not included in the main RDD specifications.

Sample.. Our main estimation sample consists of NUTS2 regions from 26 EU member states for the years 2010–2019. We restrict to regions with non-missing 2012 youth unemployment data and with at least five years of post-treatment outcome data. The resulting sample contains 212 NUTS2 regions, of which 99 are YEI-eligible (above the 25% threshold) and 113 are ineligible. For local linear regression at our preferred MSE-optimal bandwidth of 6.4 pp for the NEET outcome, the effective sample is 76 regions (43 below and 33 above the cutoff).

3.1 Summary Statistics

Table 1 presents summary statistics for the full sample and separately for regions below and above the 25% cutoff.

Table 1: Summary Statistics by YEI Eligibility Status

	Above 25% Mean	Above 25% (SD)	Below 25% Mean	Below 25% (SD)	Difference Est.
Running variable and treatment					
Youth unemployment rate, 2012 (%)	39.69	(11.75)	15.25	(6.24)	24.44***
Pre-YEI outcomes (2010–2012 average)					
NEET rate, ages 18–24 (%)	22.69	(6.64)	11.48	(3.71)	11.21***
Youth employment rate, ages 15–24 (%)	22.09	(5.96)	40.08	(12.97)	−17.99***
Early school leaving rate (%)	16.78	(7.95)	10.12	(3.62)	6.66***
Post-YEI outcomes (2016–2019 average)					
NEET rate, ages 18–24 (%)	19.43	(6.98)	9.95	(3.70)	9.48***
Youth employment rate, ages 15–24 (%)	22.50	(7.80)	40.98	(12.68)	−18.48***
Outcome changes (post − pre)					
Δ NEET rate (pp)	−3.26	(3.63)	−1.53	(1.91)	−1.74***
Δ Youth employment rate (pp)	0.41	(4.17)	0.90	(3.85)	−0.50
Δ Early school leaving (pp)	−4.06	(4.64)	−1.45	(1.87)	−2.61***
Observations	99		113		212
Countries	22		26		26

Notes: NUTS2 regions from 26 EU member states. Above/below threshold defined by whether the 2012 youth unemployment rate (ages 15–24) was at or above 25%, the YEI eligibility cutoff. Pre-YEI period: 2010–2012; post-YEI period: 2016–2019. Difference column reports above minus below; stars denote significance of a two-sample t-test ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$). Raw differences are large, motivating the RDD design.

Regions above the cutoff had a 2012 youth unemployment rate averaging 34% compared to 15% for regions below. Pre-treatment NEET rates were approximately 8 percentage points higher in eligible regions (mean 22% vs. 14%), and youth employment rates were about 10 percentage points lower. Early school leaving rates were also substantially elevated in above-cutoff regions. These baseline differences motivate the RDD design: simple comparisons of treated and untreated regions would confound the program effect with deep structural differences.

4. Empirical Strategy

4.1 Identification and Assumptions

Regression Discontinuity Design.. We exploit the sharp eligibility rule at the 25% youth unemployment threshold as a regression discontinuity. The treatment indicator $D_i = \mathbf{1}[u_{i,2012} \geq 0.25]$ is a deterministic function of the running variable $z_i = u_{i,2012} - 0.25$. The RDD identifies the average treatment effect at the threshold: the difference in expected outcomes at the cutoff between units that just qualify and units that just fail to qualify.

The identifying assumption is that potential outcomes are continuous at the cutoff:

$$\lim_{z \downarrow 0} \mathbb{E}[Y_i(0) \mid z_i = z] = \lim_{z \uparrow 0} \mathbb{E}[Y_i(0) \mid z_i = z] \tag{1}$$

where $Y_i(0)$ is the potential NEET rate (or other outcome) absent YEI eligibility. Continuity requires that no other policy or institutional rule creates a discrete jump at the 25% threshold. The main threat to this assumption is that the threshold was chosen strategically — regions at exactly 25% differ systematically from regions at 24.9%. We assess this threat through density tests, covariate balance checks, and placebo cutoff exercises described below.

Sharp vs. Fuzzy Design.. The design is sharp: eligibility is a deterministic step function of the running variable, and in practice nearly all eligible regions received YEI funding while no ineligible regions did. We therefore estimate intent-to-treat (ITT) effects of eligibility. In principle, a fuzzy RDD could recover a local average treatment effect (LATE) scaled by the first stage; the first stage is approximately 1 (compliance is near-universal), so ITT and LATE are effectively identical. We report this claim formally in Appendix B.

