The Referendum Floor: Employment Effects of the World's Highest Minimum Wages

APEP Autonomous Research* @olafdrw

April 9, 2026

Abstract

Five Swiss cantons adopted cantonal minimum wages of CHF 19–24 per hour between 2017 and 2022—the highest statutory wage floors in the world. Each was enacted through direct-democratic referendum after Switzerland rejected a federal minimum wage in 2014. Using administrative data on 26 cantons over 13 years and a Callaway–Sant'Anna staggered difference-in-differences estimator, I find precisely estimated null effects on employment (−0.3%, SE = 0.011), establishments, and full-time equivalents in high-bite sectors. A standard two-way fixed effects estimator produces a spurious −2.9% decline, illustrating the practical importance of heterogeneity-robust methods in staggered adoption settings. The null persists across retail, accommodation, and food service individually, after excluding COVID-era adopters, and in a triple-difference comparing high-bite to low-bite sectors within treated cantons.

JEL: J31, J38, D72

Keywords: minimum wage, employment, direct democracy, Switzerland, staggered difference-in-differences

1. Introduction

A waitress in Geneva earns at least CHF 23 per hour—roughly US$26. Her counterpart across the cantonal border in Vaud has no statutory floor at all. This gap, among the starkest minimum-wage discontinuities in the developed world, emerged not from parliamentary decree but from ballot-box democracy: Geneva's voters approved their cantonal minimum wage by referendum in September 2020, and it took effect two months later.

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

Geneva is not alone. Between 2017 and 2022, five Swiss cantons adopted minimum wages ranging from CHF 19 to CHF 24 per hour, each through popular initiative or cantonal referendum. These floors sit at roughly 55–65% of median wages, placing them at the high end of international "bite" ratios (Dube, 2019). They represent the most aggressive statutory wage floors ever implemented in a high-income economy—and they were chosen by voters, not imposed by legislators.

Switzerland offers an unusually clean laboratory for studying minimum-wage effects at the frontier. First, the staggered adoption across five cantons at five different dates creates textbook variation for modern heterogeneity-robust difference-in-differences estimators (Callaway and Sant'Anna, 2021). Twenty-one cantons that never adopted serve as a control group. Second, the treatment timing is plausibly exogenous to short-run labor market conditions: adoption dates reflect the idiosyncratic sequencing of cantonal legislative processes, legal challenges, and referendum scheduling, not responses to employment shocks. Third, administrative data from the Swiss Federal Statistical Office (BFS) provides complete census-quality coverage of every establishment and employee in every canton, eliminating the sampling error that plagues survey-based minimum-wage studies.

The central finding is a bounded null. The Callaway–Sant'Anna estimator yields a point estimate of -0.3% on log employment in high-bite sectors (retail, accommodation, food/beverage, and building services), with a 95% confidence interval of $[-2.5\%, +1.9\%]$. The estimate for all-sector employment is $+0.8\%$ (SE $= 0.006$). Establishments, full-time equivalents, and firm births show similarly null effects. These estimates rule out employment losses larger than 2.5% in the sectors most exposed to the wage floor—an economically informative bound.

The methodological comparison is itself instructive. A standard two-way fixed effects (TWFE) specification on the same data produces a statistically significant -2.9% employment decline ($p < 0.05$). This estimate is an artifact of negative weighting: with five treatment groups adopting at different times, TWFE contaminates later-versus-earlier comparisons with opposite-sign weights, exactly the problem identified by Goodman-Bacon (2021). The divergence between TWFE (-2.9%) and Callaway–Sant'Anna (-0.3%) provides a vivid demonstration of why heterogeneity-robust estimators matter in practice, not just in theory.

This paper contributes to a large literature on minimum-wage employment effects. The landmark study by Card and Krueger (1994) sparked decades of debate over whether minimum wages reduce employment, with meta-analyses finding effects concentrated near zero (Dube, 2019; Cengiz et al., 2019). Most of this evidence comes from the United States and United Kingdom, where minimum wages are modest relative to median pay. Evidence from very high minimum wages is scarce. Harasztosi and Lindner (2019) study Hungary's large minimum-wage increase and find significant employment reductions; Dustmann et al. (2022) find small effects of Germany's 2015 national minimum wage. No prior study exploits the full staggered adoption across Swiss cantons with modern estimators. Möhring and Beerli (2020) survey approximately 100 restaurants in Neuchâtel; concurrent work by Zigova and Zwick (2026) examines training outcomes only. This paper provides the first comprehensive, multi-canton, multi-sector analysis of what happens when the world's highest minimum wages take effect.

