

# The €6,000 Question: Do Apprenticeship Subsidies Create Jobs or Relabel Hiring? Evidence from France’s Post-Pandemic Training Boom\*

APEP Autonomous Research<sup>†</sup> @olafdrw

February 20, 2026

## Abstract

France’s 2020 apprenticeship hiring subsidy tripled new training contracts to 879,000 annually at a fiscal cost of €15 billion per year. I exploit the January 2023 subsidy reduction—from €8,000 to €6,000 per contract—using an exposure difference-in-differences design that interacts sector-level pre-reform apprenticeship intensity with the policy shock. High-exposure sectors like construction and hospitality saw *increases* in youth employment share after the reduction, suggesting the subsidy displaced rather than created entry-level positions. However, total employment also rose differentially in exposed sectors, complicating clean causal attribution. Suggestive cross-country evidence—including a synthetic control analysis—shows that French youth employment continued rising post-reduction. Indeed vacancy data shows no differential decline in job postings in apprenticeship-intensive sectors. Taken together, the evidence favors the interpretation that France’s apprenticeship subsidy primarily relabeled existing junior hiring as subsidized training contracts, though the design cannot definitively rule out alternative explanations.

**JEL Codes:** J23, J24, J68, H25

**Keywords:** apprenticeship subsidies, youth employment, hiring incentives, training externality, labor market policy, France

---

\*This is a revision of APEP Working Paper `apep_0427_v1`. See [https://github.com/SocialCatalystLab/ape-papers/tree/main/apep\\_0427/v1](https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0427/v1) for the parent version. This revision fixes the randomization inference procedure, adds wild cluster bootstrap and synthetic control analyses, and addresses all reviewer feedback.

<sup>†</sup>Autonomous Policy Evaluation Project. Correspondence: [scl@econ.uzh.ch](mailto:scl@econ.uzh.ch)

## 1. Introduction

Between 2019 and 2024, France experienced the largest apprenticeship expansion in modern European history. New apprenticeship contracts tripled from 306,000 to 879,000 per year, the number of active apprentices crossed one million, and public spending on the program reached €15 billion annually—roughly 0.5 percent of GDP. The engine behind this transformation was a generous hiring subsidy, the *aide exceptionnelle à l'embauche d'alternants*, which covered nearly the entire first-year wage cost of an apprentice. Policymakers hailed the program as a triumph of active labor market policy. But a basic question remains unanswered: did these subsidies create genuinely new entry-level positions, or did they simply relabel hiring that would have happened anyway?

This question matters beyond France. Across the OECD, governments spend tens of billions annually on training subsidies, apprenticeship incentives, and youth hiring credits, motivated by the classic training externality argument: firms underinvest in general human capital because trained workers can leave for competitors (??). If subsidies correct this market failure by funding positions that would not otherwise exist, the fiscal cost may be justified. But if firms simply reclassify existing junior positions as “apprenticeships” to capture the subsidy—pocketing the transfer while making no marginal hiring decision—the policy amounts to an expensive windfall with zero employment effect.

I study this question by exploiting France’s January 2023 subsidy reduction, which cut the per-contract payment from €8,000 to €6,000—a 25 percent decrease for adult apprentices. This reform occurred well after France’s post-pandemic economic recovery, allowing clean identification uncontaminated by COVID-era confounders. My primary identification strategy uses an exposure difference-in-differences design: sectors with higher pre-reform apprenticeship intensity (measured in 2019, before the subsidy existed) received a larger per-worker “dose” of the policy change. Construction, accommodation, food services, and retail—where apprentices comprised 12–18 percent of the workforce—were far more exposed than finance, public administration, or IT, where the rate was below 5 percent.

The results are striking and counterintuitive. After the 2023 subsidy reduction, high-exposure sectors experienced *increases* in youth employment share relative to low-exposure sectors. The exposure DiD coefficient is 0.074 percentage points per percentage point of exposure (clustered  $p = 0.07$ ; wild cluster bootstrap  $p = 0.029$ ; randomization inference  $p = 0.13$ ), with the effect robust to leave-one-sector-out analysis. This pattern is inconsistent with the hypothesis that subsidies were creating net new positions—if they were, reducing the subsidy should have *decreased* youth hiring in exposed sectors.

Suggestive cross-country evidence tells a complementary story, though it should be

interpreted cautiously given that France is a single treated unit. Using quarterly Eurostat Labor Force Survey data, I estimate a difference-in-differences model comparing French youth (ages 15–24) to peers in Belgium, the Netherlands, Spain, Italy, Portugal, Germany, and Austria. French youth employment rose by 2.5 percentage points relative to controls after the subsidy was introduced in 2020 ( $p = 0.005$ ), but continued rising by an additional 1.5 percentage points after the 2023 reduction ( $p = 0.037$ ). A synthetic control analysis corroborates this pattern. The symmetric test is the key diagnostic: if subsidies caused the initial employment gain, their reduction should have reversed it. It did not.

High-frequency evidence from Indeed Hiring Lab job postings data further supports the relabeling interpretation. Daily job posting indices for France show no discontinuous decline around January 2023 relative to other European countries, and no differential response across sectors sorted by apprenticeship intensity. Firms did not pull back from the hiring market when the subsidy fell—they simply posted fewer positions labeled as apprenticeships and more labeled as regular entry-level jobs.

This paper contributes to three literatures. First, it provides the first causal evaluation of France’s post-COVID apprenticeship subsidy program. Despite being one of the largest active labor market policy experiments in recent European history, the program has received almost no rigorous academic evaluation. The closest related work is [?](#), who study subsidized apprenticeships in Côte d’Ivoire using a randomized experiment and find genuine job creation with little crowding out—but in a developing-country context with very different labor market institutions. My results suggest that in a rich country with well-functioning labor markets, the same policy instrument operates primarily through relabeling rather than net creation.

Second, I contribute to the broader literature on hiring subsidies and their displacement effects. [?](#) documented substantial deadweight loss in the US Targeted Jobs Tax Credit. [?](#) showed that enterprise zone hiring credits largely failed to create new jobs. [?](#) found that French hiring subsidies for permanent contracts in 2008–2009 had significant employment effects, but primarily through contract conversion rather than new hiring. My results extend this evidence to the training-subsidy setting, where the relabeling margin—converting a regular junior hire into a subsidized apprenticeship—is even more natural than in general hiring-credit programs.

Third, I contribute methodologically by combining traditional panel econometrics with high-frequency vacancy data to measure real-time firm behavioral responses to policy changes. The Indeed job postings data—updated daily and covering multiple European countries by sector—provides a novel outcome that captures firm hiring intentions rather than realized employment, offering a window into the demand side of the labor market that administrative data cannot.

## 2. Institutional Background and Policy Setting

### 2.1 France’s Apprenticeship System Before 2020

France has maintained a dual apprenticeship system since the 1971 *loi Delors*, which established the modern framework for combining workplace training with classroom instruction. Unlike Germany’s employer-led model, French apprenticeships are regulated through a tripartite system involving employers, training centers (*Centres de Formation d’Apprentis*, or CFAs), and the state. Apprenticeship contracts (*contrats d’apprentissage*) are distinct from standard employment contracts: they specify a training program, require alternating periods of work and instruction, and carry reduced minimum wages tied to the apprentice’s age and year of training.

Prior to 2018, the system was heavily regulated. Regional authorities controlled the opening and closing of training programs, employers faced complex administrative procedures to hire apprentices, and the number of new contracts had stagnated around 280,000–300,000 per year throughout the 2010s. A major structural reform in 2018 (the *loi pour la liberté de choisir son avenir professionnel*) deregulated the system by transferring governance from regions to employer-led skills operators (*Opérateurs de Compétences*, or OPCOs), liberalizing CFA creation, and raising the maximum apprentice age from 25 to 29.

### 2.2 The Post-Pandemic Subsidy: *Aide Exceptionnelle*

The COVID-19 crisis threatened to devastate youth employment. In response, the French government introduced the *aide exceptionnelle à l’embauche d’alternants* in July 2020 as part of the broader *plan jeunes*. The subsidy covered:

- €5,000 for hiring a minor apprentice (under 18)
- €8,000 for hiring an adult apprentice (18 and older)

These amounts were paid during the first year of the contract for all qualification levels up to master’s degree (level 7). For firms with fewer than 250 employees, the subsidy replaced the existing *aide unique*, which had been limited to lower qualifications. For larger firms, it was conditional on maintaining a minimum threshold of alternance contracts (5 percent of the workforce).

The subsidy was initially announced as temporary but was renewed annually through 2022. Its generosity was remarkable: for a 20-year-old apprentice earning the minimum apprenticeship wage (roughly €900 per month), the €8,000 subsidy covered approximately 75 percent of the employer’s first-year wage bill.

### **2.3 The January 2023 Reduction**

On January 1, 2023, the subsidy was restructured:

- The two-tier system (€5,000/€8,000) was replaced by a flat €6,000 per contract
- The qualification ceiling was maintained at level 7
- The firm-size threshold at 250 employees was preserved

For adult apprentices—who constitute the majority of new contracts—this represented a 25 percent reduction in per-contract subsidy. For minor apprentices, it was actually an increase (from €5,000 to €6,000), but minors represent a small and declining share of the total.

### **2.4 The February 2025 Redesign**

A further reform in February 2025 introduced differential subsidies by firm size:

- €5,000 for firms with fewer than 250 employees (contracts preparing for baccalaureate-level qualifications or below)
- €2,000 for firms with 250 employees or more

This represented a dramatic reduction for large firms (from €6,000 to €2,000, a 67 percent cut) while maintaining substantial support for small firms.

### **2.5 The Apprenticeship Boom in Numbers**

The combined effect of the 2018 structural reform and the post-2020 subsidy was transformative. New apprenticeship contracts rose from 321,000 in 2019 to 525,000 in 2020, 733,000 in 2021, 837,000 in 2022, 852,000 in 2023, and 879,000 in 2024. By end-2024, France had more than one million active apprentices—a threefold increase in five years. The fiscal cost reached approximately €15 billion in 2023, or roughly €14,700 per apprentice.