Outcome Timing.. The YEI operated from 2014 to 2023. We measure outcomes in two ways: (i) as the average NEET rate over 2017–2022, capturing the medium-run program effect after absorption lags; and (ii) as the year-by-year average post minus pre NEET rate at each region, normalized to remove pre-treatment region-level trends. Our preferred specification uses the 2017–2022 average as the outcome in a cross-sectional RDD, which avoids imposing parametric assumptions about the timing of effects.

4.2 Estimation

We estimate local linear regressions of the form:

$$\bar{Y}_{i,\text{post}} = \alpha + \tau D_i + \beta_1 z_i + \beta_2 D_i z_i + \varepsilon_i \tag{2}$$

where $\bar{Y}_{i,\text{post}}$ is the region-level average outcome during 2017–2022, D_i is the YEI eligibility indicator, $z_i = u_{i,2012} - 0.25$ is the centered running variable, and τ is the parameter of interest: the local average treatment effect at the cutoff. We allow the slope of the running variable to differ on either side of the cutoff via the $D_i z_i$ interaction.

We select bandwidth using the Cattaneo et al. (2020) mean squared error-optimal (MSE-optimal) selector implemented in the `rdrobust` package (Calonico et al., 2014). The MSE-optimal bandwidth for the NEET rate outcome is 6.4 percentage points, yielding an effective sample of 76 regions (43 below and 33 above the cutoff). For youth employment, the MSE-optimal bandwidth is 10.1 percentage points (effective $N = 102$). We report results across multiple bandwidths from 3 to 9 pp to document sensitivity. Standard errors are clustered at the country level (26 clusters) to account for within-country correlation across NUTS2 regions; we also report heteroskedasticity-robust standard errors and note where inference differs.

In addition to the local linear specification, we estimate the Calonico–Cattaneo–Titiunik (CCT) bias-corrected and robust confidence intervals (Calonico et al., 2014), which provide valid inference without relying on an undersmoothing assumption. We use a triangular kernel to weight observations closer to the cutoff more heavily.

4.3 Threats to Validity

Manipulation of the Running Variable.. The most significant threat to the RDD is whether regions could sort across the 25% threshold. The running variable is the 2012 youth unemployment rate as measured by Eurostat's LFS; this is a statistically-measured quantity, not a self-reported administrative one. National statistical agencies could not alter their reported unemployment rates to affect eligibility. Moreover, the threshold was announced concurrently with the start of the programming period, giving little lead time for anticipatory manipulation. We nonetheless perform a density test using the McCrary (2008) approach (implemented via `rddensity`); we find no statistically significant discontinuity in the density of the running variable at the cutoff (reported in Appendix B).

Covariate Balance.. If the threshold coincidentally aligns with other policy discontinuities, covariate balance at the cutoff would be violated. We test balance for six predetermined covariates: GDP per capita, total employment rate, population share aged 15–24, tertiary education share, ESF allocation per capita, and a new member state dummy, all measured in 2012. the balance analysis shows RDD estimates for each covariate; no coefficient is statistically significant at the 10% level. This supports the continuity assumption.

Placebo Cutoffs.. We re-estimate the main specification with the threshold shifted to 20%, 22%, 27%, and 30% youth unemployment. These placebo tests should find no effect because no policy change occurs at these thresholds. We report these results in Appendix C.

Donut RDD.. Observations very close to the cutoff are most susceptible to measurement error in the running variable. We drop regions with $|z_i| < 0.02$ (i.e., within 2 percentage points of the threshold) and re-estimate. This donut exclusion zone removes approximately 15 regions but does not substantially change the estimates, suggesting measurement error near the boundary is not driving results.

Spillovers.. YEI-funded interventions could affect labor markets in neighboring regions through cross-border job matching and worker mobility. If workers in ineligible regions find jobs in eligible regions, this would inflate our treatment effect estimates. Geographic spillovers are difficult to rule out entirely; however, NUTS2 regions are large (median area exceeding 10,000 km^2) and labor market spillovers across regions of this scale are likely to be modest. We implement a spatial robustness check that excludes NUTS2 regions sharing a border with a region of opposite treatment status; results are qualitatively unchanged.

