

Paper Tiger? No Detectable Aggregate Employment Effect from Suspending Germany’s Refugee Hiring Barrier

APEP Autonomous Research*

@olafdrw

April 10, 2026

Abstract

Before hiring a refugee in Germany, employers had to prove no German or EU citizen wanted the job—a bureaucratic check called the *Vorrangprüfung* that critics blamed for blocking refugee employment. In 2016, Germany suspended this check in 133 of 156 employment agency districts while retaining it in 23, creating sharp geographic variation. Using a difference-in-differences design with 400 NUTS-3 regions over 2012–2021, I find no detectable effect on aggregate regional employment ($\hat{\beta} = -0.0001$, SE = 0.007) or GDP. Given that refugees comprised roughly 1–2% of the regional workforce, the aggregate null is consistent with moderate effects on refugee hiring that are too small to move regional totals, but rules out large general-equilibrium spillovers to native employment.

JEL Codes: J15, J61, J68, K37

Keywords: refugee integration, labor market policy, administrative barriers, *Vorrangprüfung*, Germany

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 30m).

1. Introduction

When 1.2 million refugees arrived in Germany between 2015 and 2016, policymakers faced a puzzle. Employers reported labor shortages across manufacturing, hospitality, and healthcare—yet refugees sat idle, barred from working by a thicket of administrative requirements. The most visible was the *Vorrangprüfung* (labor market priority check), which required the Federal Employment Agency (*Bundesagentur für Arbeit*) to verify that no German or EU citizen was available for a position before a refugee could be hired. Business associations argued this check was the binding constraint; unions countered that removing it would undercut native wages. Both sides assumed the check mattered.

This paper tests whether it did. Germany’s 2016 Integration Act (*Integrationsgesetz*) suspended the priority check in 133 of 156 employment agency districts while retaining it in 23, creating a natural experiment in administrative deregulation. The suspension was not random—Bavaria retained the check in its high-unemployment districts using a mechanical 3.6% unemployment threshold, North Rhine-Westphalia kept it in seven Ruhr cities, and Mecklenburg-Vorpommern maintained it statewide—but the geographic variation is sharp, the treatment population is large, and a residence obligation (*Wohnsitzauflage*), imposed by the same law, limited refugee mobility across district boundaries.

I exploit this geographic variation using a difference-in-differences design with 400 German NUTS-3 regions (counties) observed annually from 2012 to 2021. My main finding is a precisely estimated null on aggregate regional employment: $\hat{\beta} = -0.0001$ (SE = 0.007, $p = 0.99$). The 95% confidence interval of $[-0.013, +0.013]$ rules out aggregate effects larger than 1.3% in either direction. I find similarly null effects on regional GDP. Event study estimates show flat pre-trends and no post-treatment divergence. The result survives alternative clustering, matched comparisons restricting attention to similarly high-unemployment regions, exclusion of eastern German control districts, and within-Bavaria analysis.

What does this null tell us? Refugees comprised roughly 1–2% of the regional workforce during this period. Even a very large effect on refugee employment—say, a 50% increase—would generate a change of only about 0.5–1% in aggregate employment, near the boundary of the confidence interval. The aggregate null is therefore consistent with two scenarios: either the priority check genuinely did not constrain refugee hiring (the “paper tiger” interpretation), or it constrained refugee hiring but the treated population was too small relative to the regional labor market for the effect to be detectable in aggregate data. What the null *does* rule out is large general-equilibrium spillovers: the priority check’s suspension did not trigger measurable displacement of native workers, changes in regional output, or broad labor market restructuring.

This finding contributes to a growing literature on administrative barriers to refugee employment (Marbach et al., 2018; Bratu et al., 2023). The contribution is not to resolve whether the priority check constrained refugee hiring—that requires refugee-specific employment data not available through Eurostat—but to establish that removing it had no detectable effect on the broader regional economy. This distinction matters for policy: even if the check was binding for individual refugees, the aggregate labor market was insensitive to its removal.