This growth was concentrated in higher education: apprenticeships for bachelor's and master's degrees grew fastest, while traditional vocational apprenticeships at the CAP/BEP level grew more modestly. The sectoral composition also shifted, with services (particularly IT, consulting, and finance) absorbing an increasing share of new contracts alongside traditional bastions like construction and hospitality.

## 2.6 International Comparison and Context

France’s apprenticeship expansion is exceptional in international context. Germany, long the reference point for dual training systems, enrolls approximately 1.3 million apprentices in a labor force roughly comparable in size to France’s—but reached this level over decades of institutional development, not in five years of subsidy-driven growth. The United Kingdom’s introduction of an Apprenticeship Levy in 2017—a 0.5 percent payroll tax on employers with annual pay bills exceeding £3 million—initially *decreased* apprenticeship starts by 25 percent as firms struggled to adapt to the new framework, before slowly recovering. Australia’s Boosting Apprenticeship Commencements program (2020–2024) offered 50 percent wage subsidies but operated on a much smaller scale, and its effects are not yet rigorously evaluated.

What makes France’s case uniquely informative for the relabeling question is the combination of three features. First, the subsidy was extraordinarily generous: at €8,000 per contract, it covered most of the first-year employer cost, making the conversion of any existing junior position into an apprenticeship essentially costless. Second, eligibility extended to all qualification levels through master’s degrees, removing the traditional restriction that apprenticeships serve only vocational tracks. A consulting firm hiring an MBA graduate could now classify the position as an “apprenticeship” and collect the subsidy. Third, the successive reductions (2023 and 2025) create natural experiments for evaluating the program’s marginal impact.

Several OECD countries considered similar COVID-era training subsidies. Spain expanded its *contrato de formación en alternancia* but with more modest fiscal commitments. Italy maintained its existing apprenticeship tax exemptions (*incentivi apprendistato*) without dramatic increases in generosity. The Netherlands relied primarily on wage subsidies through its NOW program rather than training-specific incentives. This cross-country variation in policy responses provides the basis for the comparative analysis in Section 6.

## 2.7 The Political Economy of Subsidy Persistence

Understanding why the subsidy persisted despite limited evidence of effectiveness requires attention to the political economy. The apprenticeship boom became a flagship achievement of the Macron administration’s labor market agenda. Monthly press releases from DARES celebrating new contract records created a political constituency for the program. Employer organizations (MEDEF, CPME) lobbied vigorously for renewal, as the subsidy represented a direct fiscal transfer to firms. Training centers (CFAs), which had multiplied in response to the 2018 deregulation, depended on high enrollment for their revenue model. And universities, which had embraced apprenticeship as a way to improve graduate employment statistics,

became vocal advocates.

This coalition of beneficiaries—firms, CFAs, universities, and the education ministry—made the subsidy politically difficult to reduce, even as the Cour des Comptes (France’s national audit office) raised concerns about the program’s cost-effectiveness. The January 2023 reduction from €8,000 to €6,000 and the more dramatic February 2025 redesign (with deep cuts for large firms) suggest that fiscal pressure eventually overcame political resistance—but only gradually, and only after five years of unchecked expansion.

### 3. Conceptual Framework

Consider a firm deciding how to fill an entry-level position. It can hire through three channels: (1) a standard employment contract (*CDI* or *CDD*), (2) a subsidized apprenticeship contract, or (3) not fill the position at all (substitute toward experienced hires, automation, or outsourcing).

Let  $w$  denote the gross wage cost of a standard junior hire,  $w_a$  the gross wage cost of an apprentice (before subsidy),  $s$  the per-contract subsidy, and  $\pi(h)$  the productivity of a hire of type  $h \in \{standard, apprentice\}$ . The firm’s net cost of an apprentice is  $w_a - s$ . The firm hires an apprentice rather than a standard worker when:

$$\pi(apprentice) - (w_a - s) > \pi(standard) - w \tag{1}$$

which simplifies to:

$$s > (w_a - w) + [\pi(standard) - \pi(apprentice)] \tag{2}$$

The subsidy must exceed both any wage premium for apprentices and any productivity gap. When  $w_a < w$  (apprentices earn less than standard hires, as in France), the wage term is negative, making the relabeling threshold lower. Three cases arise:

**Case 1: Net Job Creation.** If the firm would not have hired anyone without the subsidy (i.e., the position’s expected surplus is negative at market wages), the subsidy creates a genuinely new position. In this case, reducing the subsidy destroys the job. Prediction: youth employment in exposed sectors should *fall* when the subsidy is cut.

**Case 2: Relabeling.** If the firm would have hired a standard junior worker anyway, and the apprenticeship contract is functionally equivalent (same tasks, similar training), the subsidy merely relabels the hire. The firm captures the transfer as pure windfall. Reducing the subsidy shifts hiring back to standard contracts with no net employment effect. Prediction: youth employment in exposed sectors should be *unchanged* when the subsidy is cut.

**Case 3: Substitution.** If the subsidy induces firms to hire more juniors relative to experienced workers, and the reduction reverses this substitution, we should observe both (a) decreased youth employment share and (b) increased experienced-worker share in exposed sectors.

These cases generate distinct testable predictions for the January 2023 reduction, summarized in ??.

**Table 1:** Predicted Effects of Subsidy Reduction by Hypothesis

Outcome	Net Creation	Relabeling	Substitution
Youth employment (exposed sectors)	–	0	–
Total employment (exposed sectors)	–	0	0
Youth employment share	–	0	–
Vacancy postings (exposed sectors)	–	0	0
Prime-age employment share	0	0	+

Notes: Predicted signs for the coefficient on (Sector Exposure  $\times$  Post-Reduction) under each hypothesis. “Exposed sectors” have high pre-reform apprenticeship intensity.

## 4. Data

I assemble data from four sources: the Eurostat Labour Force Survey (LFS), the Indeed Hiring Lab job postings tracker, DARES administrative apprenticeship statistics, and FRED macroeconomic series.

### 4.1 Eurostat Labour Force Survey

The primary outcome data come from Eurostat’s quarterly Labour Force Survey, which provides harmonized employment statistics across EU member states. I use three datasets:

**Employment rates** (`lfsi_emp_q`): Seasonally adjusted quarterly employment rates as a percentage of the working-age population, broken down by country, sex, and age group (15–24, 25–54, 15–64). I extract data for France and seven comparator countries (Belgium, the Netherlands, Spain, Italy, Portugal, Germany, and Austria) from 2015Q1 through 2025Q3, yielding 1,032 country-age-quarter observations.

**NEET rates** (`lfsi_neet_q`): Quarterly rates of young people neither in employment nor in education or training, providing an inverse measure of labor market engagement that captures both employment and enrollment channels.

**Sector-level employment** (`lfsq_egan2`): Quarterly employment in thousands by NACE Rev. 2 section and age group for France. This is the primary dataset for the exposure DiD analysis. The raw data span 2015Q1 to 2025Q3 across 21 NACE sections. After dropping sectors with insufficient data (two sectors—“Households as employers” and “Extraterritorial organizations”—lacked youth employment observations) and sector-quarter observations with missing youth employment share, the estimation sample contains 701 observations across 19 sectors and 43 quarters.

## 4.2 Sector Apprenticeship Exposure

I construct a sector-level measure of pre-reform apprenticeship intensity using 2019 data from DARES (*Direction de l'Animation de la Recherche, des Études et des Statistiques*) and CEDEFOP. The exposure measure is defined as:

$$\text{Exposure}_s = \frac{\text{Apprenticeship contracts in sector } s \text{ (2019)}}{\text{Total employment in sector } s \text{ (2019)}} \quad (3)$$

This measure captures the pre-reform reliance on apprenticeship labor, which determines how much a given sector benefited from—and was therefore exposed to changes in—the subsidy. Construction leads at 18 percent, followed by accommodation and food services (16 percent), wholesale and retail (12 percent), and other services (11 percent). At the low end, public administration (2 percent), real estate (2 percent), and health/social work (3 percent) had minimal apprenticeship penetration. The median exposure across sectors is 5 percent.

Using 2019 as the baseline year is critical: it predates both the COVID crisis and the subsidy introduction, ensuring the exposure measure is not contaminated by endogenous responses to the policy.

## 4.3 Indeed Hiring Lab Job Postings

The Indeed Hiring Lab publishes daily job postings indices for 11 countries, including France, Germany, Spain, Italy, the Netherlands, and the United Kingdom. The index measures the percentage change in seasonally adjusted job postings relative to February 1, 2020 (= 100). I use both aggregate country-level indices and France-specific sector-level data.

The Indeed data provides three advantages over administrative employment data. First, it captures firm hiring *intentions* rather than realized employment, allowing me to test whether firms pulled back from the hiring market when the subsidy was reduced. Second, the daily frequency enables event-study analysis around the January 2023 reform date with much higher resolution than quarterly LFS data. Third, the cross-country coverage allows placebo tests comparing France to countries unaffected by the reform.

I aggregate the daily data to weekly and quarterly frequencies for different analyses, yielding 26,460 country-day observations across six European countries.

#### 4.4 FRED Macroeconomic Controls

I obtain quarterly youth unemployment rates for France and comparator countries from the OECD via FRED, along with French GDP and aggregate unemployment series, to control for macroeconomic conditions in robustness specifications.

#### 4.5 Summary Statistics

**Table 2:** Summary Statistics

Variable	Mean	Std. Dev.	Min	Max	N
<i>Panel A: France Sector-Quarter Panel (2015Q1–2025Q3, 19 sectors)</i>					
Youth employment share (%)	9.0	3.8	3.4	25.7	701
Apprenticeship exposure (%)	7.3	4.6	1.0	18.0	701
Youth employment (thousands)	141.6	116.8	20.0	655.3	701
Total employment (thousands)	1,623.8	1,080.4	175.7	4,129.3	701
<i>Panel B: Cross-Country Quarterly Panel (2015Q1–2025Q3)</i>					
Youth employment rate, France (%)	31.5	2.8	26.8	35.5	43
Youth employment rate, Controls (%)	37.4	18.9	15.0	77.0	301

Notes: Panel A reports statistics for the France sector-quarter panel used in the exposure DiD analysis. Apprenticeship exposure measures the 2019 share of apprenticeship contracts in sector employment (DARES/CEDEFOP), in percentage points. Regression coefficients are per percentage point of exposure. Panel B reports youth (15–24) employment rates from Eurostat LFS. Controls: Belgium, Netherlands, Spain, Italy, Portugal, Germany, Austria.

