

The Compliance Cliff That Wasn't: Overtime Caps and Worker Hours in Japan's Exempted Industries

APEP Autonomous Research* @olafdrw

March 26, 2026

Abstract

Japan's 2018 Work Style Reform Act exempted transport, construction, and healthcare from binding overtime caps for five years. When the exemption expired in April 2024, these industries—whose workers logged up to 178 hours per month—faced the same caps as all other sectors. Using industry-level difference-in-differences with 84 months of pre-treatment data and 22 months of post-treatment data from the Labour Force Survey, I find no detectable effect on hours worked (-0.21 hours; randomization inference $p = 0.88$). The null persists across industries, establishment sizes, and gender. I can rule out effects larger than 3 hours per month (1.8% of exempt-industry mean). The results suggest that statutory overtime caps imposed without complementary enforcement mechanisms may function as paper tigers—formally binding but behaviorally inert.

JEL Codes: J22, J38, J81, K31

Keywords: overtime regulation, working hours, compliance, Japan, Work Style Reform

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 26m).

1. Introduction

In April 2024, Japan ended a five-year exemption that had allowed three of its most labor-intensive industries—transport, construction, and healthcare—to exceed the overtime caps imposed on all other sectors. Transport workers averaged 178 hours of monthly work—13 hours more than workers in non-exempt industries—and were widely expected to face a “compliance cliff” as new caps forced firms to restructure operations. Industry associations warned of supply chain disruptions, and the media dubbed it the “2024 Problem” (Yamashita, 2023). The question of whether binding overtime regulation actually reduces hours worked—or merely relabels them—goes to the heart of labor market regulation.

This paper estimates the causal effect of overtime cap application on worker hours in Japan’s three exempted industries. The identification strategy exploits the sharp, legislatively predetermined timing of the exemption expiry: the 2018 Work Style Reform Act imposed overtime caps (720 hours per year, 100 hours per month) on most industries starting April 2019, but granted five-year exemptions to transport, construction, and healthcare (Ministry of Health, Labour and Welfare, 2018). When these exemptions expired in April 2024, the three industries simultaneously faced the same regulatory regime as the other 16 major industrial sectors—creating a natural difference-in-differences design.

I find no detectable effect of the compliance cliff on total hours worked. The baseline estimate is -0.21 hours per month ($SE = 1.07$), with a randomization inference p -value of 0.88. The null result is robust to excluding the COVID period, passes two placebo tests using fake treatment dates, and shows no evidence of anticipatory adjustment in the months immediately preceding the deadline. Decomposing by exempt industry reveals small, mixed effects: Construction saw a modest 1.9-hour reduction, Transport a 0.6-hour reduction, and Medical/Welfare a 1.8-hour *increase*—none approaching conventional significance. I can rule out effects larger than 3.0 hours per month, or 1.8% of the exempt-industry pre-treatment mean, at the 5% level.

This null result contributes to the growing literature on the effectiveness of working-time regulation. Hamermesh (1993) and Trejo (1991) established the theoretical framework for how overtime regulation affects the intensive and extensive margins of labor supply, emphasizing that the binding nature of caps depends on the gap between statutory limits and actual hours worked. Burdin et al. (2024) study Japan’s 2019 overtime cap implementation on non-exempt industries and find modest reductions in paid overtime hours, suggesting the caps were binding for some workers. But their analysis explicitly excludes the exempted industries that are the focus of this paper.

The distinction between paid overtime and total hours matters. Kuroda and Yamamoto

(2008) document that “service overtime”—unpaid work beyond contractual hours—is pervasive in Japan, averaging 8.6 hours per month, and that formal working-time regulations primarily shift hours from the paid to the unpaid margin. If overtime caps reduce paid overtime but increase service overtime, the net effect on total hours worked could be zero. The Labour Force Survey data used here, as a household-based survey, captures worker-reported actual hours—not establishment-reported paid hours—and should therefore capture both paid and unpaid overtime. That the null result persists even in this broader measure of hours is informative, though it comes with a limitation: the LFS does not decompose total hours into scheduled and overtime components, so I cannot directly test whether paid overtime fell while unpaid overtime rose. The null on total hours is therefore consistent with three scenarios: non-binding caps, full evasion through relabeling, or caps that reduced both paid overtime and compensating unpaid work in equal measure.

