

The Dual Labour Market and the Motherhood Employment Penalty in Japan

Beatriz Gietner

UCD School of Economics

March 22, 2026

The Puzzle

Japan's childcare leave can be extended up to 24 months under eligibility conditions, and Japan has one of the longest father-specific entitlements in the OECD in *full-rate equivalent* terms (about 31 weeks) (OECD, 2024c).

Yet maternal employment when the youngest child is aged 0–2 remains below the OECD average in the latest OECD comparison (OECD, 2024a).

If the policy infrastructure exists, why do women still leave?

In this paper, I show that the key margin is job type: formal eligibility exists, but continuity through childbirth is much weaker in non-regular tracks.

Regular = open-ended, internal career track jobs; non-regular = fixed-term/part-time/dispatch jobs.

What This Paper Is and Isn't

What it is:

- Descriptive event study of 662 first births (KHPS/JHPS, 2004–2022)
- Predictive risk stratification: which jobs predict exit?
- Adds job-type and firm-type margins that earnings-only administrative data cannot observe directly

What it is not:

- Not a causal identification design
- Not administrative data (household panel, N=662)
- Not nationally weighted (sample-average trajectories)

Positioning: Complements Fukai and Kondo (2025), who use local tax records to estimate the child penalty with administrative precision. This paper adds the *job-type channel* that administrative earnings data cannot observe directly.

If you take one thing away: the penalty is concentrated in jobs with weaker contractual protections.

Japan's Dual Labour Market

Regular (seishain)

- Open-ended contracts
- Seniority wages, firm training
- Strong de facto job protection
- Leave is exercisable

Non-regular (hi-seiki)

- Fixed-term, part-time, dispatch
- Flat wages, limited progression
- Contracts can lapse at/around leave
- Leave is formally available but fragile

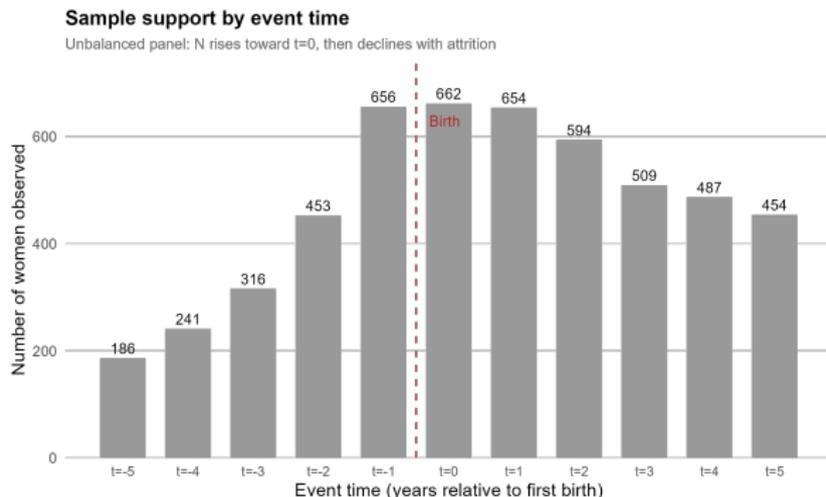
- Over 50% of employed women hold non-regular contracts (OECD, 2024b)
- Small firms face tighter staffing and higher adjustment costs around leave (OECD, 2024b)
- Secondary-earner tax/social-insurance thresholds (e.g., 103/130 man-yen) shape hours and re-entry incentives (OECD, 2024b; Nagase, 2012)
- Men's long working hours remain common (e.g., a sizeable share work >48 hours/week) (OECD, 2024b)

Data: Harmonised KHPS/JHPS household panels, 2004–2022.

Sample: 1,184 first births → 662 in event window → 5,212 person-years

- KHPS (since 2004) and JHPS (since 2009) are nationally fielded household panels with harmonised core labour variables; I pool them using the official harmonisation framework and verify robustness to survey source (KHPS vs JHPS split).
- Panel support is unbalanced over event time (women observed: 186 at $t = -5$, 662 at $t = 0$, 454 at $t = +5$).

Sample Support and Subsamples



Analytic subsamples (nested within the 662):

- Childbirth-margin risk set: 330 employed at $t = -1$, observed at $t = 0$ (124 exits)
- Pre-birth mechanism: 185 complete-case at $t = -2$ (13 exit events)
- Observed at $t = +5$: 454 (IPW sensitivity sample)

Specification and Estimand I

Specification:

$$Y_{it} = \sum_{k \neq -1} \beta_k \mathbf{1}\{\text{event_time}_{it} = k\} + \gamma_t + \varepsilon_{it}$$

- Y_{it} : outcome for woman i in survey year t ; β_k : event-time coefficients relative to $t = -1$; γ_t : calendar-year fixed effects.
- Calendar-year FE, heteroskedasticity-robust SEs; woman-clustered SEs give the same qualitative inference.
- No individual FE in the main specification: individual FE + year FE = collinearity and inflated SEs in this short unbalanced panel. Point estimates are similar, but inference becomes uninformative.
- Timing matters: $t = 0$ is the birth-report wave, and $t = -1$ can overlap late pregnancy for some women, so the estimated birth-year break is if anything conservative.

Specification and Estimand II

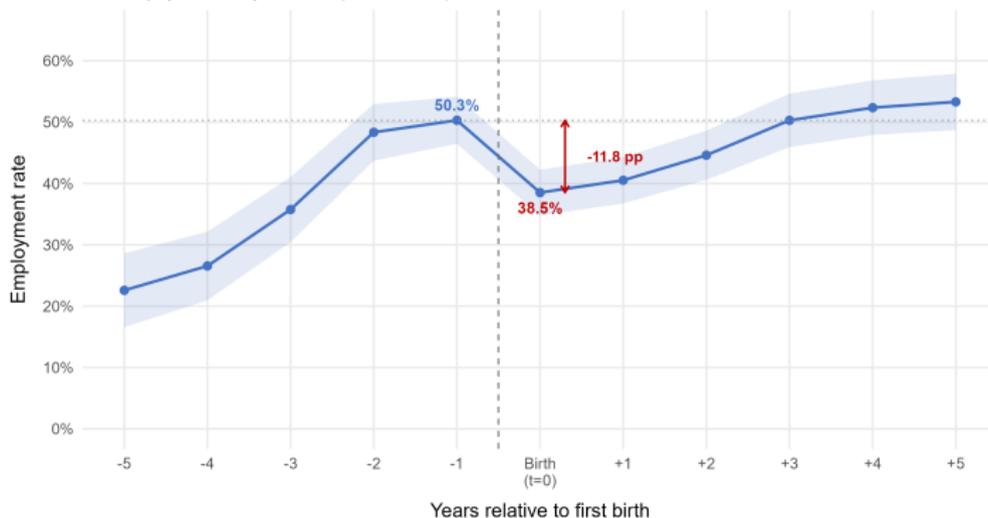
Estimand: Descriptive sample-average event-time trajectories in this analytic panel (not within-person causal treatment effects).

Interpretation strategy: Focus on $[-2, +2]$, where support is strongest; treat far leads/lags as composition-sensitive diagnostics.