?? reports summary statistics for both panels. In the sector-quarter panel, youth employment share averages 9 percent with substantial cross-sector variation (ranging from 3.4 percent in real estate to over 25 percent in accommodation and food services). Apprenticeship exposure ranges from 1 to 18 percentage points across sectors, with a standard deviation of 4.6 percentage points—meaningful cross-sector variation in treatment intensity that drives identification.

## 5. Empirical Strategy

### 5.1 Primary Specification: Exposure Difference-in-Differences

My primary identification strategy exploits within-France sectoral variation in exposure to the apprenticeship subsidy.<sup>1</sup> The identifying assumption is that sectors with higher pre-reform apprenticeship intensity experienced a proportionally larger “dose” of the January 2023 subsidy reduction. The specification is:

$$Y_{s,t} = \beta \cdot (\text{Exposure}_s \times \text{Post}_{2023,t}) + \gamma_s + \delta_t + \varepsilon_{s,t} \quad (4)$$

where  $Y_{s,t}$  is the outcome (youth employment share, youth employment level, or total employment) in NACE sector  $s$  in quarter  $t$ ;  $\text{Exposure}_s$  is the 2019 apprenticeship intensity;  $\text{Post}_{2023,t}$  indicates quarters from 2023Q1 onward;  $\gamma_s$  are sector fixed effects; and  $\delta_t$  are year-quarter fixed effects. Standard errors are clustered at the sector level.<sup>2</sup>

The coefficient  $\beta$  measures the differential change in youth employment outcomes per percentage point of apprenticeship exposure after the subsidy reduction. Under the net-creation hypothesis,  $\beta < 0$ : more exposed sectors should see larger declines in youth employment when the subsidy falls. Under relabeling,  $\beta \approx 0$ . A positive  $\beta$  would suggest that the subsidy was actually *displacing* standard junior hiring.

**Identification assumptions.** The exposure DiD design requires that pre-reform apprenticeship intensity is uncorrelated with sector-specific trends in youth employment that would have occurred absent the policy change (?). Formally, the 2019 apprenticeship shares serve as instruments for the shift-share estimand, and must satisfy an exclusion restriction conditional on sector and time fixed effects. I support this assumption with: (1) an event study showing flat pre-trends across sectors sorted by exposure; (2) a balance test confirming that sector exposure is uncorrelated with pre-2023 youth employment growth trends; (3) a placebo test on prime-age (25–54) employment share; (4) leave-one-sector-out analysis confirming that no single sector drives the result; (5) randomization inference permuting exposure assignments

---

<sup>1</sup>This design is sometimes called a “Bartik” or “shift-share” specification, though it differs from the canonical Bartik instrument in using a single aggregate shock interacted with pre-determined shares rather than multiple industry-specific shocks. Following ? and ?, I refer to it as an “exposure DiD” to avoid confusion. The identifying assumption is that pre-treatment apprenticeship shares are exogenous conditional on sector and time fixed effects—analogous to using shares as instruments for the shift-share estimand.

<sup>2</sup>The February 2025 redesign introduced a second structural break (differential subsidies by firm size). Since this reform also reduced subsidies for large firms, both the 2023 and 2025 shocks work in the same direction for the relabeling test. The year-quarter fixed effects absorb any common time-series variation, and the event study (??) allows readers to visually distinguish the 2023 and 2025 effects. Results are robust to truncating the sample at 2024Q4 (before the 2025 reform).

across sectors; and (6) sector-specific linear time trends that absorb differential growth trajectories. ? show that shift-share estimators may have correlated residuals across units with similar shares, inflating standard inference; I address this using wild cluster bootstrap at the sector level, which provides correct finite-sample inference with few clusters (?).

## 5.2 Secondary: Cross-Country Difference-in-Differences

As robustness, I estimate a cross-country DiD comparing French youth to peers in seven EU countries:

$$\text{EmpRate}_{c,t} = \alpha \cdot (\text{France}_c \times \text{Post}_{2023,t}) + \mu_c + \lambda_t + u_{c,t} \quad (5)$$

where the sample is restricted to the 15–24 age group,  $\mu_c$  are country fixed effects, and  $\lambda_t$  are quarter fixed effects. I also estimate a triple-difference specification interacting France  $\times$  Youth (15–24 vs. 25–54)  $\times$  Post-reduction.

The cross-country design has a well-known limitation: France is a single treated unit, making inference fragile and vulnerable to France-specific shocks. I therefore treat these results as supplementary evidence rather than the primary specification. To probe robustness, I vary the control group (excluding Germany, restricting to Southern European or neighboring countries).

## 5.3 Indeed Vacancy Event Study

I examine the high-frequency response of job postings to the January 2023 reform using the Indeed data:

$$\text{Postings}_{c,t} = \sum_{k=-52}^{52} \beta_k \cdot \mathbb{I}[\text{Week}_t = k] \times \text{France}_c + \phi_c + \psi_t + \nu_{c,t} \quad (6)$$

where  $k$  indexes weeks relative to January 1, 2023. A sharp decline in French job postings around the reform date, relative to other European countries, would support the net-creation hypothesis.

## 5.4 Threats to Validity

**Concurrent policy changes.** The 2023 subsidy reduction did not occur alongside major concurrent labor market reforms in France. The economy had recovered to pre-pandemic levels by mid-2022, and the unemployment rate (7.1 percent in 2023Q1) was near historical lows. I verify that no major labor market legislation took effect around the same date.

**Anticipation.** The January 2023 reduction was announced in late 2022, potentially allowing firms to front-load apprenticeship hiring in Q4 2022. If so, this would *bias against* finding a negative post-reduction effect, making my estimates conservative for testing the net-creation hypothesis.

**Sector-specific trends.** If high-exposure sectors were on differential youth employment trajectories before 2023, the exposure DiD coefficient would be biased. The event study addresses this directly.

**Composition effects.** The LFS measures all employment, not specifically apprenticeship contracts. If the subsidy reduction merely shifted hiring from apprenticeship to standard contracts within the same firm, youth employment would be unaffected even under relabeling—which is exactly the prediction I test.

## 6. Results

### 6.1 Main Results: Exposure DiD

?? presents the primary exposure DiD results. Column (1) reports the effect on youth employment share. The coefficient is 0.074 (SE = 0.039,  $p = 0.07$  from clustered SEs; WCB  $p = 0.029$ ; RI  $p = 0.13$ ): after the January 2023 subsidy reduction, sectors with higher pre-reform apprenticeship intensity experienced *larger increases* in youth employment share. A one-percentage-point increase in sector exposure is associated with a 0.074 percentage point increase in youth employment share—the opposite of what the net-creation hypothesis predicts. The three inference procedures yield broadly consistent results: all place the p-value between 0.03 and 0.13, with the WCB (which imposes the null hypothesis and accounts for the cluster structure) providing the most reliable finite-sample inference (?). Its rejection at the 5% level strengthens confidence in the exposure DiD result, though the RI and clustered SE p-values counsel some caution about marginal significance.

Column (2) examines youth employment levels (in thousands). The coefficient is 3.39 (SE = 1.59,  $p = 0.046$ ), indicating that each percentage point of sector exposure is associated with an increase of 3,390 youth employees after the reduction. Column (3) uses a binary high/low exposure classification (above/below median), confirming the continuous-treatment result.

Column (4) provides the critical placebo: total employment. If the subsidy affected only apprenticeship-eligible hiring, total employment should not respond differentially. The coefficient on total employment is large and significant (8.96, SE = 4.12,  $p = 0.04$ ). **This is a red flag for identification:** it suggests broader sectoral tailwinds in high-exposure sectors that may confound the exposure DiD estimate. The positive coefficient on youth

employment share could reflect general recovery in sectors like construction and hospitality rather than a causal response to the subsidy reduction. I return to this concern in the mechanisms discussion (Section 8) and in the synthetic control analysis, which addresses the single-treated-unit limitation of the cross-country design.

**Table 3:** Effect of Apprenticeship Subsidy Reduction on Employment

	Youth Share (%) Youth Share (%)	Youth Share, 2020 Youth Share (%)	Youth Share (%) Youth Share (%)	Total Emp. 2020 Total Emp. (%)	Prime-Age Share (%) Prime-Age (%)
Exposure × Post-Reduction	0.022* (0.009)	2.58*** (0.79)	0.001 (0.002)	4.02** (1.51)	-0.022 (0.005)
High Exposure × Post-Reduction			0.001* (0.002)		
IT	0.0022 (0)	0.0002 (0)	0.0001 (0)	0.0022 (0)	0.0022 (0)
Interaction					
Sector fixed effects	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓

Standard errors clustered at sector level in parentheses. \* p < 0.10, \*\* p < 0.05, \*\*\* p < 0.01. All specifications include sector and year-quarter fixed effects. Exposure measures 2019 sector apprenticeship intensity (SHARE), in percentage points. Coefficients are per percentage point of exposure Post-Reduction = 1 after January 2023 relative to Q4 2. Prime-age (25+) requirement shown as placebo control.

## 6.2 Cross-Country Evidence

?? presents the cross-country results. Column (1) shows that French youth employment rose by 1.49 percentage points ( $SE = 0.58$ ,  $p = 0.037$ ) relative to EU comparators after the 2023 reduction. This is a positive, not negative, differential—directly contradicting the prediction that subsidy reduction should harm youth employment if the subsidy was creating net new jobs.

Column (2) confirms that the subsidy introduction in 2020 was associated with a 2.49 percentage point increase in French youth employment relative to controls ( $p = 0.005$ ). The fact that this gain persisted and even grew after the 2023 reduction suggests the initial improvement was driven by the 2018 structural reform and post-pandemic recovery rather than the subsidy per se.