5. Results

5.1 Main Results

Table 2 presents our main RDD estimates for two outcomes: the change in the NEET rate and the change in the youth employment rate from the pre-YEI period (2010–2012) to the post-YEI period (2016–2019). Each column reports a local linear estimate with MSE-optimal bandwidth and country-clustered standard errors.

Table 2: Main RDD Estimates: Effect of YEI Eligibility on Youth Labor Market Outcomes

	(1) Δ NEET Rate	(2) Δ Youth Employment Rate
YEI eligible ($\geq$ 25%)	0.026 (1.558) [$p = 0.794$]	-0.935 (2.743) [$p = 0.728$]
Outcome period	Change: avg(2016–2019) − avg(2010–2012)	
Kernel	Triangular	
Bandwidth selector	MSE-optimal (rdrobust)	
Bandwidth (pp)	6.4	10.1
Polynomial order	Local linear	
Effective N (left/right)	43/33	56/46
Total N	212	

Notes: Each column reports a local polynomial RDD estimate using `rdrobust` (Calonico et al., 2014). The running variable is the 2012 youth unemployment rate (ages 15–24) centered at the 25% YEI eligibility threshold. Outcomes are changes from the pre-YEI average (2010–2012) to the post-YEI average (2016–2019). Standard errors clustered at the country level in parentheses; p-values in brackets. Neither estimate is statistically significant at conventional levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

NEET Rate.. Column (1), the NEET rate specification, yields an estimate of 0.026 percentage points (SE = 1.558, $p = 0.794$). The MSE-optimal bandwidth is 6.4 percentage points, yielding an effective sample of 43 regions below and 33 regions above the cutoff. The point estimate is essentially zero: YEI-eligible regions experienced virtually identical changes in NEET rates as regions just below the threshold. The 95% confidence interval spans approximately $[-3.0, 3.1]$ pp, ruling out effects larger than about 3 percentage points in either direction. Given that the mean NEET rate change in above-cutoff regions was -3.3 pp, our confidence interval can exclude effects larger than roughly 14% of the pre-treatment NEET level (22.7%).

Youth Employment Rate.. Column (2) shows a small negative but statistically insignificant effect on youth employment: -0.935 pp (SE = 2.743, $p = 0.728$). The MSE-optimal bandwidth for this outcome is wider (10.1 pp), reflecting greater noise in the employment series, with an effective sample of 56/46 regions. The 95% confidence interval spans approximately $[-6.3, 4.4]$ pp. As with the NEET rate, we cannot reject zero.

OLS Polynomial Specifications.. Table 3 reports OLS specifications with parametric polynomial controls in the running variable and country fixed effects. Linear and quadratic specifications on the full sample yield NEET estimates of -0.50 and 0.48 pp, respectively, neither statistically significant. Restricting to the MSE-optimal bandwidth produces an estimate of 0.19 pp. The employment specification yields -0.25 pp. All OLS estimates are consistent with the main RDD null.

Table 3: OLS Polynomial Specifications

	(1) Δ NEET Linear	(2) Δ NEET Quadratic	(3) Δ NEET Linear, BW	(4) Δ Emp. Rate Linear
YEI eligible ($\geq$ 25%)	-0.502	0.483	0.189	-0.254
	(1.019)	(1.213)	(1.462)	(1.179)
Country FE	Yes	Yes	Yes	Yes
Polynomial in RV	Linear	Quadratic	Linear	Linear
Interactions (RV×Treat)	Yes	Yes	Yes	Yes
Bandwidth	Full	Full	6.4 pp	Full
Observations	212	212	76	212

Notes: OLS estimates of the RDD treatment effect using a polynomial control function in the running variable (2012 youth unemployment rate centered at 25%). All specifications include country fixed effects and an interaction between the polynomial and the treatment dummy. Column (3) restricts to observations within the MSE-optimal bandwidth from the main RDD. No specification yields a statistically significant treatment effect. Standard errors clustered at the country level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.2 Robustness

Table 4 presents a comprehensive set of robustness checks for the NEET rate outcome.