Result 1: The Birth-Year Break

Maternal Employment Rate Around First Birth

Mean employment rate by event time (N=662 women). Dashed line marks childbirth.



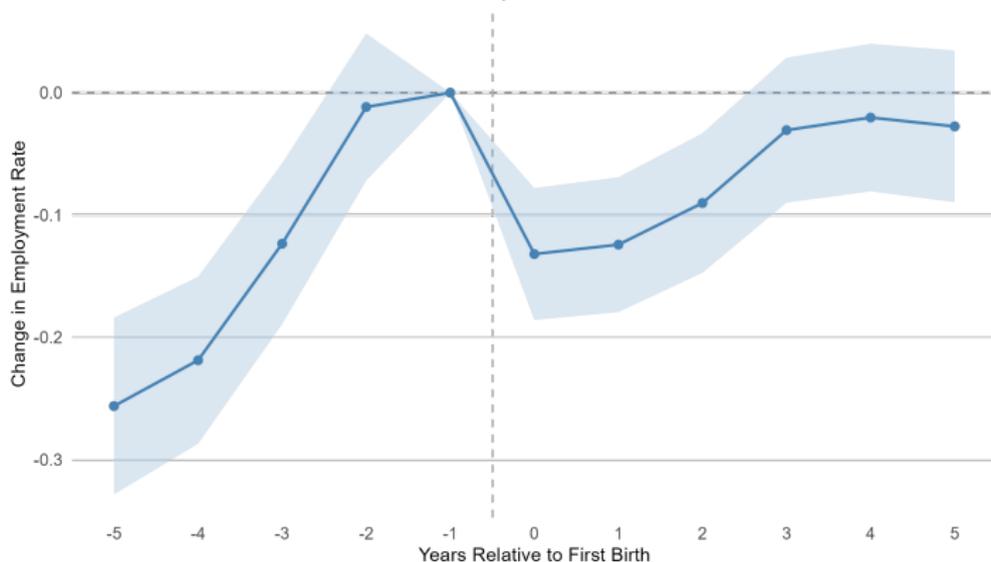
Raw means with 95% CI. Raw $t=-1$ to $t=0$ decline is -11.8 pp; year-FE-adjusted event-study estimate at $t=0$ is -13.2 pp.

Raw employment levels shown; the year-FE-adjusted drop at $t = 0$ is -13.2 pp. Employment is flat at $t = -2$. Then it drops from 50.3% to 38.5% at birth. The pre-birth rise is life-cycle entry, not a pre-trend. The far leads are composition, not anticipation.

Result 1: Coefficient Event Study

Figure 1: Employment Rate Around First Birth

Coefficients relative to $t = -1$. Year FE, heteroskedasticity-robust SEs.



The coefficient version shows the same result in year-FE-adjusted form: near lead at $t = -2$ is close to zero, and the birth-year break at $t = 0$ is -13.2 pp. I treat $[-2, +2]$ as the main interpretation window.

Identification Support: Childless Placebo

Test: Assign matched pseudo-birth years to childless married women (matched by dataset, age bin, and cohort bin) and re-estimate the event study.

Results:

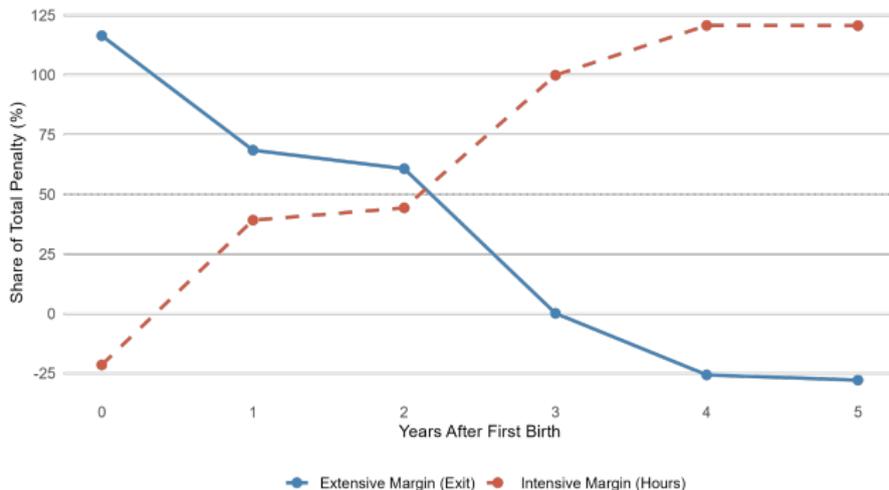
- Pseudo-birth coefficient at $t = 0$: $+0.016$ (SE 0.031, $p = 0.608$)
- Near lead at $t = -2$: -0.003 ($p = 0.937$)
- Joint far-lead test ($t = -5, -4, -3$): $p = 0.749$
- Match quality: 91.7% exact matches

Interpretation: No pseudo-birth break in the childless sample. The sharp drop in the mother sample is birth-related, not generic lifecycle drift.

But the Recovery Is an Illusion

Figure 8: Decomposition Shares Over Time

Share of total penalty attributable to each margin. Negative = employment recovery exceeds baseline.

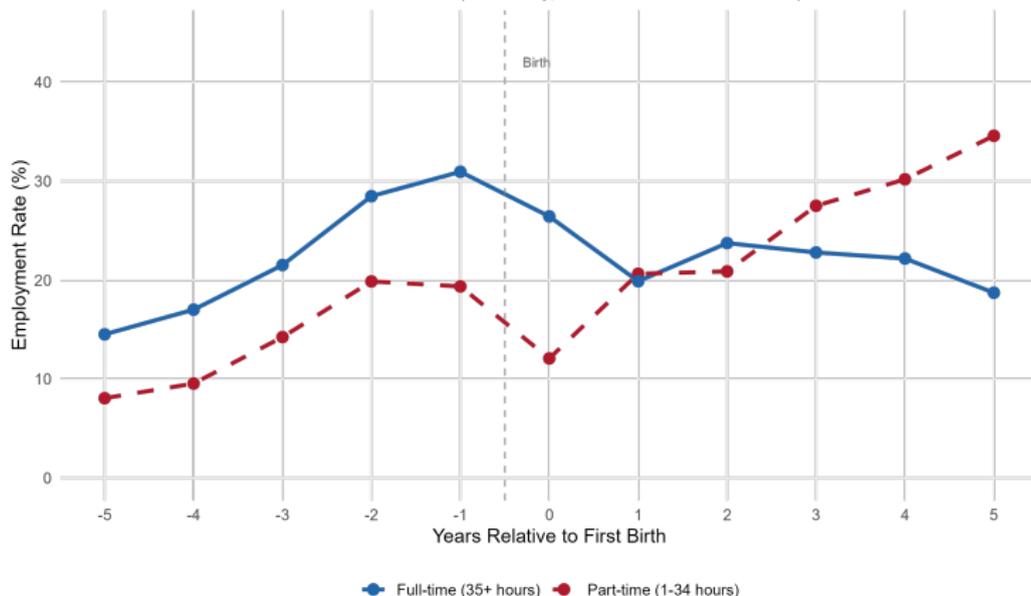


The composition of the penalty flips over time: the extensive margin (employment exit) explains most of the gap right after birth, but by $t = +3$ the intensive margin (hours/wage losses among returners) becomes dominant and remains so through $t = +5$.