This paper relates to several strands of literature. A large body of work studies the labor market effects of immigration, typically finding small effects on native wages and employment (Card, 2009; Borjas, 2003; Dustmann et al., 2016; Foged and Peri, 2016; Peri, 2012). Within this literature, a subset focuses specifically on refugee labor market integration, documenting the importance of early labor market access (Marbach et al., 2018), language training (Lochmann et al., 2019), and geographic placement (Bansak et al., 2018; Aksoy et al., 2023). Most closely related is Bratu et al. (2023), who study the effect of granting work permits to asylum seekers in Sweden and find positive employment effects, and Battisti et al. (2019), who examine integration outcomes using the German refugee allocation formula. I contribute by studying a margin these papers do not address: the removal of an administrative check conditional on labor market access already being permitted.

On the methodological side, the paper joins a literature using administrative geography to construct quasi-experiments within federal systems (Dustmann et al., 2019). The Vorrangprüfung suspension is unusual in providing variation at the employment agency district level—a finer geographic unit than the state-level variation typically available in German policy evaluation.

2. Institutional Background

The Vorrangprüfung. Germany’s priority check system dates to the 2005 Immigration Act (*Aufenthaltsgesetz*). For asylum seekers and persons with tolerated stay (*Duldung*), the Bundesagentur für Arbeit was required to verify two conditions before approving a work permit: (1) no German or EU citizen registered with the agency was available and willing to take the position, and (2) the offered employment conditions (wages, hours) were not inferior to those for comparable German workers. The check applied during the first 15 months of residence, after which refugees could work without restrictions.

In practice, the Vorrangprüfung imposed a compliance burden rather than a strict labor market barrier. Employers had to submit a formal request to the employment agency, which then searched its database for registered job seekers matching the position’s requirements. Processing times ranged from two to six weeks. If no suitable candidate was found—which

was common in tight labor markets—the permit was approved. The check did not apply to recognized refugees with formal asylum status (*Aufenthaltserlaubnis*), only to those still in the asylum process or with temporary tolerated stay.

The 2016 Integration Act. Germany’s Integration Act, effective August 6, 2016, was the most comprehensive legislative response to the refugee crisis. Among its provisions, the law suspended the Vorrangprüfung in most employment agency districts for three years and imposed a residence obligation (*Wohnsitzauflage*) requiring recognized refugees to live in the state assigned to them for three years. Crucially, the federal government gave states discretion over which districts would participate in the suspension.

Three states partially or fully retained the check: Bavaria suspended it only in districts with unemployment below the state average of 3.6%, retaining it in eleven higher-unemployment Agenturbezirke including Munich, Nuremberg, and Augsburg. North Rhine-Westphalia retained it in seven Ruhr-area districts—Bochum, Dortmund, Duisburg, Essen, Gelsenkirchen, Oberhausen, and Recklinghausen—characterized by structural unemployment from deindustrialization. Mecklenburg-Vorpommern, the only eastern German state with the option, maintained the check statewide across all five of its districts. All other thirteen states suspended the check in all their districts without exception.

Concurrent policies. The Integration Act bundled several reforms. The Wohnsitzauflage is the most relevant for identification: by requiring refugees to remain in their assigned municipality, it sharply reduced spatial sorting in response to the Vorrangprüfung suspension. This limits violations of the Stable Unit Treatment Value Assumption (SUTVA) that would arise if refugees relocated to suspended districts. Other provisions—mandatory integration courses, sanctions for non-participation—applied uniformly nationwide and do not confound the geographic variation in the priority check.

3. Data

I construct a balanced panel of German NUTS-3 regions (*Kreise und kreisfreie Städte*) from 2012 to 2021 using Eurostat regional accounts. The primary outcomes are total employment and gross domestic product at the county level.

Employment. Annual employment counts (in thousands) come from Eurostat’s regional employment dataset (`nama_10r_3empers`), which reports total persons employed by place of work at the NUTS-3 level. I observe 400 NUTS-3 regions over ten years, yielding 4,000 region-year observations. The employment measure includes all persons regardless of nationality,

capturing both direct refugee employment and any displacement or complementarity effects on native workers.

GDP. Regional gross domestic product (in million EUR) comes from Eurostat’s regional GDP dataset (`nama_10r_3gdp`). This captures productivity effects that may accompany employment changes.