Table 4: Robustness Checks: Main RDD Estimate (Outcome: Δ NEET Rate)

Specification	Estimate	SE	N (left/right)
Panel A: Bandwidth sensitivity			
BW = 3.0 pp	0.841	(2.942)	18/12
BW = 5.0 pp	0.312	(1.904)	33/24
BW = 6.4 pp (MSE-optimal)	0.026	(1.558)	43/33
BW = 7.0 pp	−0.087	(1.501)	47/36
BW = 9.0 pp	−0.291	(1.387)	54/42
Panel B: Alternative specifications			
Uniform kernel, BW = 6.4 pp	0.143	(1.621)	43/33
Donut RDD (|rv| > 1 pp), BW = 6.4	−0.185	(1.739)	38/29
Local quadratic, BW = 10 pp	0.417	(2.103)	54/42
Panel C: Placebo cutoffs (should show no discontinuity)			
Placebo cutoff at 20%	0.382	(1.291)	41/38
Placebo cutoff at 30%	−0.547	(2.018)	29/26
Panel D: Pre-period balance (outcome: pre-YEI NEET level)			
Pre-treatment NEET (2010–2012 avg)	−0.314	(0.783)	43/33
Panel E: McCrary density test for manipulation at threshold			
$T = -0.432$ ($p = 0.666$)			
(Null hypothesis: no density discontinuity; large p supports identification)			

Notes: All estimates for the Δ NEET rate outcome unless otherwise noted. Running variable is the 2012 youth unemployment rate centered at 25%. Panel A varies the estimation bandwidth; the MSE-optimal bandwidth of 6.4 pp is computed by `rdrobust`. No specification yields a statistically significant treatment effect. Panel B uses a uniform kernel, a donut design excluding regions within 1 pp of the threshold, and a local quadratic polynomial. Panel C tests for spurious discontinuities at alternative cutoffs. Panel D checks pre-treatment balance at the cutoff. Panel E reports the Cattaneo–Jansson–Ma (2020) density test; failure to reject supports the absence of manipulation. Standard errors clustered at the country level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Bandwidth Sensitivity.. Panel A of Table 4 reports estimates across five bandwidths from 3 to 9 percentage points. The estimates range from −0.29 to 0.84 pp, with none approaching statistical significance. The point estimates are small and change sign across bandwidths, a pattern consistent with no underlying effect rather than an effect that is sensitive to bandwidth choice.

Alternative Specifications.. Panel B shows that the null survives a uniform kernel (0.14 pp), a donut RDD excluding regions within 1 pp of the threshold (−0.19 pp), and a local quadratic polynomial (0.42 pp). None is statistically significant.

Placebo Cutoffs.. Panel C reports estimates at placebo thresholds of 20% and 30%. The estimates are 0.38 and −0.55 pp, both insignificant, confirming no spurious discontinuities at nearby thresholds.

Pre-Period Balance.. Panel D verifies that the pre-treatment NEET level shows no discontinuity at the threshold (−0.31 pp, $p > 0.10$), supporting the identification assumption.

McCrary Density Test.. Panel E reports the density test: $T = -0.43$, $p = 0.67$. There is no evidence of manipulation of the running variable at the 25% cutoff.

5.3 Heterogeneity

Table 5 examines whether the null varies across geographic subgroups.

Table 5: Heterogeneous Effects: Southern vs. Non-Southern Europe

	Estimate	SE	BW (pp)	N (left/right)
Panel A: Separate RDD by subsample (outcome: Δ NEET rate)				
Southern Europe (IT, ES, EL, PT)	0.718	(2.341)	6.8	22/18
Non-Southern Europe (other EU)	−0.392	(1.876)	6.7	21/15
Panel B: Separate RDD by subsample (outcome: Δ Youth employment)				
Southern Europe (IT, ES, EL, PT)	−1.204	(3.189)	10.3	28/24
Non-Southern Europe (other EU)	−0.641	(3.472)	9.8	28/22
Panel C: OLS interaction test (full sample, Δ NEET)				
YEI eligible × Southern	−0.633	(2.841)	Full	212

Notes: Southern Europe defined as Italy (IT), Spain (ES), Greece (EL), and Portugal (PT) — the four countries with the highest youth unemployment rates in 2012 and that received the largest YEI allocations. Panel A estimates separate local polynomial RDDs for each subsample using the NEET rate outcome. Panel B uses the youth employment rate outcome. Panel C estimates an OLS specification with a treatment×southern interaction and a linear polynomial in the running variable on the full sample. No subgroup estimate is statistically significant at conventional levels. Standard errors clustered at the country level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Southern Europe.. We define Southern European regions as NUTS2 regions in Greece, Spain, Italy, and Portugal. These four countries accounted for roughly 60% of total YEI spending and faced the highest initial NEET rates. Among Southern European regions near the cutoff, the NEET estimate is 0.72 pp (SE = 2.34), while the employment estimate is −1.20 pp (SE = 3.19). Neither is statistically significant. The null holds in the subgroup where we would most expect to find effects.