Where Do Women End Up? The Transition Matrix

Motherhood Penalty: Full-time vs Part-time Employment

Full-time work declines permanently; Part-time recovers after initial drop



Source: KHPS/JHPS Panel Data. N=662 women.

Employment status flows from $t = -1$ baseline. Full-time workers shift toward part-time; part-time workers shift toward non-employment. The “recovery” is largely into lower-quality jobs.

Result 2a: Who Exits at Childbirth? (Risk Profile)

Risk set: women employed at $t = -1$, observed at $t = 0$ (N=330; 124 exits).

Short answer: women in non-regular jobs.

- Exit risk is much lower in protected jobs and much higher in precarious jobs.
- Regular workers in large firms: **13.8%** exit.
- Non-regular workers in small firms: **64.8%** exit.
- Household-side variables add little once job type is included.

Result 2b: Model Summary and Robustness

Model summary

- Non-regular status: strong predictor of childbirth-margin exit (OR = 7.50, $p < 0.001$).
- Small-firm effect: positive but less precise (OR = 1.55, $p = 0.126$).

Robustness to missing-data handling

- Complete-case OR: 7.50
- Missing-indicator OR: 7.48
- High-information OR: 7.39

Takeaway: the non-regular gradient is stable across specifications, so this is not a missing-data artifact.

Heterogeneity by Contract Type

Figure 9a: Employment Trajectory by Job Type

Change relative to $t = -1$. Non-regular workers show larger, earlier penalty.

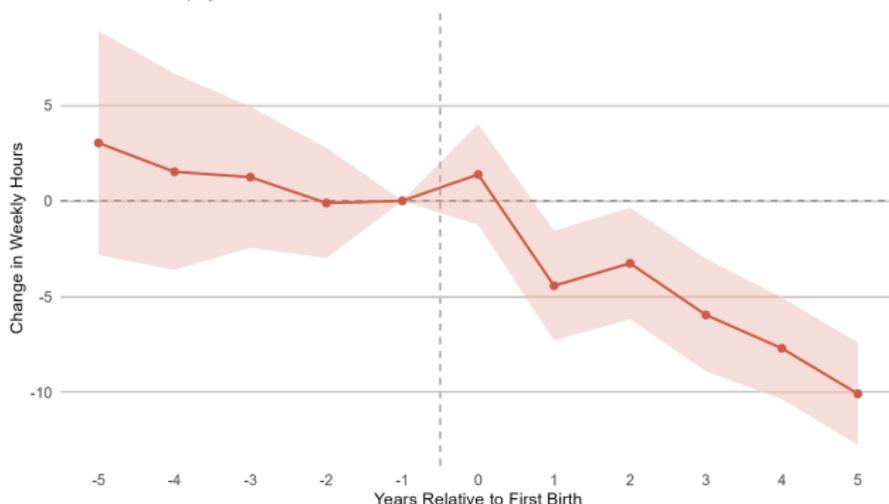


Separate event studies by job type measured at $t = -2$. Non-regular workers: -9.8 pp at $t = 0$. Regular workers: -1.6 pp. A five-fold difference in the birth-year break.

Result 3a: Persistent Hours Losses I

Figure 2: Weekly Work Hours Around First Birth

Conditional on employment. Coefficients relative to $t = -1$.



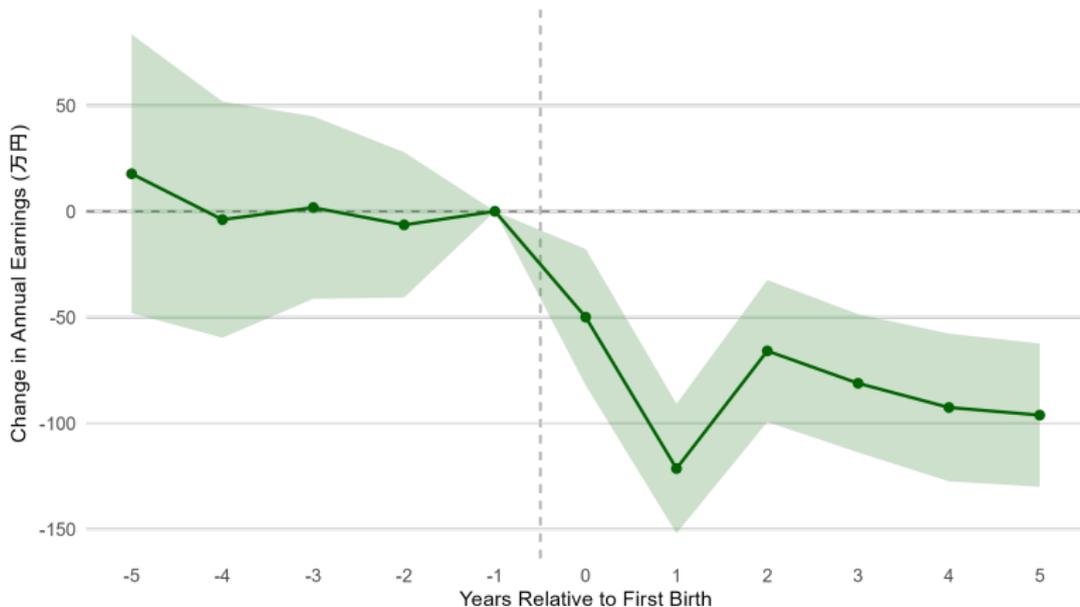
Conditional on employment (relative to $t = -1$):

- Coefficients are changes in weekly hours versus the pre-birth reference year.
- By $t = +5$: -10.1 hours/week, about a 30% decline from the $t = -1$ baseline of 34.1.

Result 3b: Persistent Earnings Losses

Figure 3: Annual Earnings Around First Birth

Conditional on employment. Units: 万円 (10,000 yen). Relative to $t = -1$.