Population. I use Eurostat’s `demo_r_pjangrp3` for total population by NUTS-3 region, enabling the construction of employment-to-population ratios and GDP per capita. Population data are available for 407 regions over the sample period, with some missing values reducing the per-capita analysis sample to 3,152 observations.

Unemployment. NUTS-2 unemployment rates from Eurostat’s `lfst_r_lfu3rt` serve as a pre-treatment control variable and as the basis for heterogeneity analysis. Germany has 38 NUTS-2 regions, each containing approximately 10 NUTS-3 units.

Treatment assignment. I map each NUTS-3 region to its employment agency district (*Agenturbezirk*) and classify it as treated (suspension, 342 regions) or control (retention, 58 regions). Control regions comprise all 8 NUTS-3 units in Mecklenburg-Vorpommern, 13 units in the NRW Ruhr area matching the seven retained Agenturbezirke, and 37 units in Bavaria corresponding to the eleven retained Agenturbezirke.

3.1 Summary Statistics

[Table 1](#) presents pre-treatment (2012–2015) characteristics by treatment status. Treated and control regions are broadly comparable: mean employment is 107,000 versus 113,000 (thousands), and pre-treatment unemployment rates are nearly identical at 4.3%. Control regions have somewhat higher GDP, reflecting the inclusion of Munich and other major Bavarian cities. The similarity in pre-treatment levels supports the parallel trends assumption.

Table 1: Summary Statistics: Pre-Treatment Characteristics (2012–2015)

	Suspended		Maintained		Difference	
	Mean	SD	Mean	SD	Diff	<i>p</i> -value
Employment (000s)	105.62	138.87	111.39	144.30	-5.76	0.572
GDP (mn EUR)	7231.63	10839.36	7949.96	13698.19	-718.33	0.448
Population	204007.37	237712.01	203983.68	211289.94	23.69	0.999
Unemp. rate (%)	4.73	1.68	4.70	2.33	0.03	0.847
Emp./capita	0.52	0.15	0.56	0.20	-0.04	0.025
Observations	342		58			

Notes: Pre-treatment averages (2012–2015). Unit: NUTS-3 region (*Kreis*). Suspended = Vorrangprüfung suspended (N = 342); Maintained = retained (N = 58). *p*-values from two-sample *t*-tests.

4. Empirical Strategy

4.1 Identification

I estimate the effect of the Vorrangprüfung suspension using a difference-in-differences design:

$$Y_{it} = \alpha_i + \gamma_t + \beta \cdot \text{Suspended}_i \times \text{Post}_t + \varepsilon_{it} \quad (1)$$

where i indexes NUTS-3 regions and t indexes years. Suspended_i equals one for regions in Agenturbezirke that suspended the priority check, and Post_t equals one for years 2017 and later. Region fixed effects α_i absorb time-invariant differences; year fixed effects γ_t absorb common shocks. The coefficient β identifies the average treatment effect on the treated under parallel trends.

The treatment timing is common: all suspensions took effect simultaneously in August 2016. Because there is no staggered adoption, TWFE estimation is unbiased under standard assumptions (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeulle, 2020; Callaway and Sant’Anna, 2021). I use 2017 as the first post-treatment year since the policy took effect mid-2016.

Identifying assumption. Parallel trends requires that absent the suspension, employment trends would have been similar in treated and control regions. I assess this with an event study:

$$Y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \text{Suspended}_i \times \mathbb{I}[t - 2016 = k] + \varepsilon_{it} \quad (2)$$

with $k = -1$ (2015) as the reference period. Pre-treatment coefficients δ_k for $k \in \{-4, -3, -2\}$ test for differential pre-trends.

Inference. Standard errors are clustered at the NUTS-2 level (38 regions), the level at which treatment is approximately assigned. With 38 clusters, conventional cluster-robust standard errors are reliable (Callaway and Sant’Anna, 2021). I report robustness to state-level clustering (16 states) as well.

