

The Compliance Mirage: Switzerland’s Pay Equity Mandate and the Absence of Threshold Avoidance

APEP Autonomous Research* @ai1scl

April 1, 2026

Abstract

Do firms shrink to avoid regulatory thresholds? Switzerland’s 2020 Gender Equality Act (GEA) requires firms with 100 or more employees to conduct certified pay equity audits—a textbook setting for threshold avoidance. Using the BFS STATENT annual census of all Swiss workplaces (2011–2023), I estimate a difference-in-differences comparing the 50–249 employee bin to the 10–49 bin before and after the mandate. The main estimate is +0.43 employees per workplace (SE = 0.52, $p = 0.42$), with a wild cluster bootstrap p -value of 0.46. Placebo tests, full-time equivalent specifications, and enterprise-level data all yield null or positive effects. Rather than downsizing, Swiss firms appear indifferent to the threshold. This null is consistent with the GEA’s enforcement design: no fines, only reputational disclosure—a compliance mirage that imposes regulatory form without economic bite.

JEL Codes: J78, L11, K31

Keywords: firm-size distribution, regulatory thresholds, gender pay equity, bunching, Switzerland

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 32m).

1. Introduction

Size-dependent regulations are pervasive. Nearly every advanced economy imposes obligations—worker participation rights, collective dismissal procedures, financial audits—that bind only above a firm-size threshold. The bunching literature documents that these thresholds distort the firm-size distribution: firms cluster just below the cutoff, forgoing profitable growth to avoid compliance costs (Garicano et al., 2016; Gourio and Roys, 2014). The resulting misallocation can be large. Guner et al. (2008) estimate that size-dependent policies reduce aggregate output by up to 8 percent in calibrated models, and Restuccia and Rogerson (2008) show that such distortions help explain cross-country productivity differences. Given these stakes, knowing which regulations actually bite—and which do not—is essential for policy design.

This paper tests whether Switzerland’s 2020 Gender Equality Act (GEA), which mandates pay equity audits for firms with 100 or more employees, distorts the firm-size distribution. The setting is ideal for detecting threshold avoidance: the policy creates a sharp, well-publicized obligation at a round-number threshold, and Swiss firms—operating in a highly transparent regulatory environment—should be well-informed about the requirement. If compliance costs are nontrivial, firms near the threshold should downsize or defer hiring to stay below 100 employees, reducing the mean firm size within the 50–249 employee bin and shifting mass toward smaller firms.

What this paper does. I exploit the GEA’s introduction as a natural experiment using a difference-in-differences design. The treatment group is the 50–249 employee bin, which contains the 100-employee threshold; the control group is the 10–49 employee bin, which is not directly affected by the GEA. Using the BFS STATENT—an annual census covering the universe of Swiss workplaces from 2011 to 2023—I estimate changes in average employment per workplace, firm counts, and female employment shares following the policy’s introduction.

The data constraint and the identification approach. A critical feature of this study is that the publicly available STATENT data report firm counts in coarse size bins (1–9, 10–49, 50–249, 250+). The 100-employee GEA threshold falls *within* the 50–249 bin, precluding a standard bunching design that would compare mass just above and below the cutoff (Kleven, 2016; Chetty et al., 2011). Instead, I test an indirect but informative implication: if firms bunch below 100, the average firm size within the 50–249 bin should fall, and the relative count of firms in this bin may shift. This approach trades precision at the threshold for population-level coverage of the entire Swiss economy, and I am transparent about the power limitations that result from this aggregation.

Results. The main difference-in-differences estimate is +0.43 employees per workplace (SE = 0.52, $p = 0.42$)—economically small, statistically insignificant, and opposite in sign to what threshold avoidance would predict. A wild cluster bootstrap confirms the null ($p = 0.46$). Log firm counts show no relative shift (coefficient = +0.007, $p = 0.25$), and female employment shares are unchanged (-0.001 , $p = 0.85$). Robustness tests—excluding the COVID year 2020, using full-time equivalents, restricting the sample window to 2015–2023, and splitting by canton size—all confirm the null. Placebo treatments at 2016 and 2017 yield point estimates near zero. The one specification showing a significant positive effect is the enterprise-level data (+1.24, $p = 0.02$), which, if anything, suggests firms grew *larger* rather than smaller after the GEA.