Result 3b: Persistent Earnings Losses

Conditional on employment (relative to $t = -1$):

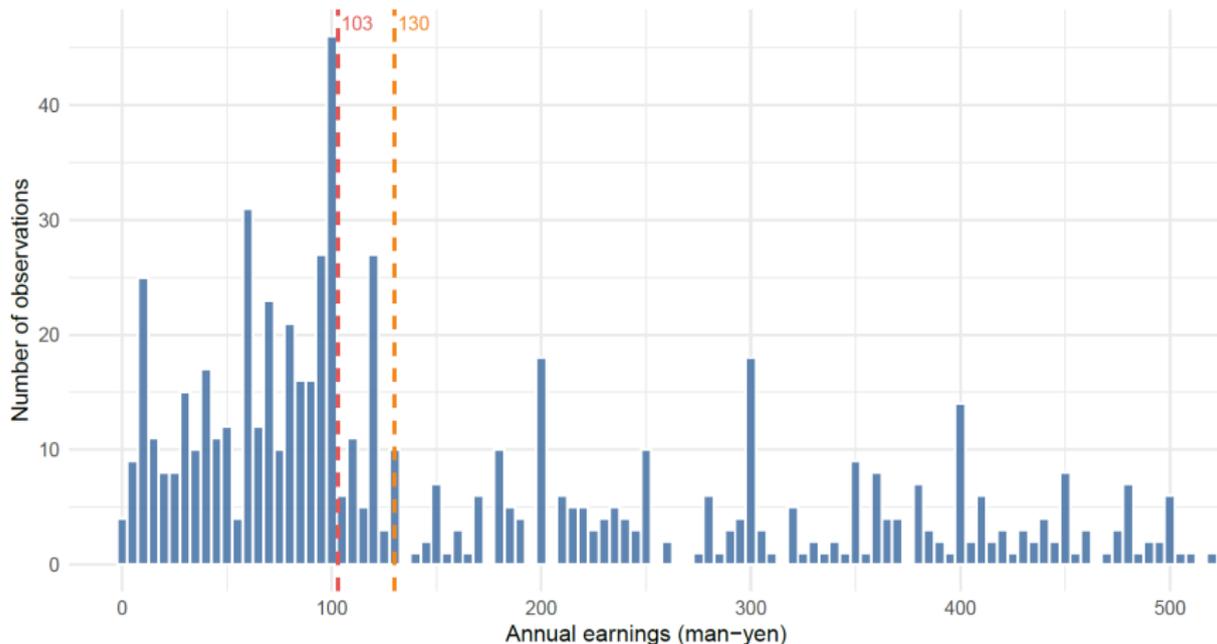
- By $t = +5$: -96.2 man-yen (about a 36% decline from baseline 268.9; implied level ≈ 172.7).
- Implied hourly wage also declines (-7.1%), consistent with lower-quality re-entry.

Unconditional outcomes (coding 0 when non-employed) show the same qualitative pattern.

Result 4: Earnings Bunching at Tax Thresholds

Annual earnings distribution for employed mothers at $t=+3$ to $t=+5$

Vertical lines mark 103 and 130 man-yen thresholds



Sample: employed observations at $t=+3,+4,+5$. Units: man-yen.

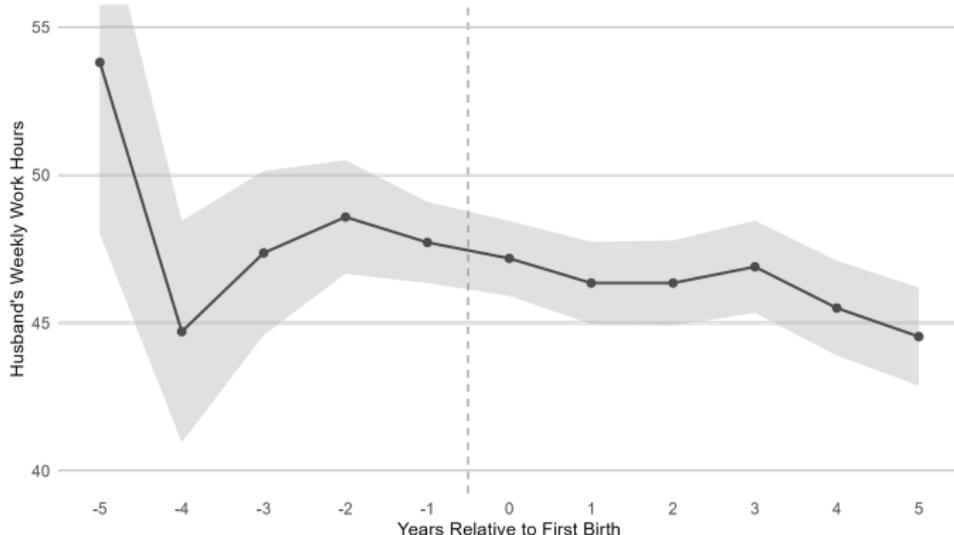
Result 4: Earnings Bunching at Tax Thresholds

Among employed mothers at $t = +3$ to $t = +5$: **47.8%** earn ≤ 103 man-yen; **55.9%** earn ≤ 130 man-yen. These are the spousal tax deduction and social insurance thresholds. Net-cost jumps around these points reduce incentives to increase hours.

The Absent Father

Figure 5: Husband's Weekly Work Hours Around First Birth

Mean hours with 95% CI. Zero adjustment = zero father penalty.



Husband's weekly hours are flat across all event times. Mean \approx 47 hours/week, no detectable adjustment at childbirth.

Only 3 of 1,183 husbands (0.3%) took childcare leave in this sample.

The Training Gap

Event time	N	Skills training (%)	Grants used (%)
-2	253	35.2	2.2
-1	475	31.2	2.3
0	568	8.6	1.0
+1	594	10.6	1.3
+3	493	15.6	0.3
+5	437	12.8	0.0

- Skills training drops from ~35% to 8.6% at birth. Partial recovery to ~13% - about 40% of baseline.
- The Education and Training Benefits System (*kyōiku kunren kyūfu seido*): <3% usage at every horizon. The system is invisible to mothers in this sample.
- A “double penalty”: childbirth reduces both labour supply *and* human capital investment.

Cohort Split: Stable Across Policy Regimes

	$t = -2$	$t = 0$	$t = +5$
Pooled	-0.012 ($p = 0.70$)	-0.132 ($p < 0.001$)	-0.028 ($p = 0.38$)
Early (≤ 2012)	+0.020 ($p = 0.63$)	-0.131 ($p < 0.001$)	-0.020 ($p = 0.66$)
Late (≥ 2013)	-0.063 ($p = 0.19$)	-0.135 ($p = 0.002$)	-0.025 ($p = 0.66$)
No-COVID	-0.004 ($p = 0.91$)	-0.145 ($p < 0.001$)	-0.028 ($p = 0.40$)

- The birth-year break is virtually identical across early and late cohorts.
- Unaffected by excluding post-COVID birth years (2020–2022).
- The 2004–2022 period spans childcare expansion, Womenomics, and multiple leave reforms - yet the penalty persists.

Evidence Hierarchy: What Is Solid vs Cautious

Most credible

- Sharp childbirth break in employment at $t = 0$ (-13.2pp).
- Near lead at $t = -2$ is close to zero.
- Childbirth-margin exit risk is strongly stratified by non-regular status (OR ≈ 7.5).

Credible but more sample-sensitive

- Long-horizon ($t = +5$) levels and recovery magnitudes.
- Small-firm effects conditional on contract type.

Interpretive context (not separately identified)

- Childcare logistics, paternal-hours narrative, and institutional mechanisms.

Full diagnostics in backup: placebo, IPW, balanced panel, leave recoding, missingness sensitivity, clustered SE, cohort split, Lee-style trimming, KHPS/JHPS split.

The penalty begins as exit and transforms into permanent downgrade.