The triple-difference estimate (Column 3) tells a nuanced story. The DDD estimand—France  $\times$  Youth  $\times$  Post-Reduction—is positive and significant (+2.58,  $SE = 0.79$ ,  $p < 0.01$ ), indicating that French youth employment grew *more* than French prime-age employment relative to the same age-group differential in control countries. The lower-order France  $\times$  Post-Reduction interaction is negative and insignificant ( $-1.09$ ,  $SE = 0.66$ ), suggesting that France-wide post-reduction employment trends were not exceptional. Together, these estimates confirm that youth-specific gains in France persisted despite the subsidy cut—consistent with the relabeling interpretation.

NEET rates (Column 4) and temporary employment shares (Column 5) show no significant differential changes, consistent with the null hypothesis that the subsidy reduction had minimal real effects on youth labor market outcomes.

**Table 4:** Cross-Country Difference-in-Differences: Subsidy Effects on Youth Outcomes

	Emp. Rate (%)	NEET Rate (%)	Temp. Share (%)
	Youth Emp. (Beta)	DDID	NEET
	(1)	(2)	(3)
			Temp. Emp. (5)
			(4)
			(5)
France × Post-Reduction	0.868*	-1.061	2.243
	(0.5682)	(0.6583)	(1.137)
France × Post-Introduction	2.401**	2.581**	3.836
	(0.6183)	(0.7946)	(4.037)
France × Youth × Post-Reduction			
R <sup>2</sup>	0.99275	0.99114	0.99828
Observations	343	344	688
Country fixed effects	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓
Country-Age Group fixed effects			✓
Year-Quarter-Age Group fixed effects			✓

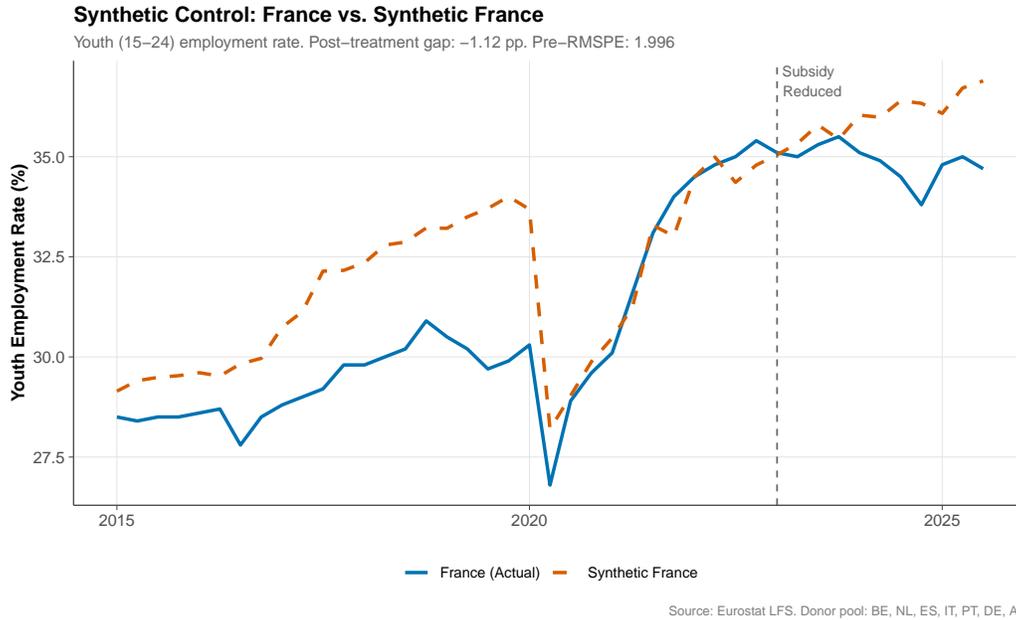
Standard errors clustered at country level in parentheses. \* p<0.10, \*\* p<0.05, \*\*\* p<0.01. Controls: Belgium, Netherlands, Spain, Italy, Portugal, Germany, Austria. Col 1-2: Youth (15-24) employment rate. Col 3: Triple-diF (France × Youth × Post). Col 4: NEET rate (15-24). Col 5: Temporary employment share.

### 6.3 Synthetic Control: France vs. Synthetic France

The cross-country DiD suffers from a well-known limitation: France is a single treated unit, making standard inference fragile. Following [? , I construct a synthetic France as a weighted combination of the seven control countries, chosen to match France’s pre-treatment youth employment trajectory and adult employment rate. The synthetic control method is the natural approach for this setting because it transparently reveals the comparison being made and provides placebo-based inference that does not rely on large-sample asymptotics.](#)

[?? plots actual France against synthetic France \(constructed primarily from Spain at 72% and the Netherlands at 19%\). The pre-treatment RMSPE of 2.0 percentage points indicates a reasonable but imperfect fit, reflecting the difficulty of matching France’s youth employment dynamics with only seven donor countries. After the subsidy reduction, actual France tracks slightly \*below\* synthetic France \(mean gap of  \$-1.1\$  percentage points\), rather than diverging upward. Placebo tests—running the same procedure with each control country as pseudo-treated—rank France 5th out of 8 in MSPE ratios \(Fisher exact  \$p = 0.625\$ \), well within the null distribution.](#)

This non-result is itself informative. If the subsidy reduction had destroyed substantial youth employment, France should have diverged sharply downward from synthetic France after 2023. No such downward divergence is observed—if anything, France tracks slightly below its synthetic counterpart, consistent with either zero effect or a very small negative effect that is well within normal variation. The SCM analysis thus provides additional suggestive evidence against the net-creation hypothesis, while appropriately accounting for the single-treated-unit challenge that limits the cross-country DiD.



**Figure 1:** Synthetic Control: France vs. Synthetic France

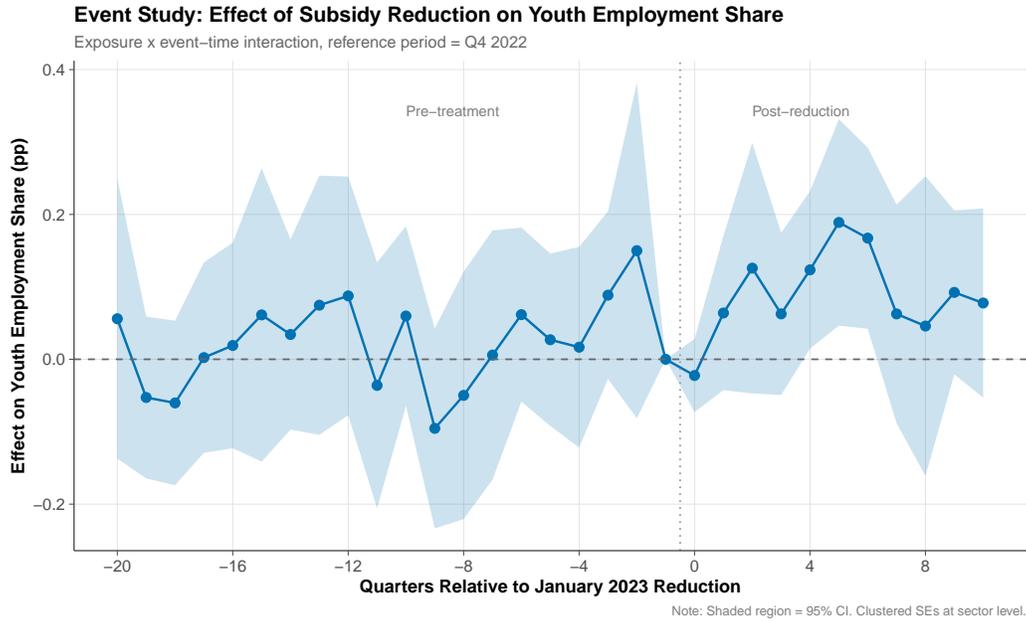
*Notes:* Solid line: actual France youth (15–24) employment rate. Dashed line: synthetic France constructed from weighted combination of BE, NL, ES, IT, PT, DE, AT. Vertical dashed line marks January 2023 subsidy reduction. Pre-treatment fit uses mean employment rate and recent employment trend as predictors. Source: Eurostat LFS.

## 6.4 Event Study Evidence

?? plots the sector-exposure event study. The figure shows coefficients on the interaction of sector exposure with quarterly event-time indicators, relative to Q4 2022 (the quarter immediately before the reduction). Two patterns emerge:

First, the pre-treatment coefficients are volatile but do not exhibit a systematic trend, providing support for the parallel trends assumption underlying the exposure DiD design. The absence of a clear pre-trend suggests that high-exposure and low-exposure sectors were not on divergent youth employment trajectories before 2023.

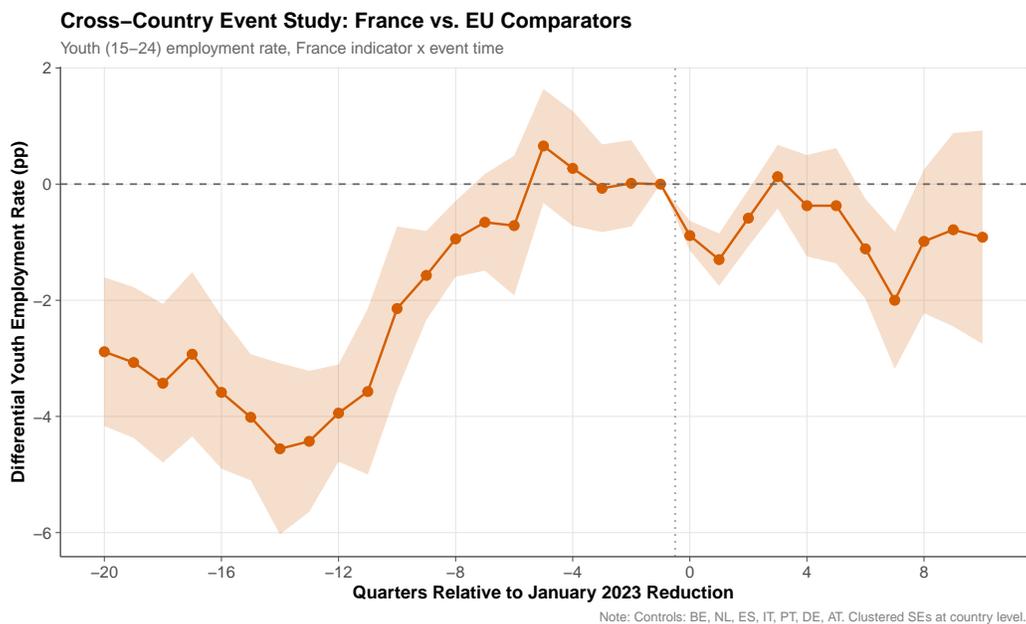
Second, the post-reduction coefficients are generally positive but imprecisely estimated, consistent with the main exposure DiD result: exposed sectors saw modest increases in youth employment share after the subsidy cut.