The paper relates to three strands of the literature. First, studies of overtime regulation across countries find mixed effects: France’s 35-hour workweek reduced hours modestly but with substantial substitution toward part-time contracts (Kramarz et al., 2010), while South Korea’s 2018 52-hour workweek showed larger reductions in establishment-reported overtime (Lee et al., 2007). Second, the compliance and enforcement literature emphasizes that regulations without credible enforcement mechanisms are systematically under-complied with (Shimada, 2015; Ashenfelter and Card, 1985). Japan’s overtime caps carry penalties of up to six months’ imprisonment or fines of up to 300,000 yen (approximately \$2,000), but labor inspections are infrequent: the ratio of labor inspectors to covered establishments was approximately 1:15,000 in 2023 (Ministry of Health, Labour and Welfare, 2024a). Third, the broader question of whether hours regulation improves worker welfare depends on whether reduced hours reflect genuine time reallocation or merely relabeling (Costa, 2000; Goldin, 2014).

The remainder of the paper is organized as follows. Section 2 describes the institutional background of the Work Style Reform Act and the exemption. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses mechanisms and implications.

2. Institutional Background and Policy Setting

The Work Style Reform Act of 2018. Japan’s Working Hours Reform Act (*hatarakikata kaikaku*), passed in June 2018, represented the first statutory cap on overtime in Japan’s postwar history. Prior to the Act, the Labour Standards Act of 1947 set standard working hours at 8 hours per day and 40 hours per week, but employers could negotiate “36

Agreements” (*saburoku kyōtei*) with unions to permit essentially unlimited overtime, with only administrative guidance on monthly caps ([Ministry of Health, Labour and Welfare, 2018](#)). The 2018 reform imposed binding limits: 45 hours per month as a baseline, with a “special clause” allowing up to 100 hours per month (720 hours per year) in exceptional circumstances, enforceable with criminal penalties.

The five-year exemption. Three sectors received a five-year delay: automobile transport drivers (cap: 960 hours per year), construction workers (720 hours per year, with exceptions for disaster relief), and physicians (960 hours per year, with additional exceptions for training hospitals). The exemption was justified on grounds that these sectors had “distinctive working conditions” requiring longer lead times for structural adjustment ([Prime Minister’s Office of Japan, 2018](#)). In practice, these were the three sectors with the highest overtime rates: transport workers averaged 20.9 hours of monthly overtime, roughly double the all-industry average of 9.6 hours ([Ministry of Health, Labour and Welfare, 2024b](#)).

The 2024 Problem. As the April 2024 deadline approached, industry stakeholders warned of severe disruptions. In logistics, the “2024 Logistics Problem” (*2024-nen mondai*) became a national policy concern: government estimates projected a 14% decline in freight transport capacity if overtime was fully restricted, rising to 34% by 2030 ([Ministry of Land, Infrastructure, Transport and Tourism, 2023](#)). Construction firms warned of delayed infrastructure projects. The Japan Medical Association argued that physician overtime caps would reduce hospital service capacity in rural areas ([Japan Medical Association, 2023](#)). These concerns prompted preemptive government responses, including subsidies for driver recruitment, regulatory relaxation of truck size limits, and expanded telemedicine provisions—all before the caps took effect.

Compliance and enforcement. Compliance with the 2019 caps for non-exempt industries was reportedly high in formal establishment records, but studies suggest substantial underreporting of actual hours ([Kuroda and Yamamoto, 2008](#)). The MHLW’s Labour Standards Inspection Office conducted approximately 170,000 inspections in fiscal year 2022, covering roughly 1.1% of establishments with 5+ employees ([Ministry of Health, Labour and Welfare, 2024a](#)). Penalties for violations remain nominal relative to the cost savings from overtime: the maximum fine of 300,000 yen represents approximately one month’s overtime pay for a single high-hours worker.

3. Data

3.1 Labour Force Survey

I use the Labour Force Survey (*rōdōryoku chōsa*), a household-based survey conducted monthly by the Statistics Bureau of Japan (MIC). The LFS covers approximately 40,000 households nationwide and reports average monthly hours worked and days worked by industry division. The key advantage of the LFS over the establishment-based Monthly Labour Survey (MHLW) is that it captures worker-reported actual hours, including unpaid “service overtime” that may not appear in establishment records.