Policy implications consistent with the descriptive evidence:

- ① **Strengthen leave continuity for fixed-term workers** - ensure leave rights survive contract renewal
- ② **Reduce accommodation costs in SMEs** - replacement-hiring subsidies and simplified leave administration
- ③ **Address men's overwork** - near-zero paternal adjustment in this sample
- ④ **Lower re-entry barriers** - training subsidies are invisible (<3% usage) and secondary-earner thresholds cap earnings

Bottom line: Expanding childcare and leave helps workers who already hold secure jobs. The binding constraint is labour-market structure - the divide between jobs that can accommodate parenthood and jobs that cannot.

What I Would Like Your Feedback On

- ① Is the *seishain/hi-seiki* distinction the right margin for this analysis, or should I also consider finer categories (*keiyaku, haken, pāto*)?
- ② Does the 103 man-yen bunching pattern match your intuition? Has the 2018 reform (threshold raised to 150 man-yen) changed behaviour in practice?
- ③ The training grants finding (<3% usage) - is this consistent with what practitioners and policymakers observe?
- ④ Are there KHPS/JHPS variables or institutional details I am missing that could sharpen the analysis?
- ⑤ How does this evidence sit alongside the administrative-data results in Fukai and Kondo (2025)?

Thank you

Beatriz Gietner
b.gietner@gmail.com

Paper, code, and all robustness tables available on request.

References I

- Casarico, A. and Lattanzio, S. (2023). Behind the child penalty: understanding what contributes to the labour market costs of motherhood. *Journal of Population Economics*, 36:1489–1511.
- de Quinto, A., Hospido, L., and Sanz, C. (2021). The child penalty: evidence from Spain. *SERIEs*, 12(4):585–615.
- Fukai, T. and Kondo, A. (2025). Parental earnings trajectories around childbirth in Japan: evidence from local tax records. Technical Report 25-E-012, RIETI Discussion Paper Series.
- Kleven, H., Landais, C., and Leite-Mariante, G. (2024). The child penalty atlas. *Review of Economic Studies*, 92(5):3174–3207.
- Nagase, N. (2012). Japanese female labor supply and the tax/social insurance system. *Japanese Journal of Labour Studies*, 54(12):12–25.
- OECD (2024a). Maternal employment rates (LMF1.2). OECD Family Database. Data for 2021 or latest available year. Accessed 2026-02-25.
- OECD (2024b). OECD economic surveys: Japan 2024. Technical report, OECD Publishing, Paris.
- OECD (2024c). Parental leave systems (PF2.1). OECD Family Database. Data as of April 2024. Accessed 2026-02-25.

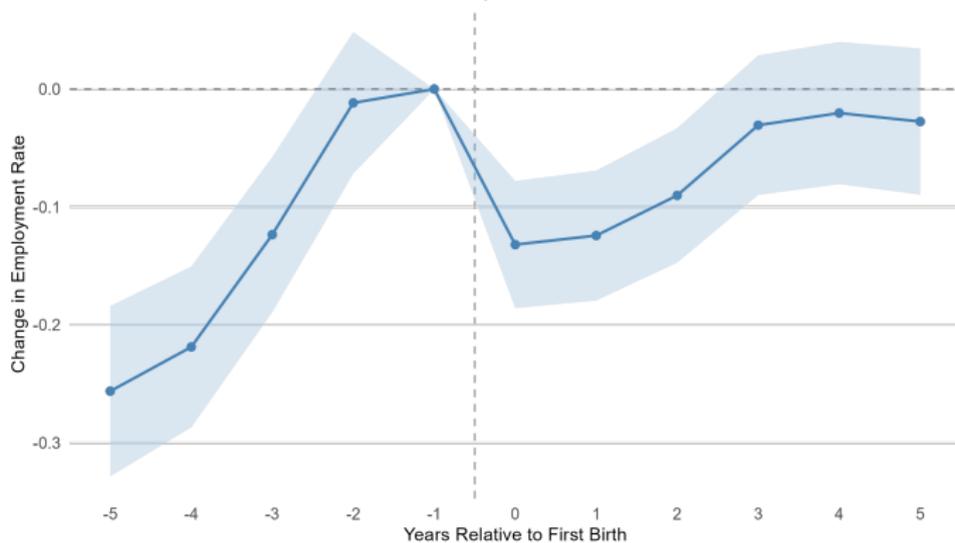
Backup: Summary Statistics at $t = -1$

	Full Sample (N=662)	Quitters (N=29)	Stayers (N=262)
Employed (%)	50.3	55.2	61.0
Full-time (%)	30.9	31.0	36.7
Part-time (%)	19.4	24.1	24.3
Hours/week (mean)	34.1	32.5	33.4
Income, man-yen (mean)	268.9	148.4	302.7

Backup: Coefficient Event Study

Figure 1: Employment Rate Around First Birth

Coefficients relative to $t = -1$. Year FE, heteroskedasticity-robust SEs.



Standard event-study coefficients relative to $t = -1$. Year FE, robust SEs.

Backup: Pre-Trend Diagnostics

The near lead at $t = -2$ is null across four specifications:

- Full sample: -0.012 , $p = 0.70$
- Trimmed pre-period (drop $t = -5, -4, -3$): $p = 0.72$
- Balanced pre-support: $p = 0.078$
- With KHPS/JHPS dataset indicator: $p = 0.71$

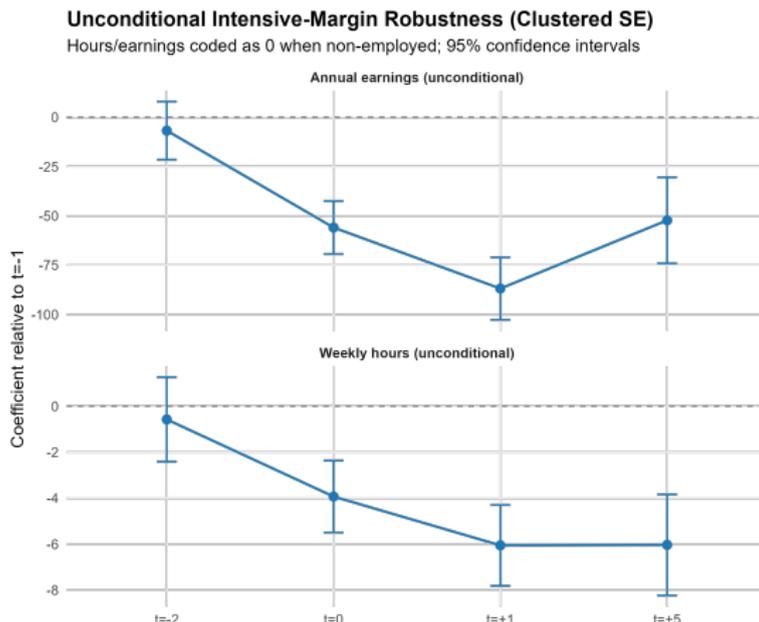
Matched childless placebo: $t = 0$ coefficient = $+0.016$, $p = 0.608$.

Joint far-lead test: $p = 0.749$.