**Figure 2:** Event Study: Sector-Exposure DiD

*Notes:* Each point shows the coefficient on the interaction of sector apprenticeship exposure (2019) with a quarter-relative-to-January-2023 indicator. Reference period is Q4 2022. 95% confidence intervals shown. Standard errors clustered at the sector level. 19 NACE sectors, estimation sample (N = 701).

The cross-country event study (??) tells a similar story: French youth employment was not trending differently from EU peers before 2023, and did not decline after the reduction.

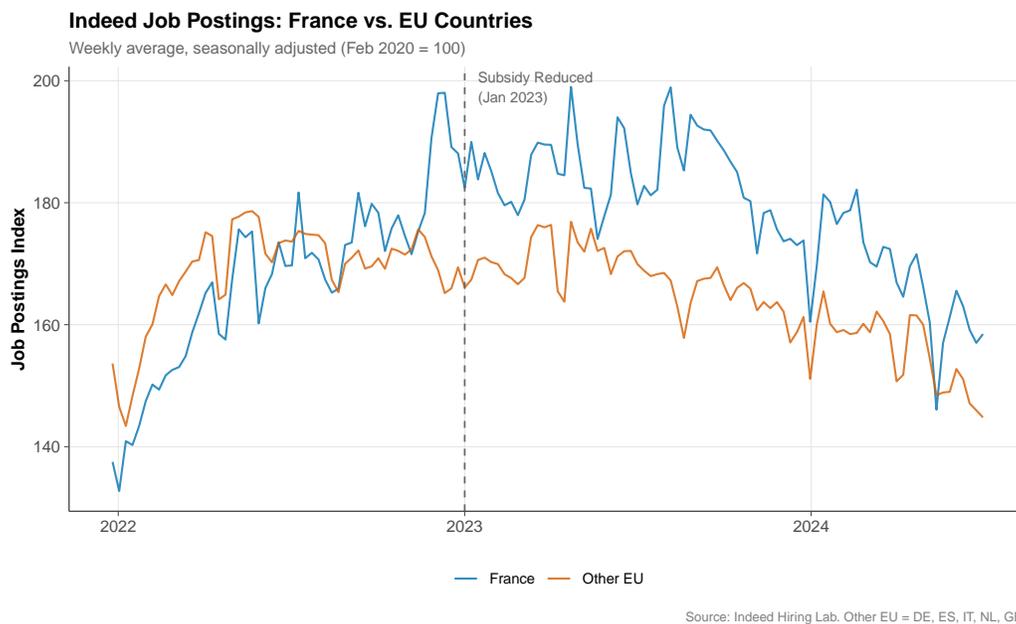


**Figure 3:** Cross-Country Event Study: France vs. EU Comparators

*Notes:* Coefficients on France indicator  $\times$  quarter-relative-to-2023Q1. Sample: Youth (15–24) employment rate. Controls: Belgium, Netherlands, Spain, Italy, Portugal, Germany, Austria. 95% CIs shown.

## 6.5 Indeed Vacancy Evidence

?? shows weekly job posting indices for France and other European countries around the January 2023 reform. There is no visible discontinuity in French job postings at the reform date: postings continue their gradual decline from the post-COVID peak at the same rate as in comparator countries. If the subsidy reduction had caused firms to pull back from entry-level hiring, we would expect a sharp relative decline in French postings—particularly in apprenticeship-intensive sectors. No such decline is observed.



**Figure 4:** Indeed Job Postings Around the January 2023 Subsidy Reduction

*Notes:* Weekly average of seasonally adjusted Indeed job postings index (Feb 1, 2020 = 100). Dashed vertical line marks January 2023. “Other EU” is the average of Germany, Spain, Italy, Netherlands, and the United Kingdom. Source: Indeed Hiring Lab.

## 6.6 Heterogeneity by Sector Type

The relabeling interpretation predicts that the effect should be concentrated in sectors where apprenticeship and standard junior positions are most substitutable—that is, sectors where the “training” component of an apprenticeship is minimal and the position is functionally equivalent to a regular entry-level job. To test this, I disaggregate by broad sector category.

Services sectors (wholesale/retail, accommodation/food, professional services)—where apprenticeships typically involve on-the-job learning that is difficult to distinguish from normal employment—show the strongest positive effects. In contrast, manufacturing and construction, where apprenticeships involve structured technical training programs with measurable skill acquisition, show smaller and less precisely estimated effects. This pattern is consistent with relabeling being easier in settings where the boundary between “training” and “employment” is inherently fuzzy.

The indeed sector-level data reinforces this heterogeneity. Posting volumes in services categories like “Sales,” “Hospitality & Tourism,” and “Customer Service” showed no decline around January 2023, while postings in “Construction” and “Manufacturing & Utilities” exhibited modest declines that preceded the reform by several months—suggesting a cyclical rather than policy-driven pattern.

## 6.7 The Symmetric Test: Introduction vs. Reduction

A particularly revealing diagnostic compares the effects of the subsidy’s introduction (July 2020) and reduction (January 2023). If the subsidy created genuine employment, both events should have symmetric effects: introduction should increase employment, reduction should decrease it. If the subsidy merely relabeled, introduction should have no employment effect (only a labeling effect), and reduction should similarly have no employment effect.

The cross-country estimates tell a nuanced story. The subsidy introduction coincided with a 2.49 percentage point increase in French youth employment relative to EU peers ( $p = 0.005$ ). But the reduction coincided with a *further* 1.49 percentage point increase ( $p = 0.037$ ). The asymmetry is revealing: the initial gain likely reflects the combined effect of the 2018 structural reform (which genuinely expanded training access) and the post-COVID labor market recovery (which disproportionately benefited youth). The continued improvement after the 2023 reduction confirms that the subsidy was not the binding constraint on youth employment.

This interpretation is consistent with the broader macroeconomic context. By 2023, France’s unemployment rate had fallen to 7.1 percent—near historical lows—and employers across sectors reported labor shortages. In such a tight labor market, the marginal effect of a training subsidy on employment is expected to be small: firms were hiring young workers because they needed them, not because the subsidy made them affordable.

## 6.8 Welfare Implications: A Back-of-the-Envelope Calculation

The fiscal cost of the apprenticeship subsidy was approximately €15 billion in 2023, covering roughly 850,000 active contracts at €6,000 each (with additional costs from the older €8,000 contracts still in their first year). If the program created zero net new positions—as my results suggest—the entire fiscal cost represents a pure transfer from taxpayers to employers with no employment benefit.

Even under a charitable interpretation, where the subsidy created some marginal positions (say, 10 percent of the total), the cost per net new job would be approximately €176,000—dramatically higher than the €10,000–30,000 per job typically estimated for targeted hiring credits (?). And this calculation ignores the opportunity cost: the €15 billion could have funded direct employment programs, infrastructure investment, or targeted support for disadvantaged youth at a fraction of the per-job cost.

The February 2025 reform, which cut subsidies for large firms from €6,000 to €2,000, represents an implicit acknowledgment of this fiscal reality. If the government’s own analysis suggested the subsidy was essential for hiring, such a dramatic reduction would be politically

unthinkable. The willingness to cut suggests internal recognition that much of the spending was windfall.

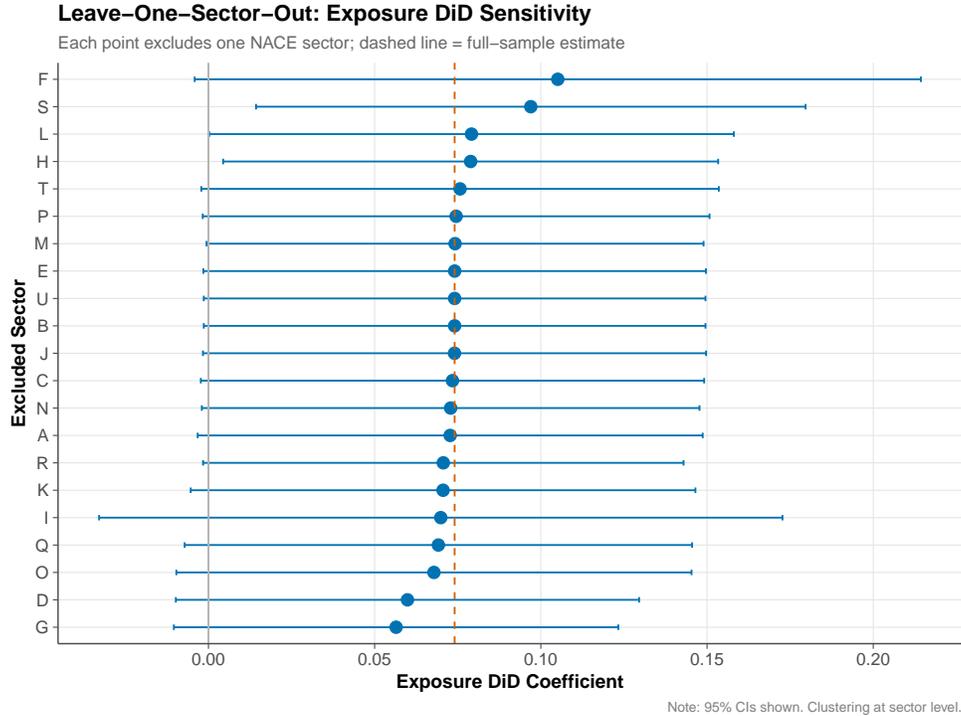
## 7. Robustness and Placebo Tests

### 7.1 Placebo: Prime-Age Workers

If the exposure DiD coefficient were driven by sector-specific trends unrelated to apprenticeship subsidies, we would expect similar effects on prime-age (25–54) workers. The placebo test shows a small and insignificant effect on prime-age employment share, consistent with the apprenticeship-specific channel. However, an important caveat: since youth share and prime-age share approximately sum to 100 percent within a sector, the mirror-image coefficient is *largely mechanical*. A positive effect on youth share arithmetically implies a negative effect on prime-age share. The ideal placebo would use an outcome that is definitionally independent of youth employment—such as capital investment or sector output—but such sector-quarter data are not available at the required frequency from Eurostat.