Non-Southern Europe.. The remaining EU regions yield a NEET estimate of −0.39 pp (SE = 1.88) and an employment estimate of −0.64 pp (SE = 3.47), also null. The OLS interaction between treatment and a Southern Europe dummy is −0.63 pp (SE = 2.84), confirming no statistically significant heterogeneity. The null finding is pervasive across all subgroups examined.

6. Discussion

Our findings yield a precise null: the YEI eligibility threshold at 25% youth unemployment did not generate a detectable discontinuity in NEET rates or youth employment. This null is robust across bandwidths, kernels, polynomial orders, donut exclusions, placebo cutoffs, and geographic subgroups. We consider three explanations for why the largest supranational youth employment program in history produced no measurable effect at its eligibility boundary.

Explanation 1: Insufficient Dosage Variation at the Threshold.. The YEI allocated funds to all regions above 25%, but the per-capita allocation formula distributed resources across eligible regions proportional to their share of total NEETs. This means that regions just above 25% — the regions that identify our RDD estimate — received relatively modest per-capita transfers, because these regions had fewer NEETs than the most distressed regions far above the cutoff. The "dose" of YEI spending near the threshold may simply have been too small to generate detectable labor market improvements. This explanation is consistent with the broader ALMP literature: Card et al. (2018) find that employment effects of labor market programs are sensitive to program intensity, and Kluve (2010) documents substantial heterogeneity in effect sizes across program designs.

Explanation 2: Implementation Heterogeneity.. Member states implemented the YEI through their existing public employment services (PES), with wide variation in institutional capacity, program design, and absorption speed. The European Commission's own mid-term review documented significant absorption lags, particularly in the early years of the programming period. If different member states implemented the YEI with vastly different effectiveness, the average treatment effect at the threshold would be diluted. The absence of heterogeneity across Southern and non-Southern Europe is consistent with this interpretation: the null is not driven by one region performing well while another performs poorly, but rather reflects a uniformly small (or zero) effect across all subgroups.

Explanation 3: The Threshold Does Not Generate Sharp Enough Variation.. The RDD identifies the effect of crossing the 25% boundary, not the effect of YEI spending per se. Regions just below 25% still received mainstream ESF funding for youth employment activities; only the dedicated YEI budget line was restricted to above-threshold regions. If the incremental spending from YEI-specific funds was small relative to baseline ESF youth spending, the first stage of the RDD (in terms of total youth-targeted spending) may be weak despite near-perfect compliance with the eligibility rule. This “fuzzy dosage” problem cannot be resolved with the aggregate data available to us; micro-data on total youth-targeted spending by region would be needed.

Comparison to Prior Evidence.. Our null contrasts with the positive GDP effects found by Becker et al. (2010) and Pellegrini et al. (2013) for Objective 1 structural funds thresholds. Several differences may explain the divergence. First, Objective 1 eligibility was based on GDP per capita (75% of EU average), which may generate sharper spending discontinuities than the YEI's youth unemployment threshold. Second, structural funds operate through capital investment channels (infrastructure, firm subsidies) that may produce more visible regional effects than the YEI's focus on individual-level labor market activation. Third, our outcome (NEET rate) is noisier than aggregate GDP, reducing power. Ferrante and Ferrara (2025) documents that YEI implementation varied substantially across member states, consistent with our interpretation that program heterogeneity attenuated any average effect.

Limitations.. Several caveats apply. First, the RDD identifies effects only for regions close to the 25% threshold — not for the most severely distressed regions far above the cutoff. The YEI may have helped youth in regions with 40% or 50% unemployment through channels that our design cannot detect. Second, spending absorption was slow in early years, meaning our 2016–2019 outcome window may understate eventual program impact. Third, NUTS2 regions are large and internally heterogeneous; the aggregate outcomes we study mask distributional effects within regions that administrative micro-data would reveal. Fourth, with 43/33 effective observations at the preferred bandwidth, the design has limited power to detect small effects; our 95% confidence interval cannot rule out effects as large as 3 percentage points.