Backup: IPW and Attrition

- Observability at $t = +5$ is weakly selective on baseline employment ($p = 0.045$)
- IPW reweighting:
 - Near lead: unchanged
 - Birth-year break: unchanged (-0.136 , $p < 0.001$)
 - $t = +5$: shifts to -0.057 ($p = 0.076$) - more negative
- Balanced panel ($t = -2$ to $t = +1$, $N=662$, 2,648 obs): core coefficients preserved

Backup: Unconditional Outcomes



Standard event-study coefficients relative to $t = -1$. Year FE, robust SEs. Coding 0 for non-employed preserves the post-birth declines and near-lead null.

Backup: Missingness Sensitivity (Profile)

Table: Missingness profile for mechanism covariates in the childbirth-margin risk set (employed at $t = -1$, observed at $t = 0$).

Covariate missingness	N	Share of risk set	Exit rate at $t = 0$
Contract type missing	36	0.109	0.444
Firm size missing	16	0.048	0.750
Either missing	48	0.145	0.500

Notes: Risk-set size is $N = 330$ women. Exit is defined as non-employment at $t = 0$ among women employed at $t = -1$.

Backup: Missingness Sensitivity (Mechanism)

Table: Childbirth-margin mechanism sensitivity to missing-data handling.

Specification	<i>N</i>	Exits	OR Non-regular	OR Small firm
Complete-case	282	100	7.505***	1.548
Missing-indicator	330	124	7.476***	1.289
High-information	237	83	7.393***	1.860**

Notes: Odds ratios reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. High-information sample requires non-missing contract type, firm size, commute indicator, and husband's overwork indicator at $t = -1$. Missing-indicator specification includes missing dummies for contract and firm size.

Non-regular OR stable: 7.50, 7.48, 7.39 across specifications.

Backup: Pre-Birth Mechanism ($t = -2, 13$ events)

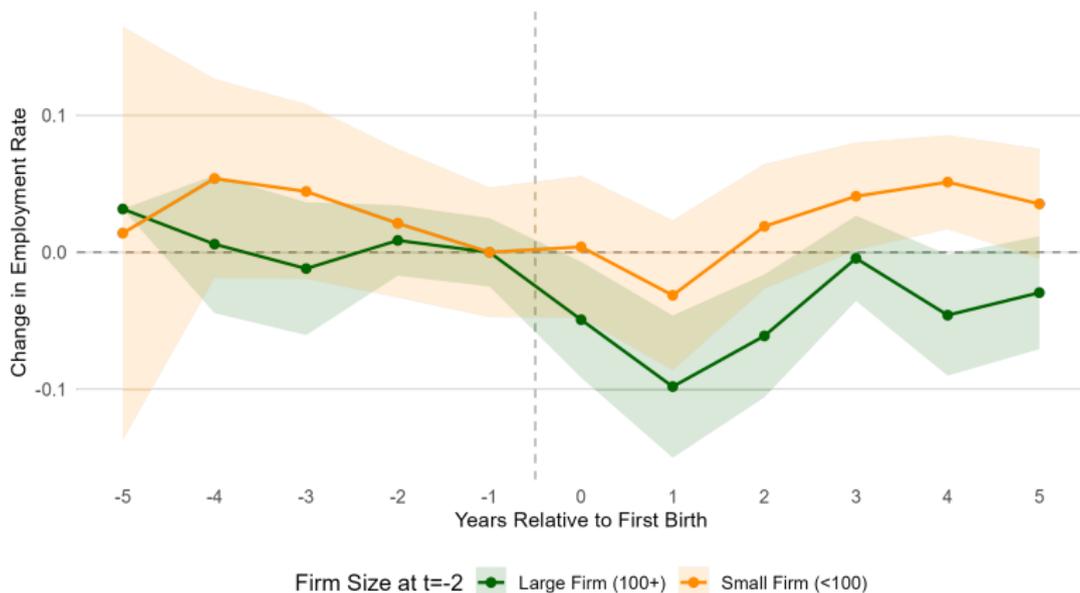
	OR	p	95% CI	Method
<i>Non-regular</i>				
Standard	6.41	0.019	-	MLE
Firth	5.32	0.009	-	Penalised
Bootstrap	-	0.008	[1.68, 19.66]	2,000
<i>Small firm</i>				
Standard	4.14	0.025	-	MLE
Firth	3.84	0.021	-	Penalised
Bootstrap	-	0.038	[1.10, 18.34]	2,000

Jackknife: non-regular OR above 1 in all 185 replications (range: 5.69–12.82).

Backup: Heterogeneity by Firm Size

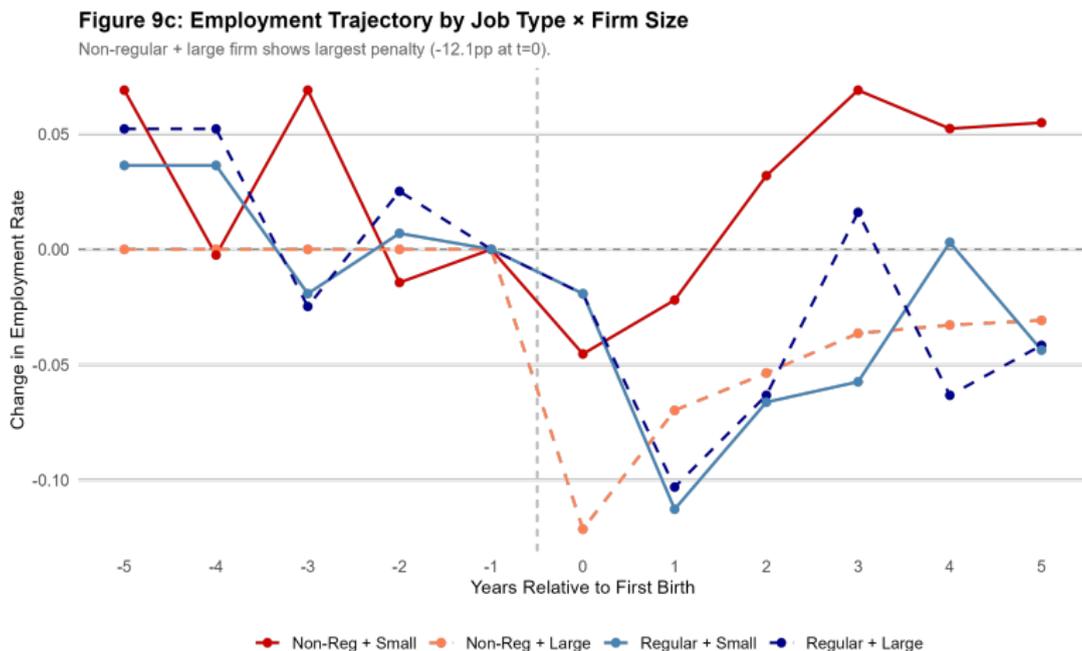
Figure 9b: Employment Trajectory by Firm Size

Change relative to $t = -1$. Penalties differ by firm size.



Separate event studies by firm size measured at $t = -2$.