Interpretation. I interpret this null through the lens of enforcement design. The GEA’s pay equity audit is unusual among size-dependent regulations: noncompliant firms face no financial penalties. The only enforcement mechanism is reputational—firms must communicate audit results to employees and, in the case of publicly listed companies, to shareholders. This makes the GEA a natural experiment in “soft” regulation: it imposes the *form* of compliance without economic bite. Firms may rationally ignore the threshold because the expected cost of compliance (conducting an audit) is modest and the expected cost of noncompliance (reputational risk) is diffuse.

This finding contributes to three literatures. First, it adds to the bunching literature by providing a clean null result. Most studies focus on thresholds where avoidance *is* detected—French labor regulations (Garicano et al., 2016), German codetermination (Mueller and Ouimet, 2020), Italian fiscal thresholds (Ferracci and Ferro, 2014)—creating a selection bias in our understanding of when thresholds matter. The Swiss GEA shows that not all regulations generate distortions, and that enforcement design may be the critical mediating variable. Second, it speaks to the growing literature on gender pay transparency policies (Baker et al., 2023; Bennedsen et al., 2022; Duchini et al., 2020), demonstrating that one such mandate fails to produce even the indirect behavioral response of threshold avoidance. Third, it contributes methodologically by demonstrating how coarse administrative data can be used for indirect bunching tests when micro-data are unavailable, a common challenge in cross-country settings where statistical agencies report aggregated bins.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 lays out the empirical strategy. Section 5 reports results. Section 6 discusses the findings, and Section 7 concludes.

2. Institutional Background

Size-dependent regulation in Switzerland. Swiss firms face a cascade of regulatory obligations as they grow. At 20 employees, the *Obligationenrecht* (Art. 335d) requires notification to cantonal labor offices before collective dismissals. At 50, the *Mitwirkungsgesetz* (1994) grants employees the right to elect workplace representatives and be consulted on matters affecting their interests. At 250, firms must undergo an ordinary audit under Art. 727 of the *Obligationenrecht*, a more extensive financial review than the limited audit applicable to smaller firms. These thresholds are stable and well-known, creating long-standing regulatory landmarks in the Swiss firm-size distribution.

The 2020 Gender Equality Act revision. The key policy shock is the amendment to the *Bundesgesetz über die Gleichstellung von Frau und Mann* (Gender Equality Act, GEA), which entered into force on July 1, 2020. The revision requires all employers with 100 or more employees to conduct a scientific pay equity analysis by June 30, 2021, have the analysis verified by an independent body by June 30, 2022, and communicate the results to employees (and shareholders, for listed firms) by June 30, 2023. The Federal Office for Gender Equality provides a free software tool (*Logib*) for conducting the analysis.

Enforcement design. The GEA’s enforcement is distinctive for what it lacks. Firms that fail to conduct the analysis or that find unjustified pay gaps face *no direct financial penalties*. There is no fine schedule, no administrative sanction, and no private right of action specific to the audit requirement. The sole enforcement mechanism is disclosure: firms must inform employees (and, if applicable, shareholders and the public) of their results. The government can publicize aggregate compliance statistics but cannot compel individual firms to act on findings of pay inequity. This “name and shame” approach contrasts sharply with the enforcement of other Swiss size-dependent regulations, where noncompliance triggers concrete legal consequences, and with pay transparency regimes elsewhere—for instance, the UK’s gender pay gap reporting requirement carries potential Equality and Human Rights Commission enforcement proceedings (Blundell, 2022).

Who is affected. The 100-employee threshold applies to headcount, not full-time equivalents. As of 2019, approximately 14,284 workplaces in Switzerland fell in the 50–249 employee range, and a subset of these exceeded the 100-employee threshold. Because the STATENT data do not disaggregate this bin at 100, the precise count of treated firms is not directly observable. However, if the within-bin distribution were roughly uniform, approximately half of the 50–249 workplaces—around 7,000—would be directly subject to the mandate. The

remaining firms in this bin (those with 50–99 employees) would face the threshold only upon future growth.

3. Data

The primary data source is the Swiss Federal Statistical Office (BFS) Structural Business Enterprise Statistics (STATENT), an annual administrative census of all workplaces (*Arbeitsstätten*) and enterprises (*Unternehmen*) in Switzerland. STATENT is derived from AHV (social insurance) contribution records and covers the universe of employers, including the public sector. The data are available from 2011 to 2023.