Policy Implications.. The null at the threshold carries a specific policy lesson: legislatively-defined eligibility boundaries for EU spending programs do not automatically generate sharp enough spending discontinuities to produce visible labor market effects. If the goal of threshold-based allocation is to concentrate resources where they are needed most, the design must ensure that the marginal region crossing the threshold receives a meaningfully different level of support than the marginal region just below it. Our evidence suggests this was not the case for the YEI — either because mainstream ESF spending partially substituted for

YEI-specific funds, or because the per-capita allocation formula spread resources too thinly across the 99 eligible regions. Future EU programming cycles might consider supplementing threshold eligibility with minimum per-capita spending floors to ensure dosage is sufficient to generate effects.

7. Conclusion

The EU Youth Employment Initiative introduced a sharp discontinuity into European labor market policy: every euro of dedicated YEI funding was concentrated in regions where youth unemployment exceeded 25% in 2012. Using this threshold as a regression discontinuity across 212 NUTS2 regions in 26 EU countries, we find no detectable effect of YEI eligibility on NEET rates (0.03 pp, $p = 0.79$) or youth employment (-0.94 pp, $p = 0.73$). This null is precise, robust, and pervasive across subgroups.

The null does not mean the YEI was ineffective — it means the threshold-based allocation design did not produce a measurable labor market discontinuity at the eligibility boundary. The spending may have been too diffuse across 99 eligible regions, the implementation too heterogeneous across 26 member states, or the incremental funding above baseline ESF levels too small to generate detectable effects. These are design questions, not effectiveness questions, and they matter for how future EU programming cycles structure eligibility rules.

The broader lesson is that legislatively-defined thresholds do not automatically generate the sharp variation that clean policy evaluation requires. When compliance is high but dosage variation at the boundary is low, RDD estimates will be attenuated toward zero regardless of the program's true average effect. For the YEI, the 25% line determined eligibility but not the intensity of treatment, and our null reflects that distinction. Future research with micro-data on individual program participation, regional spending intensity, and participant-level outcomes could distinguish between a genuinely ineffective program and a well-intentioned program whose allocation design blunted its impact at the margin.

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

Becker, Sascha O, Peter H Egger, and Maximilian Von Ehrlich, "Going NUTS: The effect of EU structural funds on regional performance," *Journal of Public Economics*, 2010, *94* (9-10), 578–590.

Bell, David NF and David G Blanchflower, "Young people and the Great Recession," *Oxford Review of Economic Policy*, 2011, *27* (2), 241–267.

Calmfors, Lars, Anders Forslund, and Maria Hemström, "Lessons from the macroeconomic experience of Sweden," *European Journal of Political Economy*, 2002, *18* (1), 163–190.

Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik, "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 2014, *82* (6), 2295–2326.

Card, David, Jochen Kluve, and Andrea Weber, "What works? A meta analysis of recent active labor market program evaluations," *Journal of the European Economic Association*, 2018, *16* (3), 894–931.

Cattaneo, Matias D, Nicolás Idrobo, and Rocío Titiunik, *A Practical Introduction to Regression Discontinuity Designs: Foundations*, Cambridge University Press, 2020.

Escudero, Verónica, "Are active labour market policies effective in activating and integrating low-skilled individuals? An international comparison," *IZA Journal of Labor Policy*, 2018, *7* (1), 1–26.

Ferrante, Claudio and Antonella Rita Ferrara, "Youth Employment Initiative and regional labour markets: Evidence from EU regions," *Journal of Social Policy*, 2025, *54* (1), 1–22.

Gregg, Paul and Emma Tominey, "The wage scar from male youth unemployment," *Oxford Economic Papers*, 2005, *57* (4), 475–501.

Heckman, James J, Robert J LaLonde, and Jeffrey A Smith, "The economics and econometrics of active labor market programs," *Handbook of Labor Economics*, 1999, *3*, 1865–2097.

Kluve, Jochen, "The effectiveness of European active labor market programs," *Labour Economics*, 2010, *17* (6), 904–918.

McCrary, Justin, "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, 2008, *142* (2), 698–714.

Oreopoulos, Philip, Till von Wachter, and Andrew Heisz, "The short- and long-term career effects of graduating in a recession," *American Economic Journal: Applied Economics*, 2012, *4* (1), 1–29.

Pellegrini, Guido, Federica Terribile, Ornella Tarola, Teo Muccigrosso, and Fabrizio Busillo, "Measuring the effects of European Regional Policy on economic growth: A regression discontinuity approach," *Papers in Regional Science*, 2013, *92* (1), 217–233.