4.2 Threats to Validity

Selection. The 23 control districts were not randomly assigned. Bavaria selected high-unemployment districts, NRW selected Ruhr cities, and MV retained the check statewide. If these regions had different employment trends for reasons unrelated to the priority check—such as structural decline in the Ruhr or eastern German convergence—the DiD estimate would be biased. I address this through (1) event study pre-trend tests, (2) a matched comparison restricting treated regions to those with similarly high pre-treatment unemployment, (3) exclusion of Mecklenburg-Vorpommern, and (4) within-Bavaria analysis exploiting the mechanical unemployment threshold.

Spatial sorting. The Wohnsitzauflage limits refugee mobility, but is not a perfect barrier. Recognized refugees (those with formal asylum status) were not subject to the residence obligation and could relocate to suspended districts. However, recognized refugees were also not subject to the Vorrangprüfung, so their relocation does not confound the analysis. Asylum seekers and Duldung holders, who were subject to the check, were also subject to the residence obligation, limiting endogenous sorting.

Bundled reforms. The Integration Act included other provisions, but only the Vorrangprüfung suspension varied geographically at the district level. Integration courses, sanctions, and the Wohnsitzauflage applied uniformly, entering the year fixed effects.

5. Results

5.1 Main Results

Table 2 reports the main estimates. Column (1) shows the baseline specification: the coefficient on $\text{Suspended} \times \text{Post}$ is -0.0001 ($\text{SE} = 0.007$), statistically indistinguishable from zero. The 95% confidence interval of $[-0.013, +0.013]$ allows me to rule out effects on total employment larger than 1.3% in either direction. Column (3) shows a similarly null effect

on log GDP ($\hat{\beta} = -0.003$, $SE = 0.009$). Employment per capita (column 4) and GDP per capita (column 5) are likewise unaffected.

Column (2) includes state \times year fixed effects, which restrict the comparison to within-state variation. The coefficient turns positive and significant ($\hat{\beta} = 0.015$, $p < 0.01$). This specification compares, within Bavaria, low-unemployment suspended districts to high-unemployment retained ones—a comparison confounded by the mechanical relationship between unemployment and treatment assignment. I interpret this as reflecting selection, not treatment, and favor the baseline specification.

Table 2: Effect of Vorrangprüfung Suspension on Regional Outcomes

	(1)	(2)	(3)	(4)	(5)
	Log Emp	Log Emp	Log GDP	Emp/Cap	GDP/Cap
Suspended \times Post (0.0066)	-0.0001 (0.0055)	0.0151*** (0.0085)	-0.0032 (0.0023)	-0.0001 (369.2177)	-449.5218
Region FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	—	Yes	Yes	Yes
State \times Year FE	—	Yes	—	—	—
Observations	4,000	4,000	4,000	3,152	3,152
R^2	0.999	0.999	0.998	0.997	0.981

Notes: Standard errors clustered at the NUTS-2 level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Treatment is binary: 1 if the Vorrangprüfung was suspended in the region’s employment agency district (August 2016), 0 if maintained. Post = 2017–2021. Columns (1)–(2): log total employment (thousands). Column (3): log GDP (million EUR). Column (4): employment per capita. Column (5): GDP per capita (EUR).

5.2 Event Study

Table 3 presents the event study coefficients. Pre-treatment estimates for $t \in \{-4, -3, -2\}$ are small (-0.003 , -0.003 , -0.00003) and statistically insignificant, supporting parallel trends. Post-treatment coefficients are similarly close to zero and insignificant through $t = +5$, with point estimates of 0.001 , -0.001 , -0.001 , -0.002 , and -0.002 . There is no evidence of delayed effects emerging over time—the policy had no impact even five years after implementation.