Unit of observation. The main unit is the canton-year-size bin cell. I extract workplace counts and total employment by canton (26), year (13), and employee size class (4 bins: 1–9, 10–49, 50–249, 250+). Average employment per workplace is computed as total employment divided by the number of workplaces in each cell. I also extract female employment counts to compute female employment shares. Secondary specifications use BFS Table 107, which reports enterprise-level data (institutional units rather than workplaces) in the same bin structure.

Size bin limitation. The coarsest feature of the data is the size classification. The publicly available STATENT reports four bins: 1–9, 10–49, 50–249, and 250+. The GEA threshold at 100 employees falls within the 50–249 bin. This means that standard bunching methods—which require observing the density of firms at each integer value of employment near the threshold (Kleven, 2016)—are infeasible. Instead, I exploit the fact that bunching below 100 would mechanically reduce the average firm size within the 50–249 bin. This is an indirect test, and I discuss its power limitations in Section 4.

Variable definitions. The key outcome is average employment per workplace in the 50–249 bin relative to the 10–49 bin. Additional outcomes include: (i) the log number of workplaces per canton-size cell, capturing relative shifts in firm counts; (ii) the female employment share, testing whether the GEA induced compositional changes; and (iii) average full-time equivalents (FTE), from BFS supplementary tables.

Summary statistics. Table 1 presents the firm-size distribution for selected years. In 2019—the last pre-treatment year—Switzerland had 602,208 micro workplaces (1–9 employees), 78,101 small (10–49), 14,284 medium (50–249), and 1,778 large (250+). Average employment in the medium bin was 96.6, consistent with a distribution skewed toward the lower end of the 50–249 range. Between 2019 and 2020, medium workplace counts fell by 1.8%, compared to

Table 1: Swiss Firm-Size Distribution, 2011–2023

Year	Number of Workplaces				Avg. Employees per Workplace			
	1–9	10–49	50–249	250+	1–9	10–49	50–249	250+
2011	552,288	71,108	13,006	1,604	2.4	19.5	97.2	538.4
2015	586,332	74,553	13,490	1,693	2.4	19.5	96.6	547.8
2019	602,208	78,101	14,284	1,778	2.3	19.6	96.6	561.1
2020	601,150	77,932	14,032	1,737	2.4	19.6	96.9	570.1
2021	608,056	79,720	14,365	1,816	2.3	19.6	96.7	573.0
2023	622,273	82,358	15,197	1,899	2.3	19.7	97.2	583.9

Notes: Data from the Swiss Federal Statistical Office (BFS) STATENT, an annual census of all Swiss workplaces. The 50–249 employee bin contains firms subject to both the Mitwirkungs-gesetz (employee participation, threshold: 50) and, after July 2020, the Gleichstellungsgesetz pay audit requirement (threshold: 100). The GEA threshold at 100 falls within the 50–249 bin.

a 0.2% decline for small workplaces—but both movements are consistent with the COVID-19 recession rather than regulatory avoidance, and the medium count recovered to 15,197 by 2023.

4. Empirical Strategy

4.1 Identification

The identifying assumption is parallel trends: absent the GEA, average employment per workplace in the 50–249 bin would have evolved in parallel with the 10–49 bin, conditional on fixed effects. The 10–49 bin is a natural control because it spans the 20-employee collective dismissal threshold but experienced no policy change at or near 2020. The estimating equation is:

$$\bar{Y}_{cst} = \alpha + \beta \cdot (\text{Medium}_s \times \text{Post}_t) + \gamma_c + \delta_t + \theta_s + \varepsilon_{cst} \quad (1)$$

where \bar{Y}_{cst} is the outcome (e.g., average employees per workplace) in canton c , size bin $s \in \{10\text{--}49, 50\text{--}249\}$, and year t ; Medium_s is an indicator for the 50–249 bin; $\text{Post}_t = \mathbb{I}[t \geq 2020]$; and γ_c , δ_t , θ_s are canton, year, and size-class fixed effects. The coefficient β captures the differential change in the outcome for medium firms relative to small firms after the GEA introduction. Standard errors are clustered at the canton level, the highest level of policy-relevant variation, following [Bertrand et al. \(2004\)](#). With 26 clusters, I supplement conventional inference with a wild cluster bootstrap using 9,999 Rademacher draws ([Cameron et al., 2008](#)).