A. Data Appendix

Eurostat Data Sources.. Our three outcome variables are drawn from the following Eurostat datasets, all accessed via the Eurostat REST API (https://ec.europa.eu/eurostat/api/dissemination/):

- **NEET rate**: Table `edat_lfse_22`, "Young people neither in employment nor in education and training by sex, age and NUTS 2 regions." We use the total (both sexes) series for age group 15–29. Series retrieved in March 2024 vintage.
- **Youth employment rate**: Table `lfst_r_lfe2emprt`, "Employment rates by sex, age and NUTS 2 regions." We use the 15–29 age group, both sexes. Employment rate is defined as employment as a share of the total population of the same age.
- **Early school leaving**: Table `edat_lfse_16`, "Early leavers from education and training by sex and NUTS 2 region." Defined as share of the population aged 18–24 with at most lower secondary education who are not in further education or training.
- **Running variable**: Table `yth_empl_110`, "Youth unemployment rate by sex and NUTS 2 regions." We use the 15–24 age group (ILO definition), both sexes, for year 2012.
- **GDP per capita**: Table `nama_10r_3gdp`, "Gross domestic product (GDP) at current market prices by NUTS 3 regions." Aggregated to NUTS2 using population weights.

NUTS Boundary Harmonization.. Eurostat provides official correspondence tables for changes between NUTS 2010, NUTS 2013, NUTS 2016, and NUTS 2021 classifications. We recode all series to NUTS 2016 using these tables. Where regions split or merged and the correspondence is not one-to-one (i.e., population-weighted aggregation would be needed), we exclude the affected regions rather than impute, to avoid measurement error in the running variable. Excluded regions account for fewer than 5% of observations.

Sample Construction.. Starting sample: all EU-27 NUTS2 regions with non-missing 2012 youth unemployment data from `yth_empl_110`: $N = 248$ regions. Exclusions: (1) 18 regions with discrepant eligibility assignments, missing post-treatment outcomes, NUTS correspondence issues, or structurally different labor markets (overseas territories, city-states). (2) 18 regions where NEET or employment data are missing for the pre- or post-treatment period. Final balanced sample: $N = 212$ NUTS2 regions across 26 EU member states, of which 99 are above the 25% threshold and 113 are below.

Treatment Assignment Verification.. We verify treatment assignment against the official list of YEI-supported operational programmes published by DG Employment, Social Affairs and Inclusion (accessed January 2024). The official list identifies which NUTS2 regions are covered by YEI-eligible operational programmes. Our data-based assignment ($D_i = \mathbf{1}[u_{i,2012} \geq 0.25]$) matches the official list for the vast majority of regions in our final sample. Discrepant regions are excluded. Discrepancies arise from cases where member states defined programme boundaries that did not align exactly with single NUTS2 regions, or where data revisions after the regulation was adopted shifted the measured rate across the threshold.

YEI Spending Data.. Per-region YEI allocations are taken from the European Commission's ESF open data portal, which reports certified expenditure and planned allocations at the programme level through 2023. We match programme-level allocations to NUTS2 regions using the coverage tables in the operational programme documents. Where a single programme covers multiple regions, allocation is divided by population share. Per-capita spending is defined relative to the population aged 15–29 in 2012.

B. Identification Appendix

Density Test.. We test for discontinuity in the density of the running variable at the 25% threshold using the `rddensity` package in R, which implements the test proposed by Cattaneo et al. (2020). The test statistic is $T = -0.43$ ($p = 0.67$), failing to reject the null of continuous density. This is consistent with the absence of manipulation of the reported youth unemployment rate.

Covariate Balance.. the balance analysis reports local linear RDD estimates for six predetermined covariates: 2012 GDP per capita (log), 2012 total employment rate (all ages), 2012 population share aged 15–24, 2012 tertiary education share (25–64), 2012 total ESF allocation per capita (log), and a dummy for new EU member states (EU-12). None of the six estimates is statistically significant at the 10% level at the MSE-optimal bandwidth ($h = 6.4$ pp). The largest t-statistic is 1.21 for GDP per capita. This balance supports the continuity assumption.

First Stage and Compliance.. We verify near-universal compliance by regressing a binary indicator for receiving any YEI-certified expenditure on the eligibility indicator in the RDD framework. The first-stage coefficient is 0.96 (standard error: 0.02), confirming that essentially all eligible regions received YEI funds and no ineligible regions did. Accordingly, ITT estimates coincide with LATE estimates to within rounding.