Backup: Combined Heterogeneity (Job Type × Firm Size)

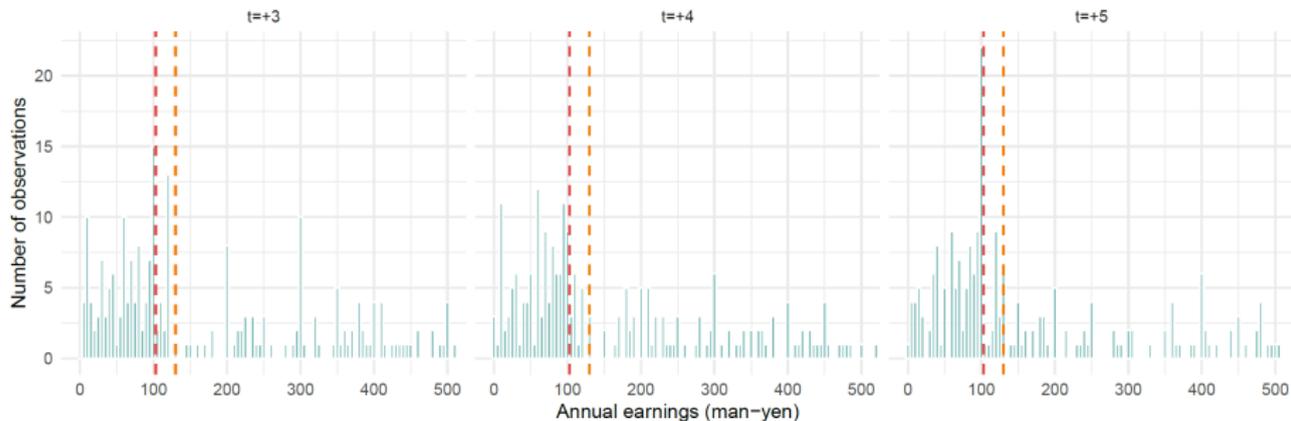


Four groups: Regular-Large, Regular-Small, Non-Regular-Large, Non-Regular-Small.
Largest penalty in Non-Regular + Large (-12.1 pp at $t = 0$).

Backup: Earnings Bunching by Event Time

Annual earnings distribution by event time (employed mothers)

Thresholds at 103 and 130 man-yen



Backup: KHPS/JHPS Split

Table: Survey-split stability check for key employment event-study coefficients (year fixed effects).

Sample	Event time	Estimate	SE	p-value	N obs	N women
Pooled	$t = -2$	-0.012	0.031	0.703	5,212	662
Pooled	$t = +0$	-0.132	0.027	0.000	5,212	662
Pooled	$t = +5$	-0.028	0.032	0.384	5,212	662
KHPS	$t = -2$	0.002	0.039	0.966	3,323	394
KHPS	$t = +0$	-0.134	0.036	0.000	3,323	394
KHPS	$t = +5$	0.010	0.040	0.796	3,323	394
JHPS	$t = -2$	-0.027	0.049	0.576	1,889	268
JHPS	$t = +0$	-0.136	0.044	0.002	1,889	268
JHPS	$t = +5$	-0.126	0.057	0.027	1,889	268

Near lead and birth-year break are stable across surveys. Long horizon diverges.

Backup: Geographic Non-Variation

- Metropolitan (Kanto/Kinki, N=355): -13.6 pp at $t = 0$
- Non-metropolitan (N=307): -12.9 pp at $t = 0$
- Near lead null in both subsamples

The dual labour market operates nationwide.

Backup: Leave-Recoding Sensitivity

Table: Leave-as-employment sensitivity at childbirth ($t = 0$).

Specification	Coef. at $t = 0$	SE	p -value	Reclassified at $t = 0$
Original (hours > 0)	-0.1319	0.0275	<0.001	0
Leave-adjusted	-0.0988	0.0276	<0.001	22

Notes: Sensitivity recodes women as employed at $t = 0$ if they are coded non-working but report leave take-up near birth. The coefficient remains strongly negative after reclassification.

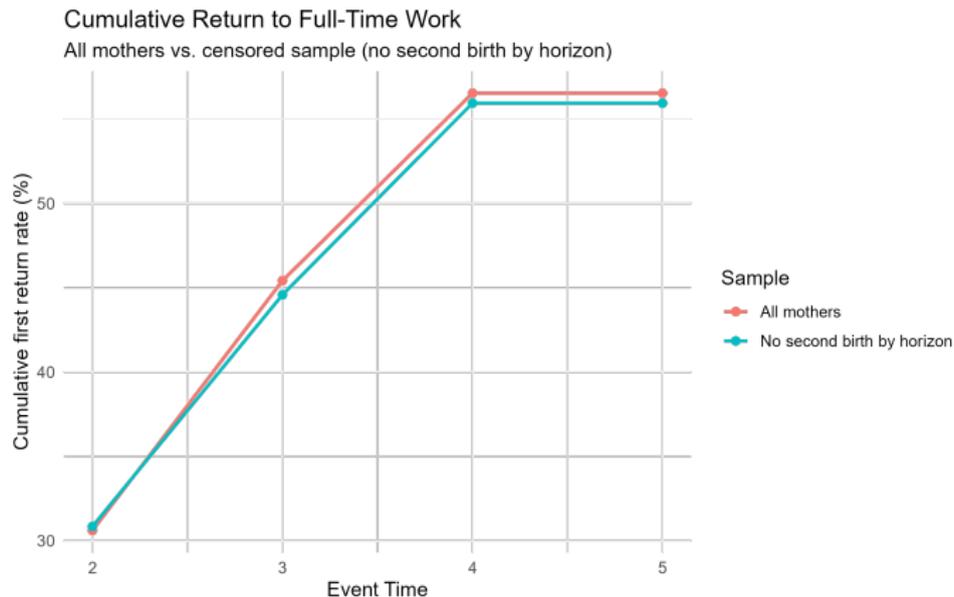
Reclassifying 22 leave-takers as employed attenuates $t = 0$ from -0.132 to -0.099 ; remains strongly significant.

Backup: Full Stability Map

Stable	Sensitive	Interpretive
Near lead null ($-0.012, p = 0.70$)	Long horizon (composition-sensitive)	Marriage-stage sequencing
Birth-year drop ($-0.132, p < 0.001$)	Far leads reflect changing support	Childcare-logistics narrative
Non-regular OR at $t=0$ ($7.50, p < 0.001$)	Pre-birth mechanism (13 events)	Household-structure channels
Cohort-split stable (early \approx late \approx no-COVID)	Industry splits selection-sensitive	Cross-country framing
Bunching at 103 man- yen (47.8%)	Hours/earnings cond. on employment	Training non-uptake causes