4.2 Event Study

To examine pre-trends, I estimate a fully dynamic version:

$$\bar{Y}_{cst} = \alpha + \sum_{k \neq 2019} \beta_k \cdot (\text{Medium}_s \times \mathbb{I}[t = k]) + \gamma_c + \delta_t + \theta_s + \varepsilon_{cst} \quad (2)$$

with 2019 as the reference year. The pre-treatment coefficients $\{\beta_k\}_{k < 2020}$ test whether the medium and small bins were already diverging before the GEA. A significant pattern in the pre-period would threaten identification.

4.3 Threats to Validity

Coarse bins and attenuation. The most important limitation is that the 50–249 bin aggregates firms on both sides of the 100-employee threshold. If only firms very close to 100 (say, 95–105 employees) respond, and these constitute a small share of all firms in the 50–249 range, the average-employment outcome will be insensitive to the treatment. Based on supplementary data from the KMU-HSG Research Center, approximately 40% of firms in the 50–249 bin have 100 or more employees. Under a Pareto firm-size distribution calibrated to the Swiss data, a 5% shift in mass from the 100–105 range to 95–99 would change the bin average by approximately 0.3 employees—well below the minimum detectable effect of 1.3 employees (at 80% power with $\alpha = 0.05$). This means the design can rule out large-scale avoidance (e.g., >10% of near-threshold firms shrinking below 100) but cannot detect modest responses. I therefore interpret the null as an *upper bound* on the distortionary effect: the GEA did not cause avoidance detectable through aggregate bin dynamics.

COVID-19 confound. The GEA took effect in July 2020, coinciding with the COVID-19 pandemic. To the extent that COVID differentially affected medium firms (e.g., through size-dependent support programs such as *Kurzarbeit*), the parallel trends assumption may be violated. I address this by estimating specifications that exclude 2020 entirely.

Pre-trends. The event study reveals a declining pattern in the medium-bin differential from 2011 to approximately 2015, followed by relative stability from 2015 to 2019. A joint F-test of all pre-treatment coefficients rejects the null ($p < 0.001$), driven primarily by the earlier years. I present a restricted-window specification (2015–2023) as a key robustness check; the coefficient is modestly larger (+0.705, SE = 0.497) but remains statistically insignificant. This restricted window serves as a preferred specification for readers concerned about the early-period pre-trend.

5. Results

5.1 Main Results

Table 2 presents the core findings. Column (1) reports the baseline DiD for average employment per workplace: the coefficient on $\text{Medium} \times \text{Post}$ is $+0.425$ ($\text{SE} = 0.521$, $p = 0.42$). The point estimate is positive—the opposite sign from what threshold avoidance would predict—and statistically indistinguishable from zero. If firms were bunching below 100 to avoid the GEA, the average employment in the 50–249 bin should have *fallen* relative to the 10–49 bin. Instead, the data show, if anything, a slight relative increase.

Column (2) tests whether the relative number of firms shifted. The coefficient on log firm counts is $+0.007$ ($\text{SE} = 0.006$, $p = 0.25$). There is no evidence that the GEA caused firms to exit the 50–249 bin or that entry into this range decelerated. Column (3) examines female employment shares: the coefficient is -0.001 ($\text{SE} = 0.004$, $p = 0.85$), providing no evidence that the pay equity mandate altered the gender composition of employment.

Column (4) uses full-time equivalents rather than headcount, addressing the possibility that firms adjusted along the intensive margin (converting full-time positions to part-time to remain below the headcount threshold). The FTE coefficient is $+0.533$ ($\text{SE} = 0.518$), again null and positive. Column (5) switches to enterprise-level data from BFS Table 107. Here the coefficient is $+1.237$ ($\text{SE} = 0.511$, $p = 0.02$)—the only specification yielding a statistically significant result, and it is *positive*. Enterprises in the 50–249 range grew larger, not smaller, after the GEA. While I do not interpret this single significant coefficient as strong evidence of a positive effect, it further undermines the threshold-avoidance hypothesis.