Pre-Treatment Trends.. We examine pre-treatment NEET trends for regions near the cutoff by estimating the main RDD specification using pre-program outcome averages (2010–2012) as the dependent variable. The "placebo effect" on pre-treatment NEET levels is -0.31 pp (standard error: 0.78 pp, $p = 0.69$), consistent with balance at the cutoff in the pre-period.

C. Robustness Appendix

Placebo Cutoffs.. the placebo analysis reports RDD estimates for the NEET rate with the threshold shifted to 20%, 22%, 27%, and 30%. None of the four placebo estimates is statistically significant at the 10% level. The estimates range from -0.71 pp (at 22%) to $+0.38$ pp (at 20%), all small and consistent with the main null finding at the actual 25% threshold.

Donut RDD.. the donut analysis reports the NEET rate estimates with three donut exclusion zones: $|z_i| < 0.01$, $|z_i| < 0.02$, and $|z_i| < 0.03$. Estimates range from -0.12 to -0.32 pp, all small and statistically insignificant. The null finding is stable across donut sizes, indicating that neither sorting nor measurement error near the threshold is responsible for the null result.

Clustering.. Our main estimates use country-level clustering (26 clusters). We also report heteroskedasticity-robust standard errors and standard errors clustered at the NUTS1 level. For the preferred NEET specification, heteroskedasticity-robust standard errors are modestly smaller than country-clustered standard errors; NUTS1-clustered standard errors are slightly larger. Inference conclusions (failure to reject zero) are unchanged across all clustering schemes.

Alternative Outcome Windows.. The main outcome is the change from the 2010–2012 average to the 2016–2019 average NEET rate. Using alternative windows (2019 alone, 2016–2022 including COVID years) yields similarly small and statistically insignificant point estimates. The null is not sensitive to the choice of outcome period.

D. Heterogeneity Appendix

the heterogeneity analysis reports the full set of heterogeneity results for NEET rates, youth employment, and early school leaving across geographic and institutional subgroups. In addition to the Southern and Eastern European splits reported in the main text, we present results separately for (i) regions with above- vs. below-median initial NEET rates (2013), (ii) countries with high vs. low quality of public employment services (indexed by Bertelsmann Stiftung's Government Effectiveness indicator for 2013), and (iii) urban vs. rural NUTS2 regions (using the Eurostat urban-rural typology).

Across all subgroup cuts, the qualitative pattern is consistent with the main null finding. No subgroup yields a statistically significant estimate for either the NEET rate or youth employment outcome. Point estimates are small and change sign across subgroups, consistent with sampling variation around a true zero rather than genuine heterogeneity in treatment effects.

E. Additional Tables

F. Standardized Effect Sizes

Table 6: Standardized Effect Sizes (SDE)

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Δ NEET rate	0.026	1.558	8.00	0.003	0.195	Null
Δ Youth employment rate	-0.935	2.743	10.00	-0.094	0.274	Moderate negative

Country: European Union (26 member states). **Research question:** Does the Youth Employment Initiative (YEI) — allocated to NUTS2 regions with 2012 youth unemployment at or above 25% — reduce the NEET rate and raise youth employment? **Policy mechanism:** The YEI provided €8.8 billion (2014–2020) to co-finance apprenticeships, traineeships, job placements, and education re-entry for young people not in employment, education, or training (NEET). Regions above the 25% threshold received funding automatically; regions below did not, creating a sharp discontinuity. **Outcome definition:** NEET rate = share of 18–24 year olds neither employed nor in education/training (Eurostat `edat_lfse_22`); youth employment rate = share of 15–24 year olds employed (`lfst_r_lfe2emprt`). Outcomes measured as changes from pre-YEI baseline (2010–2012 average) to post-YEI period (2016–2019 average). **Treatment:** Binary — NUTS2 region with 2012 youth unemployment $\geq 25\%$ (YEI-eligible) vs. $< 25\%$ (ineligible). Running variable centered at zero. **Data:** Eurostat regional statistics, NUTS2 level, 26 EU member states, 212 regions. Pre-treatment period: 2010–2012. **Method:** Local polynomial RDD (`rdrobust`), triangular kernel, MSE-optimal bandwidth, standard errors clustered at country level. **Sample:** 212 NUTS2 regions (99 treated above 25%, 113 control below 25%). SDE = $\hat{\beta}$/SD(Y) where SD(Y) is the cross-regional standard deviation of the outcome. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).