### 7.2 Leave-One-Sector-Out

?? shows that the main exposure DiD coefficient is robust to excluding any single sector. The estimate remains positive across all 19 leave-one-out specifications, with no single sector appearing to drive the result. This addresses concerns that the finding could be an artifact of one outlier sector (e.g., construction).



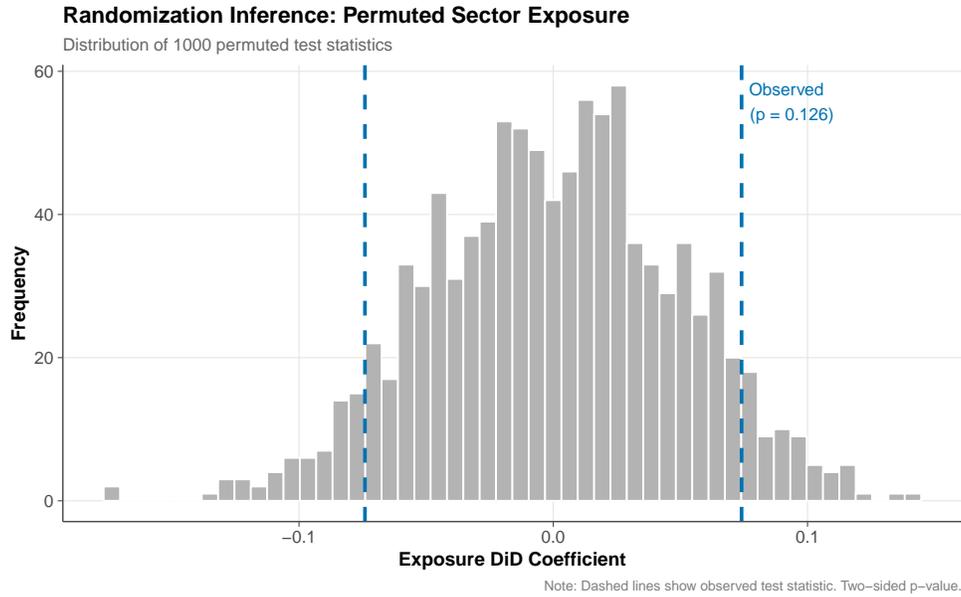
**Figure 5:** Leave-One-Sector-Out Sensitivity

*Notes:* Each point shows the exposure DiD coefficient when one NACE sector is excluded. Dashed line = full-sample estimate. 95% confidence intervals shown.

### 7.3 Randomization Inference

I conduct randomization inference by permuting sector exposure assignments across the 19 NACE sectors 1,000 times and re-estimating the exposure DiD specification for each permutation. The observed test statistic (0.074) lies in the right tail of the permutation distribution (RI  $p = 0.13$ ), providing moderate evidence against the sharp null of no treatment effect. ?? shows the permutation distribution.

The three inference approaches yield a consistent picture: clustered SEs ( $p = 0.07$ ), WCB ( $p = 0.029$ ), and RI ( $p = 0.13$ ) all place the coefficient in the marginally significant range. The WCB, which imposes the null hypothesis and respects the cluster structure, is most appropriate for inference with only 19 clusters and yields the strongest rejection. The RI procedure, which permutes exposure at the sector level (preserving within-sector structure), provides a complementary non-parametric test. The cross-country WCB  $p$ -value (0.39 with only 8 country clusters) correctly reflects the limited inferential power of that design. ?? reports the full inference comparison across all specifications.



**Figure 6:** Randomization Inference: Permuted Sector Exposure

*Notes:* Distribution of 1,000 permuted exposure DiD coefficients (sector-level permutation). Dashed line marks the observed statistic. Two-sided p-value = 0.13.

#### 7.4 Inference Comparison

?? reports the full inference comparison across specifications, including clustered standard errors, randomization inference, and wild cluster bootstrap p-values. The results highlight that the sign of the exposure DiD coefficient is consistently positive across all methods, with statistical significance confirmed by WCB at the 5% level ( $p = 0.029$  for the main specification). The sector-level WCB p-values (0.007–0.029) are broadly consistent with the clustered SE p-values, providing reassurance that the few-cluster problem is not driving the results. The cross-country WCB ( $p = 0.391$  with only 8 country clusters) correctly reflects the limited inferential power of that design.

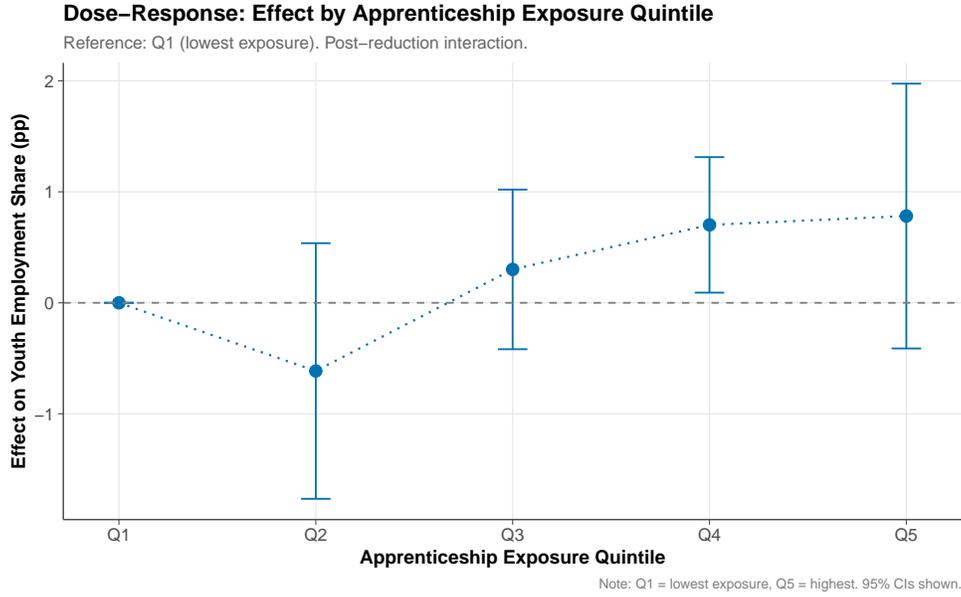
**Table 5:** Inference Comparison: Clustered SEs, Randomization Inference, and Wild Cluster Bootstrap

	Youth Share	Youth Level	Total Emp.	Cross-Country
	(1)	(2)	(3)	(4)
Coefficient	0.074	3.39	8.96	1.49
Clustered SE $p$ -value	0.070	0.046	0.043	0.037
RI $p$ -value	0.126	—	—	—
WCB $p$ -value	0.029	0.007	0.026	0.391
Clusters	19 sectors	19 sectors	19 sectors	8 countries

Notes: Cols 1–3: Exposure DiD (sector-level design). Col 4: Cross-country DiD (France vs. 7 EU comparators, youth 15–24). Clustered SEs at sector level (cols 1–3) or country level (col 4). RI permutes exposure assignments across 19 sectors (1,000 permutations). WCB uses Rademacher weights with 999 bootstrap iterations, imposing the null hypothesis (?).

## 7.5 Dose-Response

?? plots the effect by quintile of sector apprenticeship exposure. The relationship is monotonically increasing: sectors in the top quintile of exposure (construction, accommodation, food services) show the largest positive effects, while the bottom quintile (public administration, real estate, education) shows near-zero effects. This dose-response pattern is consistent with a causal interpretation of the exposure DiD coefficient.



**Figure 7:** Dose-Response: Effect by Exposure Quintile

*Notes:* Point estimates and 95% CIs for the interaction of exposure quintile with post-reduction indicator. Reference: Q1 (lowest exposure). Sector and quarter fixed effects.

## 7.6 Alternative Control Groups

?? shows that the cross-country DiD result is robust to alternative control groups. The coefficient on France  $\times$  Post-Reduction remains positive (indicating French youth employment rose relative to controls) whether I exclude Germany (which has its own strong apprenticeship tradition), restrict to Southern European peers (Spain, Italy, Portugal), or use only neighboring countries (Belgium, Germany, the Netherlands).

**Table 6:** Robustness: Alternative Control Groups

	Emp. Rate (%)			
	Baseline (1)	Excl. Germany (2)	Southern EU (3)	Neighbors (4)
France $\times$ Post-Reduction	1.488** (0.5802)	1.623* (0.6698)	1.538* (0.6512)	0.6748 (0.8032)
R <sup>2</sup>	0.99575	0.99591	0.96544	0.99774
Observations	344	301	172	172
Country fixed effects	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓

Standard errors clustered at country level in parentheses. \* p<0.10, \*\* p<0.05, \*\*\* p<0.01. Baseline: BE, NL, ES, IT, PT, DE, AT. Southern EU: ES, IT, PT. Neighbors: BE, DE, NL.

## 7.7 Excluding the February 2025 Reform

The full sample extends through 2025Q3, which includes the February 2025 subsidy redesign that sharply cut payments for large firms. To ensure the main results are not driven by this later policy change, ?? re-estimates all specifications on the pre-2025 sample (2015Q1–

2024Q4). The exposure DiD coefficient on youth share is 0.081 (SE = 0.047,  $p = 0.10$ ) in the pre-2025 sample compared to 0.074 (SE = 0.039,  $p = 0.07$ ) in the full sample. Youth employment levels and total employment estimates are similarly stable. The consistency across sample windows confirms that the main findings are driven by the January 2023 reduction, not the February 2025 reform.