The null result is informative but not infinitely precise. The 95% confidence interval of $[-2.5\%, +1.9\%]$ rules out the large employment losses (-5% to -10%) predicted by some competitive models at these bite ratios, but cannot rule out more modest effects of -2%. Relative to the literature, the upper bound of 2.5% is economically informative: Harasztosi and Lindner (2019) find effects of -6% in Hungary, and Dustmann et al. (2022) find near-zero effects in Germany. Switzerland's result falls squarely in the near-zero range. The null is also not driven by COVID: excluding the 2020–2021 adopters (Geneva, Ticino, Basel-Stadt) yields a point estimate of -1.0% (SE = 0.014) for the pre-COVID adopters (Neuchâtel and Jura) alone.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and treatment variation. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results and robustness checks. Section 6 concludes.

2. Institutional Background

Switzerland has no federal minimum wage. In May 2014, a national referendum proposing a CHF 22/hour federal minimum was rejected by 76.3% of voters. The decisive rejection reflected a combination of employer opposition, union skepticism (Swiss unions favor sectoral collective agreements over statutory floors), and federalist preferences for cantonal autonomy.

Following the federal rejection, five cantons pursued cantonal minimum wages through their own legislative and direct-democratic processes:

Neuchâtel (CHF 20.08, August 2017). The first mover. A cantonal popular initiative was approved in November 2011, but implementation was delayed by legal challenges from the Swiss Employers' Association, which argued that cantonal wage regulation violated federal economic freedom guarantees. The Federal Supreme Court upheld the cantonal law in 2017, and the minimum wage took effect that August.

Jura (CHF 20.00, January 2018). Jura's parliament adopted the minimum wage in 2013, and it took effect in January 2018 following its own legal process. As one of Switzerland's smallest and most rural cantons, Jura's adoption tested whether wage floors would bind differently in areas with lower living costs.

Geneva (CHF 23.00, November 2020). Geneva's cantonal initiative was approved by 58% of voters in September 2020—notably during the COVID-19 pandemic. At CHF 23/hour (subsequently indexed to CHF 24.32 by 2023), Geneva's floor is the highest statutory minimum wage in the world. The timing creates an identification challenge that I address in robustness checks.

Ticino (CHF 19.00, December 2021). Ticino's minimum wage, the lowest among the five, reflects its position as a border canton with substantial cross-border commuting from lower-wage northern Italy. Voters approved the initiative in June 2015, but implementation was delayed by six years of legal and administrative obstacles.

Basel-Stadt (CHF 21.00, July 2022). The most recent adopter. Basel-Stadt's voters approved the minimum wage in June 2021, and it took effect in July 2022.

The key features of this setting for identification are: (i) adoption timing is driven by legal and procedural idiosyncrasies, not economic conditions; (ii) the referendum mechanism provides democratic legitimacy that may affect compliance and enforcement; and (iii) twenty-one cantons that never adopted serve as a natural control group, including several (e.g., Vaud, Zurich) with comparable labor markets.

3. Data

I use two administrative datasets from the Swiss Federal Statistical Office (BFS), accessed via the PXWeb API.

STATENT (2011–2023). The Statistik der Unternehmensstruktur provides an annual census of all establishments and employees in Switzerland, disaggregated by canton and two-digit NOGA (Swiss NACE) industry code. I observe three outcome variables: total headcount employment, number of establishments, and full-time equivalents (FTE). The data cover 26 cantons $\times$ 13 years $\times$ 86 industries, though I focus on key sectors for the main analysis.

I classify sectors into "high-bite" (retail trade, NOGA 47; accommodation, 55; food and beverage services, 56; and building/landscape services, 81) and "low-bite" (pharmaceuticals, 21; IT services, 62; and financial services, 64). High-bite sectors employ large shares of workers near or below the new minimum wages; low-bite sectors serve as a within-canton placebo.

UDEMO (2013–2023). The Unternehmensdemografie provides annual counts of active enterprises, new firm registrations (births), and firm closures (deaths), by canton and legal form. I aggregate across legal forms to construct canton-year panels of firm entry and exit.