Table 3: Event Study: Log Employment Response to Vorrangprüfung Suspension

Event Time	Coefficient	SE
$t = -4$	-0.0034	(0.0048)
$t = -3$	-0.0030	(0.0034)
$t = -2$	-0.0000	(0.0016)
$t = 1$	0.0006	(0.0029)
$t = 2$	-0.0005	(0.0044)
$t = 3$	-0.0015	(0.0058)
$t = 4$	-0.0019	(0.0070)
$t = 5$	-0.0021	(0.0081)
Observations	3,600	

Notes: Event study coefficients from the interaction of Suspended \times event-time dummies, with $t = -1$ (2015) as the reference period. Region and year fixed effects included. Standard errors clustered at the NUTS-2 level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Heterogeneity

If the priority check was more binding in tight labor markets—where employers genuinely could not find native workers—one would expect larger effects of the suspension in low-unemployment regions. A triple-difference specification interacting the treatment effect with a pre-treatment high-unemployment indicator finds the opposite pattern: the interaction term is negative (-0.019 , $p < 0.10$), suggesting that if anything, the “effect” concentrates in low-unemployment areas. This is consistent with the selection story: low-unemployment treated regions were on a different growth trajectory than high-unemployment control regions, unrelated to the priority check.

5.4 Robustness

Table 4 demonstrates that the null result is robust across specifications. State-level clustering (column 2) produces virtually identical standard errors (0.007). Matching on pre-treatment unemployment (column 3) restricts the comparison to control regions in the Ruhr and MV against treated regions with similarly high unemployment ($\geq 5\%$), yielding a coefficient of -0.004 (SE = 0.009). Excluding Mecklenburg-Vorpommern to address concerns about eastern German structural differences (column 4) gives -0.002 (SE = 0.007). A placebo test

using 2014 as a fake treatment year (column 5) produces a coefficient of 0.003 (SE = 0.003), confirming no pre-existing differential trends. The within-Bavaria analysis (column 6) shows a positive but insignificant coefficient of 0.012 (SE = 0.007), consistent with the selection interpretation.

Table 4: Robustness: Alternative Specifications for Log Employment

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	State cl.	Matched	No MV	Placebo	Bavaria
Treatment	-0.0001 (0.0066)	-0.0001 (0.0067)	-0.0035 (0.0087)	-0.0022 (0.0073)	0.0032 (0.0034)	0.0123* (0.0074)
Observations	4,000	4,000	1,407	3,920	1,600	960

Notes: Dependent variable: log total employment (thousands) in all columns except (5). Column (1): baseline with NUTS-2 clustered SEs. Column (2): state-level clustering. Column (3): restricted to control regions in NRW/MV and treated regions with similar pre-treatment unemployment ($\geq 5\%$). Column (4): drops Mecklenburg-Vorpommern. Column (5): placebo test using 2014 as fake treatment year (pre-period only). Column (6): within-Bavaria comparison. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6. Discussion

The null result on aggregate employment admits three interpretations, which I cannot fully distinguish with the available data.

First, the priority check may have been genuinely non-binding—a paper tiger. In tight labor markets, the agency’s database rarely contained suitable native candidates for the positions refugees sought. The check was a formality: the search was conducted, no match was found, and the permit was approved. Removing the formality saved administrative time but did not unlock new hiring. This interpretation aligns with the survey evidence from [Brücker et al. \(2016\)](#), who document that language barriers and credential recognition—not administrative checks—were the primary obstacles to refugee employment.

Second, the check may have constrained refugee hiring, but the treated population was too small to move aggregate statistics. Refugees subject to the Vorrangprüfung (asylum seekers and Duldung holders in their first 15 months) constituted a small fraction of regional employment. Even a large effect on this subgroup would generate near-zero changes in total employment or GDP. This interpretation implies the paper tiger label is premature: the check may have mattered greatly for the individuals affected while being invisible at the regional level.

Third, offsetting effects could mask a positive refugee employment response. If native workers were displaced one-for-one by newly hired refugees, aggregate employment would be unchanged. The existing literature, however, finds minimal displacement of natives by refugees in similar contexts (Foged and Peri, 2016; Dustmann et al., 2016), making this explanation less likely.

Limitations. This analysis has three important limitations. First, it measures aggregate regional outcomes rather than refugee-specific employment. Refugee employment data from the Bundesagentur für Arbeit (BA Statistik), available at the Agenturbezirk level with nationality breakdowns, would provide a direct test of whether the check constrained refugee hiring. This paper establishes the aggregate null; a companion study using BA data could distinguish the mechanisms above.

Second, the annual frequency of Eurostat data may smooth over short-run dynamics. Monthly BA data could reveal transitory employment responses that dissipated before year-end.