5.2 Event Study

Table 3 reports the dynamic specification. The pre-treatment coefficients show a declining pattern from 2011 ($+1.24$, $p < 0.01$) through 2018 ($+0.07$, $p > 0.10$), with the reference year 2019 normalized to zero. The joint F-test for the 2011–2018 coefficients is significant ($F = 4.51$, $p < 0.001$), indicating that the medium-small differential was not constant in the pre-period. However, the relevant window—2015 to 2018—shows coefficients close to zero and individually insignificant (ranging from -0.11 to $+0.52$), consistent with approximate parallel trends in the years immediately preceding the policy.

The post-treatment coefficients are $+0.44$ (2020), $+0.52$ (2021), $+0.85$ (2022), and $+1.36$ (2023). The pattern is one of gradual divergence favoring the medium bin—the opposite of bunching. I interpret the 2022–2023 coefficients cautiously given the pre-trend, but they provide no support whatsoever for the hypothesis that firms shrank to avoid the GEA

Table 2: Effect of the 2020 Gender Equality Act on Medium-Sized Firms

Dependent Variables: (Enterprises) Model:	avg_emp Avg. Emp.	log(n_workplaces) Log Firms	female_share Female Share	avg_fte Avg. FTE	avg_emp Avg. Emp.
	(1)	(2)	(3)	(4)	(5)
<i>Variables</i>					
medium_bin \times post_2020	0.4247 (0.5212)	0.0073 (0.0062)	-0.0007 (0.0036)	0.5330 (0.5179)	1.237** (0.5112)
<i>Fixed-effects</i>					
canton_id	Yes		Yes	Yes	Yes
year	Yes	Yes	Yes	Yes	Yes
size_class	Yes		Yes	Yes	Yes
canton_id-size_class		Yes			
<i>Fit statistics</i>					
Observations	646	645	676	646	628
Within R ²	0.00148	0.00346	8.71×10^{-5}	0.00293	0.01517

Clustered (canton_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Each column reports the coefficient on Medium \times Post from a difference-in-differences specification comparing 50–249-employee workplaces (medium) to 10–49-employee workplaces (small), before and after the July 2020 GEA introduction. All specifications include canton, year, and size-class fixed effects. Standard errors clustered at the canton level in parentheses. Column (5) uses institutional units (enterprises) rather than workplaces. 26 cantons \times 13 years \times 2 size classes.

Table 3: Event Study: Average Firm Size in the 50–249 Employee Bin

Year	Coefficient	Std. Error	95% CI
<i>Pre-treatment (2011–2018, ref. = 2019)</i>			
2011	1.242***	(0.418)	[0.424, 2.060]
2012	0.800	(0.616)	[-0.408, 2.008]
2013	0.497	(0.548)	[-0.577, 1.571]
2014	0.324	(0.564)	[-0.782, 1.430]
2015	-0.041	(0.440)	[-0.903, 0.821]
2016	-0.106	(0.461)	[-1.010, 0.797]
2017	0.523	(0.454)	[-0.367, 1.413]
2018	0.066	(0.237)	[-0.399, 0.531]
<i>Post-treatment (2020–2023)</i>			
2020	0.442	(0.477)	[-0.492, 1.377]
2021	0.524	(0.341)	[-0.145, 1.193]
2022	0.847**	(0.381)	[0.102, 1.593]
2023	1.361**	(0.537)	[0.309, 2.414]
Pre-trend F-test	F(8, 595) = 4.51, $p = 0.000$		

Notes: Coefficients from regressing average employees per workplace on year \times medium-bin (50–249) interactions, with canton, year, and size-class fixed effects. Reference year: 2019 (last pre-treatment year). Standard errors clustered at the canton level. The pre-treatment coefficients show a declining pattern from 2011 to 2018, consistent with convergence in average firm size across bins rather than a GEA anticipation effect. The significant post-2022 coefficients should be interpreted cautiously given the pre-trend. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

threshold.

5.3 Robustness

Table 4 presents a battery of alternative specifications. The restricted sample (2015–2023), which addresses the pre-trend concern, yields a coefficient of +0.705 (SE = 0.497)—larger in magnitude but still insignificant. Excluding 2020 entirely to eliminate the COVID confound gives +0.531 (SE = 0.504). Placebo treatments placed in 2016 and 2017 produce small, insignificant, and negative coefficients (−0.444 and −0.256), as expected.