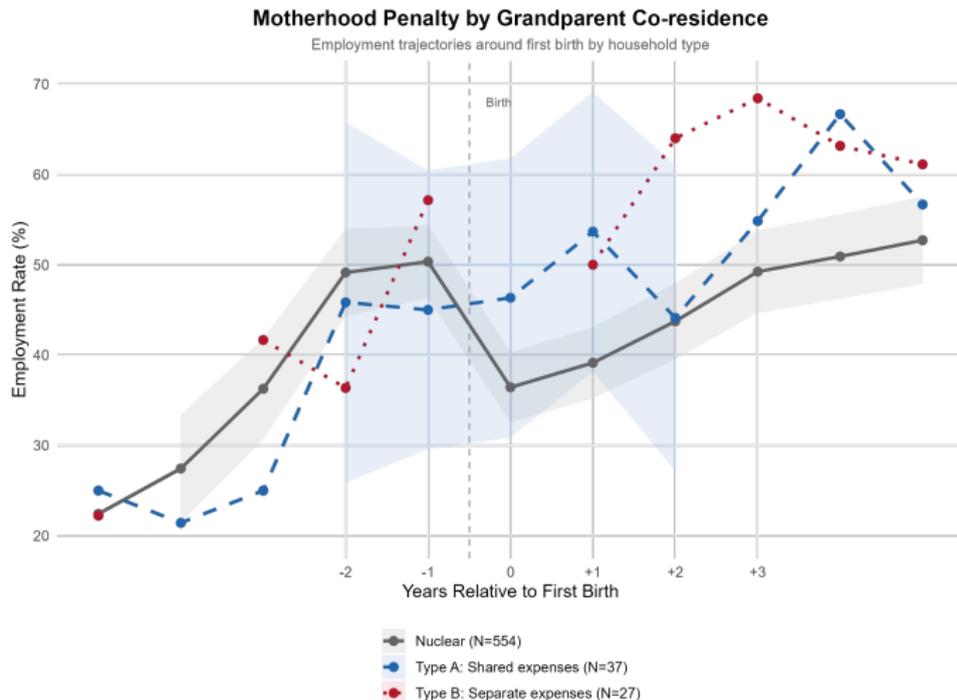
Additional checks: placebo, IPW, balanced panel, Lee bounds, clustered SEs, KHPS/JHPS split, missingness sensitivity, jackknife.

Backup: Cumulative Return Rate



Cumulative first-return proportion among women who exited employment at $t = 0$.

Backup: Grandparent Coresidence Heterogeneity



Source: KHPS/JHPS Panel Data. 95% confidence intervals shown.

Type A = Living with parents, shared household expenses. Type B = Living with parents, separate expenses.

Split by grandparent coresidence at baseline.

Q&A Backup: Identification and Sample Size

Q: Why not use individual fixed effects?

- Individual FE + year FE = near-collinearity in this short unbalanced panel. Point estimates are similar (backup figure available), but SEs inflate and inference becomes uninformative. The estimand is explicitly descriptive sample-average trajectories.

Q: The sample is only 662. Can you really draw conclusions?

- The birth-year break and non-regular OR are robust across every specification and subsample split. The evidence hierarchy is honest about what is and isn't well-powered. Jackknife confirms mechanism stability across all 185 replications.

Q: Why not use the Labour Force Survey or Employment Status Survey?

- Those are repeated cross-sections with much larger N, but they cannot track the same woman over time. The panel structure is essential for event-study design. Trade-off: smaller N but within-person trajectories.

Q: How does this compare to Fukai & Kondo (2025)?

- Their administrative tax records give much larger N and cleaner earnings measurement. They confirm a large, persistent child penalty in Japan.
- This paper adds the *job-type channel*: administrative earnings data cannot observe contract type (*seishain* vs *hi-seiki*) directly.
- The two papers are complements: Fukai & Kondo measure the penalty precisely; this paper shows *where in the labour market* it is concentrated.

Q: Is this a Japan-specific pattern or does it appear in other dual labour markets?

- The Child Penalty Atlas (Kleven et al., 2024) documents cross-country variation. Japan's aggregate penalty is moderate internationally.
- The *mechanism* – contract-type segmentation as the binding margin – may be distinctive to Japan and Korea, where the regular/non-regular divide is institutionally sharp.

- Spain (de Quinto et al., 2021) and Italy (Casarico and Lattanzio, 2023) show related temporary-contract effects but in different institutional contexts.

Q&A Backup: Non-Regular Subcategories and Leave

Q: Is hi-seiki really one category? Keiyaku, haken, and pāto are very different.

- The KHPS/JHPS classification does not always distinguish finely between non-regular subtypes, and the sample is too small to split further reliably.
- I agree this matters – *haken* workers face different constraints than *pāto* workers. A larger dataset (e.g., administrative records with contract-type flags) could explore this margin.

Q: Leave is formally available to non-regular workers since 2005/2010. Why doesn't it work?

- Eligibility required continuous employment with the same employer for 1+ year (relaxed in the 2022 revision). Contract non-renewal around pregnancy is prohibited but hard to enforce.
- The issue is not the law – it is whether the contract survives long enough to exercise the right.
- Small firms may also lack awareness of the rules or capacity to manage the absence.

Q: The 103 man-yen wall was raised to 150 man-yen in 2018. Does bunching change?

- The sample at $t = +3$ to $t = +5$ spans both pre- and post-reform periods; not enough post-reform observations to split cleanly.
- The 130 man-yen social insurance threshold was *not* raised – this may be the more binding constraint in practice.
- Bunching at 103 may also reflect employer-side coordination (firms structuring part-time hours to stay below).

Q: What about childcare availability? Isn't that the real constraint?

- Cannot observe municipality-level childcare slots in the household panel. Prefecture-level data (Appendix D.6) shows substantial geographic dispersion.
- The geographic non-variation result (metro -13.6 pp vs non-metro -12.9 pp) suggests the dual labour market operates nationwide.
- Childcare access and labour-market segmentation are likely *co-constraints*, not competing explanations.

Q: What about selection into motherhood?

- The non-regular share is actually *higher* among childless women than in the mothers' sample.
- Consistent with regular-track attachment being a precondition for in-panel family formation.
- If anything, this compresses the observed penalty relative to a broader at-risk population.

Q: You pool 2004–2022. The labour market changed enormously.

- The cohort split shows the penalty is virtually identical in early (≤ 2012) vs late (≥ 2013) cohorts. Year FE absorb aggregate trends.
- The *stability* is one of the paper's striking findings – despite Womenomics, childcare expansion, and multiple leave reforms, the structural pattern persists.

Q: What about income effects? Maybe women choose to reduce hours.

- Husband's hours are uniformly high (~ 47 hrs/week) with low variance; husband's income is not a significant predictor of exit once job type is included.
- Cannot rule out household-level optimisation. The bunching at tax thresholds is consistent with rational fiscal behaviour.

Q&A Backup: Survey Timing and Second Births I

Q: What is the timing of the survey interview relative to actual birth?

- $t = 0$ is the survey wave in which the first birth is reported, not the exact calendar date.
- Some women at $t = 0$ have had the child for a few months, others nearly a year.
- This is stated explicitly as a limitation. It means the $t = 0$ coefficient averages over heterogeneous post-birth durations.

Q: What about second births?

- Second-birth spacing and share by event time are documented in the appendix tables.
- By $t = +5$, roughly half the sample has had a second birth. Censoring at second birth does not qualitatively change the pattern.

Q: Can you separate wages from hours for returners?

- Yes. The implied hourly wage also declines (-7.1%), consistent with lower-quality re-entry rather than voluntary hours reduction alone. Figure available in backup.