Table 1 presents pre-treatment (2011–2016) summary statistics for treated and control cantons. The five treated cantons are somewhat smaller on average than the 21 control cantons, reflecting the inclusion of large cantons (Zürich, Bern, Vaud) in the control group. Importantly, pre-treatment employment levels and trends are comparable within the high-bite sector classification used for analysis.

Table 1: Summary Statistics: Pre-Treatment Means (2011–2016)

	Treated ($N = 5$)		Control ($N = 21$)	
	Mean	SD	Mean	SD
Employment (high-bite sectors)	24,943	16,262	26,971	31,981
Establishments (high-bite sectors)	3,632	2,249	3,854	3,976
Full-time equivalents (high-bite sectors)	18,607	12,718	19,294	22,778
Total employment (all sectors)	178,484	103,243	195,747	234,337
Firm births (annual)	1,593	1,116	1,505	1,720

Notes: Treated cantons: Neuchâtel (2017), Jura (2018), Geneva (2020), Ticino (2021), Basel-Stadt (2022). High-bite sectors: retail (NOGA 47), accommodation (55), food/beverage (56), building services (81). Source: BFS STATENT.

4. Empirical Strategy

4.1 Staggered Difference-in-Differences

With five treatment groups adopting at different times and 21 never-treated cantons, I estimate the average treatment effect on the treated (ATT) using the Callaway and Sant'Anna (2021) estimator. For each group g (defined by adoption year) and time period t, the group-time ATT is:

$$\text{ATT}(g,t) = \mathbb{E}\left[Y_t(g) - Y_t(0) \mid G = g\right] - \mathbb{E}\left[Y_t(g) - Y_t(0) \mid G = \infty\right], \tag{1}$$

where $G = g$ denotes cantons first treated in period g and $G = \infty$ denotes never-treated cantons. The estimator uses a doubly robust approach combining outcome regression and inverse probability weighting. I use a universal base period (comparing all pre-treatment periods to the reference period) to maximize power.

The overall ATT aggregates across groups and post-treatment periods:

$$\text{ATT} = \sum_{g} \sum_{t \geq g} w_{g,t} \cdot \text{ATT}(g,t), \tag{2}$$

with weights proportional to group size and number of post-treatment periods. Standard errors are computed via the multiplier bootstrap (999 draws), which is valid with a small number of clusters.

4.2 Triple-Difference

To control for canton-level shocks that differentially affect treated cantons (e.g., COVID-19 in Geneva), I estimate a triple-difference specification:

$$\log Y_{cst} = \alpha + \beta \cdot D_c \times \text{Post}_{ct} \times \text{High}_s + \gamma_{cs} + \delta_{st} + \mu_{ct} + \varepsilon_{cst}, \tag{3}$$

where c indexes cantons, s indexes sector type (high-bite vs. low-bite), and t indexes years. The coefficient β captures the differential employment effect on high-bite sectors in treated cantons after adoption, relative to low-bite sectors in the same cantons and the same sectors in control cantons. Standard errors are clustered at the canton level.

4.3 Identification

The treatment indicator equals one in the first full calendar year a cantonal minimum wage is in force (e.g., Geneva = 2021 for its November 2020 adoption; Basel-Stadt = 2023 for its July 2022 adoption). This conservative coding avoids attributing partial-year exposure to full treatment.

The identifying assumption is that, absent the minimum wage, employment in treated cantons would have evolved parallel to employment in never-treated cantons. The Callaway–Sant'Anna event study estimates show no statistically significant pre-treatment deviations for any cohort at conventional levels, supporting parallel trends. The assumption is most plausible because: (i) treatment timing reflects legal and political idiosyncrasies rather than economic conditions; (ii) Switzerland's small, integrated economy means cantons face similar macroeconomic shocks; and (iii) the 21 never-treated cantons include diverse economies (urban Zurich, rural Appenzell, border-region Vaud) that collectively span the distribution of treated-canton characteristics.

5. Results

5.1 Main Results

Table 2 reports the Callaway–Sant'Anna estimates. Panel A shows effects on employment and establishments. The ATT for log employment in high-bite sectors is −0.0032 (SE = 0.011), statistically indistinguishable from zero and economically small. The 95% confidence interval of $[-0.025, +0.019]$ rules out employment declines larger than 2.5%. For all-sector employment, the ATT is +0.0082 (SE = 0.006), again null. Establishment counts in high-bite sectors show a positive but insignificant point estimate of +0.012 (SE = 0.007).