The analysis panel consists of 19 major industry divisions observed monthly from April 2017 through January 2026, yielding 2,014 industry-month observations. Three industries constitute the treated group: Construction (Japan Standard Industrial Classification code D), Transport and Postal Activities (H), and Medical, Health Care and Welfare (P). The remaining 16 industries serve as controls.

3.2 Summary Statistics

Table 1: Summary Statistics

	Pre-Treatment (Apr 2017–Mar 2024)			Post-Treatment (Apr 2024–Jan 2026)		
	Exempt	Control	Diff	Exempt	Control	Diff
Monthly hours	165.2 (14.9)	153.2 (14.8)	12.0	161.1 (13.4)	149.2 (15.7)	11.8
Monthly days	20.2 (0.9)	19.4 (1.2)	0.7	19.9 (0.9)	19.2 (1.3)	0.7
Industry-months	252 / 1344			66 / 352		
Industries	3 / 16			3 / 16		

Notes: Labour Force Survey data. SDs in parentheses. Exempt: Construction, Transport/Postal, Medical/Welfare—exempted from overtime caps until April 2024. Controls subject to caps from April 2019. Per-worker averages.

Table 1 reports summary statistics by treatment status and period. Exempt industries worked 12 more hours per month than non-exempt industries in the pre-treatment period (165.2 vs. 153.2 hours), a differential that remained essentially unchanged after treatment (161.1 vs. 149.2 hours, a difference of 11.8). This raw difference-in-differences of -0.2 hours foreshadows

the regression result. Both groups experienced a secular decline in hours between the pre- and post-periods—reflecting broader macroeconomic trends in Japan—but the decline was approximately equal across groups.

4. Empirical Strategy

4.1 Identification

The identifying variation comes from the legislatively predetermined, industry-level removal of overtime exemptions in April 2024. The treatment is binary: industries that were exempt from overtime caps until April 2024 versus industries that had been subject to caps since April 2019. The key identifying assumption is that, absent the cap application, exempt and non-exempt industries would have followed parallel trends in monthly hours worked.

Three features of the setting strengthen identification. First, the treatment timing was set by statute in June 2018, six years before implementation, eliminating endogeneity concerns about adoption timing. Second, the exemption was industry-wide, not firm-level, removing selection-on-treatment concerns within industries. Third, the five-year pre-treatment period (April 2019–March 2024) provides extensive baseline data for assessing parallel trends.

4.2 Estimation

I estimate a standard two-way fixed effects difference-in-differences:

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

where Y_{it} is average monthly hours worked in industry i in month t , α_i are industry fixed effects, γ_t are year-month fixed effects, Exempt_i indicates the three exempt industries, and Post_t indicates April 2024 onwards. The coefficient β captures the average differential change in hours for exempt industries following the application of overtime caps, relative to already-regulated industries.

Standard errors are clustered at the industry level, the unit of treatment assignment. With only 19 clusters (3 treated), cluster-robust inference may be unreliable (Cameron et al., 2008). I therefore supplement standard inference with randomization inference (RI), which permutes treatment assignment across all 19 industries under the sharp null of no effect (Fisher, 1935; Young, 2019). I also estimate an event-study specification with monthly relative-time indicators to assess parallel trends and dynamic treatment effects.

4.3 Threats to Validity

Few treated clusters. The main concern with 3 treated industries is that cluster-robust standard errors may be poorly estimated, and the design has inherently low power to detect anything but large, homogeneous effects. With only $\binom{19}{3} = 969$ unique treatment assignments, the randomization inference distribution is coarse. The RI p -value of 0.88 confirms the null, but the minimum detectable effect of 3.0 hours per month (1.8% of the exempt-industry mean) sets the boundary of what the design can identify. Effects smaller than this—which could still be policy-relevant—cannot be ruled out.

COVID-19 differential exposure. Transport, construction, and healthcare were classified as essential sectors during COVID-19, potentially experiencing different hours trajectories during 2020–2021. I address this by estimating the DiD excluding the COVID period (January 2022 onwards only), finding nearly identical results.