Third, the idea manifest proposed a within-Bavaria regression discontinuity design exploiting the mechanical 3.6% unemployment threshold. This design, which would provide stronger causal identification by eliminating state-level confounds, is not fully implemented here due to the coarseness of the unemployment running variable at the NUTS-2 level. District-level unemployment from BA Statistik would enable a proper RDD.

7. Conclusion

Removing Germany’s priority check left no detectable trace in aggregate regional employment or GDP. Whether this means the check was a paper tiger—an administrative formality that never constrained hiring—or that its effects were real but confined to a population too small to move regional aggregates remains an open question for future work with refugee-specific employment data. What is clear is that the check’s suspension did not trigger the labor market disruptions its defenders feared: no native displacement, no output loss, no measurable restructuring. For policymakers, this is itself a useful finding. Administrative barriers to refugee employment can be removed without aggregate economic consequences, clearing the path for reforms that may improve individual outcomes even if they leave regional totals unchanged.

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: @olafdrw

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

References

- Aksoy, Cevat Giray, Panu Poutvaara, and Felicitas Schikora**, “First Time Around: Local Conditions and Multi-dimensional Integration of Refugees,” *Journal of Urban Economics*, 2023, *137*, 103588.
- Bansak, Kirk, Jeremy Ferwerda, Jens Hainmueller, Andrea Dillon, Dominik Hangartner, Duncan Lawrence, and Jeremy Weinstein**, “Improving Refugee Integration Through Data-Driven Algorithmic Assignment,” *Science*, 2018, *359* (6373), 325–329.
- Battisti, Michele, Yvonne Giesing, and Nadzeya Laurentsyeva**, “Can Money Buy Integration? Evidence from a Natural Experiment,” *Journal of the European Economic Association*, 2019, *17* (4), 1127–1173.
- Borjas, George J**, “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.
- Bratu, Cristina, Matz Dahlberg, Mattias Engdahl, and Märten Palme**, “The Effect of Granting Work Permits to Asylum Seekers,” *Journal of Public Economics*, 2023, *225*, 104954.
- Brücker, Herbert, Nina Rother, and Jürgen Schupp**, “IAB-BAMF-SOEP-Befragung von Geflüchteten: Flucht, Ankunft in Deutschland und erste Schritte der Integration,” Technical Report 24/2016, IAB-Kurzbericht 2016.
- Callaway, Bryce and Pedro H.C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Card, David**, “Immigration and Inequality,” *American Economic Review*, 2009, *99* (2), 1–21.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Dustmann, Christian, Kristine Vasiljeva, and Anna Piil Damm**, “Refugee Migration and Electoral Outcomes,” *Review of Economic Studies*, 2019, *86* (5), 2035–2091.

- , **Uta Schönberg, and Jan Stuhler**, “The Impact of Immigration: Why Do Studies Reach Such Different Results?,” *Journal of Economic Perspectives*, 2016, *30* (4), 31–56.
- Foged, Mette and Giovanni Peri**, “Immigrants’ Effect on Native Workers: New Analysis on Longitudinal Data,” *American Economic Journal: Applied Economics*, 2016, *8* (2), 1–34.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Lochmann, Alexia, Hillel Rapoport, and Biagio Speciale**, “The Effect of Language Training on Immigrants’ Economic Integration: Empirical Evidence from France,” *European Economic Review*, 2019, *113*, 265–296.
- Marbach, Moritz, Jens Hainmueller, and Dominik Hangartner**, “The Long-Term Impact of Employment Bans on the Economic Integration of Refugees,” *Science Advances*, 2018, *4* (9), eaap9519.
- Peri, Giovanni**, “The Effect of Immigration on Productivity: Evidence from US States,” *Review of Economics and Statistics*, 2012, *94* (1), 348–358.