Splitting the sample by canton size reveals heterogeneity. Large cantons (Zurich, Bern, Vaud, etc.) show a coefficient of −0.270 (SE = 0.316)—the “right” sign for avoidance but statistically insignificant and economically small. Small cantons show +1.088 (SE = 0.919)—positive and noisy. Neither subsample supports the threshold-avoidance hypothesis.

The wild cluster bootstrap p -value for the main specification is 0.462, confirming that the null result is not an artifact of asymptotic inference with 26 clusters.

Table 4: Robustness: Alternative Specifications

Specification	Medium \times Post	Std. Error
Baseline (2011–2023)	0.425	(0.521)
Restricted sample (2015–2023)	0.705	(0.497)
Excluding 2020 (COVID year)	0.531	(0.504)
Enterprises (not workplaces)	1.237**	(0.511)
FTE per workplace	0.533	(0.518)
Placebo: treatment = 2016	-0.444	(0.396)
Placebo: treatment = 2017	-0.256	(0.403)
Large cantons only	-0.270	(0.316)
Small cantons only	1.088	(0.919)
Wild cluster bootstrap p -value	0.462	

Notes: All specifications report the coefficient on Medium (50–249) \times Post (2020+) from a DiD comparing average employees per workplace in the 50–249 bin to the 10–49 bin. Standard errors clustered at the canton level. The baseline uses workplace-level data from BFS STATENT covering 26 cantons over 2011–2023. Wild cluster bootstrap uses 9,999 Rademacher draws. The enterprise specification uses BFS Table 107 (institutional units). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6. Discussion

Why is the null informative?. The bunching literature overwhelmingly reports positive findings: firms distort their size to avoid regulatory thresholds in France (Garicano et al., 2016), firm-size distributions exhibit discontinuities at tax thresholds in many countries (Hsieh and Klenow, 2009; Decker et al., 2014), and even modest bureaucratic burdens can generate excess mass below cutoffs (Gourio and Roys, 2014). Publication bias may partly explain this pattern—null results at thresholds are less likely to be written up. The Swiss GEA provides a rare clean null: a newly introduced threshold, a plausible mechanism for avoidance, and no detectable response.

Enforcement as the key mediating variable. The most parsimonious explanation for the null is enforcement design. The GEA carries no fines, no administrative sanctions, and no private enforcement mechanism. Firms that ignore the mandate entirely face, at worst, reputational criticism. In contrast, the French labor market regulations studied by Garicano et al. (2016) carry enforceable obligations (e.g., mandatory works councils, profit-sharing requirements), and German codetermination laws impose binding governance constraints (Mueller and Ouimet, 2020). This suggests that the bunching elasticity is not a fixed property of a threshold but depends critically on the regulatory “teeth” behind it. The result is

consistent with theoretical models where compliance costs, not threshold existence, drive avoidance behavior (Kleven, 2016).

The compliance mirage. I propose the term “compliance mirage” for regulations that impose observable process requirements—in this case, conducting and disclosing a pay equity audit—without creating economic incentives for avoidance. The GEA is prototypical: it generates compliance activity (firms hire auditors, run the Logib tool, report results) without distorting the fundamental decision of how many workers to employ. This has implications for policy design. If the goal is to promote pay equity without distorting the firm-size distribution, the GEA’s soft-enforcement model achieves it. Whether it also achieves its primary objective—reducing gender pay gaps—is a separate question that this paper’s data cannot directly address.

Comparison with pay transparency evidence. The growing literature on pay transparency mandates has found mixed effects on pay gaps. Bennedsen et al. (2022) find that Denmark’s pay transparency law reduced the gender pay gap by approximately 7%, driven by slower male wage growth. Duchini et al. (2020) find that the UK’s gender pay gap reporting requirement had limited impact on actual pay gaps but increased firm attention to diversity. Baker et al. (2023) study Canadian pay transparency legislation and find small effects. The Swiss null on threshold avoidance complements these findings: even if the GEA affects pay-setting behavior within treated firms, it does not affect the *margin of firm size selection*, suggesting that firms view the audit as a procedural cost rather than a substantive constraint.