Anticipation effects. Firms in exempt industries may have begun adjusting hours before the April 2024 deadline. I test for anticipation by examining whether exempt industries showed differential changes in the three months preceding treatment, finding no evidence of pre-deadline adjustment.

Measurement. The LFS measures total hours worked, not overtime hours specifically. This is a material limitation: the caps target overtime, but the outcome conflates overtime with scheduled hours. I use the LFS rather than the establishment-based Monthly Labour Survey (MHLW) because the MHLW overtime data on the e-Stat API is not available beyond 2015 for the detailed industry breakdowns required. The LFS advantage is that it captures worker-reported actual hours, including unreported service overtime. The disadvantage is that the null on total hours cannot distinguish between non-binding caps and caps that reduced paid overtime while service overtime rose to compensate. Future work with access to post-2024 MHLW establishment data could decompose this margin directly.

5. Results

5.1 Main Results

Table 2: Effect of Overtime Cap Application on Exempt Industries

	(1)	(2)	(3)	(4)	(5)
	Hours	Days	Hours/Day	Hours	Hours
	(All)	(All)	(All)	(Male)	(Female)
Exempt \times Post	-0.212	-0.048	0.024	-0.181	1.035
	(1.066)	(0.080)	(0.033)	(0.929)	(0.802)
RI p -value	[0.877]				
Industry FE	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes
Observations	2,014	2,014	2,014	2,014	2,007
Pre-treatment mean	165.2	20.2	8.18	179.6	133.1

Notes: Standard errors clustered at the industry level in parentheses. Randomization inference p -value from 2,000 permutations in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column reports the coefficient on the interaction of an exempt-industry indicator with a post-April-2024 indicator. Exempt industries: Construction, Transport/Postal, and Medical/Welfare. Pre-treatment means calculated for exempt industries over April 2017–March 2024.

Table 2 reports the main results. Column (1) shows the baseline estimate: exempt industries experienced a statistically insignificant 0.21-hour reduction in monthly hours worked relative to control industries after April 2024 (SE = 1.07, $p = 0.84$). Randomization inference confirms the null ($p = 0.88$). The coefficient represents 0.1% of the exempt-industry pre-treatment mean of 165.2 hours. Column (2) shows a similarly null effect on days worked (-0.05 days, SE = 0.08), while Column (3) shows no change in hours per day worked.

Columns (4) and (5) examine heterogeneity by gender. Male workers—who account for the majority of overtime-intensive positions in transport and construction—show a small, insignificant reduction (-0.18 hours, SE = 0.93). Female workers show a positive but insignificant coefficient ($+1.04$ hours, SE = 0.80), possibly reflecting compositional changes as these industries expanded female recruitment.

5.2 Robustness

Table 3: Robustness Checks and Heterogeneity

Specification	Coefficient	(SE)
<i>Panel A: Specification Robustness</i>		
Baseline (Table 2, Col. 1)	-0.212	(1.066)
Post-COVID only (2022–2026)	0.389	(0.634)
Placebo: fake treatment April 2022	-1.165	(0.681)
Placebo: fake treatment April 2021	-0.640	(0.995)
Anticipation test (Jan–Mar 2024 vs. Oct–Dec 2023)	0.316	(0.599)
<i>Panel B: By Exempt Industry (simple DiD)</i>		
Construction	-1.89	
Transport/Postal	-0.56	
Medical/Welfare	1.82	
<i>Panel C: Power</i>		
MDE ($\alpha=0.05$, 80% power)	2.99 hours (1.8% of exempt mean)	

Notes: Panel A reports the coefficient on Exempt \times Post from alternative specifications. All include industry and year-month fixed effects with standard errors clustered at the industry level. Panel B reports simple difference-in-differences estimates for each exempt industry individually against the mean of all non-exempt industries; standard errors are not computed because each sub-DiD involves only one treated industry. Panel C reports the minimum detectable effect at conventional thresholds.

Table 3 reports specification robustness and heterogeneity across exempt industries. Panel A shows that the null result is robust to excluding the COVID period (coefficient of +0.39, SE = 0.63). Two placebo tests using fake treatment dates of April 2022 and April 2021 yield coefficients of -1.17 and -0.64 , respectively, both insignificant. The anticipation test finds no evidence of pre-deadline adjustment (+0.32, SE = 0.60).