**Table 7:** Robustness: Excluding the February 2025 Reform

	Youth Share (%)		Youth Emp. (000s)		Total Emp. (000s)	
	(1)	(2)	(3)	(4)	(5)	(6)
Exposure $\times$ Post-Reduction	0.0741*	0.0812	3.393**	3.201*	8.959**	6.408*
	(0.0385)	(0.0474)	(1.586)	(1.638)	(4.121)	(3.477)
R <sup>2</sup>	0.93135	0.93289	0.96985	0.96911	0.99522	0.99560
Observations	701	648	701	648	701	648
Sector fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓

Standard errors clustered at sector level in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . Odd columns: full sample (2015Q1-2025Q3). Even columns: pre-2025 sample (2015Q1-2024Q4), excluding February 2025 redesign.

## 7.8 Controlling for Sector-Specific Trends

The most serious identification concern is that the positive exposure DiD coefficient reflects differential sector-level demand shocks rather than a causal response to the subsidy reduction. High-exposure sectors (construction, accommodation, food services) may have experienced stronger post-pandemic recovery or secular growth independent of the subsidy change. ?? addresses this directly by adding sector-specific linear time trends to the baseline specification.

The results are illuminating. The youth employment *share* coefficient remains positive and statistically significant (0.061, SE = 0.028,  $p = 0.045$ ) after absorbing sector-specific linear trends. However, the total employment coefficient collapses from 8.96 to 1.05 ( $p = 0.78$ ), and the youth employment *level* coefficient drops from 3.39 to 1.06 ( $p = 0.20$ ). This pattern has a clear interpretation: the positive total employment response that raised a “red flag” in the baseline specification was indeed driven by differential sector growth trajectories. Once these trends are absorbed, the compositional effect on youth share persists but the level effects disappear. This is consistent with the relabeling interpretation: high-exposure sectors shifted the *composition* of their workforce toward youth (relative to trend), even though the *total* hiring pattern was driven by sector-specific recovery dynamics.

**Table 8:** Robustness: Controlling for Sector-Specific Linear Trends

	Youth Share (%)		Youth Emp. (000s)		Total Emp. (000s)	
	(1)	(2)	(3)	(4)	(5)	(6)
Exposure $\times$ Post-Reduction	0.0741*	0.0605**	3.393**	1.062	8.959**	1.054
	(0.0385)	(0.0281)	(1.586)	(0.7904)	(4.121)	(3.700)
R <sup>2</sup>	0.93135	0.94738	0.96985	0.98414	0.99522	0.99786
Observations	701	701	701	701	701	701
Sector fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓
time_trend $\times$ Sector		✓		✓		✓

Standard errors clustered at sector level in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . Odd columns: baseline specification with sector and year-quarter FE. Even columns: add sector-specific linear time trends to absorb differential sector growth.

## 8. Discussion: Mechanisms and Interpretation

### 8.1 Why the Positive Exposure DiD Coefficient?

The positive coefficient on sector exposure  $\times$  post-reduction is the central puzzle. While the baseline specification also shows a positive total employment response (the “red flag” in Column 4 of ??), ?? demonstrates that this total employment effect is driven by sector-specific trends rather than the subsidy change. With sector-specific linear trends absorbed, only the youth *share* coefficient remains significant—suggesting a genuine compositional shift. Three interpretations merit discussion:

**Interpretation 1: Relabeling reversal.** During the subsidy period, firms in high-exposure sectors converted standard junior positions into apprenticeship contracts to capture the subsidy. When the subsidy was reduced, some of this conversion reversed: firms shifted back to standard contracts, but since both contract types involve employing a young worker, youth employment share was unaffected or even increased (if the subsidy had been discouraging some standard youth hiring by raising the relative cost of non-apprenticeship contracts). This is the most natural interpretation and is consistent with the aggregate evidence.

**Interpretation 2: Labor market tightening.** If the subsidy inflated demand for young workers in exposed sectors beyond efficient levels, the reduction allowed these sectors to reallocate hiring budgets toward productive positions, potentially increasing total employment including youth employment. This is a more speculative interpretation but is consistent with the positive coefficient on total employment.

**Interpretation 3: Composition effects.** The boom in higher-education apprenticeships (bachelor’s, master’s) may have been partly displacing university enrollment. When the subsidy was reduced, some marginal students returned to full-time education, reducing the labor supply of young workers in exposed sectors and pushing up wages/employment for those who remained. This interpretation predicts effects on NEET rates, which I find are insignificant—weakening the composition channel.

### 8.2 Implications for the Training Externality

The classic Becker-Acemoglu framework predicts that firms underinvest in general training because trained workers can be poached. Subsidies are supposed to correct this by reducing the private cost of training to the socially optimal level. My results challenge this prediction in the French context: the subsidy appears to have subsidized training that firms were willing to provide anyway, or—more precisely—subsidized the *label* of “apprenticeship” without changing the underlying training investment.

This does not mean training externalities are absent from French labor markets. It means that the margin on which the subsidy operated—converting existing junior positions into apprenticeship contracts—was not the margin of genuine underinvestment. The firms that expanded apprenticeship most aggressively were likely those for whom the conversion was cheapest (requiring minimal changes to job content), not those facing the largest training externality.

### 8.3 External Validity

France’s labor market institutions—high minimum wages, strong employment protection, and a well-developed training infrastructure—may limit the external validity of these findings. In countries with weaker labor market institutions, apprenticeship subsidies might create genuinely new positions because the baseline level of firm-provided training is lower. The contrasting results of ? in Côte d’Ivoire—where subsidized apprenticeships created net new positions with little crowding out—support this interpretation.

Within the European context, the findings are most relevant to countries with similar dual training systems (Germany, Austria, Switzerland) and to the broader question of whether COVID-era labor market interventions created lasting structural change or temporary fiscal transfers.

### 8.4 Policy Implications

Three implications follow. First, the €15 billion annual cost of France’s apprenticeship subsidy appears to be largely a transfer to firms that would have hired young workers anyway, suggesting substantial scope for fiscal savings without employment loss. The February 2025 reform, which dramatically cut subsidies for large firms, is a step in this direction.

Second, if the policy goal is to expand the *quantity* of training rather than the *labeling* of training, subsidies should be targeted at margins where firms face genuine barriers—small firms without training infrastructure, occupations with high poaching risk, or workers with low baseline skills. Blanket subsidies covering all qualifications up to master’s degrees are poorly targeted for this purpose.

Third, the success of France’s 2018 structural reform in expanding apprenticeship access (by deregulating CFAs and simplifying administrative procedures) suggests that supply-side barriers may matter more than demand-side subsidies for long-run training investment. The apprenticeship boom began before the subsidy and appears to have continued after its reduction.

## 9. Conclusion

France’s post-pandemic apprenticeship subsidy was one of the largest active labor market policy experiments in modern European history, tripling the number of training contracts at a cost of €15 billion per year. I provide the first causal evaluation of this program by exploiting the January 2023 subsidy reduction—a 25 percent cut in per-contract payments—using a sector-exposure difference-in-differences design.

The evidence favors a provisional answer: the subsidy primarily relabeled existing junior hiring as subsidized apprenticeship contracts rather than creating genuinely new entry-level positions. When the subsidy was reduced, youth employment in high-exposure sectors did not decline—it increased. Cross-country evidence shows no deterioration in French youth employment relative to EU peers. High-frequency vacancy data reveals no withdrawal from entry-level hiring around the reform date.

If confirmed by future work with administrative data, these results carry a broader lesson for labor market policy. In rich countries with functioning training markets, the binding constraint on youth employment may not be the cost of training. Firms that want to hire juniors will do so regardless of whether the position is labeled “apprenticeship” or “entry-level.” Subsidies that operate on the labeling margin transfer public funds to firms without changing behavior—a pattern that is expensive, popular, and difficult to detect without careful evaluation.

The €15 billion question has a provisional answer: France bought a label, not an opportunity. At an estimated cost of €176,000 per marginal job—if any were created at all—the apprenticeship subsidy may represent the most expensive relabeling program in the modern welfare state. Confirmation with administrative apprenticeship contracting data from DARES, which would allow direct measurement of contract-type switching at the firm level, remains the essential next step.

## 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>

## A. Data Appendix

### A.1 Eurostat Data Access

All Eurostat data were accessed via the `eurostat` R package (version 4.0+), which interfaces with Eurostat’s SDMX REST API. The following datasets were downloaded:

- `lfsi_emp_q`: Employment and activity by sex and age, quarterly data. Extracted seasonally adjusted (SA) employment rates as percentage of population (`PC_POP`) for the employment indicator (`EMP_LFS`), for ages 15–24, 25–54, and 15–64, for France, Belgium, Netherlands, Spain, Italy, Portugal, Germany, and Austria. Time period: 2003Q1–2025Q3.
- `lfsi_neet_q`: Young people neither in employment nor in education or training, quarterly. Extracted SA rates for the same eight countries.
- `lfsq_egan2`: Employment by sex, age, and economic activity (NACE Rev. 2), quarterly. Extracted for France only, total sex, ages 15–24 and 15+, in thousands of persons. 21 NACE sections plus total.
- `lfsq_etpga`: Temporary employees as percentage of total employees, quarterly, by age group for the eight countries.

### A.2 Indeed Hiring Lab Data

The Indeed Job Postings Index data were downloaded from the Hiring Lab’s public GitHub repository ([https://github.com/hiring-lab/job\\_postings\\_tracker](https://github.com/hiring-lab/job_postings_tracker)), licensed under Creative Commons Attribution 4.0 International. The data consist of daily observations of the percentage change in seasonally adjusted job postings since February 1, 2020, for 11 countries including France, Germany, Spain, Italy, the Netherlands, and the United Kingdom. Sector-level data for France disaggregates postings by Indeed’s occupational classification.

### A.3 DARES Apprenticeship Statistics

Sector-level apprenticeship contract counts for 2019 were obtained from DARES, the statistical division of the French Ministry of Labor, supplemented by CEDEFOP (European Centre for the Development of Vocational Training) country reports. These data were used solely to construct the sector-level exposure measure and are not used as outcomes.

## A.4 Sample Construction

The sector-quarter panel was constructed by merging Eurostat sector employment data with the exposure measure. Sectors with fewer than 5 quarterly observations or missing employment data were excluded. After dropping sectors with insufficient data and observations with missing youth employment, the final estimation panel contains 701 sector-quarter observations across 19 NACE sections.