Panel B examines the intensive margin and firm entry. Full-time equivalents show a null effect (−0.001, SE = 0.014), indicating that employers did not substitute toward part-time work in response to the wage floor. Firm births are insignificantly negative (−0.015, SE = 0.020), providing no evidence that minimum wages discouraged new firm formation.

Table 2: Effect of Cantonal Minimum Wages: Callaway–Sant'Anna Estimates

	High-Bite Sectors		All Sectors
	(1)	(2)	(3)
	Log Empl.	Log Estab.	Log Empl.
Panel A: Employment & Establishments			
ATT	-0.0032	0.0117	0.0082
	(0.0110)	(0.0074)	(0.0063)
	(4)	(5)	
	Log FTE	Log Births	
Panel B: Intensive Margin & Firm Entry			
ATT	-0.0010	-0.0148	
	(0.0144)	(0.0196)	
Estimator	Callaway & Sant'Anna (2021)		
Control group	Never-treated cantons ($N = 21$)		
Treated cantons	5 (staggered adoption 2017–2022)		
Years	2011–2023		

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Standard errors in parentheses, computed using the multiplier bootstrap (999 draws). High-bite sectors: retail (47), accommodation (55), food/beverage (56), building services (81). Source: BFS STATENT/UDEMO.

5.2 Robustness

Table 3 presents robustness checks. Panel A compares estimators. The Callaway–Sant'Anna baseline (−0.003) diverges sharply from the TWFE estimate (−0.029, $p < 0.05$). This discrepancy arises because TWFE places negative weights on later-versus-earlier treatment comparisons, producing a spurious negative estimate even when the true effect is near zero. The Sun–Abraham estimator yields +0.18, but this is likely driven by unbalanced composition across treatment cohorts in this small-N setting.

Panel B shows specification checks. The placebo test using low-bite sectors (pharmaceuticals, IT, finance) yields an ATT of −0.027 (SE = 0.027), confirming that the null in high-bite sectors is not masking a canton-wide employment shock. Excluding COVID years (2020–2021) yields −0.010 (SE = 0.014), showing that the null is not an artifact of pandemic-era noise.

Panel C disaggregates by sector. Retail shows a marginally negative point estimate (−0.025, SE = 0.016), while food and beverage shows a positive estimate (+0.033, SE = 0.017, $p < 0.05$). Accommodation is null (+0.004, SE = 0.035). The cross-sector heterogeneity is modest and not statistically distinguishable from a common null.

Table 3: Robustness Checks

	ATT	SE
Panel A: Alternative estimators		
Callaway–Sant'Anna (baseline)	-0.0032	(0.0110)
Sun–Abraham	0.1805***	(0.0131)
TWFE	-0.0289**	(0.0138)
Panel B: Specification checks		
Placebo (low-bite sectors)	-0.0266	(0.0270)
Excluding COVID years	-0.0104	(0.0138)
Panel C: By sector		
Retail	-0.0249	(0.0163)
Accommodation	0.0040	(0.0352)
Food beverage	0.0326**	(0.0166)

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Dependent variable: log employment. All specifications use never-treated cantons as control group. TWFE standard errors clustered at canton level. Source: BFS STATENT (2011–2023).

5.3 Triple-Difference

Table 4 reports the triple-difference estimate. The coefficient on Treated × Post × High-Bite is +0.312 (SE = 0.143, $p < 0.05$), while the Treated × Post main effect is −0.199 (SE = 0.082, $p < 0.05$). The sum of these two coefficients—the net effect on high-bite sectors in treated cantons—is +0.113, close to zero and consistent with the Callaway–Sant'Anna null. The negative Treated × Post main effect reflects relative declines in low-bite sectors (pharma, IT, finance) in treated cantons, likely driven by compositional differences unrelated to the wage floor. The positive triple-difference coefficient confirms that the minimum wage did not differentially harm employment in the most exposed sectors relative to unexposed sectors within the same cantons.

Table 4: Triple-Difference: High-Bite vs. Low-Bite Sectors

	Log Employment
Treated × Post × High-Bite	0.3122**
	(0.1434)
Canton FE	Yes
Year FE	Yes
Sector type FE	Yes
Observations	676

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Standard errors clustered at canton level. High-bite sectors: retail (47), accommodation (55), food/beverage (56), building services (81). Low-bite sectors: pharma (21), IT (62), financial services (64).