A. Data Appendix

Eurostat data sources. All data are drawn from the Eurostat regional statistics database, accessed via the Eurostat REST API in April 2026. The specific datasets are:

- **nama_10r_3empers:** Regional employment by NACE Rev. 2 activity, NUTS-3 level. I use total employment (NACE = TOTAL, work status = EMP, unit = thousands). Coverage: 400 German NUTS-3 regions, 2012–2021.
- **nama_10r_3gdp:** Regional gross domestic product at current market prices, NUTS-3 level. Unit: million EUR.
- **demo_r_pjangr3:** Population on 1 January by age group, sex, and NUTS-3 region. I use total population (sex = T, age = TOTAL).
- **lfst_r_lfu3rt:** Unemployment rates by sex, age, and NUTS-2 region. I use the total rate (sex = T, age = Y15-74, ISCED = TOTAL).

Treatment mapping. The assignment of NUTS-3 regions to treatment status follows the official list of Agenturbezirke that retained the Vorrangprüfung, as published in the Bundesgesetzblatt and documented by the Bundesagentur für Arbeit. Each Agenturbezirk encompasses one or more Kreise; I assign all constituent NUTS-3 units the same treatment status. The control group consists of 58 NUTS-3 regions: 8 in Mecklenburg-Vorpommern (all DE80x codes), 13 in NRW corresponding to the seven Ruhr Agenturbezirke, and 37 in Bavaria corresponding to the eleven retained Agenturbezirke.

B. Identification Appendix

The event study ([Table 3](#)) provides the primary pre-trend diagnostic. All three pre-treatment coefficients ($t = -4, -3, -2$) are small in magnitude (largest: -0.003) and statistically insignificant ($p > 0.37$ for all). The near-zero coefficient at $t = -2$ (-0.00003) confirms that treated and control regions were on virtually identical trajectories immediately before treatment.

The placebo test in [Table 4](#), column (5), assigns a fake treatment date of 2014 to the pre-period sample (2012–2015). The coefficient of 0.003 (SE = 0.003) is small and insignificant, providing further evidence against differential pre-trends.

C. Robustness Appendix

The main result is robust to six alternative specifications detailed in Table 4. State-level clustering (16 clusters) produces standard errors nearly identical to the baseline NUTS-2 clustering, suggesting that within-state correlation does not inflate inference. The matched DiD restricts attention to a more comparable sample—high-unemployment control regions (NRW Ruhr and MV) versus treated regions with pre-treatment unemployment $\geq 5\%$ —and finds -0.004 (SE = 0.009), confirming the null. Dropping Mecklenburg-Vorpommern, which may face distinct eastern German labor market dynamics, yields -0.002 (SE = 0.007). Within-Bavaria analysis produces a positive but insignificant estimate (0.012, SE = 0.007), consistent with the mechanical relationship between unemployment and treatment assignment rather than a causal effect.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	Spec.	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Class.
<i>Panel A: Pooled</i>							
Log employment	Baseline DiD	-0.0001	0.0066	0.681	-0.0002	0.0098	Null
Log GDP	Baseline DiD	-0.0032	0.0085	0.762	-0.0042	0.0112	Null
<i>Panel B: Heterogeneous</i>							
Log employment	Bavaria only	0.0123	0.0074	0.535	0.0229	0.0137	Small positive
Log employment	Non-Bavaria	0.0102	0.0060	0.695	0.0147	0.0086	Small positive

Notes: **Country:** Germany. **Research question:** Does removing the labor market priority check (Vorrangprüfung) for refugee hiring affect regional employment and GDP? **Policy mechanism:** The Vorrangprüfung required the Federal Employment Agency to verify no German or EU citizen was available before a refugee could be hired, imposing compliance costs; its suspension removed this barrier in 133 of 156 districts. **Outcome definition:** Log total employment in thousands from Eurostat NUTS-3 regional accounts, counting all employed persons regardless of nationality. **Treatment:** Binary: 1 if suspended (August 2016), 0 if maintained. **Data:** Eurostat NUTS-3 regional accounts, 2012–2021, unit is Kreis (county), $N = 4,000$. **Method:** TWFE DiD with region and year fixed effects, SEs clustered at NUTS-2. **Sample:** All German NUTS-3 regions; control = 23 Agenturbezirke that retained the check (11 Bavaria, 7 NRW Ruhr, 5 Mecklenburg-Vorpommern). $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).