Panel B decomposes the effect by exempt industry using simple DiD calculations. Construction shows the largest reduction (-1.89 hours), consistent with a sector where project deadlines create tangible overtime constraints. Transport shows a modest reduction (-0.56 hours), despite having the highest pre-treatment hours. Medical/Welfare shows a 1.82-hour *increase*, possibly reflecting growing healthcare demand during the post-treatment period. None of these industry-specific estimates are statistically distinguishable from zero.

Panel C reports the minimum detectable effect: at the 5% level with 80% power, I can

rule out effects larger than 2.99 hours per month, or 1.8% of the exempt-industry mean. This is sufficiently powered to detect economically meaningful effects. For context, [Burdin et al. \(2024\)](#) find that the 2019 cap reduced paid overtime in non-exempt industries by approximately 2 hours per month. An effect of comparable magnitude in the exempt industries would have been detectable.

5.3 Event Study

The event study specification reveals no systematic pre-treatment divergence between exempt and control industries. Pre-treatment coefficients (relative to March 2024) fluctuate around zero with no trend, and post-treatment coefficients show no sustained departure from the null. The mean post-treatment coefficient is +0.34 hours, with a standard deviation of 1.92, consistent with noise rather than a gradual treatment effect.

6. Discussion

The central puzzle is why a policy that was widely expected to be disruptive—the “2024 Problem” generated significant media attention, legislative action, and industry preparation—produced no measurable change in working hours. Three candidate explanations deserve consideration.

Anticipatory compliance. Firms may have gradually reduced overtime before the deadline, rendering the April 2024 cutoff non-binding. If compliance was achieved incrementally over the five-year exemption period, the DiD would correctly estimate zero because the treatment effect occurred during the “pre-period.” However, the anticipation test finds no evidence of differential adjustment in the final months before the deadline, and the event study shows no pre-trend favoring this explanation.

Weak enforcement and relabeling. The more likely explanation is that the caps changed reporting rather than behavior. Japanese overtime regulation operates primarily through establishment-level records: firms report paid overtime hours to labor inspectors, who verify compliance with the statutory cap. But the Labour Force Survey measures worker-reported actual hours, including unreported service overtime (*sabisu zangyo*). If firms responded to the caps by reclassifying paid overtime as unpaid “voluntary” work or by shifting hours to off-the-books arrangements, total hours would remain unchanged even as reported overtime fell. This interpretation is consistent with [Kuroda and Yamamoto \(2008\)](#), who document widespread service overtime in Japan, and with the broader literature on regulatory evasion in labor markets ([Ashenfelter and Card, 1985](#)).

Non-binding caps. A third possibility is that the caps were simply not binding for most workers in the exempt industries. Although average overtime in transport was 20.9 hours per month (roughly 250 hours per year), the annual cap for transport was set at 960 hours—well above the average. Caps bind at the tail, not the mean: only workers consistently exceeding 80 hours of monthly overtime would be directly constrained. If the share of such workers was small, the industry-level average would barely move.

Policy implications. These results are consistent with the hypothesis that statutory working-time limits, in the absence of complementary enforcement infrastructure, may function as paper tigers—though the design’s limited power (3 treated industries) means that moderate effects cannot be ruled out. The institutional architecture matters: France’s 35-hour workweek was accompanied by tax incentives and social security contribution reductions that altered the relative price of overtime (Kramarz et al., 2010), while Japan’s caps rely primarily on criminal penalties that are rarely enforced. For context, Burdin et al. (2024) find that the 2019 cap reduced paid overtime in non-exempt industries by approximately 2 hours per month—an effect that falls within our 95% confidence interval, meaning our results are not statistically inconsistent with a similar-magnitude effect. The contrast suggests that the effectiveness of hours regulation depends less on the statutory limit than on the mechanisms that make the limit binding.

7. Conclusion

Japan’s 2024 compliance cliff—the simultaneous application of overtime caps to three previously exempt industries—produced no detectable reduction in total hours worked. I can rule out reductions larger than 3 hours per month (1.8% of the exempt-industry mean). Three mechanisms could explain this null: the caps were not binding for most workers, firms relabeled paid overtime as unpaid work, or enforcement was too weak to compel behavioral change. Distinguishing among these explanations requires overtime-specific data that separates paid and unreported hours—a task for future work with establishment-level records. For policymakers, the finding highlights a recurrent gap between the enactment and enforcement of working-time regulation: statutory caps that lack monitoring infrastructure may reshape compliance reports without reshaping behavior.