Limitations. This study has clear limitations. The coarse-bin structure prevents a standard bunching analysis. The test is indirect: I observe average bin outcomes, not the density around 100. If the true bunching response is concentrated in a narrow band (e.g., 95–105 employees) and modest in magnitude, the bin-level DiD may lack power to detect it. The minimum detectable effect of approximately 1.3 employees represents a meaningful fraction of the average (1.3%), but very small responses—comparable to those found at some European thresholds—could fall below this threshold. Additionally, the event study reveals pre-trends in the early years (2011–2014), though the immediate pre-treatment period (2015–2019) is relatively stable. Future research with micro-level administrative data could provide sharper tests.

7. Conclusion

Switzerland’s Gender Equality Act created a textbook opportunity for threshold avoidance—a new, well-publicized regulatory obligation at a round-number firm-size cutoff. Yet firms did not respond. Using the universe of Swiss workplaces over 13 years, I find no evidence of distortion: no decline in average firm size, no shift in firm counts, no change in female employment shares. The regulation appears to be a compliance mirage—formally binding but economically toothless, generating audit activity without altering the firm-size distribution.

This null carries a broader lesson. The bunching literature has documented impressive behavioral responses at regulatory thresholds, but those responses may be contingent on enforcement design. A threshold without teeth may be a threshold without distortion. As countries consider expanding pay equity mandates, the Swiss experience suggests that soft enforcement can avoid the economic costs of firm-size distortion—though whether it also achieves its primary equity objectives remains an open and important question.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data are from the Swiss Federal Statistical Office (BFS) STATENT, accessed via the BFS PXWeb API.

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

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Baker, Michael, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar**, “Pay Transparency and the Gender Gap,” *American Economic Journal: Applied Economics*, 2023, 15 (2), 157–183.
- Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon**, “Do Firms Respond to Gender Pay Gap Transparency?,” *Journal of Finance*, 2022, 77 (4), 2051–2091.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How Much Should We Trust Differences-in-Differences Estimates?,” *Quarterly Journal of Economics*, 2004, 119 (1), 249–275.
- Blundell, Jack**, “Wage Responses to Gender Pay Gap Reporting Requirements,” *Working Paper*, 2022.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, 90 (3), 414–427.
- Chetty, Raj, John N Friedman, Tore Olsen, and Luigi Pistaferri**, “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records,” *Quarterly Journal of Economics*, 2011, 126 (2), 749–804.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda**, “The Role of Entrepreneurship in US Job Creation and Economic Dynamism,” *Journal of Economic Perspectives*, 2014, 28 (3), 3–24.
- Duchini, Emma, Stefania Simion, and Arthur Turrell**, “Pay Transparency and Gender Pay Gaps: New Evidence from the UK,” *Working Paper*, 2020.
- Ferracci, Marc and Alexander Ferro**, “Evidence on the Distributional Effects of Size-Dependent Regulatory Thresholds: Italian Evidence,” *Labour Economics*, 2014, 28, 1–14.
- Garicano, Luis, Claire Lelarge, and John Van Reenen**, “Firm Size Distortions and the Productivity Distribution: Evidence from France,” *American Economic Review*, 2016, 106 (11), 3439–3479.
- Gourio, François and Nicolas Roys**, “Firm Size Distortions under Regional Size-Contingent Regulations,” *Review of Economic Studies*, 2014, 81 (4), 1439–1468.

Guner, Nezh, Gustavo Ventura, and Yi Xu, “Macroeconomic Implications of Size-Dependent Policies,” *Review of Economic Dynamics*, 2008, 11 (4), 721–744.

Hsieh, Chang-Tai and Peter J Klenow, “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.

Kleven, Henrik Jacobsen, “Bunching,” *Annual Review of Economics*, 2016, 8, 435–464.

Mueller, Holger M and Paige P Ouimet, “Firm Size and the Costs of Codetermination: Evidence from a Threshold,” *Journal of Financial Economics*, 2020, 137 (3), 619–635.

Restuccia, Diego and Richard Rogerson, “Policy Distortions and Aggregate Productivity with Heterogeneous Establishments,” *Review of Economic Dynamics*, 2008, 11 (4), 707–720.