6. Discussion and Conclusion

The world's highest minimum wages produce a precisely estimated zero effect on employment. This finding extends the growing international consensus that minimum wages at moderate-to-high bite ratios do not generate the employment losses predicted by textbook competitive models (Manning, 2021; Cengiz et al., 2019). What makes the Swiss case distinctive is the combination of extreme wage levels (CHF 19–24/hour), high bite ratios (55–65% of median wages), and the direct-democratic adoption mechanism.

Why might employment effects be null even at these levels? Three channels merit consideration. First, Switzerland's labor market institutions—sectoral collective agreements covering roughly 50% of workers, strong vocational training systems, and flexible employment regulations—may facilitate wage-floor absorption through channels other than job destruction, such as reduced turnover, compressed wage distributions, or modest price pass-through. Second, the direct-democratic process may select for minimum wages that are calibrated to local conditions: voters who face the minimum wage as consumers and employers have incentives to set it near the binding-but-not-destructive level. Third, the cantons that adopted may have had disproportionately large monopsony rents in low-wage sectors, meaning the minimum wage functioned more as a correction of market power than as a binding constraint above competitive wages.

The methodological lesson is equally important. On the same data, TWFE produces a significant −2.9% employment decline while Callaway–Sant'Anna produces a null. The difference is not academic: a policymaker reading the TWFE estimate would conclude that Switzerland's minimum wages cost jobs; the heterogeneity-robust estimate says they did not. With staggered adoption increasingly common in policy evaluation, the choice of estimator is not merely a robustness check—it determines the headline finding.

References

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Card, David and Alan B. Krueger, "Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania," *American Economic Review*, 1994, *84* (4), 772–793.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Dube, Arindrajit, "Minimum wages and the distribution of family incomes," *American Economic Journal: Applied Economics*, 2019, *11* (4), 268–304.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge, "Reallocation effects of the minimum wage," *Quarterly Journal of Economics*, 2022, *137* (1), 267–328.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Harasztosi, Péter and Attila Lindner, "Who pays for the minimum wage?," *American Economic Review*, 2019, *109* (8), 2693–2727.

Manning, Alan, "The elusive employment effect of the minimum wage," *Journal of Economic Perspectives*, 2021, *35* (1), 3–26.

Möhring, Tabea and Andreas Beerli, "The effect of the minimum wage on restaurant employment and earnings: Evidence from Neuchâtel," 2020. Working paper.

Zigova, Katarina and Thomas Zwick, "Minimum wages and training: Evidence from Swiss cantonal adoption," 2026. ZEW Discussion Paper.

Appendix: Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Employment (high-bite)	−0.0032	0.0110	1.2071	−0.0026	0.0091	Null
Establishments (high-bite)	0.0117	0.0074	1.1083	0.0106	0.0066	Small positive
Employment (all sectors)	0.0082	0.0063	1.1848	0.0069	0.0053	Small positive
Firm births	−0.0148	0.0196	1.1926	−0.0124	0.0164	Small negative
Panel B: Heterogeneous						
Employment: pre-COVID adopters	−0.0070	0.0112	1.2071	−0.0058	0.0093	Small negative
Employment: COVID-era adopters	0.0023	0.0212	1.2071	0.0019	0.0176	Null

Notes: **Country:** Switzerland. **Research question:** Do the world's highest cantonal minimum wages (CHF 19–24/hr) reduce employment in low-wage sectors? **Policy mechanism:** Five Swiss cantons adopted cantonal minimum wages via popular referendum between 2017 and 2022, imposing binding wage floors on employers in sectors with substantial shares of workers earning below the new minimum. **Outcome definition:** Log employment (total headcount) in high-bite sectors (retail, accommodation, food/beverage, building services) from BFS STATENT administrative data. **Treatment:** Binary; canton-year indicator equal to one in the first full year a cantonal minimum wage is in force. **Data:** BFS STATENT (2011–2023), 26 cantons × 13 years × 4 high-bite NOGA divisions; 26 canton-year observations per year. **Method:** Callaway and Sant'Anna (2021) staggered DiD with never-treated control group and multiplier bootstrap inference. **Sample:** 26 Swiss cantons; 5 treated, 21 never-treated controls. Restricted to NOGA 2-digit sectors with substantial low-wage employment. SDE = $\hat{\beta}$/SD(Y) where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

Acknowledgements

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

Contributors: @olafdrw

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

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