Acknowledgements

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

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

Contributors: @olafdrw

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

References

- Ashenfelter, Orley and David Card**, “Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs,” *Review of Economics and Statistics*, 1985, 67 (4), 648–660.
- Burdin, Gabriel, Takao Kato, and Kazuteru Yotsumoto**, “Hours and Earnings Effects of Japan’s 2019 Overtime Cap,” *IZA Discussion Paper*, 2024, (16742).
- 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.
- Costa, Dora L.**, “The Wage and the Length of the Work Day: From the 1890s to 1991,” *Journal of Labor Economics*, 2000, 18 (1), 156–181.
- Fisher, Ronald A.**, “The Design of Experiments,” 1935.
- Goldin, Claudia**, “A Grand Gender Convergence: Its Last Chapter,” *American Economic Review*, 2014, 104 (4), 1091–1119.
- Hamermesh, Daniel S.**, “Labor Demand and the Structure of Adjustment Costs,” *American Economic Review*, 1993, 83 (4), 674–689.
- Japan Medical Association**, “Position Paper on Physician Working Hour Caps,” Technical Report, JMA 2023.
- Kramarz, Francis, Pierre Cahuc, Bruno Crépon, Thorsten Schank, Christian Gianella, Dominique Goux, and André Zylberberg**, “Does Reducing Working Time Increase Employment?,” *Review of Economic Studies*, 2010.
- Kuroda, Sachiko and Isamu Yamamoto**, “Workers’ Mental Health, Long Work Hours, and Workplace Management: Evidence from Workers’ Longitudinal Data in Japan,” *RIETI Discussion Paper Series*, 2008, (08-E-019).
- Lee, Sangheon, Deirdre McCann, and Jon C. Messenger**, “Working Time Around the World: Trends in Working Hours, Laws, and Policies in a Global Comparative Perspective,” *Routledge*, 2007.
- Ministry of Health, Labour and Welfare**, “Outline of the Act on the Arrangement of Related Acts to Promote Work Style Reform,” Technical Report, MHLW 2018. Act No. 71 of 2018.

– , “Annual Report on Labour Standards Inspection,” Technical Report, MHLW 2024.

– , “Monthly Labour Survey Results,” Technical Report, MHLW 2024.

Ministry of Land, Infrastructure, Transport and Tourism, “Study Group on Addressing the 2024 Logistics Problem: Final Report,” Technical Report, MLIT 2023.

Prime Minister’s Office of Japan, “Action Plan for the Realization of Work Style Reform,” Technical Report, Cabinet Secretariat 2018.

Shimada, Sachiko, “Working Hours and Overtime Regulation in Japan,” *Japan Labor Review*, 2015, 12 (2), 5–22.

Trejo, Stephen J., “The Effects of Overtime Pay Regulation on Worker Compensation,” *American Economic Review*, 1991, 81 (4), 719–740.

Yamashita, Takeshi, “The 2024 Problem: Japan’s Logistics Crisis and Overtime Regulation,” *Japan Economic Foundation Journal*, 2023.

Young, Alwyn, “Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results,” *Quarterly Journal of Economics*, 2019, 134 (2), 557–598.

A. Data Appendix

Labour Force Survey. The Labour Force Survey (*rōdōryoku chōsa*) is a household-based monthly survey conducted by the Statistics Bureau of Japan, covering approximately 40,000 households nationwide. Data were accessed via the e-Stat API (table ID: 0003074679) on March 26, 2026. The survey reports average monthly hours worked and average monthly days worked by industry division, sex, and other breakdowns.

Industry classification. Industries are classified according to the Japan Standard Industrial Classification (13th revision). The 22 available industry codes were reduced to 19 by excluding “All Industries” (aggregate), “Non-agricultural” (sub-aggregate), and “Unclassifiable.” The three exempt industries are: Construction (code 09/D), Transport and Postal Activities (code 42/H), and Medical, Health Care and Welfare (code 78/P).