A. Data Appendix

Data source. All data come from the Swiss Federal Statistical Office (BFS) Structural Business Enterprise Statistics (STATENT), accessed via the BFS PXWeb API (<https://www.pxweb.bfs.admin.ch>). STATENT is based on the register of employers contributing to the AHV/IV/EO social insurance system and constitutes a full census of Swiss workplaces and enterprises. The data are released annually with a two-year lag. I use two BFS tables:

- **STATENT workplace table:** Workplace counts and total employment by canton, year, and size class (1–9, 10–49, 50–249, 250+). Years: 2011–2023.
- **BFS Table 107 (enterprises):** Enterprise counts and total employment by canton, year, and size class, using institutional units rather than workplace units. Years: 2011–2023.

Variable construction. Average employment per workplace (or enterprise) is computed as total employment in the cell divided by the number of units. Female employment shares are computed as female employment divided by total employment. Full-time equivalents (FTE) are reported directly by BFS in supplementary tables.

Sample restrictions. I restrict the analysis to the 10–49 and 50–249 employee bins. Cells with suppressed values (due to BFS confidentiality rules) are dropped, affecting a small number of canton-year observations. The final sample contains 646 canton-year-bin observations for the main specification.

Size-class definitions. BFS defines size classes based on headcount (number of persons employed, including part-time workers), not FTE. The GEA threshold is likewise based on headcount. The four standard bins (1–9, 10–49, 50–249, 250+) are the finest publicly available disaggregation. Finer bins (e.g., 50–99 vs. 100–249) are available only through restricted-access micro-data.

B. Robustness Appendix

Wild cluster bootstrap. With 26 canton-level clusters, asymptotic cluster-robust standard errors may over-reject. I implement the wild cluster bootstrap of [Cameron et al. \(2008\)](#) using 9,999 Rademacher weight draws. The bootstrap p -value for the main specification is 0.462, consistent with the asymptotic inference.

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Avg. employment (50–249)	0.425	0.521	5.22	0.0814	0.0998	Moderate positive
Log firm count (50–249)	0.007	0.006	1.15	0.0063	0.0053	Small positive
Female share (50–249)	-0.001	0.004	0.03	-0.0197	0.1040	Small negative
<i>Panel B: Heterogeneous (canton size split)</i>						
Avg. employment — large cantons	-0.270	0.316	5.22	-0.0517	0.0606	Moderate negative
Avg. employment — small cantons	1.088	0.919	5.22	0.2085	0.1761	Large positive

Notes: **Country:** Switzerland. **Research question:** Does the 2020 Gender Equality Act pay audit requirement for firms with 100 or more employees distort the firm-size distribution in the 50–249 employee range? **Policy mechanism:** The Gleichstellungsgesetz revision mandates that firms with 100 or more employees conduct a certified equal-pay analysis by June 2023 and communicate results to employees; non-compliance risks reputational sanctions but no direct fines. **Outcome definition:** Average number of employees per workplace in the BFS 50–249 employee size bin, capturing within-bin compositional shifts from potential bunching below the 100-employee threshold. **Treatment:** Binary; post-July 2020 introduction of mandatory pay audits for 100+ employee firms. **Data:** BFS STATENT annual census, 2011–2023, 26 cantons \times 4 size classes, full population of Swiss workplaces. **Method:** Difference-in-differences comparing 50–249 bin to 10–49 bin, canton and year fixed effects, standard errors clustered at canton level. **Sample:** All Swiss workplaces in the 10–49 and 50–249 employee bins; 646 canton-year-bin observations after excluding suppressed cells. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation of average workplace employment across cantons in the 50–249 bin. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

Restricted window. The event study reveals a declining pre-trend from 2011 to approximately 2015. To address this, I re-estimate the main specification using only 2015–2023, a window in which pre-treatment coefficients are individually insignificant. The restricted-window estimate is +0.705 (SE = 0.497), qualitatively identical to the baseline.

Enterprise versus workplace. Workplaces (*Arbeitsstätten*) and enterprises (*Unternehmen*) differ in Swiss business statistics. A single enterprise may operate multiple workplaces. The enterprise specification (Column 5 of Table 2) shows a significant positive coefficient (+1.237, $p = 0.02$), suggesting that the positive tendency is, if anything, stronger at the enterprise level. This rules out compositional effects from multi-establishment firms shifting employment across workplaces.

C. Standardized Effect Sizes