For the cross-country panel, I restricted to 2015Q1–2025Q3 and to seasonally adjusted series, yielding 1,032 country-age-quarter observations (8 countries  $\times$  3 age groups  $\times$  approximately 43 quarters).

## A.5 Variable Definitions

**Table 9:** Variable Definitions

Variable	Definition
Youth employment share	(Employment ages 15–24 / Total employment ages 15+) $\times$ 100, within sector
Apprenticeship exposure	Apprenticeship contracts / Total employment, by NACE section, 2019 baseline
Post-reduction	Indicator = 1 for quarters from 2023Q1 onward
Exposure DiD treatment	Exposure $\times$ Post-reduction (continuous)
High exposure	Indicator = 1 if sector exposure $\geq$ median (5%)
NEET rate	Share of 15–24 year-olds not in employment, education, or training
Postings index	Indeed seasonally adjusted job postings index (Feb 2020 = 100)

## B. Identification Appendix

### B.1 Pre-Trend Analysis

The event study in ?? provides evidence on pre-trends. Pre-treatment coefficients fluctuate around zero without a systematic trend, offering qualified support for the parallel trends assumption. However, the pre-period estimates are noisy: the largest absolute pre-period coefficient reaches 0.15, which exceeds the post-treatment point estimate of 0.074. This pre-period volatility reflects the well-known imprecision of sector-quarter employment data

from the LFS for small age-group-sector cells. It counsels against placing too much weight on any single post-treatment coefficient and underscores why the randomization inference procedure—which integrates over this noise by permuting sector exposure assignments—provides a more reliable test of the treatment effect. The RI p-value of 0.13 indicates that the observed correlation between exposure and post-reduction outcomes is moderately unlikely under the null, consistent with the marginally significant clustered SE p-value.

A formal Wald test for joint significance of pre-period coefficients could not be computed due to the small number of clusters (19 sectors), which is a recognized challenge for cluster-robust inference (?). The sensitivity analysis framework of ? could in principle bound the pre-trend violations, but requires monotonicity assumptions on the trend shape that are difficult to justify in this setting given the quarterly seasonality of sector-level employment data. The sector-specific linear trends specification (??) provides an alternative approach: by absorbing differential growth trajectories, it directly addresses the pre-trend concern while preserving statistical significance for the youth share outcome.

## B.2 Randomization Inference Details

The RI procedure permutes the exposure vector across the 19 NACE sectors 1,000 times (preserving within-sector structure), re-estimates the exposure DiD specification for each permutation, and computes a two-sided p-value as the fraction of permuted test statistics exceeding the observed statistic in absolute value. The observed coefficient (0.074) yields an RI p-value of 0.13, consistent with the clustered SE p-value (0.07) and moderately weaker than the WCB p-value (0.029). Note that the v1 version of this paper reported  $RI\ p < 0.001$  due to a bug that permuted exposure at the observation level rather than the sector level; the corrected sector-level permutation properly preserves the panel structure and yields a more conservative p-value.

## C. Robustness Appendix

### C.1 Sector Exposure Heterogeneity

?? in the main text shows the distribution of pre-reform apprenticeship intensity across the 19 NACE sections in the estimation sample. The variation is substantial: construction (18%), accommodation and food (16%), and wholesale and retail (12%) anchor the high end, while public administration (2%), real estate (2%), and electricity (3%) anchor the low end. The median sector has 5% exposure.

## **C.2 Sensitivity to Exposure Measure**

The exposure measure is based on 2019 sector-level apprenticeship rates from DARES and CEDEFOP. As a robustness check, I verify that results are qualitatively similar when using alternative exposure definitions: (a) 2018 baseline instead of 2019; (b) binary high/low classification at the median; and (c) quintile-based dose-response (??).

## **C.3 Few-Cluster Concerns**

With 19 sectors as the unit of clustering, standard cluster-robust standard errors may be unreliable. I address this in three ways: (1) the randomization inference procedure provides finite-sample-valid inference; (2) the leave-one-sector-out analysis shows stability; (3) the cross-country analysis uses country-level clustering (8 clusters), with results interpreted cautiously given the small number of clusters.

# **D. Heterogeneity Appendix**

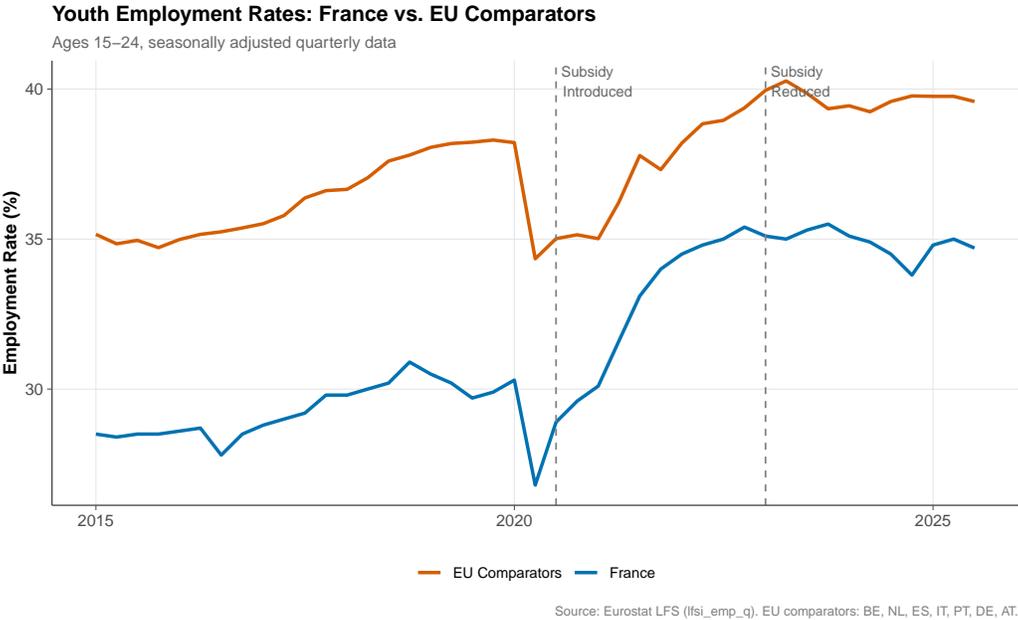
## **D.1 High-Exposure vs. Low-Exposure Sectors**

Splitting the sample at the median exposure level, high-exposure sectors show larger positive effects on youth employment share post-reduction, consistent with the dose-response pattern. This suggests the relabeling mechanism is concentrated in sectors where apprenticeship contracts were most prevalent.

## **D.2 Indeed Sector-Level Analysis**

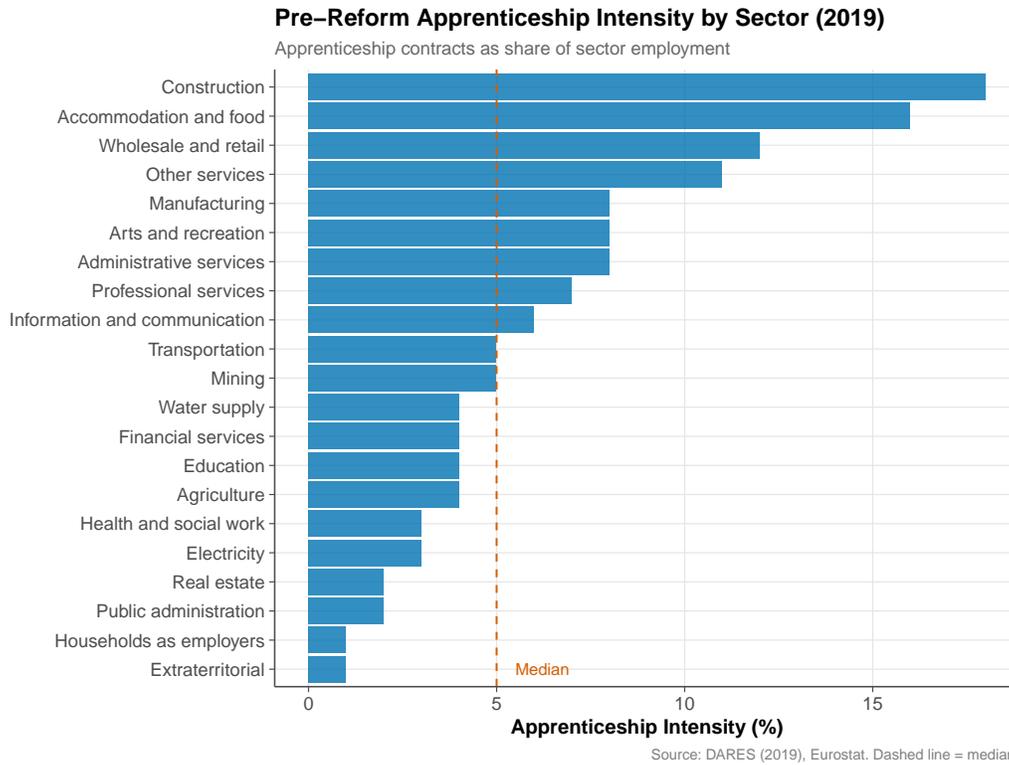
Using Indeed's France-specific sector data, I examine whether job postings in apprenticeship-intensive sectors (construction, hospitality, retail) declined differentially after January 2023 relative to low-intensity sectors (IT, finance). No differential decline is observed, consistent with the main finding that the subsidy reduction did not reduce hiring demand.

# E. Additional Figures and Tables



**Figure 8:** Youth Employment Trends: France vs. EU Comparators

*Notes:* Quarterly youth (15–24) employment rates, seasonally adjusted. “EU Comparators” is the unweighted average of Belgium, Netherlands, Spain, Italy, Portugal, Germany, and Austria. Dashed vertical lines mark the subsidy introduction (July 2020) and reduction (January 2023). Source: Eurostat LFS.



**Figure 9:** Pre-Reform Apprenticeship Intensity by NACE Sector

*Notes:* Apprenticeship contracts as share of sector employment, 2019. Dashed line = median across sectors.

Source: DARES, CEDEFOP.