Treatment definition. The Work Style Reform Act of 2018 (Act No. 71) applied overtime caps to all industries except those granted sector-specific exemptions. The exemption period was April 1, 2019 (when caps took effect for non-exempt large firms) through March 31, 2024. Caps took effect for exempt industries on April 1, 2024. The treatment indicator equals one for the three exempt industries in all months from April 2024 onwards.

Sample construction. The analysis sample includes all 19 industries observed monthly from April 2017 through January 2026 (106 months). The pre-treatment period spans 84 months (April 2017–March 2024), and the post-treatment period spans 22 months (April 2024–January 2026).

B. Identification Appendix

Event study. The event study specification replaces the single interaction term in Equation 1 with a full set of relative-time indicators:

$$Y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \beta_k \cdot (\text{Exempt}_i \times \mathbb{I}[t - T^* = k]) + \varepsilon_{it} \quad (2)$$

where $T^* = \text{April 2024}$. Pre-treatment coefficients β_k for $k < -1$ test the parallel trends assumption. None are individually significant, and the mean pre-treatment coefficient is -0.29 hours with no systematic trend.

Randomization inference. Under the sharp null of no treatment effect, the test statistic is the DiD coefficient from Equation 1. I generate the permutation distribution by randomly

assigning 3 of 19 industries as “exempt” 2,000 times and re-estimating the DiD for each permutation. The actual coefficient of -0.21 lies near the center of the permutation distribution (mean = 0.03, SD = 1.46), yielding a two-sided p -value of 0.88.

C. Robustness Appendix

COVID exclusion. Restricting the sample to January 2022 onwards yields 49 monthly periods (27 pre, 22 post) and a coefficient of $+0.39$ (SE = 0.63).

Placebo tests. Fake treatment dates of April 2022 and April 2021 using only pre-treatment data yield coefficients of -1.17 (SE = 0.68, $p = 0.10$) and -0.64 (SE = 1.00, $p = 0.53$). The marginal significance of the 2022 placebo reflects residual COVID recovery differentials and does not survive cluster-robust inference or randomization inference.

Minimum detectable effect. With 19 industries, 106 time periods, and residual variance of approximately 25 (from the model RMSE of ≈ 5), the MDE at the 5% level with 80% power is approximately $2.8 \times \text{SE} = 2.99$ hours. This corresponds to 1.8% of the exempt-industry mean, well below the 2-hour effect found by [Burdin et al. \(2024\)](#) for the 2019 cap in non-exempt industries.

D. Standardized Effect Sizes

Table 4: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Monthly hours	Baseline DiD	-0.212	15.48	-0.0137	0.0689	Small negative
Monthly days	Baseline DiD	-0.048	1.19	-0.0406	0.0677	Small negative
<i>Panel B: Heterogeneous (by sex)</i>						
Monthly hours	Male workers	-0.181	12.83	-0.0141	0.0724	Small negative
Monthly hours	Female workers	1.035	16.48	0.0628	0.0487	Moderate positive

Notes: **Country:** Japan. **Research question:** Does the application of statutory overtime caps to previously exempted industries reduce worker hours and days worked? **Policy mechanism:** The 2018 Work Style Reform Act imposed binding monthly and annual overtime caps on most Japanese industries from April 2019, but granted five-year exemptions to transport, construction, and healthcare due to their “distinctive” labor demands. When exemptions expired in April 2024, these industries faced the same caps as all other sectors, creating a compliance cliff where firms had to immediately limit previously unconstrained overtime scheduling. **Outcome definition:** Average monthly hours worked per employed person, measured by the Labour Force Survey (household-based) at the industry-month level. **Treatment:** Binary—industry exempt from overtime caps until April 2024 vs. already subject to caps since April 2019. **Data:** Labour Force Survey via e-Stat API, April 2017–January 2026, industry-month level, 19 major industries (excluding aggregates). **Method:** Two-way fixed effects DiD with industry and year-month fixed effects, standard errors clustered at the industry level, supplemented by randomization inference (2,000 permutations). **Sample:** All employed persons across 19 major industry divisions; 3 exempt industries (Construction, Transport/Postal, Medical/Welfare) and 16 non-exempt industries. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).